Energy to Solution: A New Mission for Parallel Computing

Arndt Bode

Leibniz-Supercomputing Centre, Garching bei München
and
Lehrstuhl für Rechnerorganisation und Rechnertechnik
Technische Universität München
Germany
bode@lrz.de

Abstract. For a long period in the development of computers and computing efficient applications were only characterized by computational – and memory complexity or in more practical terms elapsed computing time and required main memory capacity. The history of Euro-Par and its predecessor-organizations stands for research on the development of ever more powerful computer architectures that shorten the compute time both by faster clocking and by parallel execution as well as the development of algorithms that can exhibit these parallel architectural features. The success of enhancing architectures and algorithms is best described by exponential curves regarding the peak computing power of architectures and the efficiency of algorithms. As microprocessor parts get more and more power hungry and electricity gets more and more expensive, "energy to solution" is a new optimization criterion for large applications. This calls for energy aware solutions.

Components of Energy Aware Computing

In order to reduce the power used to run an application, four components have to be optimized, three of them relate to the computer system and the programs to be extended, one relates to the infrastructure of the computer system:

- **Energy aware infrastructure:** This parameter relates to the fact, that computers need climate, cooling, uninterruptable voltage supply, building with light, heating and additional infrastructure components that consume power. Examples for measures to reduce energy are: Use of liquid cooling, direct cooling, free cooling, waste heat reuse, adsorption machines, monitoring and optimizing of energy consumption and infrastructure control, coupling of infrastructure power requirements with behavior of computers and the application execution.
- **Energy aware system hardware:** This parameter describes all mechanisms in new hardware to reduce power in the system itself: sleep modes of inactive parts, clock control of the parts, fine grain hardware-monitoring of

F. Wolf, B. Mohr, and D. an Mey (Eds.): Euro-Par 2013, LNCS 8097, pp. 1–2, 2013.

system consumption, any hardware that relates to accumulating, processing and using consumption data for power reduction in the part and its relation to other parts of the system, autonomous optimizing and learning behavior.
- **Energy aware system software:** This parameter describes all sorts of automatic, semi-automatic, or online and offline user controlled tools, that do monitor, analyze and control the execution of application software on the system: support for optimal pinings of threads, finding the optimal clock based on previous runs and data on execution behavior looking at relations between processing, memory, interconnection and storage/I/0 boundedness.
- **Energy aware algorithms:** In most cases, the fastest algorithm consumes the minimum power. But there are exceptions, if algorithms use redundancy (sometimes in applications with super linear speedup). This is a very broad research area in its own.

Experiences with SuperMUC

Energy-aware HPC solutions are tried out at Leibniz Supercomputing Center with SuperMUC a system with 160.000 cores of XEON (IBM iDataPlex). Energy consumption is measured and controlled at many different levels starting at the course grain level with the overall infrastructure control down to very fine grain tools on the individual chip level. Some of the tools are under user control, other were controlled by the datacenter management team, some are even fully automatic. Clocking control is offered by LRZ to the general user of SuperMUC as a tool that supports measuring in detail the execution of a program at a first run, puts data into a database and uses this for subsequent runs to optimize the clocking of SuperMUC.

These measures implemented so far present a first step and will be enhanced in the future. Further development and research is needed to couple and optimize the effects of the various tools. We also advocate for better energy awareness indicators (such as PUE) that do measure the total consumption including the entire infrastructure and take into account that the percentage of peak performance used in highly parallel applications is rather poor and varies with the system architecture, in order to allow for fair evaluations and comparison.

If we want to afford the electricity for an Exascale system, energy awareness of computing has to be approved by orders of magnitude. The methods and tools that have to be developed to do so, present an interesting new field of research for the Euro-Par community.