Evaluation of Clusters based on Systems on a Chip for High-Performance Computing: A Review*

Melissa Johanna Aldana[**]
Jaime Alberto Buitrago[***]
Julián Esteban Gutiérrez[****]

Recibido: 11/10/2018 • Aceptado: 14/03/2020
https://doi.org/10.22395/rium.v19n37a4

Abstract

High-performance computing systems are the maximum expression in the field of processing for large amounts of data. However, their energy consumption is an aspect of great importance, which was not considered decades ago. Hence, software developers and hardware providers are obligated to approach new challenges to address energy consumption, and costs. Constructing a computational cluster with a large amount of systems on a chip can result in a powerful, ecologic platform, with the capacity to offer sufficient performance for different applications, as long as low costs and minimum energy consumption can be maintained. As a result, energy efficient hardware has an opportunity to impact upon the area of high-performance computing. This article presents a systematic review of the evaluations conducted on clusters of Systems on a Chip for High-Performance computing in the research setting.

Keywords: Systems on a chip; high-performance computing; clusters; benchmarks.

* This review paper stems from a research project (no. 811) called "Cluster of mobile systems on a Chip for high-performance computing", funded by the Universidad del Quindío through its Vice-rectory of Research. Start date: 2017. End date: 2018.

** Magíster en proyectos educativos mediados por TIC, e ingeniera de sistemas y computación. Profesora asistente del Programa de Ingeniería de Sistemas y Computación, Grupo de Investigación GRID, Universidad del Quindío. Correo electrónico: mjaldana@uniquindio.edu.co. Orcid: https://orcid.org/0000-0002-0257-4783.

*** Ph. D. en Ingeniería e ingeniero electrónico. Profesor asistente del Programa de Ingeniería Electrónica, Grupo de Investigación Sinfoci, Universidad del Quindío. Correo electrónico: jalbertob@uniquindio.edu.co. Orcid: https://orcid.org/0000-0002-8328-8470.

**** Ph. D en Ciencias de la Computación, Ingeniero de Sistemas, Profesor Titular Programa de Ingeniería de Sistemas y Computación, Grupo de Investigación GRID Universidad del Quindío. Correo electrónico: jugutier@uniquindio.edu.co. Orcid: https://orcid.org/0000-0002-1825-0097.

Evaluación de clústeres basados en sistemas en un chip para computación de alto desempeño: una revisión

Resumen

Los sistemas de computación de alto desempeño son la máxima expresión en el campo de procesamiento para grandes cantidades de datos. Sin embargo, su consumo de energía es un aspecto de gran importancia que no era tenido en cuenta en décadas pasadas. Por lo tanto, desarrolladores de *software* y proveedores de *hardware* están obligados a enfocarse en nuevos retos para abordar el consumo de energía y costos. Construir un clúster informático con una gran cantidad de sistemas en un chip puede dar como resultado una plataforma poderosa, ecológica y capaz de ofrecer el rendimiento suficiente para diferentes aplicaciones, siempre y cuando se puedan mantener bajos costos y el menor consumo de energía posible. Como resultado, el *hardware* eficiente en el consumo de energía tiene la oportunidad de tener un impacto en el área de la computación de alto desempeño. En este artículo se presenta una revisión sistemática para conocer las evaluaciones realizadas a clústeres de sistemas en un chip para computación de alto desempeño en el ámbito investigativo.

Palabras clave: sistemas en un chip; computación de alto desempeño; clústeres; *benchmarks.*

Avaliação de clusters baseados em sistemas em um chip para a computação de alto desempenho: uma revisão

Resumo

Os sistemas de computação de alto desempenho são a máxima expressão no campo de processamento para grandes quantidades de dados. No entanto, seu consumo de energia é um aspecto de grande importância que não era levado em consideração em décadas passadas. Portanto, desenvolvedores de software e provedores de hardware estão obrigados a focar-se em novos desafios para abordar o consumo de energia e custos. Construir um cluster informático com uma grande quantidade de sistemas em um chip pode dar como resultado uma plataforma poderosa, ecológica e capaz de oferecer o rendimento suficiente para diferentes aplicações, desde que possam ser mantidos baixos custos e o menor consumo de energia possível. Como resultado, o hardware eficiente no consumo de energia tem a oportunidade de ter um impacto na área da computação de alto desempenho. Neste artigo, apresenta-se uma revisão sistemática para conhecer as avaliações realizadas a clusters de sistemas em um chip para computação de alto desempenho no âmbito investigativo.

Palavras-chave: sistemas em um chip; computação de alto desempenho; clusters; *benchmarks*.

INTRODUCTION

High-performance computing (HPC) is the concept that encompasses the principles, methods, and techniques that allow to address problems with complex computer structures, and of high requirements. The solution to those problems involves massive data sets, a large amount of variables, and complex calculation processes, which require efficient application of modern parallel computation tools [1]. In addition, the scientific challenges arising in engineering, geophysics, bioinformatics, and other types of applications of intensive computational use, require ever-higher amounts of computational calculations [2].

