

# Energy Efficiency Evaluation of Multi-level Parallelism on Low Power Processors

Vinícius Garcia Pinto, Arthur F. Lorenzon, Antonio Carlos S. Beck,  
Nicolas Maillard and Philippe O. A. Navaux

<sup>1</sup>Federal University of Rio Grande do Sul (UFRGS)  
Informatics Institute  
Porto Alegre, Brazil

{vgpinto, aflorenzon, caco, nicolas, navaux}@inf.ufrgs.br

**Abstract.** *Energy efficiency and consumption are becoming major concerns in HPC area. One considered alternative to reach better energy efficiency has been the use of unconventional architectures in the HPC scenario, e.g., embedded and mobile processors. In this paper, we present an evaluation about the use of multi-level parallelism in two low-power architectures: Intel Atom and ARM Cortex-A9. Our results show that for all tested cases Intel Atom outperforms ARM Cortex-A9 in terms of execution time and Energy-Delay Product.*

## 1. Introduction

In the next generation of supercomputers, namely exascale systems, a major concern of High Performance Processing community is related to the energy consumption. Exascale machines will have 100 more times processing power than the best current machines. However, the energy required to maintain these systems correspond to power from a nuclear plant of medium size [Wehner et al. 2009]. Therefore, in the same way that is necessary to increase the performance, it is also mandatory to reduce the energy consumption of these supercomputers [Barroso 2005, Asanovic et al. 2009].

Currently, the High Performance Computing (HPC) systems are composed of General Purpose Processors, e.g., Intel Xeon. These processors have great processing power when compared to low-power processors (e.g., Intel Atom), but with high Thermal Design Power (TDP). While the processors present in HPC systems have similar TDP of 130 Watt, low-power processors have TDP much lower than 130 Watt, where in some cases correspond to 2% of this value. For example, the Intel Atom N2600 has maximum TDP of only 3.5 Watt and the ARM Cortex A9 has TDP of 2.5 Watt [Intel 2013, Blem et al. 2013, ARM Ltd 2013]. Therefore, the use of low-power processors is an alternative to the HPC systems that may join the exascale era.

This study aims to compare the use of multi-level parallelism for low-power architectures: Intel Atom and ARM Cortex-A9. The comparison will be performed in terms of performance, energy consumption and Energy Delay Product (EDP – Metric used to study the trade-off between energy and performance). For this, the set of NAS Parallel Benchmarks (NPB) Multi-Zone Version will be used. The main contribution of this paper is the use of multi-level parallelism for low-power systems as an alternative to the high energy consumption of HPC Systems.

This paper is structured as follows: Section 2 presents work related to our research and discusses the contributions of this work in relation to existing work. The contextualization of the set of benchmarks used in this study is done in Section 3. After, in Section 4 two low-power processors target of this work are presented. The results are discussed in terms of energy efficiency and performance in Section 5. Finally, Section 6 contains the conclusions of this study, followed by acknowledgements and our references.

## 2. Related Works

Energy consumption has been identified as one of the major challenges to achieve exascale computing [Bergman et al. 2008]. Many recent studies have considered the use of low power processors as a possible alternative to increase the energy efficiency of the new HPC systems.

Roberts-Hoffman and Hegde [Roberts-Hoffman and Hegde 2009] presents a comparison between Cortex-A series from ARM and Atom N330 by Intel. They use four sequential benchmarks, two integer-based and two floating-point-based. Their results show that the Atom N330 has better raw performance while the Cortex-A8 has significantly greater power efficiency.

The study from Stanley-Marbell and Cabezas [Stanley-Marbell and Cabezas 2011] presents a detailed characterization of three low-power processors (Intel Atom, Power Architecture e500 and ARM Cortex-A8) in terms of performance, power and thermal. In this characterization were used three benchmarks suites: Phoenix MapReduce, MiBench and SPEC CPU 2000. They demonstrated that ARM platform has the lower power dissipation and better energy-efficiency with single-core execution, but using dual core for execution, Intel Atom achieves better energy-efficiency.

