



Proceedings of the First PhD Symposium on Sustainable Ultrascale
Computing Systems (NESUS PhD 2016)
Timisoara, Romania

Jesus Carretero, Javier Garcia Blas
Dana Petcu
(Editors)

February 8-11, 2016



This work is licensed under a Creative Commons Attribution-
NonCommercial-NoDerivs 3.0 Unported License

Energy aware execution environments and algorithms on low power multi-core architectures

SANDRA CATALÁN, RAFAEL RODRÍGUEZ-SÁNCHEZ, ENRIQUE S. QUINTANA-ORTÍ

Universitat Jaume I, Spain
catalans@uji.es, rarodrig@uji.es, quintana@uji.es

Abstract

Energy consumption is a key aspect that conditions the proper functioning of nowadays data centers and high performance computing just like the launch of new services, due to its environmental negative impact and the increasing economic costs of energy.

The energy efficiency of the applications used in these data centers could be improved, especially when systems' utilization rate is low or moderate, or when targeting memory bounded applications. In this sense, energy proportionality stands for systems which power consumption is in line with the amount of work performed in each moment. As a response to these needs, the main objective of this project is to study, design, develop and analyze experimental solutions (models, programs, tools and techniques) aware of energy proportionality for scientific and engineering applications on low-power architectures. With the aim of showing the benefits of this contribution, two applications, coming from the image processing and dynamic molecular simulation fields, have been chosen.

Keywords Energy, low-power architectures, linear algebra, NESUS

I. MOTIVATION

Nowadays there is a vast variety of scientific, industrial and engineering applications that have great computing power and storage requirements, and their demand is still growing. In order to obtain more precise solutions in these applications, scientists need to build and work with sophisticated physical and mathematical models. Scientific computation (seen as the elaboration of mathematical models and the use of computers to analyze and solve scientific problems) is an efficient tool to make scientific discoveries that are complementary to the most traditional methods based on theory and experimentation [1] As a consequence, new data processing systems and high performance computing centers collapse just a few weeks later from their commissioning [2].

To face the mathematical formulation at the bottom of the physical laws advanced numerical algorithms are required: linear algebra, spectral methods (e.g, FFT), N-body methods, mesh methods to solve partial differential equations, as well as searching, classifica-

tations demanded to solve these scientific, industrial and engineering applications can be decomposed into a reduced number of well known matrix computation problems, e.g., simple operations of linear algebra, linear equation systems, minimum least square problems, eigenvalue and eigenvector problems. In this way, the efficiency of these computation problems determine as a last resort the effectiveness of the software application.

Large scale HPC (high performance computing) systems are great energy consumers, using computing resources and auxiliary systems to work [1]. This energy consumption has a direct impact on the operation costs and maintenance of the computing centers, threatening their existence and complicating the acquisition of new facilities. However, electricity cost is not the only problem; energy consumption results in carbon dioxide emissions dangerous to the environment and public health, and the heat reduces the reliability of the hardware components [3].

HPC centers' pressure forced hardware manufacturers to improve their designs to get better energy

the power supply unit) provide some energy saving strategies, based on the system transition to a low power consumption state or the dynamic adaptation of frequency and voltage (DVFS or Dynamic Voltage Frequency Scaling) [4]. On the other hand, software systems, communication libraries and, specially, computational libraries and application codes used in HPC centers have been, traditionally, unaware of power consumption. In fact, the Top500 [5] list is a good example. Computers listed in this ranking are classified depending on their sustained performance (in FLOPS) when running the Linpack test (basically, solving a dense linear system of scalable dimension). However, the numerical method behind this test, LU factorization, is far from being representative for most real scientific codes [6].

Despite the great benefits [7] that HPC energy aware solutions can provide in terms of run time optimization and energy conservation, this topic is still at an early research stage if compared with energy study in other segments. Recently, HPC community has presented energy aware metrics, e.g., Energy Delay Product (EDP), Energy To Solution (ETS), FLOPS/Watt or FTTSE [8], that are becoming more significant when evaluating algorithms and computers performance. In fact, the Green500 [9] ranking, which uses these metrics to compare and classify supercomputers all around the world regarding their energy efficiency, is becoming more considered every day.

II. RELATED WORK

Nowadays HPC linear algebra libraries make use of hardware concurrency in multi-core processors using multi-threaded implementations highly optimized for a small set of linear algebra kernels (particularly, BLAS matrix-vector product and matrix-matrix product). For years, this approach was successfully followed by the scientific community, since it provides an interface that has allowed the development of complex and architecture independent packages of numerical methods with portable performance. However, with the increasing number of cores (e.g, Intel Xeon Phi), this solution became suboptimal due to the fact that concurrency at BLAS level implies a high number of thread synchronizations, causing a high overhead.

