

Phases detection in high performance parallel applications

Keywords: High Performance Computing, Control theory, Traces and Data Analysis.

1 Description

Parallel applications are made of various computational tasks, as well as communications between computing resources: during a first set of computations, computing units produce intermediate results that are communicated to other computing units, thereby enabling new computations that will pursue the same scheme. This holds for domain-based simulations in which the computation is divided in time steps separated by global communications, but also for data-flow computations in which elementary computation tasks communicate directly with each other to transmit their intermediate results.

In this general scheme, computational units are characterized by their state : they might be busy, executing a computing task, sending their intermediate results, or idle, waiting for incoming data. From a global point of view, the evolution of the states of computing units form phases in the application, computing phases of various kinds, separated by communication phases. These phases might be synchronous, if all the computing units change their state at the same time, or spreaded within a large time interval.

Whatever their shape, knowing about these phases can be a huge benefit:

- the succession of phases is an indication of the application progress and can serve as a monitoring metric
- a phase change is an inflection point in which the application behavior will change, and, thus, this is the perfect time to adapt the resources management policies:
 - to save energy, by lowering CPU frequency when threads are idle

- to improve resources utilization, by releasing or requesting computing units depending on the workload
- to improve efficiency, by redistributing data or migrating computations

Therefore, phases detection has been widely studied in the literature. Various techniques based on data topology, spectral analysis or information theory have been designed to detect phases in application traces. This will be the focus of this internship.

2 Expectations

Despite the multiple phase detection techniques that can be found in the literature, no thorough comparison between them has been conducted. A phase detection technique is usually justified by the resulting gain at the resources management level but not by the precision and the efficiency of the detection. Needless to say that the various models found in the existing works often differ completely, even regarding their starting point: the definition of a phase. Furthermore, they mostly focus on synchronous phases detected in *post mortem* application traces. Although convenient from the detection point of view, this does not match the actual needs regarding high performance computing.

The purpose of this internship is twofold:

- compare existing approaches and propose an unifying model for phases detection in application traces
- extend existing works to detect phases online (during the collection of the application trace) and to detect asynchronous phases. These objectives are quite challenging and the candidate, after a bibliographical study, will have to precise the focus of his work.

3 Contact

- Guillaume Huard <Guillaume.Huard@imag.fr>
- Swann Perarnau <swann@anl.gov>