Rajovic et al. [Rajovic et al. 2011] built a prototype system based on ARM Cortex-A9. They evaluated the single-core performance of ARM Cortex-A9 vs Intel Core i7 with benchmarks from Dhrystone, STREAM and SPEC CPU 2006 suites. Intel Core i7 outperforms Cortex-A9 by a factor of nine (Dhrystone) and by a factor of five (STREAM). However, Cortex-A9 uses less energy to execute the benchmarks.

A comparison between ARM based cluster and Intel X86 workstation was made by Ou et al. [Ou et al. 2012]. Their tests were made using three applications: web server throughput, in-memory database, and video transcoding. They showed that the ARM cluster is more energy-efficient than the Intel workstation. However, multiple ARM processors are needed to provide comparable performance to an Intel workstation.

Padoin et al. [Padoin et al. 2012] presents a comparison between ARM and Xeon in terms of Time-to-Solution and Energy-to-Solution. They conclude that for HPC using ARM instead of Xeon is still questionable.

The architecture of Tibidabo, a HPC cluster built with ARM processors, was introduced by Rajovic et al. [Rajovic et al. 2013a]. They conclude that Tibidabo energy-efficiency can be competitive with AMD Opteron 6128 and Intel Xeon X5660-based systems. Experiences with NVIDIA Tegra 2, Tegra 3 and Quadro 1000M are reported by Rajovic et al. [Rajovic et al. 2013b]. Their evaluation showed that Tegra 3 reduces in 67% the required energy to solution in comparison with Tegra 2.

A detailed analysis comparing Intel and ARM was performed in [Blem et al. 2013].

In this work, the authors compare the impact of different microarchitectures on performance, power, energy and EDP. In the analysis, different ISA (Instruction Set Architectures) for both general purpose and low-power systems are considered. A great number of variables are evaluated, such as the number of executed instructions, cycle counts, average instruction length, number of memory accesses, and execution time so on. However, only sequential applications running on a single processor are considered, without parallelism exploitation.

An extensive study about the viability of using low power processors in HPC platforms was made by Jarus et al. [Jarus et al. 2013]. They evaluated five platforms (Intel Xeon E7, Intel Core i7, Intel Atom N2600, AMD Fusion and ARM Cortex A9) with seven benchmarks suites (Phoronix, CoreMark, Fhourstones, Whetstone, Linpack, OSU and High-Performance Linpack). Their results showed that Cortex A9 energy usage is up to 12 times lower than the rest of the CPUs, but the execution time of Intel Xeon E7 or Intel Core i7 was up to 117 times shorter.

Different from other studies presented in this section, our work presents an evaluation of low power processors (ARM and Atom) with multi-level parallelism (OpenMP and MPI) benchmarks. In this work we use benchmarks from Multi-Zone Version of the well-known suite NPB (presented in the next section), we also use an innovative metric (EDP) to express the relation between energy saving and delay in execution time.

### 3. NAS Parallel Benchmarks - Multi-Zone Version

The NAS Parallel Benchmarks is a set of programs design to help evaluate the performance of parallel supercomputers and parallelization tools, originally developed by NASA [NAS 2013]. The NPB programs are derived from computational fluid dynamics codes and are implemented with different Parallel Programming Interfaces, such as MPI [Saphir et al. 1997], OpenMP [Jin et al. 1999], Java [Frumkin et al. 2003], and High Performance Fortran (HPF) [Frumkin et al. 1998].

In this work, we use the Multi-Zone version of NPB (NPB-MZ). The NPB-MZ version is designed to exploit multiple levels of parallelism in applications and to test the effectiveness of multi-level and hybrid parallelization paradigms and tools. In this version, are implemented three (LU, BT and SP) of the eight benchmarks available in the single-zone version of NPB. All benchmarks are implemented in Fortran language and parallelized with Open Multi-Processing (OpenMP) and Message-Passing Interface (MPI).

