HPC vs HPC
High Performance Computing vs High Power Consumption

Edson Luiz Padoin[1,2], Philippe O. A. Navaux[2]
[1]Departamento de Ciencias Exatas e Engenharias (DCEEng)
Universidade Regional do Noroeste do Estado do Rio Grande do Sul (UNIJUI) – Ijui, RS – Brasil
[2]Instituto de Informatica
Grupo de Processamento Paralelo e Distribuido (GPPD)
Universidade Federal do Rio Grande do Sul (UFRGS) – Porto Alegre, RS – Brasil
{elpadoin, navaux}@inf.ufrgs.br

Abstract

The High Power Consumption(HPC) is one of biggest problems for the High Performance Computing (HPC) community and one of the major obstacles for exascale systems design. The new generations of HPC systems intend to achieve exaflop performances and will demand even more energy to processing and cooling. Nowadays, the growth of HPC systems is limited by energy issues, in this context, in this paper we present a study and an analyse of the some issues that has impact in the energy efficiency.

1. Introduction

The power consumption has been a serious problem for the High Performance Computing (HPC) community and one of the major obstacles for exascale systems design [11] [4]. In a world with limited energy resources and an rising demand for more computational power, energy consumption is limiting the scale of computers that can be deployed. Therefore, the scientific community is searching for ways to improve the power efficiency of the HPC systems [7].

In the Top500 list, published in November 2009, the average energy consumption of the top 10 computers reaches 2 MW. Among these, the most power consuming system has reached at 7 MW [4]. After only a year, in the 36^{th} list, this average increased 65%.

With this, we can conclude what the growth of systems of high performance computing is limited by energy issues. So, to build exascale systems, it is necessary to improve energy efficiency(EE) (Flops/Watt) of current systems.

This paper presents a review about High Power Consumption of High Performance Computing systems, one of the major hurdles in the path to exascale system. It is organized as follows: Section 2 shows an evolution of the processors. Section 3 introduces the power consumption, the power efficiency and rank lists. Section 4 shows the problems to build exascale systems. Section 5 reviews the related works and the Section 6 concludes the paper.

2. Evolution of the Processors Architecture

The performance of HPC systems has evolved through the addition of processors and cores, being the Intel one of the largest manufacturers of processors. The comparison of their generations of processors allows us to analyze the increased performance and a significant increase in energy consumption.

Table 1 shows information on 3 models of Intel family processors. The important changes are:"

Model	Clock (Mhz)	Transistor size(nm)	Transistor amount(M)	Voltage (V)	Current max(A)	TDP (W)
Pentium	60	800	3.1	5	2.92	14.6
Pentium 4	3800	90	125	1.25	119	115
Times	63.333	8.89	40.3	4.00	40.75	7.87
Pentium	60	800	3.1	5	2.92	14.6
Core i7	3200	32	1170	0.8	110	130
Times	53.333	25.00	377.41	6.25	37.67	8.90

Table 1. Evolution of the Intel Family Processors

- the increase of operating frequency from 60 MHz in the first Pentium model to values near to 4 GHz in some current models (63 times);

- with the decrease of size of the transistor from 800 nm to 32 nm (25 times) was a very significant increase in the number of transistors per chip, from 3.1 million to 1,170 million (377 times);
- the voltage was reduced from 5 V since the processors 8086/8088 to values less than 1V in current processors (6.25 times);
- the electric current was increased from 2.92 A to values above 100 A. As an example, see the processor i7 (37.7 times);
- all these changes have caused a significant increase in Thermal Dissipation Power (TDP) of processors. The first Pentium models had a TDP of 14 W, while the processors i7 have value of 130 W. (8.9 times);

An important feature was the significant increase of operating frequencies of processors, reaching the supported limit of cooling and the physical limits of materials used. These technological innovations have produced a considerable increase in computational performance but also increased energy consumption of processors.

