

Scalability and Energy Efficiency of HPC cluster with ARM MPSoC *

Edson L. Padoin^{1,2}, Daniel A. G. de Oliveira¹, Pedro Velho¹, Philippe O. A. Navaux¹,
Brice Videau³, Augustin Degomme³, Jean-Francois Mehaut³

¹Institute of Informatics

Federal University of Rio Grande do Sul (UFRGS) – Porto Alegre, RS – Brazil

²Department of Exact Sciences and Engineering

Regional University of Northwest of Rio Grande do Sul (UNIJUI) – Ijuí, RS – Brazil

³Laboratoire d'Informatique de Grenoble (LIG)

Grenoble University – Grenoble – France

{elpadoin, dagoliveira, pvelho, navaux}@inf.ufrgs.br

{brice.videau, augustin.degomme, jean-francois.mehaut}@imag.fr

Abstract

Power consumption has already become one of the main challenges to design new HPC architectures slowing down the conception of exascale systems. In response, researchers aim at overcoming power consumption constraints using embedded processors. State-of-the-art projects, like Mont-Blanc, bet on ARM to build energy efficient HPC architectures. Therefore, investigating ways to improve energy efficiency is one of highest priority for future HPC systems. In the present paper we evaluate scalability and energy efficiency of 3 different MPSoCs in cluster. The achieved results indicate that Tegra 2 has the greater scalability and Snowball had the best energy efficiency among the tested MPSoCs.

1. Introduction

The currently number one supercomputer on the Top500 list, the Tianhe-2, performs 33.8 PFlops spending 17.8 MW, however several applications still claim for more computing speed. Research on supercomputing aim at providing faster supercomputers to aid researchers with their huge experiments. However, computing speed costs power and today we are facing a focus change where the main challenge to

increase computing speed lies on reducing power consumption.

To build exascale systems matching the increasing demand of computing performance we need to consider energy consumption constraints [1, 16]. During the last decade the computing power of fastest supercomputers of the world, on the Top500 list, increased exponentially over time. This exponential increase in computing power also represent an exponential increase in the power consumption. Conceiving exascale supercomputers scaling the current cutting edge technology would demand over a GW of power. This is equivalent to the entire production of a medium size nuclear power plant [14]. To avoid this tremendous waste of natural resources a global research effort rises to break the exascale barrier. In 2008 specialists alerted on the official DARPA report [2, 6] that the acceptable power budget to reach exascale is 20 MW. So, the energy efficiency of futures exascale systems have a limit of 50 GFlops/W.

One possible approach to increase speed without incurring on the exponential growth of power consumption rely on using low power processors. Indeed, the handheld industry target processing speed respecting power consumption to improve device's battery life. One architecture that follows this path is ARM. Several companies embed ARM processor cores in the same chip, also called a Multiprocessor System-on-Chip (MPSoC). The first generation of ARM architectures targeted mostly low power consumption and were unsuitable for HPC for 2 main reasons. First a moderate processing speed. Second the lack of floating point unit and hardware support for SIMD instructions. In the other hand, the new ARM Cortex-A processors consid-

* This work was supported by CNPq, CAPES, FAPERGS and FINEP. This research has been partially supported by CAPES-BRAZIL under grants 3471-13-6. Work developed on the context of the associated international laboratory between UFRGS and *Université de Grenoble - LICIA* and project HPC-GA.

erably increases processing speed slightly increasing power consumption. Due to these characteristics several research towards exascale bet on ARM processors.

Project Mont-Blanc is one of the first to introduce the idea of an ARM based supercomputer [7]. The project bets on highly heterogeneous MPSoC combining ARM and GPU to achieve high processing speed at low power consumption. The project aims at decrease power consumption at least 15-fold compared to current fastest supercomputers. Precisely, the Mont-Blanc goal is to delivery 200 PFlops respecting a 10 MW power budget [7, 13]. The project current prototype, Tibidabo, look forward to reach 7 GFlops/W until the end of 2014 [10]. Tibidabo is the first supercomputer using NVIDIA Tegra 2 technology featuring ARM plus GPU on the same MPSoCs, even though this GPU is exclusively for graphic rendering.

This paper question the feasibility to achieve exascale with ARM low power processors. To address this question we analyze the scalability and energy efficiency of 3 different MPSoCs. with ARM processors using HPL benchmark. This paper presents up-to-date results of two papers submitted to international conferences [8, 9].

The rest of the paper is structured as follows. In Section 2 details the ARM architecture. In Section 3 we discuss our proposal and methodology of evaluation. Section 4 describes the tests results obtained. Section 5, we present related works on energy consumption and performance on ARM platforms. Section 6 outlines our conclusions, contributions and future work perspectives.