High-performance computing systems have shown exponential growth in computational power in recent decades, and this is due mainly to the evolution of microprocessor technology [3]. Large HPC systems are dominated by processors using mainly x86 and Power instruction sets, supplied by three providers: Intel, AMD, and IBM [3]. Performance of HPC systems has increased continuously with advances of Moore's Law and parallel processing, while energy efficiency could be considered a secondary problem [4]. By 2000, hardware manufacturers improved their processors' performance by adding optimization technology, like bifurcation prediction, predictive execution, and increased cache size, besides making the clock frequency faster. However, the downside of this situation was increased energy consumption, which obligated manufacturers to add multiple nuclei in a processor to avoid having problems due to overheating [3]. According to the aforementioned, it was clear that energy consumption was the dominant parameter in the scaling challenges to achieve better performance. It is widely accepted that future HPC systems will be limited by their energy consumption [2], and thermal problems [5]. Improvement is required in energy efficiency to achieve exascale computing (1018 FLOPS) [4-5].

Using concepts of integrated technologies, like systems on a chip (SoC, System-on-Chip), have emerged naturally to address the problem of energy consumption. The first prototypes were studied based on a vast number of microprocessors of many low-power nuclei instead of rapid complex nuclei, to comply with HPC and power consumption demands. In addition, recent progress and the availability of 64-bit Advanced RISC Machine (ARM) nuclei open new expectations for cluster development [4].

One way to address this problem of energy consumption is to replace the central processing units (CPU) of high-end servers with the low-power processors traditionally found in embedded systems. The use of integrated processors in clusters is not new: Diverse BlueGene machines use integrated PowerPC chips [7]. Likewise, low-power processors, which were originally destined for mobile devices, have had enormous improvement with respect to their computational power. Low-power modern processors,

like the most recent ARM designs, provide great performance in relation to their low energy consumption. Today, most mobile devices are themselves small supercomputers, which supply computing power due to multiple processors and sophisticated graphic processors, while maintaining the low energy consumption that is necessary for mobile applications [5].

Implementation of HPC platforms is a costly undertaking, which may be inaccessible for small and medium institutions. Projects like Mont-Blanc [http://www.montblanc-project.eu/] and COSA (Computing on SoC Architectures, http://www.cosa-project.it/) [8] have successfully demonstrated the use of SoC-based clusters for HPC systems, but the need still exists to analyze their viability. The principal deficiencies of the current evaluations of these systems are the lack of detailed information on the performance levels in distributed systems, and the comparative evaluation in large-scale applications [9].

This work analyzed publications in the area of clusters of Systems on a Chip for high-performance computing. Inspired on the PRISMA Declaration, a search and selection of investigations was conducted. Similarly, a systematic review was carried out to learn of the evaluations made of these types of systems in the research setting. This evaluation does not compare the systems to each other, since each application has its own characteristics and different configurations, which does not make it appropriate to know which is better or worse at the performance level. Likewise, it is highlighted that the objective of this article is to know the type of evaluations carried out on Systems on a Chip for high-performance computing.

1. METHODOLOGY

The search and systematic review used an adaptation of the PRISMA Declaration, which sought to help the authors to improve the presentation of the systematic reviews [10]. The PRISMA Declaration consists of a 27-item checklist and a four-phase flow diagram, which are the methodological route to conduct the search and systematic review. Additionally, PRISMA can be useful for the critical evaluation of published systematic reviews [11]. The following sections will present the results of the search process and systematic review. Based on the PRISMA Declaration, we first defined the theme or problem to analyze and then formulated three research questions, which are the base of the search criteria and selection of the bibliography review.

2. RESULTS

This work sought to identify what types of performance evaluations have been conducted on clusters based on systems on a chip for high-performance computing. Hence, it

is important to know what hardware, benchmarks, and measurement parameters have been used in implementing these types of systems.

Specifically, three questions focused on the recognition of the performance evaluations mentioned were formulated, to allow concentration in the search of the bibliography review.

- Research questions:
  - What types of devices were used in the investigations consulted?
  - What tests or benchmarks were implemented to measure cluster performance?
  - What other parameters were considered to conduct the evaluations?

2.1 Search sources

According to the research questions, a set of search sources was defined, which include IEEE Xplore Digital Library (Institute of Electrical and Electronics Engineers), Science Direct, Engineering Village, and Google Scholar, which contain published scientific bibliography according to the theme of interest.

