Language and Run–time Support for Network Parallel Computing

Peter A. Dinda David R. O'Hallaron Jaspal Subhlok* Jon A. Webb Bwolen Yang

School of Computer Science
Carnegie Mellon University
Pittsburgh PA 15213

Abstract

This paper describes NetFx, the extended Fx compiler system for network parallel computing. NetFx uses the Fx model of task parallelism to distribute and manage computations across the sequential and parallel machines of a network. The central problem in compilation for network parallel computing is that the compiler is presented with a heterogeneous and dynamic target. Our approach is based on a novel run–time system that presents a simple communication interface to the compiler, but uses compiler knowledge to customize communication between tasks on heterogeneous machines over complex networks. The compiler generates code for data parallel modules with point–to–point transfer of distributed data sets between modules. This allows the compiler to be portable and enables communication generation without knowledge of exactly how the modules will be mapped at run–time, and what low level communication primitives will be used. The compiler also generates a set of custom routines, called distribution functions, to translate between different data distributions. The run–time system performs the actual communication using a mix of generic and custom distribution functions depending on run–time parameters like the type and number of nodes assigned to the communicating modules and the data distributions of the variables being communicated. This mechanism enables the run–time system to exploit compile–time optimizations, and enables the compiler to manage foreign modules that use non–standard data distributions. We describe the NetFx programming model, the communication interface, how NetFx programs are compiled to the interface, and the issues in implementing the interface.

1 Introduction

New networking technologies like HiPPI and ATM have made it possible to implement efficient network parallel applications. Because of considerable improvements in the latency and bandwidth characteristics of networks, the domain of applications that can benefit from network parallel computing has become much larger, including diverse fields like scientific simulations, real-time applications and real-world modeling. However, standard parallel languages like High Performance Fortran (HPF) [6] offer little support for network parallel computing. Without such support, the construction of an efficient, scalable, and environment–independent network parallel application is a difficult and error–prone task in which the applications programmer must wear many hats, including that of the network specialist.

The goal of our research is to make it possible to program network parallel applications in an easy and efficient way. Compiling for a variety of machines across a network presents several new challenges. This paper shows how parallelism inside a single machine and across a set of machines can be expressed in a uniform framework. We also present the design of a compiler and a run–time system that enables efficient generation of programs for network parallel computing.

The programming model we use combines task and data parallelism. Procedures in an application are data parallel and each procedure is mapped to a homogeneous set of nodes, which may be different nodes of a large parallel machine or a cluster of workstations. Different procedures are placed on different sets of nodes on the network and the compiler manages the communication between them.

The most challenging part of compilation for a network is generating efficient communication. The most efficient communication interface across a network is typically different from the communication interface between the nodes of a parallel machine, hence the compiler has to manage multiple communication paradigms. More important,