The multi-zone benchmarks lower-upper symmetric Gauss-Seidel (LU), scalar penta-diagonal (SP), and block tri-diagonal (BT) stress the need to exploit multiple levels of parallelism for efficiency and to balance the computational load. All three benchmarks compute discrete solutions of the unsteady, compressible Navier-Stokes equations in spatial dimensions  $x$ ,  $y$  and  $z$  [Jin and der Wijngaart 2006]. In the case of LU multi-zone, the mesh is divided in zones with identical size which makes relatively better the load balancing of the parallelized code. The SP multi-zone code follows a similar strategy, but in this case the number of zones in each of the two horizontal dimensions grows with the problem size. In BT multi-zone, as in the case of SP the number of zones grows with the problem size, but in this benchmark the mesh is not divided in identical size zones. Due to that, in BT is harder to balance the load than for SP and LU, which makes it a more

Class	Memory (approx.)
S	1 MB
W	6 MB
A	50 MB
B	200 MB
C	800 MB
D	12.8 GB

**Table 1. Memory requirements for each problem class**

realistic case. More details about mesh division and boundary communication can be found at [Van der Wijngaart and Jin 2003] and [Jin and der Wijngaart 2006]. The memory requirements for each problem class for the three benchmarks are presented in Table 1.

#### **4. Low Power Platforms**

This section contextualize about the two low power platforms used in this work: ARM and Atom.

##### **4.1. ARM**

ARM is a company that designs processors, architectures and licenses them to manufacturers as Nvidia, Samsung, STMicroelectronics and Texas Instruments. The ARM processors are the world's leading in market of embedded processors.

The last generations of ARM processors includes supports to Single Instruction, Multiple Data (SIMD) instructions, multiprocessing and better floating point support. The ARM Cortex-A family is optimized for low power and high performance applications, and it is used in several modern embedded systems, like digital TVs and smartphones [ARM Ltd 2013]. Because of this characteristics, these processors are considered a possible alternative over traditional CPUs to build supercomputers [Rajovic et al. 2013b, Jarus et al. 2013, Ou et al. 2012, Rajovic et al. 2011].

##### **4.2. Atom**

Intel Atom is the brand name for a line of x86 microprocessors from Intel, designed for ultra portable computers, smartphones and other portable devices with low power consumption. The first Atom processors had a single core. With the evolution of its development, it is now possible to find devices with dual-core technology with Simultaneous Multi-Threading technology [Intel 2013].

Although developed for the purpose of acting in embedded systems, the Atom inherited different characteristics of x86 architectures. The main characteristic is the compatibility with the x86 instruction set. Thus, programs compiled for general-purpose architectures can run without changes in Atom. Because it is a processor to be used in battery-powered systems, has very low thermal dissipation.

#### **5. Experimental Results**

This section discusses the results, provides a comparison between ARM and Atom in terms of performance and energy-efficiency and describes our testing environment.

	ARM Cortex-A9	Atom
Processor	OMAP4430	N2600
Frequency	1 GHz	1.66 GHz
# Cores	2	2
# Threads	2	4
L1 Data Cache	32 KB	24 KB
L1 Inst Cache	32 KB	32 KB
L2 Cache	1 MB/chip	512 KB
Memory RAM	1 GB	2 GB
Platform	Pandaboard	Dev Board
Processor Power (avg)	1.25 W	2.42 W
Memory RAM Power	400 mW	576 mW

**Table 2. Platform Summary**

### 5.1. Execution Environment

To allow multi-level parallelism comparison between the platforms, two clusters were used, one with processors ARM Cortex-A9 and other with processors Atom N2600. The main features of both processors can be found in Table 2 while the complete comparison in [Blem et al. 2013]. In both clusters, the network interconnection between nodes was 100 MB/s and the operating system in use was Linux Debian on its stable version.

The tests were performed using the Class A of NPB-MZ (Section 3) for a total of five nodes in Atom cluster and seven nodes in ARM Cluster. The results presented in the next two sections are an average of 10 executions and the standard deviation is shown together with the results.