2.2 Search criteria

In line with the search sources, the study created the permitted criteria to carry out the bibliography review according to terminology employed by experts in high-performance computing. General criteria, like date range, expressions of interest, languages selected, and expressions not permitted are listed ahead:

- Study period: 2010 onwards.
- Languages: Spanish and English.
- Expressions:
  - Measurement in system on chip and cluster
  - Measurement in cluster system on chip high performance computing
  - Benchmark in system on chip and cluster
- Expressions not permitted (the word SoC captures too much information that does not belong to the study theme):
  - Assessment System on Chip

- Metrics System on Chip

- Measurement in SoC
- Benchmark in SoC

With these terms defined, a search chain has been specified that includes adequate terminology according to the functioning of the databases (figure 1).

```
{
        {
         "System on Chip" AND "messurment" AND "cluster"
        }
   OR
        {
         "System on Chip" AND "benchmark" AND "cluster"
        }
   OR
        {
         "System on Chip" AND "cluster" AND "HPC"
        }
   OR
        {
         "System on Chip" AND "cluster" AND "High
         Performance Computing"
        }
}
```

Figure 1. Chain of Selected Search

Source: Prepared by the authors.

Additionally, the study defined the exclusion criteria of the works related:

- Publications prior to 2010.
- Duplicate publications.
- Studies in languages other than those expressed.
- Publications without metrics, or evaluations of clusters of mobile systems on a chip.

2.3 Systematic Review Results

Specifically, the information of interest, upon obtaining the publications selected, was based on evaluation strategies and tools, types of tests, and results of evaluations for clusters based on systems on a chip for high-performance computing. Table 1 shows the amount of works found, excluded, and selected from each source consulted to be analyzed in this review.

Table 1. Amount of Publications per Each Research Source

Source	Amount of Works	Excluded	Final amount of works
IEEE	36	25	11
Science Direct	17	14	3
Engineering Village	61	59	2
Google Scholar	30	27	3
Total	144	125	19

Source: Prepared by the authors.

According to the PRISMA Declaration, 144 studies were identified in the "*Identification*" phase, of which 40 publications duplicated in the different databases were manually excluded, thus, leaving 104 studies, which make up the "*Screening*" phase of the Prisma model. The "*Eligibility*" phase excluded 75 more publications and studies that did not contain the requirements according to the inclusion and exclusion criteria initially expressed. This analysis was first conducted by the title, followed by the abstract, and the contents of the evaluation methodologies. Finally, these 29 documents were analyzed in depth, finding that 10 of the publications contained evaluations of devices with emphasis on computer networks, which is not part of the respective analysis, thereby, considering the 19 studies that finally complied with all the criteria and the emphasis required, and which were selected to be part of this systematic review. The phases of the methodology, as well as the number of publications worked in each of them, can be observed in figure 2.

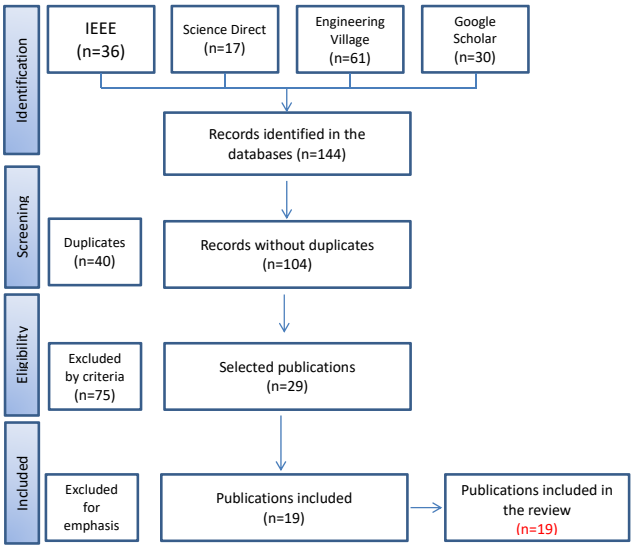

Figure 2. Flow Diagram, PRISMA Declaration

Source: Prepared by the authors.

2.4 Statistical information

The following presents the statistical data on the studies selected, from a general vision of the information analyzed in the evaluation of clusters based on systems on a chip for high-performance computing. Table 2 summarizes the data collected.