3. Power Consumption of the HPC systems

Twenty years ago, HPC systems consumed less than a megawatt. The Earth Simulator was the first system to exceed 10 MW [6]. Today, power consumption is one of the main issues in researches for construction of large scale systems. Given the current energy consumption of petascale systems, and considering the limited supply of energy, it will not be possible to build systems of higher computing power.

The average consumption of the first 10 systems of the 34^{th} TOP500 list was 2 MW [4]. After only one year, in the 36^{th} list, this average was increased by 65%, where the system of higher consumption is the Jaguar, from Oak Ridge National Laboratories with almost 7 MW of power consumption.

Power is proportional to the product of the Capacitance, Frequency and Voltage squared [11] [3]. For integrated circuits implemented in CMOS technology, the average power consumption is given by:

$$P = f.C_L.Vdd^2.\alpha$$

where f is the frequency of the system clock, C_L is the load capacitance, Vdd is the voltage and α is the transition activity.

In this scenario, an important issue is the measurement and analysis of power consumption. Some systems have hardware for this measurement, an example is the Cray XT, that provides interfaces in each node to measure the power consumption [11]. On the other hand, in other systems, the measurement can be a problem. In these cases, the solution may be to simulate the power consumption, where simulators as Wattch, PowerScope or others are used.

In other studies, the energy consumption is often measured via an external power meter, such as the Yokogawa Digital Power Analyzers and Dranetz Power Platform.

3.1. Lists of Performance and Power Consumption

HPC systems are often evaluated by its computing power (flops). Since 1993, the Top500 list offers a ranking of the 500 fastest supercomputers. Running the Linpack benchmark, the supercomputers are compared and ranked according to their performance.

However, due to the large increase in the energy consumption of these systems, interest in this area has increased and the analysis of power efficiency has also been used to evaluate systems [9].

The green computing area is becoming increasingly important in a world with limited energy resources and the HPC (High Power Consumption) has been one constraint on HPC systems designs. With this concern, the Green500 list was created to provide a ranking of the most energy-efficient supercomputers.

So, the Top500 List offers a ranking of the 500 fastest supercomputers (Flops) and the Green500 List, ranks the top 500 supercomputers in the world by energy efficiency (Flops/Watt).

The first sistem in the Green list, published in november 2010, is the prototype IBM Blue Gene/Q system at the TJ Watson Research Center, with 1684.2 Mflops/W. Followed by HP ProLiant SL 390s, with 958.35 Mflops/W and by a hybrid cluster core at the Nacional Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign with 933.86 Mflops/W [13].

The performance increase of HPC systems has been achieved with the increase of processors/cores and, recently, by adding accelerators such as graphic processor units (GPU). This organization, with different resources, is called heterogeneous architecture and in newer versions of the Top 500 list can be seen the their increasing number. Aaccelerator-based systems take 8 of the top ten places on 36^{th} Top500 list [13]. This class of computers appears in the Green500 Lists with an average efficiency of 756Mflops/W, while others have 211Mflops/W [13].

For the construction of exascale systems, much research has been conducted in several centers of high-performance processing, all considering the power consumption as the primary metric to be reduced [1].

Supercomputer performance has doubled more than 3000 times in the past 15 to 20 years, the performance per watt has increased 300 times and performance per square foot has only increased 65 times in the same pe-

riod of time [16]. This relationship can be seen in Figure 1.

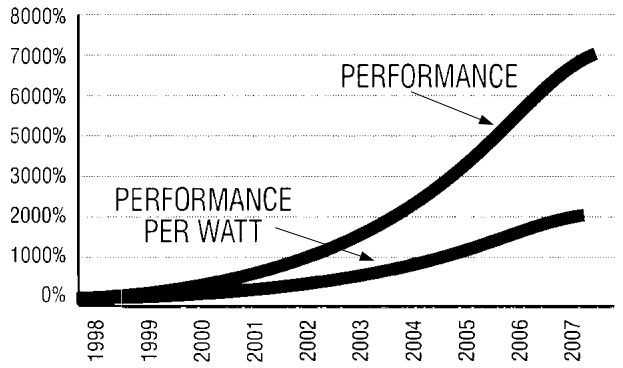

Figure 1. Increased performance vs Increased performance/watt [16] [15].

