Workpackage 1: Pseudo Application Set

Objectives

The objective of this workpackage is to construct a collection of informal descriptions of parallel application programs which embody our intended range of conceptual structures and their varied instantiation. This is necessary to ensure that the standard has excellent coverage in terms of the end-user applications.

This is a crucial and difficult task for many reasons. Firstly, parallel applications programming is an emerging and rapidly evolving discipline, both in terms of language base and conceptual structures. Secondly, the existing usage of structured approaches is quite diverse and often informal. Apparently simple concepts such as "pipeline" and "divide & conquer" have already been used in may ways. For example, pipelines may or may not have internal state and may or may not choose to vary the relationship between the cardinalities of input and output sets. As soon as one starts to formalise such concepts, there is a danger of making specifications which exclude certain well established and useful applications. While this may sometimes be inevitable, such an outcome should be a result of conscious engineering rather than an unforseen side-effect.

Description of Work

The work has two phases: survey and distillation. The survey phase involves a substantial literature review. This will be based on:

- the detailed personal knowledge of the participants, who are all experienced in the area;
- the resources (bibliographies, workshop proceedings) which have emerged from previous informal collaborations;
- the wider parallel applications literature (conference proceedings, annual reports of activities in supercomputing centres).

The survey activity will be built around a pre-conceived set structural classes, consisting of those which are already common currency in the area (pipeline, farm, divide & conquer and so on). For each candidate application, investigators will report briefly on its structural aspects in terms of these classes, highlighting any novel factors. Informality will be encouraged and applications may be classified in several ways as appropriate. The set of application descriptions will be stored on the project web-site.

In the second phase we will distill from this survey the key concepts and variations which our standard must encompass. To do this, we will select a subset of the applications which collectively span the space investigated. These will collected into a report in which each section will describe one application with pseudo-code and a commentary highlighting use of the conceptual structures and nuances of their instantiation.

Deliverables

A report, as outlined above, describing our pseudo applications in terms of their abstractable parallel structure.

Milestones and Expected Results

The milestones are simply the completion of the survey, midway through the workpackage, and the production of the report. The result is that we will have refined our understanding of the application universe and demonstrated that the standard will be representative of the needs of users.