Table 2. Publications Selected for Systematic Review

Publication	Hardware	Benchmarks
A study of big data processing constraints on a low-power Hadoop cluster [12]	Cubieboard	BigDataBench. micro-benchmarks: • wordcount • sort • grep
Evaluating Performance and Energy on ARM-based Clusters for High Performance Computing [13]	BeagleBoard PandaBoard	High-Performance Linpack (HPL)
Evaluation of Mobile ARM-Based SoCs for High Performance Computing [14]	NVIDIA Jetson TK1 Odroid XU3-Lite The Parallella Board Wandboard Quad	OpenBLAS STREAM benchmark for measuring memory bandwidth EPCC OpenMP microbenchmark suite.
Performance measurements and comparison between Cruz Cluster I and Cruz Cluster II [15]	Raspberry Pi 2 Raspberry Pi B+	High Performance Linpack (HLP)
Optimizing performance and power consumption for an ARM-based big data cluster [16]	Cubieboard	JVM, CRC32, data compression
Performance and Energy evaluation of Spark applications on low-power SoCs [17]	Intel E5-2650 Intel i5-430M Raspberry Pi 3 Snapdragon410	Spark applications benchmark (micro-benchmarks)
A Low-Cost Energy-Efficient Raspberry Pi Cluster for Data Mining Algorithms [6]	Raspberry Pi 2 Intel Xeon Phi	Data mining algorithms: Association Rule Learning (A priori) and K-Means.
A Raspberry Pi Cluster Instrumented for Fine-Grained Power Measurement [7]	Raspberry Pi Model 2B	Cluster benchmarking: Linpack STREAM
Efficiency Modeling and Analysis of 64-bit ARM Clusters for HPC [4]	Juno ARM AMD Seattle Micro X-Gene	SGEMM, DGEMM [10] and HPL [11].
Energy Efficiency of a Low Power Hardware Cluster for High Performance Computing [5]	Odroid-C2 Odroid-XU4 Raspberry Pi 3 ASRock J3160TM-ITX	Himeno Benchmark. NASA NAS Parallel Benchmarks. MPI-based k-means Algorithm. Distributed Video Encoding.
Evaluating ARM HPC Clusters for Scientific Workloads [9]	ODROID-X	Benchmarks (STREAM, Sysbench, and PARSEC) and benchmarks (High Performance Linpack (HPL), NASA Advanced Supercomputing (NAS) Parallel Benchmark (NPB), and Gadget-2)

Publication	Hardware	Benchmarks
Evaluating Systems on Chip through HPC bioinformatics and astrophysics applications [8]	NVIDIA Jetson K1 ODROID-XU3 ODROID-XU-E ARNDALE OCTA Freescale IMX.6	High Performance Linpack (HPL) benchmark GROMACS, NAMD and GADGET2
Study of Raspberry Pi 2 Quad-core Cortex-A7 CPU Cluster as a Mini Supercomputer [18]	Raspberry Pi 2 model B	High Performance Linpack Benchmark
Supercomputing with Commodity CPUs Are Mobile SoCs Ready for HPC [19]	NVIDIA Tegra 2 NVIDIA Tegra 3 Samsung Exynos 5250	11 Micro-kernels benchmarks: Vector operation Dense matrix-matrix multiplication 3D volume stencil computation 2D convolution One-dimensional Fast Fourier Transform Reduction operation Histogram calculation Generic merge sort N-body calculation Markov Chain Monte Carlo method xi. Sparce Vector-Matrix Multiplication
Tibidabo - Making the case for an ARM-based HPC system [2]	NVIDIA Tegra 2	SPEC CPU2006 benchmarks STREAM benchmark High-Performance Linpack (HPL)
Tiny GPU Cluster for Big Spatial Data A Preliminary Performance Evaluation [20]especially for geo-referenced spatial (or geospatial	NVIDIA Tegra K1	Micro-Kernels benchmarks: taxi-nycb g10m-wwf
Towards an energy efficient SoC computing cluster [21]	Cubieboard 2	NAS Parallel Benchmarks
Building a low consumption cluster using SBC technology [22]	Raspberry PI B+	NAS Benchmarks
Building and benchmarking a low power ARM cluster [3]	Pandaboard ES Raspberry Pi	CoreMark STREAM benchmark Linpack High Performance Linpack (HPL) ping pong benchmark NAS parallel benchmark

Source: Prepared by the authors.

According to the information illustrated in figure 3, the date range of the publications selected is between 2012 and 2017; bear in mind that the study consultation began in 2010.

2012	2013	2014	2015	2016	2017
10.5 %	5.3 %	26.3%	10.5 %	36.8 %	10.5 %

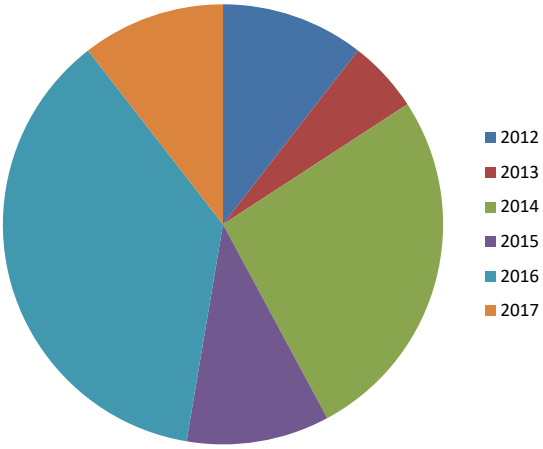

Figure 3. Percentage of Studies per Year of Publication

Source: Prepared by the authors.

