Workpackage 7: Benchmarks

Objectives

We will show the applicability of the skeleton library to a range of practical application problems. We will consider data-parallel problems as well as problems that require task parallelism. The goal is to demonstrate that these problems can be handled more easily and more quickly with our skeletons than with hand-crafted MPI. This should be particularly easy for problems requiring task parallelism and the interaction of task and data parallelism because the system will automatically take care of the technical details of mapping MPMD-style computations to SPMD-style MPI.

Moreover, the efficiency of the code produced using skeletons should be similar to that of hand-crafted MPI code. The latter will be demonstrated by experimental results across a range of parallel architectures.

Description of Work

First of all, a representative set of application problems must be selected. As far as possible, they will be taken from benchmark suites that are widely accepted in the high-performance community, like NAS, SPLASH, and PARKBENCH. The main novelty of our benchmarks will be that they are not restricted to raw performance but also take into account the ease of programming. Once the application problems have been selected, the problems and corresponding algorithms must be described in a report. Then these problems have to be solved using skeletons as well as using MPI directly. For most of these problems, MPI-based implementations are already available. Both approaches will be compared with others wrt. the time needed for the implementation and run-time across different hardware platforms.

Deliverables

We will develop:

- a document describing the selected application problems, the algorithms used to tackle them, the implementation using skeletons, and corresponding experimental results,
- the implementations of the applications using skeletons.

We will make both freely available via the Web.

Milestones and Expected Results

The above-mentioned document will be produced in several steps. The sections on application problems and algorithms are not dependent on other milestones and can be written within the first 6 months. The section on the implementation can only be written once the implementation is finished, and this in turn requires that the prototype implementation of the skeleton library (see Workpackage 4) is available. We therefore expect to be able to finish the implementation of the application problems and the corresponding description within 18 months. The experimental results and their description in the report will be ready within approximately 21 months.