Optimizing Communication and Synchronization in CAF Applications

Alessandro FANFARILLO [a] Davide DEL VENTO [a] Patrick NICHOLS [a]
[a] *National Center for Atmospheric Research*

Abstract. Since the beginning of distributed computing, overlapping communication with computation has always been an attractive technique used to mask high communication costs. Although easy to detect by a human being, communication/computation overlapping requires knowledge about architectural and network details in order to be performed effectively.

When low level details influence performance and productivity, compilers and run-time libraries play the critical role of translating the high level statements understandable by humans into efficient commands suitable for machines.

With the advent of PGAS languages, parallelism becomes part of the programming language and communication can be expressed with simple variable assignments. As for serial programs, PGAS compilers should be able to optimize all aspects of the language. That would include communication, but unfortunately this is not yet the case.

In this work we consider parallel scientific programs written in Coarray Fortran and we focus on how to build a PGAS compiler capable to optimize the communication, in particular by automatically exploiting opportunities for communication/computation overlapping. We also sketch an extension for the Fortran language that allows one to express the relation between data and synchronization events; we finally show how this relation can be used by the compiler to perform additional communication optimizations.

Keywords. MPI, Coarray, Fortran, PGAS, Compilers

Introduction

Communication and synchronization are two of the most important characteristics in parallel computing. The future computing architectures, equipped with billions of heterogeneous devices and different memory levels, will make these factors even more critical than they are today. In fact, communication (in general data movement) will dominate performance and energy consumption because of technological and physical limitations, whereas the high level of parallelism will suffer due to inefficiencies of coarse grain synchronization methods (e.g. full barriers).

In scientific computing, in particular in atmospheric science, the most used approach for parallel programming is usually based on the data-parallel model.