The origin of the publications included in the review can be seen in figure 4. These publications are concentrated in three continents, America, Asia, and Europe, with Europe having the highest percentage of publications on the theme selected with 52 % of all the studies. It should be noted that Spain was the European country with the highest number of publications. Likewise, in the other continents, The United States, Brazil, and Thailand had the highest participation (figure 5).

Figure 4. Publications per Country

Source: Prepared by the authors.

According to the country of origin of the publications studied, three institutions stand out in Europe: The Supercomputing Center in Barcelona, the Department of Computer Architecture at Universidad Politécnica de Catalunya, and Cambridge University in the United Kingdom, which had a higher number of studies on the theme of the systematic review.

Consequently, from the northern region of the American continent the study obtained greater information from institutions like the University of Maine in Orono, the City University of New York Graduate Center and the City College of the City of New York, and the University of Oklahoma in The United States. Additionally, in the southern region of the American continent some of the investigations were presented by Pontificia Universidade Católica Minas Gerais, and Universidade Federal de Santa Catarina in Brazil. Furthermore, the Asian continent was represented by institutions like the College of Computer Engineering, and the Suranaree University of Technology in Nakhon Ratchasima Province in Thailand.

Figure 5. Institutions Where the Studies Selected were Conducted

Source: Prepared by the authors.

2.5 Answering the Research Questions

According to the questions initially made, which were focused on recognizing the performance evaluations of the clusters based on systems on a chip for high-performance computing, it was possible to extract information that answers the three research questions.

The initial question corresponded to the inquiry: What type of devices were used in the investigations consulted? Regarding the technology used in the studies analyzed, it was found that the Raspberry, Odroid and Nvidia systems are the most representative,

according to the selection made in the bibliography review. These technologies represent 58.2 % of all the technologies used, according to figure 6.

Each of these systems represented 19.4 %. The remaining percentage focused on systems like Juno ARM, AMD SeattleMicro X-Gene, Snapdragon 410, Arndale Octa, Freescale IMX.6, The Parallella Board, Wandboard Quad, Pandaboard, and Cubieboard.

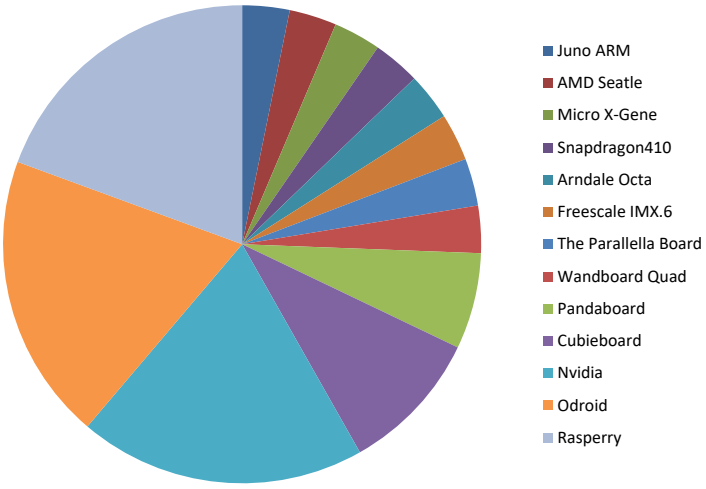

Juno ARM	AMD Seatle	Micro X-Gene	Snapdragon 410	Arndale Octa	Freeascale IMX.6	The Parallela Board
3.2 %	3.2 %	3.2 %	3.2 %	3.2 %	3.2 %	3.2 %
Wandboard Quad	Pandaboard	Cubieboard	Nvidia	Odroid	Rasperry	
3.2 %	6.5 %	9.7 %	19.4 %	19.4 %	19.4 %	

Figure 6. Technology Used in the Publications of Clusters of Systems on a Chip

Source: Prepared by the authors.

The second question referred to What tests or benchmarks were implemented to measure cluster performance. As a result of the review, approximately 25 different performance tests were found (figure 7), highlighting the High Performance Linpack (HPL), Stream, and NAS Parallel Benchmarks tests.

Finally, the third question asked What other parameters were kept in mind to conduct the evaluations. To answer this question, it was established that the benchmarks currently used are standardized, and the measurement parameters are already identified for the comparisons made to be measured under the same conditions. Hence, no additional parameters were evidenced from those specified in the performance tests.