### 5.2. Timing and Scalability

The first executed test consists on run the benchmarks in only one node. This test was made using one MPI process and two OpenMP threads. The graphs from Figure 1 show the execution time obtained in ARM cluster with BT, LU and SP benchmarks. Similarly, graphs from Figure 2 present the results from Atom cluster. For all benchmarks, with one thread, ARM cluster is slower than Atom cluster (on average 54% slower for BT, 80% for LU and 88% for SP). With two threads, ARM cluster is also slower than Atom cluster (on average 38% slower with BT, 61% with LU and 79% with SP). However, with BT and LU, ARM cluster achieves higher speedups than Atom cluster (see Table 3).

Benchmark	Threads	ARM-based Speed up	Atom-based Speed up
BT-MZ Class A	1	1.000	1.000
BT-MZ Class A	2	1.847	1.691
LU-MZ Class A	1	1.000	1.000
LU-MZ Class A	2	1.708	1.644
SP-MZ Class A	1	1.000	1.000
SP-MZ Class A	2	1.605	1.642

**Table 3. One-node speed up**

Benchmark	Threads	ARM-based Speed up	Atom-based Speed up
BT-MZ Class A	1	1.000	1.000
BT-MZ Class A	2	1.880	1.960
BT-MZ Class A	4	3.467	3.737
BT-MZ Class A	6	4.505	4.785
BT-MZ Class A	8	5.203	6.121
BT-MZ Class A	10	6.015	7.591
BT-MZ Class A	12	7.541	-
BT-MZ Class A	14	8.014	-
LU-MZ Class A	1	1.000	1.000
LU-MZ Class A	2	1.757	1.951
LU-MZ Class A	4	3.129	3.786
LU-MZ Class A	6	4.031	4.191
LU-MZ Class A	8	4.829	6.219
LU-MZ Class A	10	5.988	5.953
LU-MZ Class A	12	6.456	-
LU-MZ Class A	14	6.928	-
SP-MZ Class A	1	1.000	1.000
SP-MZ Class A	2	1.723	1.983
SP-MZ Class A	4	2.791	3.824
SP-MZ Class A	6	3.291	4.168
SP-MZ Class A	8	4.002	6.020
SP-MZ Class A	10	4.830	5.465
SP-MZ Class A	12	5.250	-
SP-MZ Class A	14	5.693	-

**Table 4. Multi-node speed up**

The second executed test consists on run the benchmarks in distributed nodes, up to 7 nodes for ARM cluster and up to 5 nodes for Atom cluster. This test was made using multiple MPI processes and multiple OpenMP threads. The results are presented in Figures 3 and 4. As in the first test, for all benchmarks, Atom-based cluster outperforms ARM-based cluster. In the case of BT benchmark 12 ARM cores are needed to provide comparable performance to 6 Atom cores. Similar behavior is observed in the LU (14 ARM cores to 6 Atom cores) and SP (2 ARM cores to 1 Atom cores) benchmarks. For all number of threads used Atom cluster achieves higher speed ups than ARM cluster, except for LU with 10 threads when speed up is equal (see Table 4).

### 5.3. Power Consumption and Energy Delay Product Analysis

This section presents and discusses the impact of the aforementioned results on energy consumption. We also consider the Energy Delay Product, so one can have a better idea of the tradeoff considering energy and performance. To calculate energy, it was considered that the cost per core shown in Table 2. The memory power data were gathered from the Cacti Tool [cac 2013, Muralimanohar et al. 2009].

Benchmark	Threads	ARM Energy (J)	Atom Energy (J)	ARM EDP	Atom EDP
BT-MZ Class A	1	1907	2194	2078396	154354
BT-MZ Class A	2	1695	1987	958414	712855
BT-MZ Class A	4	1686	1960	517043	370313
BT-MZ Class A	6	1887	2237	445439	328799
BT-MZ Class A	8	2145	2304	438305	264809
BT-MZ Class A	10	2297	2306	406053	213742
BT-MZ Class A	12	2184	-	307971	-
BT-MZ Class A	14	2387	-	316661	-
LU-MZ Class A	1	1437	1410	1180112	638049
LU-MZ Class A	2	1289	1283	553874	297484
LU-MZ Class A	4	1327	1239	320238	148030
LU-MZ Class A	6	1498	1610	280614	170431
LU-MZ Class A	8	1641	1458	256685	106037
LU-MZ Class A	10	1638	1911	211998	146713
LU-MZ Class A	12	1812	-	211998	-
LU-MZ Class A	14	1961	-	213791	-
SP-MZ Class A	1	1350	1282	1041759	526828
SP-MZ Class A	2	1247	1146	518586	237161
SP-MZ Class A	4	1412	1113	362673	119552
SP-MZ Class A	6	1741	1498	379273	147491
SP-MZ Class A	8	1880	1498	336783	93205
SP-MZ Class A	10	1929	1869	286235	140391
SP-MZ Class A	12	2116	-	288905	-
SP-MZ Class A	14	2266	-	285310	-