Table 2 compares the performance, the energy consumption, the energy efficiency (EE) of the supercomputers with performance above of 1 petaflops of the 36^{th} Top500 list. Also shows the order of these system in energy efficiency.

Rank	System	Rmax	Power	EE	Rank
TOP500	Name	(Pflops)	(MW)	(Mflops/W)	EE
1	Tianhe-1A	2.57	4.04	636	2
2	Jaguar	1.76	6.95	251	6
3	Nebulae	1.27	2.58	493	3
4	Tsubame 2.0	1.19	1.39	850	1
5	Hopper	1.054	2.91	362	5
6	Tera-10	1.05	4.59	229	7
7	Roadrunner	1.04	2.34	446	4

Table 2. Petaflops Supercomputers [5]

The Tianhe-1A supercomputer, with 2.57 petaflops of performance, was the current Top500 leader, but the second in energy efficiency and the eleventh in Green500 list, with 636Mflops/W.

4. Exascale Systems

The petascale systems of today have accelerated studies that were not possible some years ago. However, indications from researchers are that they would need far more powerful computing tools to meet the ever increasing challenges of an increasingly complex world [12].

According to the 36^{th} Top500 list, the fastest supercomputer is the China's Tianhe-1A with performance of 2.57 petaflops and power consumption of 4 MW. To get an exaflop[1] system we have a factor of almost 400 times and the consumption of this system would be equivalent to 1.6 GW.

The same relationship can be made with the Blue Waters system, the petascale supercomputer being designed and built in the National Center for Supercomputing Applications (NCSA) of the University of Illinois at Urbana-Champaign. This machine is going to consume 15 MW and will have performance of 10 petaflops. To achieve exascale, with a factor of 100 times, this system would consume 1.5 GW [10].

Therefore, with todays technology, exascale system would consume over a gigawatt of power, making it economically and ecologically impracticable [2] [10] [12]. In this context, the study of energy efficiency becomes important. The challenge is to build systems with computational power of exaflops, without exponentially increasing the power consumption.

5. Research to building Exascale

The energy consumption of computer systems has been the subject of research, especially in high-performance environments. Many researches have evaluated performance and power consumption of applications using heterogeneous systems. However, other studies also show the evaluation results of energy efficiency in homogeneous systems.

Peter in [10] considers the energy expended per flop. At the time, computation circuits required about 70 picojoules for each operation. An approach could be able to get the energy requirement of a flop down to about 5 to 10 pJ. Also assuming that the microprocessors would support voltages lower than 1 volt, and a peak of 0.5 V to circuitry manufactures by 2015.

The researcher Mateo Valero, from the Barcelona Supercomputing Center (BSC) is coordinating a project for the construction of the first HPC system using ARM processors. Using ARM Cortex-A9 processors, which require only 0.25 watts, the project called the Mont Blanc Zero aims to a performance of 200 petaflops and power consumption of 10 MW [14].

Other system, the Mont Blanc, aims to have efficiency of 20 gigaflops/W, higher than the current Green500 leader, the prototype IBM Blue Gene/Q system with 1.7 gigaflops/W [14].

As power consumption is proportional to the CPU frequency, one technique that is being explored is the use of Dynamic Voltage and Frequency Scaling (DVFS) within clusters and supercomputers [16]. It enables the above scaling of frequency and voltage to reduce the operational costs of powering and cooling [9].

1 10^{18} floating point operations per second

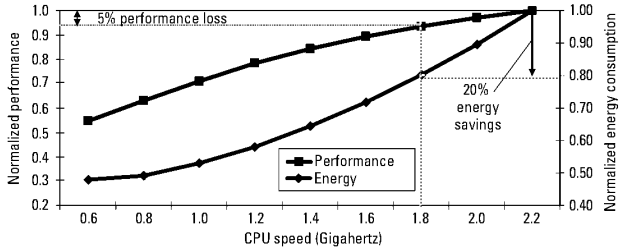

Figure 2. Performance x Frequency [8] [16].