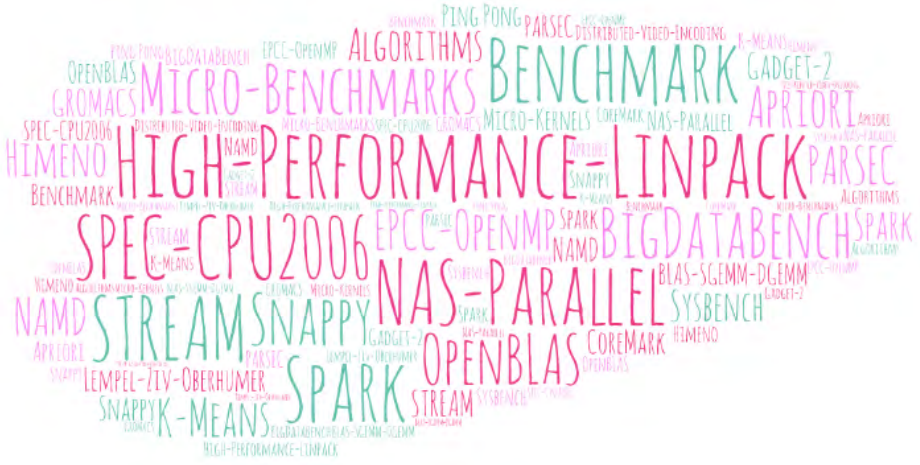

Figure 7. Tests or Benchmarks Employed in the Studies Selected

Source: Prepared by the authors.

DISCUSSION

As stated, cutting-edge HPC systems have high-energy consumption, which has motivated the community to seek strategies aimed at reducing energy consumption while offering high performance. These strategies use diverse forms of evaluating these systems to demonstrate their performance and reduction of energy consumption. These evaluations, methodologically speaking, are known as performance tests or benchmarks.

In the field of HPC, benchmarks appear during the stages of development of the clusters, including the design, implementation, and maintenance, seeking to estimate, evaluate, and ensure the effectiveness of the system [23]. Benchmarks help to quantify the capacity of HPC systems. The result of the benchmarks is that of finding the suitability of a given architecture for a selection of computer applications. Said measurements must be based on real computer applications, and on a comparative and consistent evaluation process [24].

Because of the review, it was possible to identify that the most relevant performance tests in the field of HPC are High-Performance Linpack (HPL), Stream, and NAS Parallel Benchmarks.

The HPL is an implementation of the Linpack benchmark, developed specifically to be executed in architectures that provide support to the parallel-distributed processing. Likewise, it is used to measure the performance of a high-performance computer and is currently used by the Top500 and Green500 to compile their rankings. The HPL is based

on the resolution of dense linear systems (generated randomly) from double-precision equations (64 bits), over a system with distributed memory [25]. The HPL provides programs to evaluate the numerical precision of the results and the time of resolution. The reported performance value depends on diverse factors, but under certain efficiency assumptions of the intercommunication network; the HPL resolution algorithm may be considered scalable, and its efficiency is maintained constant with respect to the memory used by each processor. Execution of HPL requires an implementation of MPI distributed memory and the BLAS library. The unit to present the HPL results is Flops [15].

STREAM is a benchmark that permits evaluating the efficiency of access to the principal memory through four operations over vectors and it is the standard industry benchmark to measure a system's memory bandwidth and the computation rhythm of simple vectors. The program runs four copy, scale, add, and triad operations of floating-point values over large matrices to calculate the memory bandwidth obtained in each operation in MB/s. Table 3 shows the number of flops and bytes counted every loop iteration [3].

Table 3. Operations of the STREAM benchmark

Operation	Instructions	Per iteration	
		bytes	Flops
COPY	*a(i)= b(i)*	*16*	*0*
SCALE	*a(i)=q* b(i)*	*16*	*1*
SUM	*a(i)=b(i)+ c(i)*	*24*	*1*
TRIAD	*a(i)=b(i)+ q*c(i)*	*24*	*3*

Source: [3].

The "NAS Parallel Benchmarks" suite is a project developed by the NASA Advanced Supercomputing (NAS) Division to evaluate the performance of highly parallel computing in supercomputers. The benchmark derives from Computational Fluid Dynamics (CFD) applications, and consists of five kernels and three pseudo-applications. Currently, NAS Benchmarks comprise 11 benchmarks, of which eight were originally available in the NPB 1 version: MultiGrid (MG), Conjugate Gradient (CG), Fast Fourier Transform (FT), Integer Sort (IS), Embarrassingly Parallel (EP), Block Tridiagonal (BT), Scalar Pentadiagonal (SP), and Lower-Upper symmetric Gauss-Seidel (LU); and three other benchmarks available from version NPB 3.1 and NPB 3.2: Unstructured Adaptive (UA), Data Cube Operator (DC), and Data Traffic (DT) [3].

According to the benchmarks presented and the works studied, these three performance tests are robust tools to evaluate cluster systems based on SoC for HPC.

To implement these tests, it is necessary to conduct a specific study in each of these tools to obtain standardized results that are comparable with the related literature.

In most of the works analyzed, it was found that the evaluation conducted made possible to know the performance and power consumed by the clusters. Performance was mainly evaluated by using performance tests already named. Power consumption is a determinant factor in HPC systems based on SoC, given that the number of instructions executed per watt consumed are related, with this being a currently recognized parameter, which was previously not relevant for the construction of clusters for HPC.

