HPC Standards

Esprit project 21111
Final Management Report

May1998

Contents

Executive Summary

The purpose of this project was to provide a mechanism to ensure coherent European input to the development of international standards in the area of HPCN. In order for European industry to be well placed to take advantage of emerging standards it is vital that Europe continues to be involved in the standardisation activities. EU funding for these standardisation activities is critical if we wish to have a vote on the content of a standard. Indicative of this importance is the fact that HPF-2 voting rights were in fact lost for GMD due to delays in funding HPC-Standards. Regular attendance at meetings is important not only to maintain voting rights but also to keep abreast of developments.

Due to the state of development of HPF when the project started few standardisation meetings were attended by European representatives. Nevertheless, attendance at the last meeting and especially at the User Group meeting, was very valuable in terms of obtaining up-to-date knowledge of the status of the standard. The Vienna workshop and the strong presence at the User Group meeting demonstrated European strength in HPF activities. Their experience with the technology has enabled Europeans to have significant impact through the interaction between end-users and compiler vendors. One benefit of the perceived European strength in this area is that the next User-Group meeting will take place in Europe. NA Software, one of Europe's few HPF compiler developers derived direct benefit from the meetings, as did academic researchers. The project brought together European workers (some of whom were previously not in contact) and raised the profile of this field internationally.

The main European input to MPI-2 has been at the level related to applications, and indeed the MPI-Forum has taken more application-oriented decisions as a result of attendance. In particular the following parts of the standard were influenced as a result of European voting; Fortran 90 bindings, dynamic process creation, external interfaces, parallel I/O and interoperability. The direct consequence of this input has been to support some European applications (for example, VAMPIR from PALLAS benefited from the choice of external interface) and to represent the interests of other Esprit Projects with regard to language interoperability. The indirect benefits of attendance, knowing the precise state of the standard and the directions it was moving, were probably even more important. For example, PALLAS has been able to provide MPI-2 related services to customers at a very early stage, and have produced one of the first MPI-2 implementations for an industrial customer. This understanding of the developing standard has also been beneficial to GMD dissuading them from relying excessively on the dynamic process management functionality contained in MPI-2.

HPC-Standards funding has allowed significant European-US discussions in the area of the PARKBENCH low-level benchmarks. European contribution was made in the following areas; the new "COMMS" benchmarks, a Java applet tool (PICT) for postprocessing and parameterisation of PARKBENCH results, run and reporting rules, Java benchmarks and parallel I/O benchmarks. PARKBENCH awareness and dissemination of results has been increased through the Performance Evaluation and Modelling for Computer Systems (PEMCS) electronic journal which was initiated in early 1997 at the suggestion of the HPCnet Network of Excellence. HPCnet provided some start-up funding for this activity which promotes responsible performance reporting. Therefore, although not funded directly by HPC-Standards or the PARKBENCH group, the journal has active involvement of EU and US experts and this can be seen to be a result of EU support for PARKBENCH in HPC-Standards.

The dissemination of information from standardisation activities consisted most importantly of "European Information Workshops". The project was responsible for holding three such workshops, on HPF-2, MPI-2 and PARKBENCH, and also funded a fourth workshop, on parallel tools, arranged in conjunction with Smith Systems Engineering and the Parallel Applications Centre (PAC) at the University of Southampton.

1 Introduction

The objective of the HPC Standards Project has been to ensure European involvement in the development of High Performance Computing standards and to disseminate information on these standards. The three major standards considered by the project were High Performance Fortran (HPF-2), the Message Passing Interface (MPI-2) and the Parallel Kernels and Benchmarks (PARKBENCH) standard. Support was also provided for a Parallel Tools Workshop. Funding was directed towards enabling European representatives to regularly attend meetings of the standards bodies held in the USA in order to obtain and preserve voting rights. The dissemination of information was managed from the University of Southampton, and consisted most importantly of ‘‘European Information Workshops”. The project was responsible for holding three such workshops, on HPF-2, MPI-2 and PARKBENCH, and also funded a fourth workshop, on parallel tools, arranged in conjunction with Smith Systems Engineering and the Parallel Applications Centre (PAC) at the University of Southampton. Besides the workshops, general dissemination to the community by normal academic means such as a project Web page, mailing lists, and conference talks took place.

This final report summarises the activities in each of these areas over the duration of the project, which lasted from 1st June 1996 until 1st December 1997.

The rest of the report is structured as follows. Section 2 describes the project Web page (which contains more detailed information than is possible in this report). Sections 3 to 5 discuss each standards area separately, emphasising the impact of the project and the benefits construing to Europe. Section 6 contains details of the “European Information Workshops”. Appendix 1 contains copies of feedback on the overall project solicited directly from the representatives. Appendix 2 lists tables of the standards meetings attended along with the names and affiliations of the European representatives. Appendix 3 contains some of the reports provided by the representatives on the meetings and workshops they attended. Finally, appendices 4 and 5 contain additional details on the new PARKBENCH COMMS1 benchmark and PICT tool mentioned in the main body of the text.

2 Project Web page

An "HPC Standards" project Web page has been set up at URL:

`http://www.ccg.ecs.soton.ac.uk/Projects/hpc-stds`

During the course of the project, this Web page was extended to provide up-to-date information on the status of the standards and on project events as they occurred. It contains full information about the project and links to pages about the HPF-2, MPI-2 and PARKBENCH standardisation activities, and also about the project-sponsored ‘European information workshops’. The information on these pages is more thorough than the details contained in this report.

The HPF-2, MPI-2 and PARKBENCH pages contain meeting reports and sets of minutes from project participants attending standardisation meetings, as well as the ‘official’ meeting minutes, information on the official email discussion lists and how to join them, and links to the latest draft standard and to other sites with relevant information.

The project Web page also points to information about the “European information workshops" including their final programs and slides of many of the presentations. The HPCN software tools survey undertaken by Smith System Engineering and the Parallel Applications Centre at the University of Southampton is also available from the Web page.

3 HPF-2

3.1 Background

The High Performance Fortran Forum (HPFF) met from March 1992 to March 1993 to define a set of extensions to Fortran called High Performance Fortran (HPF). The aim was to address the problems of writing portable data parallel programs for architectures where the distribution of data has a significant impact on performance, and to achieve high efficiency across different parallel machines with the same high level HPF program. Some of the fundamental factors affecting the performance of a parallel program are the degree of available parallelism, exploitation of data locality, and choice of appropriate task granularity. HPF provides mechanisms for the programmer to guide the compiler with respect to these factors through HPF directives, which appear as structured comments suggesting implementation strategies or which assert facts about a program to the compiler.

HPF-2 was an effort to develop extensions to, and provide corrections, clarifications and interpretations of the HPF language definition in the light of user and compiler writer experience. Specifically, HPF-2 includes features that broaden the applicability of HPF, and whose implementation can exploit successful research experience. HPF-2 began in January 1995 with broad participation from vendors, users and software developers. The standard was finalised in January 1997.

3.2 Project

Because the standard was finalised soon after the HPC standards project started, European representatives were only able to attend one standardisation meeting in September 1996. For this reason, six attendees drawn from the European consortium were funded to attend the first annual HPF User Group meeting in February 1997.

The workshop, “Summer of HPF in Vienna” in July 1997, was very successful, attracting 70 participants from across Europe, Canada and the USA. The event enabled many compiler developers, both commercial and research, to come together with potential and actual users of High Performance Fortran in order to discuss their problems and needs. Compiler writers need guidance from users in order to understand how best to improve their products; application developers need to find out how to write their codes in ways that help the compiler generate fast object code. Additional tutorial sessions introduced new users to HPF with “hands-on” sessions. Details of the workshop may be found in the Section 6 and on the Web page.

3.3 European Input and Benefits

Due to the state of development of HPF when the project started few standardisation meetings were attended by European representatives. Voting rights were in fact lost for GMD and this was due to delays in funding HPC-Standards. This is indicative of the importance of EU funding for these standardisation activities if we wish to have a vote on the content of a standard. Nevertheless, attendance at the last meeting and especially at the User Group meeting, was very valuable in terms of obtaining up-to-date knowledge of the status of the standard. The Vienna workshop and the strong presence at the User Group meeting demonstrated European strength in HPF activities. Their experience with the technology has enabled Europeans to have significant impact through the interaction between end-users and compiler vendors.

One benefit of the perceived European strength in this area is that the next User-Group meeting will take place in Europe. NA Software, one of Europe's few HPF compiler developers derived direct benefit from the meetings, as did academic researchers.

The project brought together European workers (some of whom were previously not in contact) and raised the profile of this field internationally.