**Table 5. Multi-node energy consumption and EDP**

The results presented in the Table 5 correspond to the total energy consumption for the ARM-based and Atom-based clusters. We note that for BT-MZ benchmark, the total energy consumption for execution of an application of the ARM was always less than the Atom. In the best case, ARM has saved 18% of energy with 6 threads. However, increasing the number of threads (8 and 10), this difference decreases. As for the LU-MZ, only two cases ARM consumed less energy than Atom, at 6 and 10 threads. This behavior does not occur for SP-MZ, in that for all test cases, Atom consumed less energy.

The same table also shows the results for the EDP, which corresponds to the relationship between performance and the total energy consumption of the application. In this case, when it is taken into consideration the runtime, the Atom-based cluster obtained the best results. This occurs because although the Atom has consumed more energy to run an application in full, he got better performance than ARM.

## **6. Conclusion**

This work shows an evaluation of the energy efficiency of low-power processors using multi-level parallelism. For this, the set of benchmarks NAS-MZ was used, which contains three applications parallelized with MPI and OpenMP. A comparison in terms of execution time, power consumption and EDP was performed between two different low-power processors: ARM Cortex and Intel Atom.


The results obtained shown that Atom got the best results. Although in some cases, executions in ARM possessed lower power consumption, this factor is not the same when it relates to performance. Thus, for all test cases, the Atom-based cluster proved to be the best option for use of multi-level parallelism at low power processors.

As part of our future work, we plan to repeat executions in general purpose architectures, e.g., Intel Xeon trying to find a ratio between the energy efficiency of low-power systems and general purpose.

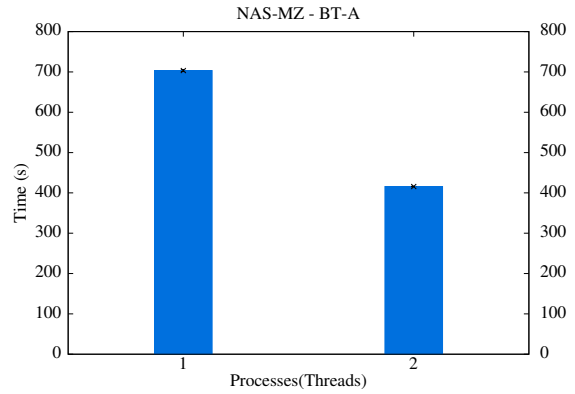
## **Acknowledgment**

The authors would like to thank GridRS Project (<http://gridrs.lad.pucrs.br/>) for providing access to ARM-based cluster, Victor Abaunza for technical support with ARM-based cluster and the Embedded System Laboratory of UFRGS for providing access to Atom-based cluster. This work was partially supported by the Brazilian agencies CAPES and CNPq.







