

High-Level Data Mapping for Clusters of SMPs

Siegfried Benkner
Institute for Software Science
University of Vienna, Austria

Thomas Brandes
SCAI - Institute for Algorithms and Scientific Computing
GMD - German National Research Center for Information Technology
St. Augustin, Germany

e-mail: sigi@ieee.org

Clusters of shared-memory multiprocessors (SMPs) have become the most promising parallel computing platforms for scientific computing. However, SMP clusters significantly increase the complexity of user application development when using the low-level application programming interfaces MPI and OpenMP, forcing users to deal with both distributed-memory and shared-memory parallelization details. In this paper we present extensions of High Performance Fortran for SMP clusters which enable the compiler to adopt a hybrid parallelization strategy, efficiently combining distributed-memory with shared-memory parallelism. By means of a small set of new language features, the hierarchical structure of SMP clusters may be specified. This information is utilized by the compiler to derive inter-node data mappings for controlling distributed-memory parallelization across the nodes of a cluster, and intra-node data mappings for extracting shared-memory parallelism within nodes. Additional mechanisms are proposed for specifying inter- and intra-node data mappings explicitly, for controlling specific SM parallelization issues, and for integrating OpenMP routines in HPF applications. The proposed features are being realized within the ADAPTOR and VFC compiler. The parallelization strategy for clusters of SMPs adopted by these compilers is discussed as well as a hybrid-parallel execution model based on a combination of MPI and OpenMP. Early experimental results indicate the effectiveness of the proposed features.