3.4 New Features of HPF-2

The HPF 2.0 standard differs from the HPF 1.1 standard in a number of ways:

- Repartitioning of the Language:

  The new document describes two components: the HPF 2.0 language (which is expected to be widely and relatively rapidly implemented) and the set of Approved Extensions (which are not part of HPF 2.0 but may be included in future implementations in response to user demand, as the compilation technology matures.)

- Features Now in Standard Fortran:

  Fortran, instead of Fortran 90 is now defined as the base language for extensions; this implies that HPF includes all features added to Fortran at the 1995 revision. With this revision, a few HPF 1.1 features are now part of the Fortran standard (the FORALL statement and construct; the PURE attribute for procedures; and extensions to the MINLOC and MAXLOC intrinsics to include an optional DIM argument), and hence no longer appear as HPF extensions to Fortran.

- Features Removed or Restricted in HPF 2.0:

  - Sequential arrays may no longer be explicitly mapped.
  - In any procedure call in which distributed data may require redistribution, the procedure must now have an explicit interface.
  - The treatment of the INHERIT directive has been simplified in that it is no longer possible to specify both INHERIT and DISTRIBUTE together.
  - The treatment of pointers has been simplified.

- Elimination of the HPF Subset:

  Unlike HPF 1.1, HPF 2.0 no longer has a recommended minimal subset for faster implementation (i.e. Subset HPF), although the original HPF 1.1 Subset is documented in an annex.

- Features Moved to Approved Extensions:

  The DYNAMIC attribute and the REDISTRIBUTE and REALIGN statements have been moved to the Approved Extensions.

- New Features of HPF 2.0:

  The following new constructs have been introduced in HPF 2.0:

  - The REDUCTION clause for INDEPENDENT loops;
  - The new HPF_LIBRARY procedures SORT_DOWN, SORT_UP.

- New Approved Extensions:

  The Approved Extensions include the following features not part of HPF 1.1:

  - Mapping of objects to processor subsets;
  - Explicit mapping of pointers and components of derived types;
  - New distribution formats: GEN_BLOCK and INDIRECT;
  - New directives: RANGE, SHADOW, ON, RESIDENT, TASK_REGION;
  - Additional intrinsic procedures: ACTIVE_NUM_PROCS,ACTIVE_PROCS_SHAPE, and a generalized TRANSPOSE intrinsic;
  - New HPF_LIBRARY procedures: HPF_MAP_ARRAY and HPF_NUMBER_MAPPED; revision of procedures HPF_ALIGNMENT, HPF_DISTRIBUTION and HPF_TEMPLATE;
  - Support for asynchronous I/O with a new statement WAIT, and an additional I/O control parameter in the Fortran READ/WRITE statement;
  - Extensions to the EXTRINSIC facilities to support interoperability with C and FORTRAN 77.

- Recognized Externally-Supported HPF Extrinsics:

  Finally, the document acknowledges a new category, HPF-related EXTRINSIC interfaces, that are recognized as meeting appropriate standards for such interfaces, but are not included as Approved Extensions. Responsibility for the content of each such interface is assumed by the organization proposing it rather than by the HPF Forum.

The full HPF-2 Language Specification is available at the following URL:

`http://www.crpc.rice.edu/HPFF/hpf2/index.html`

3.5 Current Status

HPF-2 was a response to the inadequacy of HPF-1 directives in particular with regard to irregular data structures. A complementary activity, funded by the LTR unit of Esprit, is the HPF+ project. This has explored further structures and directives outside of HPF-1 and HPF-2. The partners in HPF+ are well represented in the HPC-Standards project and the two projects have been mutually beneficial. HPF+, in particular, devoted time to trials of new features in real industrial applications. Their experience has contributed to EU feedback via HPF workshops. HPF+ has now finished (April 1998). There is expected to be a period of stability for the standard while compiler vendors implement the new HPF-2 features.

4 MPI-2

4.1 Background

The message-passing paradigm for programming parallel computers has been developing over many years and has led to a very successful standard, MPI-1, which was agreed in May 1994 and is now supported by all major vendors. One reason for the success of this standard was its clearly defined scope: less mature areas of functionality such as parallel I/O were not included. Beginning in March 1995, these omissions were addressed by the MPI-2 standard, which was eventually published in July 1997. The MPI-2 standard contains extensions and clarifications of MPI-1, additions providing new functionality and new language bindings. The bulk of the work has been in determining appropriate standards for the new areas of functionality, comprising: dynamic processes, one-sided communication and parallel I/O.

4.2 Project

The HPC standards project funded European representatives to attend 9 meetings of the MPI Forum and held a workshop in February 1997. The meetings were held almost bimonthly with an average of between 2 and 3 representatives at each meeting, and considerable consistency of attendees. The appendix contains details of the meetings and lists of the representatives. Further information is available via the HPC standards Web page.

Attendees have been active in disseminating news from the meetings, both informally and though talks at other European meetings.

The Third European MPI Workshop was hosted by the Edinburgh Parallel Computing Centre February 13-14 1997 and was funded by the project. It formed the latest in a series of workshops on MPI that started with the 1st European MPI Workshop held in January 1994. The topic was strictly MPI-2 and the workshop provided an excellent opportunity for users and developers to learn more about MPI-2 and meet members of the MPI-2 Forum and the MPI developer community. The Developer sessions of the workshop consisted of presentations from vendor and third party developers, in which MPI and MPI related products and development plans were described, while the User sessions were concerned with the application oriented experiences and perspectives of MPI from end users. The meeting generated significant interest in MPI-2 with some 56 delegates, representing 36 organisations across 11 countries, attending. A full report on the meeting is available on the Edinburgh Parallel Computing Centre Web page and can be accessed directly from the HPC standards Web page.

4.3 European Input and Benefits

Since vendor-specific architectural issues have significant effect on the new functionality contained within MPI-2, vendors have had most input in determining this part of the standard. The main European input has been at the levels more directly related to applications, and indeed the MPI-Forum has taken more application-oriented decisions as a result of attendance. In particular the following parts of the standard were influenced as a result of European voting:

- Fortran 90 bindings
- Dynamic Process Creation
- External Interfaces
- Parallel I/O
- Interoperability

The direct consequence of this input has been to support some European applications (for example, VAMPIR from PALLAS benefited from the choice of external interface) and also to represent the interests of other Esprit Projects with regard to language interoperability.

The indirect benefits of attendance, knowing the precise state of the standard and the directions it was moving, were probably even more important. For example, PALLAS have been able to provide MPI-2 related services to customers at a very early stage, and have produced one of the first MPI-2 implementations for an industrial customer. This understanding of the developing standard has also been beneficial, though in the opposite direction, though dissuading GMD from relying excessively on the dynamic process management functionality contained in MPI-2.

4.4 Current Status

Although it is too early to judge the success of MPI-2, it is widely supported and all major vendors have committed themselves to providing MPI-2 implementations within a year (reported at 12th Real Applications on Parallel Systems (RAPS) meeting, December 1997). The Japanese Vendors seem particularly enthusiastic and much of their work in this area is performed in Europe.

One of the most important aspects of MPI-2 for HPC is parallel I/O since this provides a route to much more easily portable code. As implementations of MPI-2 appear and are made efficient, this is the area in which significant progress may be expected.

5 PARKBENCH

5.1 Background

The PARKBENCH (PARallel Kernels and BENCHmarks) committee, originally called the Parallel Benchmarking Working Group, was founded at Supercomputing '92 in Minneapolis, when a group of about 50 people interested in computer benchmarking met under the joint initiative of Tony Hey and Jack Dongarra. Its objectives are:

1. To establish a comprehensive set of parallel benchmarks that is generally accepted by both users and vendors of parallel systems.

2. To provide a focus for parallel benchmark activities and avoid unnecessary duplication of effort and proliferation of benchmarks.

3. To set standards for benchmarking methodology and result reporting together with a control database/repository for both benchmarks and the results.

4. To make the benchmarks and results freely available in the public domain.

Involvement in PARKBENCH is open without charge to all members of the high performance computing community, and operates similarly to the High Performance Fortran Forum.

The committee divides it work between 5 groups, namely:

- Methodology (chair: David Bailey, NASA, USA).
- Low-level benchmarks (chair: Roger Hockney, visiting Professor at Southampton and Westminster Universities).
- Kernel benchmarks (chair: Tony Hey, University of Southampton, UK).
- Compact application benchmarks (chair: David Walker, University of Cardiff, UK).
- HPF compiler benchmarks (chair: Tom Haupt, NPAC, Syracuse University, USA).

