Workshop 23
ESPRIT Projects

Ron Perrott and Colin Upstill

Co-chairmen

The ESPRIT Programme has now been in existence for over a decade. One of the areas to benefit substantially from the injection of ESPRIT research and development funding has been the area of parallel computing or high performance computing. It is therefore appropriate that at this conference there should be two sessions dealing with HPC projects. A wide range of papers were submitted, and those selected represent state-of-the-art projects in various areas of interest. The projects break down into applications which can benefit from the use of HPC and those which deal with the software required for making high performance computers more efficient.

In the former category are papers like Parallel Crew Scheduling in Pharos. This is an application which considers the problem of scheduling airline crews and the use of a network of workstations. It is based on the widely known Carmen system which is used by most European airlines to schedule their crews among aeroplanes. The objective is to produce a parallel version of the Carmen system and execute it on a network of workstations and demonstrate the improvements and benefits which can be obtained. Essentially there are two critical parts in the application dealing with pairing generation and optimisation. The pairing generator distributes the numeration of pairings over the processors and the optimiser is based on an iterative Lagrangian heuristic. Initial results indicate the benefits the speed-up which can be obtained with this application on a network of workstations.

The second paper deals with CORBA (Common Object Request Broker Architecture) and develops a compliant programming environment for HPC. This is part of an ESPRIT Project known as PACHA. The main objective is to help the design of HPC applications using independent software components through the use of distributed objects. It proposes to harness the benefits of distributive and parallel programming using a combination of two standards, namely CORBA and MPI. CORBA provides transparent remote method invocations which are handled by an object request broker that provides a communication infrastructure independent of the underlying network. This project introduces a new kind of object which is referred to as a CORBA parallel object. It then allows the aggregation of computing resources to speed up the execution of a software component. The interface is described using an extended IDL to manage data distribution among the objects of the collection. CORBA is being implemented using Orbix from Iona Technologies and has already been tested using a signal processing application based on a client server approach.

The OCEANS project deals with optimising compilers for embedded applications. The objective of OCEANS is to design and implement an optimising

compiler that utilises aggressive analysis techniques and integrates source level restructuring transformations with low level machine-dependent optimisations. This will then provide a prototype framework for iterative compilation where feedback from the low level is used to guide the selection of a suitable sequence of source level transformations and vice versa. The OCEANS compiler is centred around two major components, a high level restructuring system known as MT1 and a low level system for supporting language transformations and optimisations known as Salto. In order to validate the compiler, four public domain multimedia codes have been selected. Initial results indicate that the results using the initial prototype are satisfactory in comparison with the production compiler but their implementation still needs to be refined at both the high and low levels.

The fourth paper deals with industrial stochastic simulations on large-scale meta-computers, and reports the acheivements of the PROMENVIR (Probabilistic Mechanical Design Environment) project. This project culminated in a pan-European meta-computer demonstration of the use of PROMENVIR running a multibody simulation application. PROMENVIR is an advanced computing tool for performing stochastic analysis of generic physical systems. The tool provides the user with a framework for running a stochastic Monte Carlo simulation using a preferred deterministic solver and then provides sophisticated means to analyse the results. The European meta-computer consisted of a total of 102 processors geographically distributed among the members of the consortium. The paper details the problems that were encountered in establishing the meta-computer and carrying out the execution of the chosen application. It makes useful contributions on the feasibility and reliability of such an approach.

HIPEC stands for High Performance Computing Visualisation System Supporting Network Electronic Commerce applications. This is an ESPRIT Project whose main objective is to integrate advanced high performance technologies to form a generic electronic commerce application. The application is aimed at giving a large number of SMEs a configuration and visualisation tool to support the selling process within the show room and to enlarge their business using the same tool over the Internet. Computer graphics technology plays a primary role in the project as realistic images are required in such commerce. The particular experiment was based on the generic technology and applied to bathroom furniture application. The project was shown to be of benefit in this particular application area.

The final paper is the porting of the SEMC3D electromagnetics code to HPF. This was a project within the ESPRIT PHAROS Project, in which four industrial simulation codes are ported to HPF. The electromagnetic simulation code was ported onto two machines, namely the Meiko CS2 and the IBM SP2. The paper describes the application and the numerical computational features which are important and then analyses the porting of the code to a parallel machine. A series of phases such as code cleaning, translation to Fortran 90, and inserting HPF directives are described. This is followed by a series of measurements. The first stage in the port was the conversion of the Fortran 77 code to Fortran

90, for which converted code single processor times on the IBM SP2 where obtained. This was then followed by conversion to HPF code and performance times over a range of processors documented. Finally, a comparison with message passing code was carried out. The speed-ups obtained in the benchmarking of this particular code are described as encouraging by the users, and give evidence of the benefits of using HPF.