(a) BT-MZ Class A




(a) BT-MZ Class A



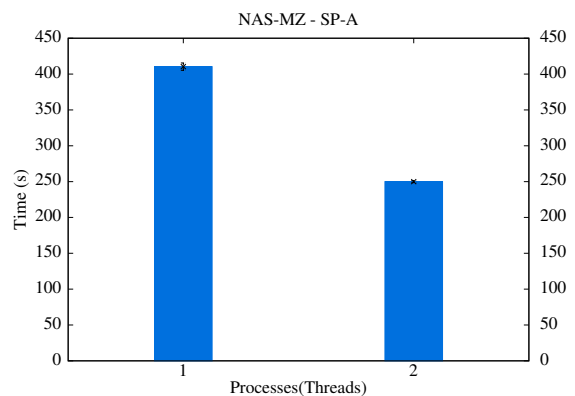
(b) LU-MZ Class A



(b) LU-MZ Class A



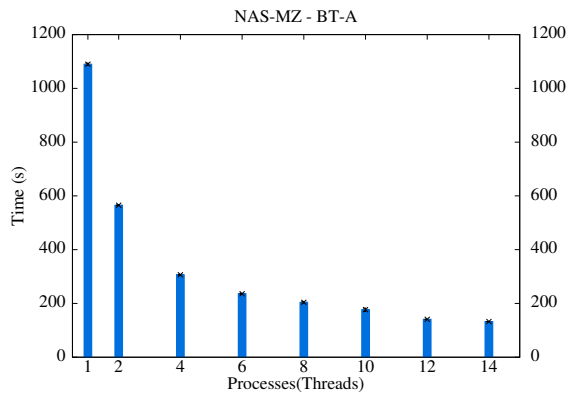
(c) SP-MZ Class A




(c) SP-MZ Class A

**Figure 1. Time for one-node execution on ARM-based Cluster.**

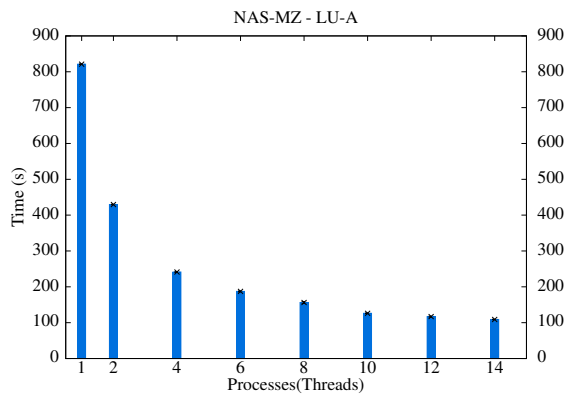
**Figure 2. Time for one-node execution on Atom-based Cluster.**