The committee usually meets 2 or 3 times a year in Knoxville, Tennessee in the USA and at the Supercomputing conference.

To facilitate discussion and exchange of information the following email list exists:

`PARKBENCH-comm@cs.utk.edu`

5.2 European Input and Benefits

HPC standards funding has been used to maintain collaborative links with the US and has helped to initiate some new projects described in more detail below. Over the duration of the HPC-Standards project PARKBENCH has been an excellent forum for discussions particularly in the area of low-level benchmarking (notably COMMS1) and especially the sub-set covering the message-passing performance. Prof Roger Hockney (the chairman of the low-level group) provides comprehensive details on the following Web page:

http://www.minnow.demon.co.uk/Pbench/index.htm

5.2.1 New "COMMS" Benchmarks

The COMMS1 (or pingpong) benchmark from the PARKBENCH low-level suite measures the basic communication properties of a message-passing MIMD computer (further details are given in Appendix 4). The existing COMMS1 benchmark, which was developed and tested for message lengths up to 4*10^4, has been found to produce invalid performance parameters when applied to some recent data with very long message lengths reaching 10^7 or even 10^8 Byte. This is because the least-squares fitting procedure used minimises the sum of the squares of the absolute error between the fitted curve and the data. This leads to the fitted curve being "tied" to the values for long messages and virtually ignoring values for short messages. This can lead to some very strange and spurious results.

The requirement for new COMMS benchmarks has led to significant discussion at recent PARKBENCH committee meetings as well as presentations at the PARKBENCH Workshop. The European involvement here has been strong with contributions from Southampton (Roger Hockney and Tony Hey), Portsmouth (Mark Baker) and Westminster Universities (Vladimir Getov) in the U.K.; Pallas (Hans Plum) and NEC Europe (Rolf Hempel) in Germany; VCPC (Ian Glendinning) in Austria; Utrecht University (Aad van der Steen) in The Netherlands, and others. Europe has a solid background in this area as the low-level benchmarks originated from the European project GENESIS and the EuroBen initiative. There has also been significant exchange of ideas in this area with PARKBENCH members in the USA, including, Pat Worley (ORNL), Bodo Parody (Sun), Charles Grassl (SGI/Cray Research), Adolfy Hosie (LANL), and Ron Sercely (HP/CXTC). Two new codes have been proposed and after the latest e-mail discussions in this area the research community and the vendors are very close to an agreement about a general specification of the new low-level COMMS benchmarks.

Charles Grassl has developed new FORTRAN/C and C/Winsock low-level pingpong benchmarks, whilst Roger Hockney (visiting Professor at Southampton and Westminster Universities) has produced a modified version of the existing COMMS1 benchmark. This modified COMMS1 benchmark solves the existing problem by minimising the relative rather than absolute error between fitted curve and data, and also provides a three-parameter fit which can be tied to three data points (normally the shortest and longest message and a point in between).

This new COMMS1 is available at URL:

http://www.minnow.demon.co.uk/Pbench/comms1/

5.2.2 PARKBENCH Interactive Curve Fitting Tool (PICT)

The PARKBENCH Committee recommended that an interactive curve-fitting tool be produced for the postprocessing and parametrisation of PARKBENCH results using the latest

Internet Web technology. Roger Hockney has produced a Java applet tool called PICT (PARKBENCH Interactive Curve-fitting Tool). This can be used to fit curves to raw performance data (time t, message length n) or output files from the new COMMS1 benchmark.

The PICT tool can be used on or off-line to fit performance curves to experimental data. Interactive fitting is achieved by dragging the performance curve around the graph with the mouse until a "best" fit is obtained. The corresponding performance parameters and relative error of the fit to the data are displayed as this takes place.

Two and three-parameter performance curves (or models) are provided. The two-parameter (rinf, nhalf) model assumes a linear relation between time and length, and a pipe-function dependence of performance on length. In the three-parameter model, time is a non-linear function of length and gives extra flexibility when the two-parameter model is inadequate.

Data may be input from any URL on the Web, or from a local file. Input filters are provided. A number of examples are included dating from 1996/7 from the COMMS1 benchmark on the Cray T3D &T3E, IBM SP2, SGI Origin 2000, and Convex.

The PICT applet can be run from the following URL:

`http://www.minnow.demon.co.uk/pict/source/pict2a.html`

Appendix 5 gives further details of the new 3-parameter model and lists some of the features of the PICT tool**.**

5.2.3 Run and Reporting Rules

Roger Hockney (chair of the low-level group) has produced a set of run and reporting rules for the PARKBENCH codes. These cover:

- Benchmark Compilation and Measurement
    - Baseline Runs
    - Optimized Runs
    - Limitations of Optimization
- Disclosing and Publishing PARKBENCH Results
- Submitting Results to PARKBENCH Committee

5.2.4 The Electronic Journal of Performance Evaluation and Modelling for Computer Systems

The PARKBENCH Committee discussed ways of increasing awareness of performance and modelling issues associated with HPC and it was decided that a new peer-reviewed electronic journal should be devised. One focus for such a journal would be dissemination and analysis of PARKBENCH results. The Performance Evaluation and Modelling for Computer Systems (PEMCS) electronic journal was initiated in early 1997 at the suggestion of the HPCnet Network of

Excellence. HPCnet provided some start-up funding for this activity which promotes responsible performance reporting. Therefore, although not funded directly by HPC-Standards or the PARKBENCH group, the journal has active involvement of EU and US experts and this can be seen to be a result of EU support for PARKBENCH in HPC-Standards.

Four of the major factors that determine the success of a suite of benchmark codes are:

- Usefulness of the results obtained
- Ease of use
- Availability of the benchmark codes
- Availability of the results obtained

PARKBENCH has addressed availability issues through the use of the Web. All codes are freely downloadable from the Web and PARKBENCH performance results are available through an interactive graphical tool on the Web, The Graphical Benchmark Information Service (GBIS), which was developed by Mark Papiani at the University of Southampton, UK. Ease of use is currently being addressed in the USA (see current status). The usefulness of the codes has been recently enhanced through the PICT tool. Increased awareness and dissemination of results is being achieved through activities such as the HPC-Standards PARKBENCH Workshop, through conference contributions (see Section 5.2.7) and now through the PEMCS electronic journal.

Details of the PEMCS journal are given below, and it should be noted that the editors and associate editors are drawn from the PARKBENCH committee.

Editors:

Prof. Tony Hey (University of Southampton, UK)
Prof. Jack Dongarra (University of Tennessee)

Associate Editors:

Dr Mark Baker (University of Portsmouth, UK)
Dr Vladimir Getov (Westminster University, UK)
Dr Erich Strohmaier (University of Tennessee, USA)
Dr Subhash Saini (NASA)

Aims and Scope

Starting at the beginning of 1997, the journal aims to publish on the Web high quality, peer reviewed, original scientific papers alongside review articles and short notes in the rapidly developing area of performance evaluation and modelling of computer systems with special emphasis on high performance computing.

The rush for higher and higher performance has always been one of the main goals of electronic computers. Currently, high performance computing is moving rapidly from an era of ‘Big Iron’ to a future that will be dominated by systems built from commodity components. Very soon, users will be able to construct high-performance systems by clustering off-the-shelf processing modules using widely available high-speed communication switches. Alternatively, the Web itself represents the largest available ‘parallel’ computer, with more than 20 million potential nodes worldwide. This makes the innovative Web technologies particularly attractive for distributed computing. Equally exciting is the goal to achieving Petaflop computing rates on real production codes.

All this makes the performance evaluation and modelling of emerging hybrid shared/distributed memory parallel architectures with complex memory hierarchies and corresponding applications a natural area of priority for science, research and development.

The main objectives of this journal are, therefore, to provide a focus for performance evaluation activities and to establish a flexible environment on the Web for reporting and discussing performance modelling and Petaflop computing. The Electronic Journal will concentrate on computer performance evaluation and modelling and will includes such as:

- Evaluation methodologies
- Analysis of scalability and efficiency
- Modelling, estimation and optimization
- Evaluation results and comparisons
- Low-level benchmarks and parameters
- Full-applications’ performance
- Tools for performance, visualization and tuning
- Performance issues for Petaflop computing
- Performance programming
- Performance portability

The journal is published simultaneously at:

`http://hpc-journals.ecs.soton.ac.uk/`

and

`http://www.netlib.org/utk/papers/PEMCS/`

5.2.5 JAVA Benchmarks

The Performance Engineering Group Software at the University of Westminster, UK (led by Dr Vladimir Getov) has produced a Java version of the PARKBENCH single-processor low-level benchmarks. They are also developing high-performance Java environments (JavaMPI: a Java binding for MPI, and Java bindings for numerical libraries such as BLAS, PBLAS, LAPACK, ScaLAPACK). They now participate in the Java Grande Forum. This involvement would not have been possible without initial work instigated by PARKBENCH and the low-level benchmarks in Java.

