Taking advantage of Data Analytics to improve energy consumption

D. Trystram
Univ. Grenoble Alpes, Grenoble INP
Georges Da Costa, Jean-Marc Pierson, Patricia Stolf
Univ. Toulouse

October 3, 2020

Energy issues in large-scale distributed platforms

Energy has always been a concern in large scale distributed platforms. Nowadays, it becomes more critical due to the transition to the next generation extreme scale High Performance Computing platforms (HPC).

In this spirit, the ENERGUMEN project, funded by the ANR in 2018, aims at studying efficient and practical tools for managing the allocation of jobs to the various components of a large/extreme scale HPC platforms.

There exist various mechanisms for reducing the energy in such large scale HPC platforms. First, it is possible to decrease the clock frequency of the computing units (known as *Dynamic Voltage and Frequency Scaling* – DVFS). Another way is to switch-off some nodes (*shutdown*) or to put them into a sleep mode for some time. Sometimes, both mechanisms can even be used in a coordinated manner. Alternatively, saving energy can be obtained as a consequence of reducing data movements by adequate communication-aware allocations of the jobs (resulting to internal communications with high locality, close to I/O nodes, etc.). Today, some approaches already use this leverage, but they are limited since they assume that they do not have any influence on the applications themselves (in other words, there are interferences due to communications inside nodes and between nodes).

In the ENERGUMEN project, we proposed two new complementary mechanisms for addressing the energy/performance trade-off in HPC platforms. First, by revisiting the classical speed-scaling and power down mechanisms by using a *malleable model*, which allows to shrink or stretch dynamically the execution time of the jobs according to the current energy profile and by designing new optimized policies for energy-aware data allocations at the software level. These mechanisms aim at introducing more flexibility into the management of the heterogeneous resources of extreme scale HPC platforms.

Most existing performance studies in HPC platforms target the reduction of the execution times, very few are considering the energy dimension (sometimes, the energy resource is considered as a second hand objective in multi-criteria approaches).

We propose in this subject to focus on energy as a major factor of performance in link with ENERGUMEN.

Analysis of actual execution traces

Today, the power consumption becomes a major constraint, in particular by the impact of the digital activities on the climate changes and by the always increasing demand of computing power for managing huge data. Being able to increase the *required* computing performances within the same order of magnitude in energy as today is a great challenge (or with less power consumption). This clearly shows that reaching the two previous targets needs a revolution in the way of handling resource management problems.

Recently, we also developed an analysis of execution logs on large scale platforms in Grenoble and Toulouse (several years of data recording). The objective was to put a focus on the induced energy consumption. A software

has been designed and implemented for recording fine-grained data linked with power consumption of the computing units. This analysis leads to a classification of the energetic profile of the jobs, but it did not study the links with the job allocation strategies.

Links between energetic traces and allocation strategies

The subject of this internship is to study how to use the existing classification of the traces to improve the allocation of the jobs (both sides of pure performances and energy should be considered).

First, we would like to strengthen the existing analysis, for instance by selecting some specific jobs submitted by the same users. Then, the work will focus on coupling some jobs with complementary energetic profiles.

Skills:

- Knowledge of times series data analisys
- Knowledge and usage of HPC plateforms
- Knowledge of optimization techniques

Contacts

Denis Trystram, professor at University Grenoble Alpes
Georges da Costa, Jean-Marc Pierson and Patricia Stolf, Toulouse