2. Low Power Processors

Several processors are known as power hungry because their design main focus is speed assuming an interrupted power source. Current supercomputers mainly rely CPUs that follow this last characteristic. However, power consumption is the main concern to build faster supercomputers. Due to this, power hungry CPUs are leaving space to new technologies in supercomputers design. Embedded systems historically targeted low power consumption to improve battery life, presenting a possible alternative over CPUs that are conventional to supercomputers. The ARM Cortex processor family target high performance and low energy consumption [15]. Several innovations that incorporate ARM MPSoCs take energy consumption into account. Through this section we discuss characteristics of several ARM processors which are low power alternatives to build the next generation of supercomputers.

ARM was made with power consumption constraints in mind. In this context, speed was also important but nonetheless a secondary goal. Early versions of the ARM ISA lack of single-precision FP or double-precision support. This made ARM less adapted to scientific applications. Recently,

the ARMv7 ISA have native support to single-precision FP and double-precision. Another improvement that aim speed in this new ISA is the support for SIMD instructions with the NEON unit.

3. Experimental Method

This section describes the methodology for our study. We present the execution environment, and then discuss the Benchmark and Workload methodology.

3.1. Execution Environment

The execution environment is composed of 3 MPSoCs with ARM processors, PandaBoard, Snowball e Tegra 2. Table 1 shows the main characteristics of each test machine.

MPSoC	PandaBoard	Snowball	Tegra 2
Processor ARM Cortex	A9	A9	A9
Manufacturer	Texas Inst.	ST-Ericsson	Nvidia
Instruction Set Architecture (ISA)	ARMv7	ARMv7	ARMv7-a
Processor Model	OMAP4430	A9500	Q7
Processor Technology (nm)	45	45	40
Clock Frequency (GHz)	1	1	1
Cores/Processor (#)	2	2	2
Multi-Threading	Yes	Yes	Yes
TDP (W)	0.25	0.25	0.25
Low Power Memory (GB)	1	1	1
Cache L1/Core (KB)	64	64	64
Cache L2 (KB)	1024	512	1024
Cache L3 Shared (MB)	0	0	0
Advanced SIMD	NEON	NEON	-
Floating Point Unit (FPU)	VFPv3	VFPv3	VFPv3
Graphics Processing Unit (GPU)	-	Mali 400 MP1	GeForce
NIC Ethernet (Mbits)	10/100	10/100	1000
Maximum Power of MPSoC (W)	8.2	2.5	5.7

Table 1. Detailed configuration for each ARM MPSoC.

The operating system used in our tests was the provided by each manufacturer, a modified version of Ubuntu GNU/Linux distribution with kernel version 3.0. The application compiles using arm-linux-gnueabi-gcc version 4.5.

3.2. Benchmark and Workload

To analyze the scalability and energy efficiency, we use High-Performance Linpack (HPL) benchmark parallelized with MPICH2 interface version 1.4 in a cluster with two MPSoCs. Was measured the power consumption of the whole systems, because this determine the final constraint of electrical power on the supplier.

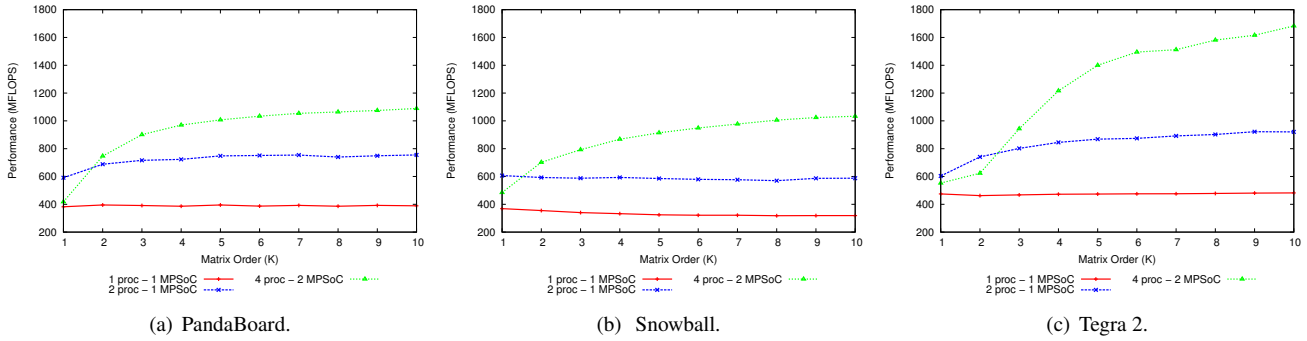


Figure 1. Maximum performance of each MPSoC with different processes.