Low-level arithmetic Java benchmarks and further details can be found on the Web at:

http://perun.hscs.wmin.ac.uk/CSPE/software.html

5.2.6 Parallel I/O Benchmarks and EU Input to PARKBENCH

A group composed of E. J. Zaluska, D. Lancaster and P. Melas at the University of Southampton have been working on parallel I/O benchmarks with the aim of their eventual inclusion in the PARKBENCH Suite. The development of the MPI-2 and associated MPI-I/O standard described elsewhere in this project report has been a strong motivating factor since MPI-I/O is presently the best way of portably implementing parallel I/O at the application programming level. MPI-1 has been widely taken up in the community and provided MPI-2 is similarly adopted it will be possible, for the first time, to make comparisons of parallel I/O between widely dissimilar systems. HPF-2, by contrast, chose not to include any parallel I/O standard. The parallel I/O benchmarks therefore remain firmly within the context of MPI and besides testing the intrinsic I/O performance they implicitly test aspects of the MPI implementation such as the MPI library and the compilers. The set of benchmark programs is intended to test parallel I/O at several different levels organised using the standard approach of low-level, kernel and compact application classes.

Firstly, a parallel I/O system needs carefully instrumented, low-level benchmarks to facilitate improvements at both the MPI-I/O implementation level and the file system level. The programs in this class are known as low-level tests and the overriding design criterion is that they are simple and make clearly defined measurements. Most applications will make more complex use of I/O and to reflect this there are two additional classes of test called kernel and compact applications.

Kernel programs identify a set of characteristic I/O behaviours that recur in many typical applications. These essential parts are abstracted as free-standing programs and allow the characteristic behaviours to be investigated in detail. This allows the relative efficiency of different typical characteristic I/O behaviours to be measured and provides an evaluation framework against which future I/O intensive applications can be assessed. These kernel programs are still artificial, concentrating only on one characteristic I/O type, so there is still need for a final level of “compact application” programs that combine I/O with computation to provide a more realistic class of test.

The tests generate clearly defined timing measurements, which are analyzed by a separate set of tools. These measurements and the results of analysis address the parallel I/O at levels running from file system, through MPI implementation, to the application level. The levels broadly correspond to the three classes of test. The questions the test suite sets out to answer and the kinds of problems it should help solve are listed below:

- Performance tuning and debugging of file systems and implementations.

- Comparison between different systems and implementations (especially in the context of the characteristic kernel types)
- Performance tuning of applications within the MPI context.
- Measuring parameters needed to predict performance.

In summary, the analysis of the lower classes of test will lead to better understanding of I/O at file system and MPI implementation levels, thus leading the way to performance improvements. The analysis of the higher classes of test aids tuning at the application level.

In order to best address the goals, tools rather than benchmarks are emphasised so the suite is more under the control of the user, and thus implicitly less “automatic”. With this in mind, the analysis requirements are greater than for a simple benchmark so data gathering is separated from analysis in order to simplify the tests themselves, yet allow analysis at varying degrees of sophistication. Learning from criticisms of past benchmarking activity, we insist on simplicity both of the tests themselves and of the run procedure.

The present status is that the low-level and kernel classes of benchmark have been completed but there has so far been little opportunity for access to suitable machines for testing purposes.

5.2.7 PARKBENCH Dissemination at the Euro-Par ’98 Conference Tutorial: Performance Analysis, Evaluation and Optimisation

PARKBENCH dissemination to a European audience will take place at the ACM and IFIP sponsored Euro-Par ’98 Conference being held at University of Southampton from 1st-4th September 1998. One of the six tutorials at this conference is entitled, “Performance Analysis, Evaluation and Optimisation” and is being given by three members of the PARKBENCH committee, Vladimir Getov (University of Westminster), Tony Hey (University of Southampton) and Roger Hockney (Visiting Professor at the Universities of Southampton and Westminster)

Many issues that provide the focus of the PARKBENCH suite will be discussed including, parametrisation of performance measurements, characterisation of applications and architectures, and interpretation of results. This tutorial will give a comprehensive introduction to performance evaluation methodology as well as an overview of performance analysis and optimisation techniques. The methods of analytical performance modelling and estimation will be introduced in detail, and their application will be shown in case studies. These issues will be illustrated with latest benchmark results including PARKBENCH and NPB codes for the state-of-the-art parallel computers - Origin-2000, SP2, T3E, TERA, and others.

Further details are available at URL:

`http://www.europar98.ecs.soton.ac.uk/`

5.3 Current Status of PARKBENCH

The Web site for PARKBENCH is hosted and administered by the University of Tennessee in the USA. The results database, GBIS, is hosted and updated by the University of Southampton in the UK and is mirrored at the University of Tennessee. Prof Roger Hockney in the UK has made a new interactive curve-fitting tool (PICT) available. Recent effort in the USA have concentrated on making the benchmarks easier to run with improved makefiles, the inclusion of all necessary libraries and improved documentation.

The current version of the PARKBENCH is at release 2.1.1. This contains mainly bug fixes to release 2.0, the main features of which are itemised below.

- Full Integration of Kernel and Compact Application Benchmarks
- All Benchmarks are now available through PVM and MPI
- Inclusion and integration of all necessary libraries
- New Set of Run Rules: HTML or PostScript
- More efficient NPB Benchmarks included
- Maximal message sizes extended to 10MB
- Improved Makefiles and Documentation

The PARKBENCH suite currently contains:

- 5 Low Level Sequential codes
- 5 Low Level Communication codes
- 5 Parallel Linear Algebra Kernels
- 2 NASA Ames Parallel Benchmark Kernels
- 3 NASA Ames Compact Application codes
- PSTSWM - Parallel Spectral Transform Shallow Water Model code
- Support of both PVM and MPI as Message Passing APIs.

Overall, PARKBENCH suffers from lack of resources and in particular funded manpower. This does not mean that PARKBENCH is not important for industry. On the contrary, the industrial interest in PARKBENCH has always been high (see for example the recent discussion on the low-level COMMS benchmarks above).

Due to the lack of manpower, it was decided at the HPC-Standards funded PARKBENCH workshop, that PARKBENCH effort should focus on the low-level benchmark suite and associated tools in the near future. One area under discussion is the integration of the PICT tool with a new Java version of the GBIS performance database.
In the future, as a result of the HPF and MPI standards supported by the project, there will be need for new low-level benchmarks in these areas. PARKBENCH has already heard from projects in HPF and MPI-IO that have indicated plans to integrate the benchmark codes they produce into the PARKBENCH Suite. The HPF work is organised by Chuck Keobel at Rice University, but the MPI-IO work is European. At the Southampton PARKBENCH Workshop in September 1997, Dave Snelling of Fujitsu reported on a test suite for parallel I/O based on the MPI-2 interface. The suite is to have the following characteristics:

- The tests will be low-level and the timings will have clear significance.
- Data gathering will be separated from analysis and no model for the behaviour of the data is assumed.
- Curves of data, rather than single numbers will be produced.
- The tests are to be written in Fortran 90.

6 Workshops

6.1 "Summer of HPF" Workshop, Vienna, July 1-4 1996

An HPC Standards information workshop on HPF-2 was held in Vienna on 3-4 July 1996, preceded by two days of tutorials: an HPF-2 tutorial on 1 July, and a tutorial on the use of NA Software's HPF+ compiler and debugger on 2 July. The workshop was divided into sessions on:

- Commercial HPF compilers (with presentations by Thinking Machines, IBM, ACE, Portland Group, NA Software and Applied Parallel Research)
- Research HPF compilers (shpf, VFCS, and KeLP from UCSD)
- Migrating programs to HPF
- HPF applications
- HPF benchmarking.

These were accompanied by an exhibition of commercial and public domain HPF compilers and tools. The workshop attracted about 70 participants from across Europe and the USA and Canada. The workshop proceedings, with copies of many of the presentations, can be found on the WWW at:

`http://www.vcpc.univie.ac.at/news/summer-of-hpf`

The HPC Standards project Web page has a link to this page.

6.2 Parallel Tools Workshop, Brussels, 5 February 1997

A workshop on parallel tools was held in Brussels, 5 February 1997, organised jointly by Smith Systems Engineering and the PAC, University of Southampton. The aim of the meeting was to present the results of a survey of the current status of HPCN tools undertaken by SSE and PAC, and to assess future developments and trends in the area. The timing was to enable the findings of the workshop to be shared with the tool software developer community in advance of the EC's March 1997 call for proposals.