Additionally, it must be highlighted that referent measurements exist to evaluate this technology, like the Top500 and Green500 lists [26]. The former classifies supercomputers according to the performance of instructions in floating point, while the latter classifies the most energy efficient supercomputers.

The Top500 list does not have as criterion to classify big machines according to energy consumption, although it is a parameter that has always been measured. The Green500 list calculates the performance of the cluster, and captures the power consumed by the cluster during the execution of the performance tests. Thereafter, performance per watt is determined by using the following formula: Performance per watt (PPW) = Performance/power [3].

It should be stressed that when designing an HPC cluster, the main goal is to obtain a hardware configuration that offers good results in the application or sets of applications. This hardware will be evaluated through performance tests that will allow to know the efficiency of these systems. These tests must be conducted with standard benchmarks for their results to be comparable.

4. CONCLUSIONS

This work presented a systematic review of the evaluations conducted on clusters of Systems on a chip for high-performance computing, using elements from the Prisma Declaration. The Declaration made possible to establish a methodological route to carry out the search and systematic review of the bibliography consulted.

As a result of the review, it was found that the performance tests most used were: High-Performance Linpack (HPL), NAS Parallel Benchmarks, and STREAM benchmark. These tests are the best known in HPC, given that they introduce standard results, which permit using comparatives. Likewise, the hardware most used were the Raspberry, Odroid, and Nvidia systems.

REFERENCES

[1] Schadt, E., Linderman, M., Sorenson, J. *et al.* "Computational solutions to large-scale data management and analysis," *Nat Rev Genet,* no. 11, pp. 647–657, 2010. DOI: https://doi.org/10.1038/nrg2857

[2] N. Rajovic, A. Rico, N. Puzovic, C. Adeniyi-Jones, and A. Ramírez, "Tibidabo: Making the case for an ARM-based HPC system*," Future Generation Computer Systems,* no. 36, pp. 322–334, 2014. DOI: https://doi.org/10.1016/j.future.2013.07.013

[3] N. Balakrishnan, *Building and benchmarking a low power ARM cluster*, M.S. Thesis, EPCC Edinburgh Parallel Computing Center, The University of Edinburgh, 2012. Available: http://static.epcc.ed.ac.uk/dissertations/hpc-msc/2011-2012/Submission-1126390.pdf

[4] J. W. Weloli, S. Bilavarn, S. Derradji, C. Belleudy and S. Lesmanne, "Efficiency Modeling and Analysis of 64-bit ARM Clusters for HPC," 2016 Euromicro Conference on Digital System Design (DSD), Limassol, pp. 342-347, 2016. DOI: https://doi.org/10.1109/DSD.2016.74

[5] M. Görtz, R. Kühn, O. Zietek, R. Bernhard, M. Bulinski, D. Duman, B. Freisen, U. Jentsch, T. Klöppner, D. Popovic, and L. Xu, "Energy Efficiency of a Low Power Hardware Cluster for High Performance Computing," Eibl, M. & Gaedke, M. (Hrsg.), INFORMATIK 2017. Gesellschaft für Informatik, Bonn, pp. 2537-2548, 2017. DOI: https://doi.org/10.18420/in2017_256

[6] J. Saffran *et al.*, "A Low-Cost Energy-Efficient Raspberry Pi Cluster for Data Mining Algorithms," in Desprez F. et al. (eds) Euro-Par 2016: Parallel Processing Workshops. Euro-Par 2016. Lecture Notes in Computer Science, vol 10104. Springer, Cham. 2017. DOI: https://doi.org/10.1007/978-3-319-58943-5_63

[7] M. Cloutier, C. Paradis, and V. Weaver, "A Raspberry Pi Cluster Instrumented for Fine-Grained Power Measurement," *Electronics,* vol. 5, no. 4, p. 61, 2016. DOI: https://doi.org/10.3390/electronics5040061

[8] L. Morganti, D. Cesini, and A. Ferraro, "Evaluating Systems on Chip through HPC Bioinformatic and Astrophysic Applications," in *2016 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP),* Heraklion, pp. 541-544, 2016. DOI: https://doi.org/10.1109/PDP.2016.82

[9] J. Maqbool, S. Oh, and G. C. Fox, "Evaluating ARM HPC clusters for scientific workloads," *Concurrency Computation*, vol. 27, no. 17, pp. 5390–5410, 2015. DOI: https://doi.org/10.1002/cpe.3602

[10] R. Manchado Garabito, S. Tamames Gómez, M. López González, L. Mohedano Macías, M. D'Agostino, and J. Veiga de Cabo, "Revisiones Sistemáticas Exploratorias," *Medicina y Seguridad del Trabajo*, vol. 55, no. 215, pp. 28–51, 2009.