4. Results

Scalability is one of the big challenges in HPC systems and is related to problems of all levels in the system. In the context of this work, the goal is to analyze scalability and energy efficiency of 3 different MPSoCs in cluster. In this section we present the test results of each MPSoC.

The Figure 1 shows the maximum performance obtained with the best settings of HPL benchmark varying the number of processes and the amount of MPSoCs (n proc - n MP-SoC). To analyze scalability the tests were performed according to the number of processors of each MPSoC.

Table 2 presents the maximum amount of operations executed per second (MFlops) performing 2 processes in a MP-SoC and 4 processes in a cluster with two MPSoCs. Also highlights the scalability achieved.

Performance x MPSoC	PandaBoard	Snowball	Tegra 2
1 MPSoC - 2 processes (MFlops)	755.2	587.6	920.6
2 MPSoC - 4 processes (MFlops)	1089.0	1033.3	1683.0
Scalability	44.2%	75.8%	82.8%

Table 2. Scalability of MPSoCs evaluated with matrix order 10K in HPL

With 4 processes running in the cluster, Tegra 2 had a peak performance of 1.683 MFlops while PandaBoard and Snowball reached only 1.089 MFlops and 1.033 MFlops. That is, Tegra 2 was 1.54 times faster than the PandaBoard MPSoC and 1.62 times faster than the Snowball MPSoC.

Tegra 2 has the greater scalability among the tested MP-SoCs. The Tegra 2 scalability, 82.8%, was practically twice higher than the scalability of PandaBoard.

Tables 3 presents the amount of operations executed per second (MFlops), the maximum instantaneous power (W) and the amount of operations executed per Watt spent (MFlops/W) of each MPSoC.

The best performance achieved was 748.6 MFlops on PandaBoard, 587.6 MFlops on Snowball and 920.6

MPSoC	PandaBoard	Snowball	Tegra 2
Performance (MFlops)	748.6	587.6	920.6
Maximum Instantaneous Power (W)	8.2	2.5	5.7
Energy Efficiency (MFlops/W)	91.3	235.0	161.5

Table 3. Energy Efficiency of each MPSoC with matrix order 5K

MFlops on Tegra 2. However, the maximum instantaneous power of PandaBoard, measuring the entire system power consumption, was 8.2 W against 2.5 W on Snowball [3] and 5.7 W on Tegra 2 [10]. According to these results, the energy efficiency of the MPSoCs PandaBoard, Snowball e Tegra 2 were 91.3 Mflops/W, 235.0 Mflops/W and 161.5 Mflops/W respectively.

Tables 4 is similar to Table 3, but presents the energy efficiency of two MPSoC in cluster. Was adopted as maximum instantaneous power the multiplying the power of a MPSoCs by the amount of MPSoC used. However, as the clusters have the same type of interconnection equipment, in these tests we disregard their power consumption.

MPSoC	PandaBoard	Snowball	Tegra 2
Performance (MFlops)	1089.0	1033.0	1683.0
Maximum Instantaneous Power (W)	16.4	5.0	11.4
Energy Efficiency (MFlops/W)	66.4	206.6	147.6

Table 4. Energy Efficiency of two MPSoCs in cluster with matrix order 10K

Tegra 2 had the best performance. On the other hand, when related performance with power, the MPSoC Snowball had the best energy efficiency.

5. Related Work

The energy consumption is a central question on the quest for the next generation supercomputers. Many papers evaluate the power consumption with a variety of ap-

plications and architectures. Unconventional solutions appear in the care of power consumption on the supercomputing scene. An approach is to use ARM MPSoCs.

Jack Dongarra and Piotr Luszczek [4] present a landscape to analyze ARM. The authors make a comparison of energy efficiency with ARM and several Intel CPUs. Their results point that ARM has better efficiency (about 4 GFlops/W). Valero *et al.* [13] present similar results. The Cortex A9 have an efficiency of 4 GFlops/W and the future Cortex A15 have 8 GFlops/W. These values depict better efficiency than INTEL and IBM processors.

Roberts-Hoffman *et al.* [11] describes a comparison between ARM Cortex A8 and Intel Atom N330. Furlinger *et al.* [5] analyzes the energy efficiency of parallel and distributed computing commodity devices. Their studies compare the performance and energy consumption of an AppleTV cluster using ARM Cortex A8 processors.

Stanley-Marbell *et al.* [12] discuss thermal constraints on low power consumption processors. They establish a connection between energy consumption and the processor architecture, either, ARM, Power and Intel Atom. The ARM platform has minor energy consumption, presenting a better efficiency with light-weight workloads. However, the Intel platform, which has more power consumption rates, had the best energy efficiency for heavy workloads.