Details of the survey can be found on the WWW at:

`http://www.smithsys.co.uk/smith/public/tech/market.htm`

The project helped in promoting this workshop via Usenet announcements and postings to relevant mailing lists.

6.3 Third European MPI Workshop, Edinburgh, February 13-14 1997

An HPC Standards workshop on MPI was held in Edinburgh on 13-14 February 1997, organised by Lyndon Clarke and the EPCC. The aims of the workshop were:

- To disseminate information concerning the latest version of MPI, MPI-2
- To dissemination of various implementations and intended future amendments of MPI-2
- To provide a forum for feedback from users to developers
- To provide a platform for general and open discussion on matters relating to MPI-2
- To provide feedback to members of the MPI-2 forum prior to standardisation

Some 56 delegates, representing 36 organisations across 11 countries, attended. Full details of attendees, together with technical details of the proceedings of the workshop, can be found on the WWW at:

`http://www.epcc.ed.ac.uk/mpi2euro`

The HPC Standards project Web page has a link to this page. Apart from providing basic funding enabling the workshop to proceed, several of the named HPC-Standards beneficiaries also received funding for attendance:

- James Cownie (DIS)
- Klaus Wolf (GMD)
- Hans-Christian Hoppe (PALLAS)

The following delegates from the USA also received funding:

- Steve Huss-Lederman (U. Wisconsin)
- Bill Gropp (ANL)

Two further Europeans were funded:

- Ed Zaluska (U. Southampton)
- Panagiotis Melas (U. Southampton)

A detailed report of the meeting by Panagiotis Melas appears in the appendix to this report.

The MPI-2 standard was finalised in July 1997. Information concerning the standard can be reached from the project Web page.

6.4 Fall-97 PARKBENCH Workshop, Southampton, 11-12 September 1997

A two-day PARKBENCH workshop was held in Southampton from 11-12 September 1997 and funded by HPC Standards. It was organised by Mark Baker (University of Portsmouth) and Vladimir Getov (University of Southampton). The aims of the workshop were

- To disseminate information about the present status of the PARKBENCH benchmarks.
- To encourage discussion of approaches to performance evaluation and modelling of HPC systems.
- To bring together US experts, European academics and vendors.

The event also featured discussion of parallel I/O and data-parallel issues in HPF.

A workshop programme with abstracts and slides of the presentations (in postscript or HTML format) can be found on the WWW, at the URL

http://hpc-journals.ecs.soton.ac.uk/Workshops/PEMCS/fall-97/

The presentations are made available through the auspices of the electronic *Journal of Performance Evaluation and Modelling of Computer Systems*, in the founding of which several HPC Standards representatives were involved. A link to the above site exists on the project Web page.

7 Project Coordination and Finances

The first coordinator of the project was Dr John Merlin, who was at its helm from its inception until the end of October 1996. Upon his departure to Vienna, the duties of coordination were split between two individuals, Drs Robin Allen (Administrative Coordinator, 20% of full-time) and David Lancaster (Technical Coordination, 80%). Dr Allen was involved with the project from November 1996 until its cessation, Dr Lancaster from March 1997 until its cessation.

7.1 Administrative Coordination

The duties of the Administrative Coordinator were as stipulated in the TA. He was the point of contact for submission of expenses claims and requests for funding, and was also responsible for extracting feedback from attendees of HPC-Standards-funded meetings. A project webpage was set up describing the aims of the project, and it was regularly updated with reports and minutes of standardisation meetings attended by HPC Standards participants. He was also responsible for initiating and overseeing the holding of several workshops detailed in the TA, relevant to the standards being funded (although the parallel tools workshop was arranged directly through DG3). These meetings were advertised through mailing lists and Usenet announcements, and details (such as programmes) were made available on the World Wide Web. As such, the duties of the admin coordinator were largely confined to email and 'phone communication with project representatives, and to website maintenance. This required comparatively little work on a day to day basis, but intensified at the times of standards meetings and significantly increased when workshops were being organised.

7.2 Technical Coordination

The work of the Technical Coordinator has been concerned with MPI-2. An important component of the MPI-2 standard is MPI-I/O which provides the first general standard for a parallel I/O interface. All major vendors have stated their commitment to this standard and some preliminary implementations are already available. This presents, for the first time, the opportunity to directly compare different underlying parallel file systems and thereby judge their relative merits. Such a study is important because although parallel I/O is widely regarded as the cure for the I/O performance bottleneck that is expected in many future HPC applications, the technique is still at a research stage and is not yet well developed. The Technical Coordinator has devoted most of his time to developing tools based on the MPI-2 interface that can be used to test the performance of a variety of parallel file systems. Besides testing the intrinsic I/O performance these tools also test aspects of the MPI implementation. They can thus be used to compare different strategies used to build parallel file systems, can help in performance tuning and debugging of filesystems and implementations and can also be used to aid tuning at the application level. In addition to aiding the dissemination of MPI-2 by providing useful tests, it is expected that this work will be integrated into the PARKBENCH Suite.

7.3 Project Finances

A breakdown of the financial status of the project is shown in the table below:

Cost Cat	ECU	ECU Awd	ECU Underspend
Labour	57 275	59 500	2 225
Equipment	0	0	0
Travel & Subsistence	78471	91 392	12921
Cons & Comp	15 058	14 675	-383
Other Significant			
CIC	30 161	33 113	2 952
TOTAL	180 965	198 680	17 715

(This is a corrected version of an earlier table, in which part of the travel expenditure had been miscoded by the finance department as an equipment item.)

It can be seen that the largest underspend was in the travel budget. 37 claims were made against a budgeted 40. The shortfall is largely due to many claims being only for partial expenses or less than the budgeted amount.

In understanding this table it is important to bear in mind the varying rate of exchange between the ECU and Sterling. Esprit project 21111 was a fixed price contract, set at 198,680 ECU. At the commencement of the project, the exchange rate was 1.19 ECU to the pound; however, within a few months the pound had increased in strength, soon reaching 1.4 ECU to the pound, and has crept up since. This meant that, in Sterling terms (with which the Finance Department works), the budget available for the project actually decreased over time: one consequence of this was that the salary of the coordinator/s, while fixed in Sterling, effectively increased in ECUs, perhaps making the extremely small underspend in this category in the above table misleading. Another was that three of the four workshops were held at a time when Sterling was strong, resulting in a slight overspend in that category. It also accounts for the Adminstrative Coordinator's mistaken belief at the cessation of the project that the total underspend was significantly greater than was actually the case.

Another consequence of the varying rate of exchange is that, had the project been significantly more proactive, it could have run significantly over budget - at the original exchange rate, use of the full ECU budget would have resulted in a £30,000 overspend.

Appendix 1: HPC-Standards Benefits: Summary Comments from Funded Representatives

A1.1 HPF-2

A1.1.1 John Merlin

Subject: Re: HPC Standards
From: **John Merlin** <jhm@vcpc.univie.ac.at>

HPC Standards was a bit too late to influence the HPF-2 standard much. That's probably why most of the HPCS HPF people didn't take up the offer of funding to go to the meetings.

Nevertheless, although it may not have influenced HPF-2, I believe the fact that NAS could attend a couple of HPFF meetings was very valuable for them, since:

-- they got to meet many of the important people in the field and established a presence in the HPFF community;

-- they can claim that they were involved in the HPFF language definition process, which can be useful for marketing;

-- they got detailed, in depth knowledge of what's in HPF-2, which is valuable for the further development of their compiler.

I think a good case can be made that this was very useful for NAS, which is one of Europe's few HPF compiler developers.

Likewise I think that Thomas Brandes -- who has developed the ADAPTOR public domain HPF compiler but who hadn't been to any previous HPFF meetings as far as I'm aware -- benefited from attending the HPFF meetings for pretty much the same reasons as above.

Probably the main benefit of the project in HPF terms was its support for the 'Summer of HPF' in Vienna' meeting and attendance at the 1st HPF User Group (HUG) meeting, which were both enormously useful.

Europe had a strong presence at the 1st HUG meeting, thanks in large part to HPCS. In fact in the discussion afterwards it was stated that Europe was very strong in HPF activities (i.e. R&D and use of the language for code development); this would not have been apparent were it not for HPCS. **This is one of the reasons why it was decided to hold the next HUG meeting in Europe** (in Porto, see http://www.vcpc.univie.ac.at/activities/news/HUG98).