This technique is used by authors in [16] to scale down the frequency and save energy. Test results are shown in Figure 2. In this work, the authors concluded that, by reducing frequency in 18%, the performance is reduced by only 5%.

Power-aware algorithms can provide detailed hints to hardware about its resource requirements. One important aspect of this effort is the correlation of high level application behavioral metrics with power consumption. It is necessary algorithms and resource scheduling systems that are capable of using information from the energy monitoring system to improve energy efficiency, without sacrificing computational performance [7].

6. Conclusions

The high energy costs for processors, data storage, network interconnection and cooling will be impractical in exascale systems. Therefore, it is essential to study of energy consumption of parallel architectures under different workloads and organizations, trying to find new ways to increase the energy efficiency.

If we assume the use of todays technology to build an exascale system, it would consume over a gigawatt of power. Then, to build exascale systems, it is necessary to analyse other issues such as software scalability, memory, IO, and storage bandwidth, and system resiliency to improve energy efficiency.

As future works we intend to analyze the performance and energy consumption of some applications in different configurations of hardware.

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*, 2011.

[3] A. Chandrakasan, S. Sheng, and R. Brodersen. Low-power cmos digital design. *Solid-State Circuits, IEEE Journal of*, 27(4):473–484, 1992.

[4] Y. Dong, J. Chen, and T. Tang. Power measurements and analyses of massive object storage system. In *2010 10th IEEE International Conference on Computer and Information Technology (CIT 2010)*, pages 1317–1322. IEEE, 2010.

[5] J. Dongarra. Architecture-aware algorithms and software for peta and exascale computing. In *9th Annual Workshop on Charm ++ and Its Applications 2011*. 2011.

[6] J. Dongarra, P. Beckman, T. Moore, P. Aerts, G. Aloisio, J. Andre, D. Barkai, J. Berthou, T. Boku, B. Braunschweig, et al. The international exascale software project roadmap. *International Journal of High Performance Computing Applications*, 25(1):3, 2011.

[7] J. Evans. On performance and energy management in high performance computing systems. In *2010 39th International Conference on Parallel Processing Workshops*, pages 445–452. IEEE, 2010.

[8] W. Feng, X. Feng, and R. Ce. Green supercomputing comes of age. *IT Professional*, 10(1):17–23, 2008.

[9] C. Hsu and W. Feng. A power-aware run-time system for high-performance computing. pages 1–9. IEEE Computer Society, 2005.

[10] P. Kogge. The tops in flops. *Spectrum, IEEE*, 48(2):44–50, 2011.

[11] J. Laros, K. Pedretti, S. Kelly, J. Vandyke, K. Ferreira, C. Vaughan, and M. Swan. Topics on measuring real power usage on high performance computing platforms. In *Cluster Computing and Workshops, 2009. CLUSTER'09. IEEE International Conference on*, pages 1–8. IEEE, 2009.

[12] S. S. Pawlowski. Exascale science: the next frontier in high performance computing. In *The 24th International Conference on Supercomputing (ICS),2010*, page 1, 2010.

[13] P. Schreier. How cool are supercomputer? *Scientific Computing World*, pages 22–24, 2011.

[14] M. Valero. Towards exaflop supercomputers. *Conference Center of the University of Patras - High Perfomance Computing Academic Research Network (HPC-net)*, 2011.

[15] M. Wagner. The efficiency challenge: Balancing total cost of ownership with real estate demands. Technical report, Cherokee International, Tech. Rep, 2008.

[16] A. Younge, G. von Laszewski, L. Wang, S. Lopez-Alarcon, and W. Carithers. Efficient resource management for cloud computing environments. In *International Conference on Green Computing*, pages 357–364. IEEE, 2010.