Using Simulation to Evaluate and Tune the Performance of Dynamic Load Balancing of an Over-decomposed Geophysics Application

Rafael Keller Tesser
Informatics Institute UFRGS
rktesser@inf.ufrgs.br

Lucas Mello Schnorr
Informatics Institute UFRGS
schnorr@inf.ufrgs.br

Arnaud Legrand
Univ. Grenoble Alpes
CNRS, Inria, Grenoble INP
LIG, F-38000, Grenoble, France
arnaud.legrand@imag.fr

Fabrice Dupros
BRGM, Orlans, France
f.dupros@brgm.fr

Philippe O. A. Navaux
Informatics Institute UFRGS

Abstract

Finite difference methods are commonplace in scientific computing. Despite their apparent regularity, they often exhibit load imbalance that damages their efficiency. We characterize the spatial and temporal load imbalance of Ondes3D, a seismic wave propagation simulator. We reveal that this imbalance originates from the nature of the input data and from low-level CPU optimizations. Such dynamic imbalance should therefore be quite common and is intractable by any static approach or classical code reorganization. An effective solution, with few code modifications, combines domain over-decomposition and dynamic load balancing (e.g., with AMPI), migrating data and computation at the granularity of an MPI rank. It generally requires a careful tuning of the over-decomposition level, the load balancing heuristic and frequency. These choices are quite dependent on application and platform characteristics. In this paper, we propose a methodology that leverages the capabilities of the SimGrid framework to conduct such study at low experimental cost. It combines emulation, simulation, and application modeling that requires minimal code modification and yet manages to capture both spatial and temporal load imbalance, faithfully predicting its overall performance. We compare simulation and real executions results and show how our strategy can be used to determine the best load balancing configuration for a given application/hardware configuration.