Support for the 'Summer of HPF' workshop was also of course very valuable. The workshop disseminated info about HPF activities and the emerging HPF-2 standard to European industry and academia. HPCS supported the attendance of several US invited participants, whose participation was invaluable since they are among the key people in the HPF field, and it also funded some European participants.

A1.1.2 Barbara Chapman

Subject: Re: HPC Standards
From: **Barbara Chapman** <chapman@mcs.anl.gov>

I have just seen John Merlin's reply to your message and think he has got it pretty well summed up. development **As you are aware, these funds were too late to enable active participation in the HPF standard**. However, the task for compiler developers is anything but over. In particular, vendors must judge for themselves the relative importance of the advanced features, which are currently an optional part of the implementation. In order to do so, they need contact with end users -- who are in the best position to state what they feel they need most in additional support -- and this is what we are trying to provide them by holding regular meetings with participation by both groups.

The HPC Standards funding was very helpful on that front.

As John pointed out, it also highlighted European efforts to the US audience. However, I feel the interaction between end users and vendors to be the single most important contribution from the project.

A1.1.3 Cecile Germain

Subject: Re: HPC Standards, Santa Fe '97
From: **Cecile Germain** <Cecile.Germain@lri.fr>
Organization: LRI - Universite Paris-Sud

My attendance at this meeting has definitively been worthwhile for me : I got in touch with people from PARALLAB (Research lab, Norway), and we have built a joint proposal (LRI = me, PARALLAB and the oil company Norsk Hydro in Norway, and Ecole des Mines (France)) for a grant from FNS. FNS is "Fondation Franco-Norvegienne pour la recherche scientifique et technique et le developpement industriel". The project deals with the design of a parallel version of a 3d pre-stack migration code, and includes an HPF version. The project is currently under review.

I think that the significant European input is proved by the fact that the next meeting will take place in Europe. In a more informal way, it seems that the quite important French participation at the first meeting was not expected, and was good news.

A1.1.4 Thomas Brandes

Subject: Re: HPC Standards
From: **Thomas Brandes** <Thomas.Brandes@gmd.de>

Influence of GMD for HPF 2 within the HPF Forum

Unfortunately, GMD could not be very helpful to achieve the aim of bringing substantial input to the Forum and playing a leadership role for important subgroups.

As the HPC Standards project started rather late, we could only participate in one of the last meetings. We were also not able to get the voting rights. But together with VCPC we brought forward our input that we have got within the common projects, mainly PPPE. And we followed the discussions of the HPF Forum and the state of the language definition during the whole time to become a valuable discussion partner.

It was indeed very important to participate in the one meeting. Members of the HPF Forum were very open for discussions and clarifications and took our comments into account. The detailed discussions within the Forum were also very helpful to understand the background of the language definition in detail. **And our interest in HPF showed the European interest.**

The participation was also very helpful for the PHAROS project. Though PHAROS did not use HPF 2 for porting the industrial applications, it was very helpful to convince the partners that this new language definition will avoid many of the problems that they have had for porting their applications to HPF 1. It also increased the motivation that investing in HPF as a technology might become more interesting in the future.

GMD made also an active contribution in the first HPF Users Group meeting together with other project partners of PHAROS. It became obvious at the meeting that Europe has already more experience in using HPF as a technology than the USA. And that has certainly increased the visibility of Europe in the HPF community.

A1.1.5 Alistair Ewing

Organisation: EPCC

The first HPF User Group Meeting was held in Santa Fe to provide an opportunity for users of HPF to meet each other, share ideas and experience and get up-to-date information about implementations and future directions. In each of these areas, I found the meeting extremely useful and stimulating.

My presentation at the meeting was "Optimising HPF Codes: A User’s Experience" and aimed at outlining the types of techniques available to users to improve the performance of their HPF codes. These techniques range from single processor optimisations to making best use of locality, to exploiting tuned third party libraries to including explicit message-passing calls in the HPF framework. My conclusion was that there is much scope for optimising HPF codes and I outlined a strategy for such optimisation and other further work of interest.

From the feedback I received I feel that the talk was well-received, by both users and developers alike, and although not providing general solutions, demonstrated what is possible given a standard set of optimisation techniques by a user of HPF and which is widely applicable. Moreover, it demonstrated that leading edge work in this field is being undertaken in Europe.

The funding from the Standards Project was invaluable for making attendance of this meeting possible, and I was very pleased to mention this at the start of my talk. Not only was I able to promote European participation and input into Standards meetings, but I was also able to demonstrate the research contributing to the European HPF effort. Ironically, it took a trip to the US to meet many other European researchers involved in HPF, however, I feel that these contacts will significantly increase the strength of the European involvement in future HPF developments. In addition to this, the experience I gained from other HPF developments very naturally feeds into EPCC's HPF training manuals and will enhance the technical application support for the users of our facility.

A1.1.6 Henk Sips

Subject: Re: HPC standards
From: **Henk Sips** <sips@cs.tudelft.nl>

As I can recall we did only once use the HPC standards fund. Not for myself, but for Will Denissen of TNO to attend the "last" HPF meeting in Santa fe last year.

Before that time there was an urgent need, but no money, so most of the trips the HPFF were financed by the PREPARE (Esprit) project. This was always a problem, because you needed permission to go to the US for European projects. This is really a silly rule, since costs for travel are not much different for Europe or the US.

In short: HPC standards was a good thing, however, as all things in Europe: Too Late !!

For the Java thing, history seems to repeat itself (if you want I can send you (and Agnes Bradier) a short summary of the meeting).

regards,

- Henk

A1.1.7 Bob Boland

Subject: Re: HPC Standards/Santa Fe UGM, Feb '97
From: **Bob Boland** <wrb@lanl.gov>
Organization: Los Alamos National Laboratory

Re: European input at the 1st HPF UGM. Several of the Europeans were session chairs. That included Brandes and Chapman. I talked with most of them too. A couple of observations. All were very active participants; they interfaced well with those on "this side of the pond", our little joke about the Atlantic. (Some of the visitors from the Far East had a more difficult time due to language problems.) There was excellent exchange between the European delegates and the Americans; sometimes it was spirited, but always professional.

Sometimes a "first time" meeting is not very successful. This one, in my opinion, was extremely successful and those whom you supported were very vital to the success. We

acknowledged the support received from Esprit at the meeting and I am still very grateful for your generosity.

A1.2 MPI-2

A1.2.1 James Cownie

Subject: Re: HPC Standards
From: **James Cownie** <jcownie@dolphinics.com>

HPC Standards participant feedback

Jim Cownie, Dolphin Interconnect Solutions.

I believe that the HPC Standards project significantly improved the participation of European industry in the MPI-2 process, by providing funding for the attendance of a number of people (myself included).

The information gathered by attendees was made available to a wider European audience through the successful, HPC Standards funded, "European MPI meeting" in Edinburgh. This allowed many other people with an interest in MPI to be briefed on MPI-2, and to meet some of the important US participants in the process (such as Rusty Lusk, the MPI-2 convener).

It is clear that the MPI process has been very successful (for instance all of the message passing codes in the DoE ASCI program are being written using MPI), and **it was therefore important that Europe continued to participate in the MPI process. HPC Standards was able to ensure that this happened, allowing people who had already attended MPI-1 through the PPPE project to continue to attend MPI-2. The continuity of attendance of this European contingent between MPI-1 and MPI-2 was also beneficial to the whole MPI-2 process, since many of the US participants changed between MPI-1 and MPI-2.**

A1.2.2 Lyndon Clarke

Subject: Re: HPC Standards
From: **Lyndon J. Clarke** <lyndon@makespan.co.uk>

Support from the ESPRIT Project "HPC Standards" (2111) was an essential component of European representation in the MPI-2 standards forum. Without this support EPCC could not have represented the views and, through discussion in working subcommittees and formal procedures, supported the requirements of its numerous European academic and industrial partners.

The 3rd European MPI Workshop, held at the John MacIntyre Centre, Pollock Halls, The University of Edinburgh on February 13-4 1998, and organised by EPCC, was an important part of the project. 56 delegates attended the workshop. The 49 European delegates came from 38 organisations across 11 European nations. The 38 European organisations comprised 8 companies, 4 government service agencies and 24 academic institutions. The workshop also included seven US delegates, five of whom were prominent members of the MPI-2 Forum, and six European members of the MPI-2 Forum participated. **The European delegates received a clear and concise description of the, at the time, proposed MPI-2 standard.**

The MPI-2 Forum representatives received useful feedback from the European developer and user community.

