Using MPI to Implement Scalable Libraries

Ewing Lusk

Mathematics and Computer Science Division
Argonne National Laboratory, USA

MPI is an instantiation of a general-purpose programming model, and high-performance implementations of the MPI standard have provided scalability for a wide range of applications. Ease of use was not an explicit goal of the MPI design process, which emphasized completeness, portability, and performance. Thus it is not surprising that MPI is occasionally criticized for being inconvenient to use and thus a drag on software developer productivity. One approach to the productivity issue is to use MPI to implement simpler programming models. Such models may limit the range of parallel algorithms that can be expressed, yet provide sufficient generality to benefit a significant number of applications, even from different domains. We illustrate this concept with the ADLB (Asynchronous, Dynamic Load-Balancing) library, which can be used to express manager/worker algorithms in such a way that their execution is scalable, even on the largest machines. ADLB makes sophisticated use of MPI functionality while providing an extremely simple API for the application programmer. We will describe it in the context of solving Sudoku puzzles and a nuclear physics Monte Carlo application currently running on tens of thousands of processors.