Recently, many projects demonstrated the benefits of applying parallelism at a higher level, in both dense

of-order execution by means of an scheduling aware of the tasks' dependencies. Examples of this successful solution are libflame (SuperMatrix [10]), PLASMA (Quark) [11], SMPs [12], StarPU [13], etc., based on the ideas/techniques firstly proposed by the project Cilk [14] of MIT. These execution frameworks aim at the gross performance as final value for the user. However, they are completely energy unaware. Initial research efforts showed the possibility of keeping isoeficiency/isoscalability in a parallel solver while getting low power consumption, and the benefits derived from this approach. This can be done, for instance, scheduling non-critical tasks to less powerful and low power consumption cores (on heterogeneous environments) or through processor frequency adjustment, and promoting idle cores to low power states [15, 16, 17].

Previously mentioned solutions try to efficiently identify and make use of task parallelism in software applications. To this end, they provide the user with an explicit or implicit mechanism to identify tasks and dependencies among them. There is a part in this framework that builds a Directed Acyclic Graph (DAG) that gathers all the dependencies, and this information is used by the scheduler, which in turn issues tasks to execution when their dependencies are solved and there are enough free computational resources. Some of these frameworks also tackle the existence of multiple address spaces, providing the programmer with an explicit transfer mechanism or, alternatively, a memory control mechanism built in the scheduler that performs transparent transfers for the programmer. Scheduling algorithms at the bottom of these execution frameworks aim at optimizing performance, but generally, they do not consider energy as a variable to make decisions. However, for some operations, it is possible to improve energy efficiency during the dynamic execution of a DAG if some non-critical tasks are executed at a lower speed (via, e.g., the frequency reduction of cores applying DVFS).

On the way towards the construction of exaflop supercomputers, some research lines stand for the utilization of highly heterogeneous systems, composed of some nodes, with a huge amount of simple and low-power multi-core processors, combined with some other nodes, featuring hardware accelerators [18]. In the same vein, some recent works reveal energy advantages when using low-power processors, such as Intel Atom, ARM A-15, or more specialized systems, like

Signal Processors (DSP) of Texas Instruments [19, 20].

III. THESIS IDEA

The main objective of the research proposal is to study, design, develop and experimentally analyze solutions that are aware of the energy proportionality (models, programs, tools and techniques) of scientific and engineering applications running on low-power architectures. This objective is composed of two specific targets:

- Studying, characterizing and modeling low-power architectures' performance and energy efficiency, which include, Intel Atom, ARM Cortex-A15, Texas Instruments DSP C66x, among others.
- Designing, developing and evaluating energy proportional solutions for scientific applications in the field of hyperspectral image processing and macromolecular simulations.

So far the improvement of these kind of applications was focused on increasing their performance, through traditional parallel systems that were to a large extent energy proportionality oblivious. The novelty of this proposal is founded on the study of specific HPC techniques for low-power architectures, capable of making the best of the greater energy proportionality of these systems.

To achieve the proposed goal, the first stage of the work will consist of analyzing, modeling and optimizing basic kernels on low-power architectures. To this end, a representative number of low-power architectures will be selected in order to build experimental energy models with an appropriate collection of parameters and to determine computing and memory access costs in terms of energy. In addition, the same basic kernels will be used to characterize the energy consumption of the different components in a given architecture. After this initial study, the improvement of hyperspectral image processing problems and macromolecular simulations will be tackled. In both cases, the exploitation of parallelism at different levels (fine grain, gross grain and task parallelism) and the use of the MPI paradigm will be key to get the best of these applications on low-power architectures.

IV. CONCLUSION AND FUTURE WORK

ciety point of view, this proposal is also part of the greenhouse gas reduction challenge and the energy efficiency goal. Moreover, this project is strongly connected with the climate change action and the use of raw materials and natural resources. On the other hand, the macromolecular simulations, and to a large extent also the hyperspectral image processing, make use of and produce huge amounts of data/results. Consequently, these two kind of applications belong to the "big data" category, being also characterized as a priority topic by the economical and digital society challenges.

As future work, the improvement of dense linear algebra operations (focused on the BLIS library [21]) on low-power architectures has to be completed and the improvement of hyperspectral image processing problems and macromolecular simulations need to be performed.

Acknowledgment

This work is partially supported by EU under the COST Program Action IC1305: Network for Sustainable Ultrascale Computing (NESUS) and the FPU program of MECED.