A1.2.3 Hans-Christian Hoppe

Subject: Re: HPC-Standards
From: **Hans-Christian Hoppe** <hch@pallas.de>

With the HPC-Standards funding, PALLAS has been able to attend most of the MPI-2 standardisation meetings, in addition to the MPI workshop in Edinburgh and the IMPI meeting in November 1997. **We could attain voting member status in the MPI Forum, which has enabled us to influence the MPI-2 standard in a number of key issues:**

- Attachment of separate MPI-2 applications
- MPI-I/O
- Support for Fortran 90
- Profiling interface for MPI-2
- External interface for profiling/debugging tools to access MPI-2 internals

During the MPI-2 discussions, we've been able to provide input derived from our experience with HPC applications, both from within other Esprit projects and from our "normal" business. **This has helped the MPI Forum to take more application-oriented decisions. Cooperation with the EPCC representative has helped us a lot in these issues. In addition to actively influence the outcome of discussions, our presence at the MPI-2 meetings did give us first-hand knowledge of the emerging standard, which in turn has helped us a lot improving our MPI-related tools and extending them to MPI-2. In addition, we've been able to offer customers MPI-2 related services.**

The other European participants (EPCC, GMD, University of Stuttgart) have likewise helped to improve the MPI-2 standard and to bring European issues to bear.

The Edinburgh workshop early in 1997 was very useful to exchange information between key MPI-Forum members and the European MPI user community, as well as promoting the upcoming MPI-2 standard in Europe. Both sides - the MPI Forum and the MPI users - certainly got considerable benefits from this event.

A1.2.4 Klaus Wolf

Subject: Re: HPC Standards
From: **Klaus Wolf** <klaus.wolf@gmd.de>
Organization: GMD, Forschungszentrum Informationstechnik

As already described in the progress and management report from December 1996 my primary role as a MPI-2 user in attending the MPI-2 meetings was to keep track of the ongoing standardisation and to transfer the actual information into CEC projects like CISPAR or WINPAR (additionally into national funded projects). My work and interests in MPI-2 were clearly described in three publications and presentations: the first at the MPI-2 developers conference in Notre-Dame/Illinois (July 96), the second at the 3rd European MPI-Workshop in Edinburgh (February 97) and the third on the European PVM-MPI conference in Krakow/Poland (November 97).

To summarize my (personal) impression on MPI-2 is that the time-gap between (theoretical) standardisation and (practical) implementation of MPI2 libraries seems too big and is still growing. Full and comparable implementations (!!!) of MPI2 are still missing on typical machines like from IBM, Cray/SGI or NEC. Besides that - and that seems to be the bigger problem - interoperability between different MPI-2 implementations will not be available in the near future.

To formulate the resulting relation between GMD as partner in the European HPC-STD project and the (nearly American) MPI standard I would say the following: simply the fact that we know better about the (technical and social/business) internals and the gap between theory and implementation preserved us from relying too early on MPI-2 in ongoing projects. As GMD is still playing an key role as an expert in parallelisation in lots of CEC projects it was and is important to know about those internals and realistic time scales.

A1.3 PARKBENCH

A1.3.1 Tony Hey

Subject: PARKBENCH activity in HPCstds
From: **Tony Hey** <ajgh@ecs.soton.ac.uk>

The HPC-Standards project funded European participation in the PARKBENCH activity. Meetings took place about three times a year - twice in Knoxville with a Birds-of-a-Feather session organized at the US Supercomputing conference each year. In addition a workshop on Benchmarking activity and a European PARKBENCH meeting was held in Southampton in September 1997.

The PARKBENCH initiative started from the European Genesis Distributed Memory Benchmarks and has adopted much of their approach. The meetings have been regularly attended by major players in the US benchmark scene including NASA Ames and ORNL, and have been hosted in Knoxville by Professor Jack Dongarra. **At all meetings there have several European attendees and the activity remains firmly a joint Europe-US activity which is supported by most of the major vendors.** There has been dialogue with the SPEC-HPC group and I attended a meeting of this group with Jack Dongarra before a PARKBENCH meeting in Knoxville.

The focus of the group's activities has been sharpened at the last few meetings and at the Southampton Workshop it was agreed the group should concentrate on the low-level benchmarks originally introduced by the Genesis project. While there has been quite wide acceptance and dissemination of the benchmarks, in my view the activity is still sub-critical and is likely to remain so until some dedicated effort can be directed at the benchmark codes. The group remain optimistic that such funding will become available in the near future.

Tony Hey
Chairman of the PARKBENCH Group

A1.3.2 Mark Baker

Subject: HPC Standards
From: **Mark Baker** <mab@sis.port.ac.uk>

Here is how HPC standards has helped me:

* Meant that I've been able to meet (face to face) people I have collaborated with in the past but never actually met.
* Enabled me meet up at fairly regular intervals with colleagues interested in Benchmarking.
* Helped me build up a network research partners - currently I am working on low-level benchmarks with people in SGI-Cray (USA), a University in Holland and a University in Adelaide.

* Made the administration of the PEMCS easier - again meeting up with people instead of email contacts.
* I have been able to solicit papers for PEMCS more easily by meeting researchers in-person.

Appendix 2: Tables of Standardisation Meetings Attended

During 1996 a total of seven meetings, in addition to a project-funded workshop, were attended by representatives. In 1997, representatives attended a total of five standardisation meetings, in addition to three project-funded workshops. This section details the meetings by date and location, and lists those who attended.

Summaries of what took place at each meeting have been solicited from European participants, and appear in Appendix 3 of this report. Some longer accounts can be found on the HPC Standards project Web page.

A2.1 HPF-2

Date	Location	HPC Standards Participants
18-20 Sept 1996	San Francisco	Thomas Brandes (GMD) Mike Delves (NAS)

There were no standardisation meetings for HPF-2 in 1997; the standard was finalised in January 1997. However, several delegates were funded to attend the first Annual HPF User Group Meeting (UGM) in Santa Fe in February 1997

Date	Location	HPC Standards Participants
24-26 Feb 1997	Santa Fe	Thomas Brandes (GMD) Will Dennissen (TNO) John Merlin (VCPC) Christian Borel (MATRA) Cecile Germain (LRI) Alistair Ewing (EPCC)

The HPC Standards project generously funded several European representatives to attend the foregoing meeting, over and above those allowed for as standard project beneficiaries. These extra attendees were drawn chiefly from industry. The contribution of the project to the meeting was acknowledged in its promotional materials, and in the talks of individual speakers.

A2.2 MPI-2

Date	Location	HPC Standards Participants
5-7 June 1996	Chicago	Lyndon Clarke* Hans-Christian Hoppe (PALLAS) Klaus Wolf (GMD)
17-19 July 1996	Chicago	James Cownie (DIS) Hans-Christian Hoppe (PALLAS) Klaus Wolf (GMD)
3-6 Sept 1996	Chicago	Lyndon Clarke (EPCC) James Cownie (DIS) Hans-Christian Hoppe (PALLAS) Klaus Wolf (GMD)
8-11 Oct 1996	Chicago	Lyndon Clarke (EPCC) Hans-Christian Hoppe (PALLAS) Klaus Wolf (GMD)
17-22 Nov 1996	Pittsburgh	Hans-Christian Hoppe (PALLAS)

* This meeting attendance was not funded by HPC Standards.

Date	Location	HPC Standards Participants
21-23 Jan 1997	Chicago	Hans-Christian Hoppe (PALLAS) James Cownie (DIS)
5-7 Mar 1997	Chicago	Lyndon Clarke (EPCC) Hans-Christian Hoppe (PALLAS)
23-25 April 1997	Chicago	Lyndon Clarke (EPCC) Hans-Christian Hoppe (PALLAS)
15-21 Nov 1997	IMPI Meeting, SC '97, Pittsburgh	Hans-Christian Hoppe (PALLAS)

A2.3 PARKBENCH

Date	Location	HPC Standards Participants
31 Oct 1996	Knoxville	Mark Baker (U. Portsmouth) Erik Riedel (Genias)
17-22 Nov 1996	Pittsburgh	Tony Hey (U. Southampton)

Date	Location	HPC Standards Participants
9 May 1997	Knoxville	Mark Baker (U. Portsmouth) Vladimir Getov (U. Southampton)
15-21 Nov 1997	SC '97 Pittsburgh	Tony Hey (U. Southampton) Vladimir Getov (U. Southampton)

Appendix 3: Meetings Summaries

Several participants funded to attend standardisation meetings produced reports or minutes of these meetings. Lengthier contributions can be found at the project Web site. Briefer examples are reproduced here.