[11] G. Urrútia, and X. Bonfill, "Declaración PRISMA: una propuesta para mejorar la publicación de revisiones sistemáticas y metaanálisis," *Med. Clin. (Barc),* vol. 135, no. 11, pp. 507–511, 2010. DOI: https://doi.org/10.1016/j.medcli.2010.01.015

[12] C. Kaewkasi, and W. Srisuruk, "A study of big data processing constraints on a low-power Hadoop cluster," 2014 International Computer Science and Engineering Conference (ICSEC), Khon Kaen, pp. 267-272, 2014. DOI: https://doi.org/10.1109/ICSEC.2014.6978206

[13] E. L. Padoin, D, P. Velho, and P. O. A. Navaux, "Evaluating Performance and Energy on ARM-based Clusters for High Performance Computing," in *41st International Conference on Parallel Processing Workshops, Pittsburgh,* 2012. DOI: https://doi.org/10.1109/ICPPW.2012.21

[14] A. Selinger, K. Rupp, and S. Selberherr, "Evaluation of Mobile ARM-Based SoCs for High Performance Computing," in *Proceedings of the 24th High Performance Computing Symposium (HPC '16). Society for Computer Simulation International,* pp. 1–7, 2016. DOI: https://doi.org/10.22360/SpringSim.2016.HPC.022

[15] C. Salazar, "Medidas de rendimiento y comparación entre el Clúster Cruz I y el Clúster Cruz II," *Revista de la Facultad de Ciencias de la UNI, Revciuni,* vol. 17, no. 1, pp. 9–16, 2014.

[16] C. Kaewkasi and W. Srisuruk, "Optimizing performance and power consumption for an ARM-based big data cluster," *TENCON 2014 - IEEE Region 10 Conference, Bangkok,* pp. 1-6, 2014. DOI: https://doi.org/10.1109/TENCON.2014.7022399

[18] I. Stamelos, D. Soudris, and C. Kachris, "Performance and energy evaluation of spark applications on low-power SoCs," in *2016 International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (SAMOS), Agios Konstantinos,* pp. 300-305, 2016. DOI: https://doi.org/10.1109/SAMOS.2016.7818362

[19] A. Mappuji, N. Effendy, M. Mustaghfirin, F. Sondok, R. P. Yuniar and S. P. Pangesti, "Study of Raspberry Pi 2 quad-core Cortex-A7 CPU cluster as a mini supercomputer," in *8th International Conference on Information Technology and Electrical Engineering (ICITEE), Yogyakarta,* 2016, pp. 1-4, 2016. DOI: https://doi.org/10.1109/ICITEED.2016.7863250

[20] N. Rajovic, P. M. Carpenter, I. Gelado, N. Puzovic, A. Ramirez and M. Valero, "Supercomputing with commodity CPUs: Are mobile SoCs ready for HPC?," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis,* pp. 1-12, 2013. DOI: https://doi.org/10.1145/2503210.2503281

[21] J. Zhang, S. You and L. Gruenwald, "Tiny GPU Cluster for Big Spatial Data: A Preliminary Performance Evaluation," *in IEEE 35th International Conference on Distributed Computing Systems Workshops*, pp. 142-147, 2015. DOI: https://doi.org/10.1109/ICDCSW.2015.33

[22] Z. Krpić, G. Horvat, D. Žagar and G. Martinović, "Towards an energy efficient SoC computing cluster," 37th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), Opatija, pp. 178-182, 2014. DOI: https://doi.org/10.1109/MIPRO.2014.6859556

[23] L. O. Salvador, *Building a low consumption cluster using SBC technology*, B.Sc. Thesis, Ingeniería Informática, Universidad de Cantabria, 2016. Available: http://hdl.handle.net/10902/9383

[24] M. Tsuji, W. T. C. Kramer and M. Sato, "A Performance Projection of Mini-Applications onto Benchmarks Toward the Performance Projection of Real-Applications," 2017 IEEE International Conference on Cluster Computing (Cluster), Honolulu, HI, pp. 826-833, 2017. DOI: https://doi.org/10.1109/CLUSTER.2017.123

[25] M. Sayeed, H. Bae, Y. Zheng, B. Armstrong, R. Eigenmann and F. Saied, "Measuring High-Performance Computing with Real Applications," *Computing in Science & Engineering,* vol. 10, no. 4, pp. 60-70, 2008. DOI: https://doi.org/10.1109/MCSE.2008.98

[26] A. Remy. *Solving dense linear systems on accelerated multicore architectures*, PhD thesis, Hardware Architecture, Université Paris Sud - Paris XI, 2015. Available: https://tel.archives-ouvertes.fr/tel-01225745/document

[27] Top 500.org, "Top 500 The list," 2018. [Online]. Available: https://www.top500.org/ [Accessed: 28-Jan-2018].