REFERENCES

- [1] J. Dongarra, *et al*, The international ExaScale software project roadmap, *Int. J. of High Performance Computing & Applications* 25 (1) (2011) 3–60.
- [2] International technology roadmap for semiconductors, <http://www.itrs.net/> (2013).
- [3] W.-c. Feng, X. Feng, R. Ge, Green supercomputing comes of age, *IT Professional* 10 (1) (2008) 17–23. doi:10.1109/MITP.2008.8.
- [4] W. Y. Lee, Energy-saving DVFS scheduling of multiple periodic real-time tasks on multi-core processors, in: *Distributed Simulation and Real Time Applications, 2009. DS-RT '09. 13th IEEE/ACM International Symposium on*, 2009, pp. 216–223. doi:10.1109/DS-RT.2009.12.
- [5] The Top 500 list, <http://www.top500.org/> (2014).
- [6] P. Kogge, K. Bergman, S. Borkar, D. Campbell,

- R. Lucas, M. Richards, A. Scarpelli, S. Scott, A. Snively, T. Sterling, R. S. Williams, K. Yelick, ExaScale computing study: Technology challenges in achieving ExaScale systems (2008).
- [7] S. Albers, Energy-efficient algorithms, *Commun. ACM* 53 (2010) 86–96.
- [8] C. Bekas, A. Curioni, A new energy aware performance metric, *Computer Science - Research and Development* 25 (2010) 187–195.
- [9] The Green 500 list, <http://www.green500.org/> (2014).
- [10] E. Chan, F. G. Van Zee, P. Bientinesi, E. S. Quintana-Ortí, G. Quintana-Ortí, R. van de Geijn, SuperMatrix: A multithreaded runtime scheduling system for algorithms-by-blocks, in: *ACM SIGPLAN 2008 symposium on Principles and practices of parallel programming (PPoPP'08)*, 2008, to appear.
- [11] F. Song, S. Tomov, J. Dongarra, Enabling and scaling matrix computations on heterogeneous multi-core and multi-gpu systems, in: *Proceedings of the 26th ACM International Conference on Supercomputing, ICS '12*, ACM, New York, NY, USA, 2012, pp. 365–376, <http://doi.acm.org/10.1145/2304576.2304625>. doi:10.1145/2304576.2304625.
- [12] R. M. Badia, J. R. Herrero, J. Labarta, J. M. Pérez, E. S. Quintana-Ortí, G. Quintana-Ortí, Parallelizing dense and banded linear algebra libraries using SMPs, *Conc. and Comp.: Pract. and Exper.* 21 (2009) 2438–2456.
- [13] R. M. Badia, J. R. Herrero, J. Labarta, J. M. Pérez, E. S. Quintana-Ortí, G. Quintana-Ortí, Parallelizing dense and banded linear algebra libraries using smpss, *Concurrency and Computation: Practice and Experience* 21 (18) (2009) 2438–2456.
- [14] R. D. Blumofe, C. F. Joerg, B. C. Kuszmaul, C. E. Leiserson, K. H. Randall, Y. Zhou, Cilk: An efficient multithreaded runtime system, Vol. 30, *ACM*, 1995.
- [15] P. Alonso, M. F. Dolz, F. D. Igual, R. Mayo, E. S. Quintana-Ortí, DVFS-control techniques for Development 1–10 <http://dx.doi.org/10.1007/s00450-011-0188-7>.
- [16] P. Alonso, M. F. Dolz, R. Mayo, E. S. Quintana-Ortí, Saving energy in the LU factorization with partial pivoting on multi-core processors, 2012, to appear.
- [17] P. Alonso, M. Dolz, R. Mayo, E. Quintana-Ortí, Improving power efficiency of dense linear algebra algorithms on multi-core processors via slack control, *Proc. Int. Conf. High Performance Computing & Simulation–HPCS* (2011) 463–470.
- [18] Mont-blanc project, <http://www.montblanc-project.eu/> (2013).
- [19] J. I. Aliaga, H. Anzt, M. Castillo, J. C. Fernández, G. León, J. Pérez, E. S. Quintana-Ortí, Performance and energy analysis of the iterative solution of sparse linear systems on multicore and manycore architectures, *Springer*, 2014.
- [20] M. Castillo, J. Fernández, F. Igual, A. Plaza, E. Quintana-Ortí, A. Remón, Hyperspectral unmixing on multicore DSPs: Trading off performance for energy, *Selected Topics in Applied Earth Observations and Remote Sensing*, *IEEE Journal of DOI:10.1109/JSTARS.2013.2266927*.
- [21] F. G. Van Zee, R. A. van de Geijn, BLIS: A framework for generating BLAS-like libraries, *ACM Trans. Math. Soft.* To appear. <http://www.cs.utexas.edu/>.