(a) BT-MZ Class A




(a) BT-MZ Class A




(b) LU-MZ Class A



(b) LU-MZ Class A



(c) SP-MZ Class A



(c) SP-MZ Class A

**Figure 3. Time for multi-node execution on ARM-based Cluster.**

**Figure 4. Time for multi-node execution on Atom-based Cluster.**

## References

- (2013). CACTI 6.0. <http://www.cs.utah.edu/~rajeev/cacti6/>.
- ARM Ltd (2013). Cortex-A Series. <http://www.arm.com/products/processors/cortex-a/index.php>.
- Asanovic, K., Bodik, R., Demmel, J., Keaveny, T., Keutzer, K., Kubiawicz, J., Morgan, N., Patterson, D., Sen, K., Wawrzynek, J., et al. (2009). A view of the parallel computing landscape. *Communications of the ACM*, 52(10):56–67.
- Barroso, L. A. (2005). The price of performance. *Queue*, 3(7):48–53.
- Bergman, K., Borkar, S., Campbell, D., Carlson, W., Dally, W., Denneau, M., Franzone, P., Harrod, W., Hiller, J., Karp, S., Keckler, S., Klein, D., Lucas, R., Richards, M., Scarpelli, A., Scott, S., Snively, A., Sterling, T., Williams, R. S., and Yelick, K. (2008). Exascale computing study: Technology challenges in achieving exascale systems. Technical report.
- Blem, E., Menon, J., and Sankaralingam, K. (2013). Power Struggles: Revisiting the RISC vs. CISC Debate on Contemporary ARM and x86 Architectures. In *2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, pages 1–12. IEEE Computer Society.
- Frumkin, M., Jin, H., and Yan, J. (1998). Implementation of NAS Parallel Benchmarks in High Performance Fortran. *NAS Technical Report NAS-98-009*.
- Frumkin, M., Schultz, M., Jin, H., and Yan, J. (2003). Performance and Scalability of the NAS Parallel Benchmarks in Java. In *Parallel and Distributed Processing Symposium (IPDPS'03), 2003. Proceedings of the International*, pages 1–6.
- Intel (2013). Intel Atom Processor. <http://www.intel.com/content/www/us/en/processors/atom/atom-processor.html>.
- Jarus, M., Varrette, S., Oleksiak, A., and Bouvry, P. (2013). Performance Evaluation and Energy Efficiency of High-Density HPC Platforms Based on Intel, AMD and ARM Processors. In Pierson, J.-M., Da Costa, G., and Dittmann, L., editors, *Energy Efficiency in Large Scale Distributed Systems*, Lecture Notes in Computer Science, pages 182–200. Springer Berlin Heidelberg.
- Jin, H. and der Wijngaart, R. F. V. (2006). Performance Characteristics of the Multi-Zone NAS Parallel Benchmarks. *Journal of Parallel and Distributed Computing*, 66(5):674–685. IPDPS '04 Special Issue - 18th International Parallel and Distributed Processing Symposium.
- Jin, H., Frumkin, M., and Yan, J. (1999). The OpenMP implementation of NAS Parallel Benchmarks and its performance. Technical report, Technical Report NAS-99-011, NASA Ames Research Center.
- Muralimanohar, N., Balasubramonian, R., and Jouppi, N. P. (2009). CACTI 6.0: A Tool to Understand Large Caches. Technical report. <http://www.cs.utah.edu/~rajeev/cacti6/cacti6-tr.pdf>.
- NAS (2013). NAS Parallel Benchmarks. <http://www.nas.nasa.gov/Software/NPB/>.

- Ou, Z., Pang, B., Deng, Y., Nurminen, J. K., Ylä-Jääski, A., and Hui, P. (2012). Energy- and Cost-Efficiency Analysis of ARM-Based Clusters. In *2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2012)*, pages 115–123. IEEE.
- Padoin, E. L., de Oliveira, D. A., Velho, P., and Navaux, P. O. (2012). Time-to-Solution and Energy-to-Solution: A Comparison between ARM and Xeon. In *2012 Third Workshop on Applications for Multi-Core Architecture (WAMCA)*, pages 48–53. IEEE Computer Society.
- Rajovic, N., Puzovic, N., Vilanova, L., Villavieja, C., and Ramirez, A. (2011). The Low-Power Architecture Approach Towards Exascale Computing. In *Proceedings of the second workshop on Scalable algorithms for large-scale systems - ScalA '11*, page 1, New York, New York, USA. ACM Press.
- Rajovic, N., Rico, A., Puzovic, N., Adeniyi-Jones, C., and Ramirez, A. (2013a). Tibidabo1: Making the case for an ARM-based HPC system. *Future Generation Computer Systems*.
- Rajovic, N., Rico, A., Vipond, J., Gelado, I., Puzovic, N., and Ramirez, A. (2013b). Experiences with Mobile Processors for Energy Efficient HPC. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2013*, pages 464–468, New Jersey. IEEE Conference Publications.
- Roberts-Hoffman, K. and Hegde, P. (2009). ARM Cortex-A8 vs. Intel Atom: Architectural and Benchmark Comparisons. Technical report.
- Saphir, W., Van der Wijngaart, R. F., Woo, A., and Yarrow, M. (1997). New Implementations and Results for the NAS Parallel Benchmarks 2. In *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing, PPSC 1997*.
- Stanley-Marbell, P. and Cabezas, V. C. (2011). Performance, Power, and Thermal Analysis of Low-Power Processors for Scale-Out Systems. In *2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*, pages 863–870. IEEE.
- Van der Wijngaart, R. F. and Jin, H. (2003). NAS Parallel Benchmarks, Multi-Zone versions. *NASA Ames Research Center, Tech. Rep. NAS-03-010*.
- Wehner, M., Oliner, L., and Shalf, J. (2009). A Real Cloud Computer. *IEEE Spectrum*, 46(10):24–29.