6. Conclusion

This paper presented an analysis of 3 different MPSoCs to analyze the feasibility of building cluster with ARM processors. Our contributions, using the HPL benchmark, include: (i) evaluation of the performance and scalability and (ii) analysis of the energy efficiency.

We evaluated 3 MPSoCs based on ARM, varying the number of processes and the number of MPSoC. In our experiments, we saw that Tegra 2 was 1.54 times faster than the PandaBoard MPSoC and 1.62 times faster than the Snowball MPSoC. Tegra 2 also has the greater scalability among the tested MPSoCs. On the other hand, the MPSoC Snowball had the best energy efficiency.

Our future works will focus on investigating news MPSoC and evaluate energy efficiency using a cluster with greater amount of MPSoCs.

References

- [1] K. Barker, K. Davis, A. Hoisie, D. Kerbyson, M. Lang, S. Pakin, and J. Sancho. Using Performance Modeling to Design Large-Scale Systems. *IEEE Computer*, 42(11):42–49, 2009.
- [2] P. Beckman, B. Dally, G. Shainer, T. Dunning, S. C. Ahalt, and M. Bernhardt. On the Road to Exascale. *Scientific Computing World*, (116):26–28, 2011.
- [3] Calao-Systems. *Snowball Home Page*. 2012. <http://www.calao-systems.com/articles.php?lng=fr&pg=6186>.
- [4] J. Dongarra and P. Luszczek. Anatomy of a Globally Recursive Embedded Linpack Benchmark. In *16th IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, MA, USA, 2012.
- [5] K. Furlinger, C. Klausecker, and D. Kranzlmüller. Towards Energy Efficient Parallel Computing on Consumer Electronic Devices. In *1st International Conference on Information and Communication on Technology*, 2011.
- [6] P. Kogge, K. Bergman, S. Borkar, D. Campbell, W. Carson, W. Dally, M. Denneau, P. Franzon, W. Harrod, and K. Hill. Exascale Computing Study: Technology Challenges in achieving Exascale Systems. *Defense Advanced Research Projects Agency (DARPA IPTO)*, pages 1–297, 2008.
- [7] Mont-Blanc. *Mont-Blanc Project Home Page*. Last visited January 2013. <http://www.montblanc-project.eu/>.
- [8] E. L. Padoin, D. A. G. Oliveira, P. Velho, and P. O. A. Navaux. Evaluating the performance and energy of ARM processors for High Performance Computing. In *41st International Conference on Parallel Processing Workshop (ICPPW - UCAA)*.
- [9] E. L. Padoin, P. Velho, D. A. G. de Oliveira, P. O. A. Navaux, B. Videau, A. Degomme, and J.-F. Mehaut. Análise de desempenho, escalabilidade e eficiência energética de MPSoCs com processadores ARM. In *Conferencia Latino Americana de Computación de Alto Rendimiento (CLCAR)*, pages 1–10, San José, Costa Rica, 2013.
- [10] N. Rajovic, N. Puzovic, and A. Ramirez. Tibidabo: Making the Case for an ARM Based HPC System. In *Mont-Blanc Publications*, pages 1–11, BSC, Spain, 2012.
- [11] K. Roberts-Hoffman and P. Hegde. ARM Cortex-A8 vs. Intel Atom: Architectural and Benchmark Comparisons. In *Report at University of Texas*, pages 1–8, Dallas, USA, 2009.
- [12] P. Stanley-Marbell and V. Cabezas. Performance, Power, and Thermal Analysis of Low-Power Processors for Scale-Out Systems. In *25th IEEE Parallel & Distributed Processing, Workshops and Phd Forum (IPDPSW)*, pages 863–870, Anchorage, Alaska, USA, 2011.
- [13] M. Valero. Towards Exaflop Supercomputers. In *High Performance Computing Academic Research Network (HPC-net)*, pages 1–117, Rio Patras, Greece, 2011.
- [14] M. Wehner, L. Olikar, and J. Shalf. A real cloud computer. *Spectrum, IEEE*, 46(10):24–29, 2009.
- [15] G. Yeung, P. Hoxey, P. Prabhat, Y. Chong, D. O’Driscoll, and C. Hawkins. Low Power Memory Implementation for a GHz+ Dual Core ARM Cortex A9 Processor on a High-K Metal Gate 32nm Low Power Process. In *IEEE Design, Automation and Test (VLSI-DAT)*.
- [16] A. Younge, G. von Laszewski, L. Wang, S. Lopez-Alarcon, and W. Carithers. Efficient Resource Management for Cloud Computing Environments. In *IEEE International Green Computing Conference*, Chicago, IL, USA, 2010.