A3.1 HPF-2 Meetings in 1996

A3.1.1 GMD (Thomas Brandes) and NAS (Mike Delves), HPF-2 Meeting, 18-20 Sept 1996, San Francisco

Introduction: HPF-2 and the dedicated Needs in European Software-Projects

The High Performance Fortran (HPF) language is designed as a set of extensions and modifications to the established International Standard for Fortran. It is now the new de facto standard language for writing data parallel programs for shared and distributed memory parallel architectures. As HPF provides a global address space at the user level, HPF programs are much easier to write than conventional message-passing programs.

The new standard HPF-2.0 is intended to overcome the difficulties of the previous standard HPF-1.1. The new 'Approved Extensions' include advanced features that meet specific needs, but are not likely to be supported in all initial compiler implementations. These include:

- New data distribution patterns (including INDIRECT mappings and SHADOW regions) and DYNAMIC data distributions (REALIGN and REDISTRIBUTE)
- New parallel control mechanisms (including the ON clause and TASK_REGION directive)
- Asynchronous I/O operations
- Several predefined EXTRINSIC types (some, like HPF_LOCAL, supported by the HPF Forum; others, like HPF_CRAFT, supported by other groups).

In the CEC-funded project PHAROS, industrial applications are ported to HPF. Although the current ports will be ported to HPF-1.1, the need for the new proposed extensions has been identified. The availability of HPF-2.0 may improve the willingness to port other industrial applications to HPF. The feedback of the PHAROS project will give certain priorities for the implementation of approved extensions in future HPF compilers.

We have attended just one HPC Standards-funded HPF-2 meeting to date, on 18-20 Sep 1996.

Report on HPF-2 Meeting, San Francisco, Hyatt Regency Hotel, 18-20 Sep. 1996

Report by Thomas Brandes and Mike Delves

Primary aims of meeting:

- To finalize the HPF 2.0 Document for the Public Review
- To draw up a schedule for the Document
- Establish HPF Users Group

European Participants: Mike Delves, NA Software, UK; Thomas Brandes, GMD, Germany.

Other Participants:

Robert Babb, University of Denver
Alok Choudhary, Syracuse University
Ken Kennedy, Rice University/CRPC
Charles Koebel, Rice University
David Loveman, DEC
Piyush Mehrotra, ICASE
Carol Munroe, TMC

Bob Roland, LANL
P. Sadayappan, Ohio State U.
Jaspal Subhlock, CMU
Carl Weiss, DEC
Joel Williamson, HP
Mary Zosel, LLNL
Henry Zongaro, IBM, Canada

Main issues addressed:

- Restructuring of Chapter 4(INDEPENDENT) (by Mike Delves)
- New section about differences HPF-2.0 and HPF-1.1 (by Mike Delvs)
- ON HOME, especially interprocedural (Koelbel, Subhlok)
- Shadow width for distributed arrays
- Pointer and related distribution directives
- Syntax for extended attributes
- Restrictions for providing explicit interfaces
- Asynchronous I/O (many related problems due to availability of the values read in)
- Restructuring of chapters related to EXTRINSIC
- HPF First Users Group Meeting
- New structure of the document
- Schedule of draft
- Letter of invitation for comments

Work in subgroup 'Mapping':

Members: Mehrotra, Weiss, Brandes, Zongara, Sadayappan

- Definition of 'Is specialization of'

- Shadow edges
- Remapping of pointers (giving up of current restrictions)
- INHERIT directive for pointers
- Prescriptive/descriptive directives

Users Group Meeting

The first annual HPF Users Group Meeting will take place in Santa Fe, New Mexico, Feb 24-26, 1997, Eldorado Hotel. The EC PHAROS project will participate with an invited presentation.

Participation of EC members

Thomas Brandes worked within the subgroup 'Mapping'. His participation was useful to identify current deficiencies and to clarify some definitions. Mike Delves did some editorial work on different chapters of the final document.

A presentation of HPF-2.0 features has been given within the PHAROS project at the last porting workshop (10-11 Oct, 1996)

Results

The _High Performance Fortran Language Specification, version 2.0.delta_ is now available (since 29th Oct). This is a public comment version of the document. The HPFF Forum will value any and all reactions that users care to make. These will be taken into consideration in preparing the final draft in December.

A3.2 HPF-2 Meetings in 1997

A3.2.1 TNO - Institute of Applied Physics (Will J.A. Dennissen), First Annual HPF User Group Meeting, Santa Fe, 24-26 Feb 1997

HPF User Group

At the first annual HPF User Group meeting there were 69 participants. Most of them (49) were from the United States, 17 from Europe and 3 from Japan. Given the travelling distance to New Mexico this is not a bad contribution for Europe or Japan. Besides the familiar HPF Forum faces a lot of new people were present. Most of the representatives were from universities or national research labs, or hardware/tool vendors. But from an industrial perspective more interesting, the first signs of industrial end-user interest showed up in representatives of Boeing, Mobil and Amoco. Another 20 real industrial HPF users were indirectly visible as part of the Japan Association for High Performance Fortran, which brings together compiler developers and advanced HPC users.

Contribution TNO-TPD

We have contributed to the HPFUG meeting with one presentation and a poster session. In our presentation "Migrating an industrial CFD code to high-performance Fortran" we outlined the results of a pilot project. It focussed on the end user reactions to a new language. It turned out that real industrial end users with large strategic codes do not have the highest performance on the top of there priority list. Things that are much more important are portability to other platforms, standardization of the language in a ISO/ANSI standard and the learning curve of the new language for their programmers. How to protect their current investments in application code, working practices, tools and experiences. This view was confirmed by the other presentations as well as the feedback obtained on our presentation. In particular the elegant high-level parallelization contructs that are offered by HPF matches much better with the industrial end-user requirements than the explicit message passing programming model. Therefore we believe HPF is a better alternative for the parallelization market.

Status of HPF

At the HUG Future and Wrap-up meeting, our impression was that there is a clear consensus that HPF as a language is a better alternative than explicit message passing, that the HPF compilers and associated tools are not mature enough and it is just a matter of time to solve these issues. For example: We tried the kernel CFD code (about 500 lines) on a cluster of workstations using our PREPARE research compiler, and compared the results with the same code compiled by the commercially available PGI compiler on the Cray T3E. The figures showed a nice scaling on the cluster of workstations but no scaling at all on the Cray T3E. The timings on one processor were more or less the same. This kind of compiler dependent behaviour is unacceptable to industrial end-users that expect portability of code AND performance. Much broader tooling support is needed while parallelizing real applications, like analysis tools, profilers and debuggers. A survey amongst the participants revealed a strong interest in first completing the core HPF 2.0 functionality, before going into the approved extensions.

Education/Promotion/Support

Besides developing HPF tooling, a lot of new initiatives showed up in teaching/explaining HPF to new uses. Some excellent course material is available on the Internet. To help promote HPF the HPF user group showed strong interest in making HPF WEB pages which can give new users a clear overview about the language standard, the available courses, tool vendors supported hardware and last but not least an overview of the success stories in migrating real (large) industrial codes to HPF. Also in this respect our contribution was welcomed.

Conclusion

Although each new language has it's own time scale to full acceptance, we think HPF is very promising, the user interest is growing and also real industrial success stories are already showing up.

A3.2.2 VCPC (John Merlin), First Annual HPF User Group Meeting, Santa Fe, 24-26 Feb 1997

I attended the First HPF Users' Group (HUG) meeting in Santa Fe on 24-26 February, as well as the final HPF2.0 Forum meeting that was held immediately afterwards at the same venue.

At the HUG meeting I presented my recent work on the design and development of a SPMD-HPF programming system. The purpose of this work is to extend HPF to support dynamic, irregular, block-structured applications, which are of considerable industrial importance in fields such as CFD. My implementation is based on *shpf*, a public domain HPF compilation system that I developed partly in a former Esprit project, PUMA, and this presentation also provided a good opportunity to publicise shpf.

Attendance at the meeting also allowed me to learn about the latest trends and techniques that have arisen in HPF programming, which will be valuable for the VCPC's HPF tutorial and dissemination activities. A report of the meeting was given to the Esprit Pharos project, in which we are partners. In particular, I reported back on proposals to set up an HPF applications survey, and to publish a special issue of a journal on the subject of HPF applications and performance, both of which will provide good avenues for Pharos to publicise the results of its activities, namely the porting of 4 industrial codes to HPF.

Also at this meeting the future organisation, activities and leadership of the HPF User Group (HUG) were decided, with subgroups being set-up in the USA, Europe and Japan. Barbara Chapman of the VCPC (whose attendance was not funded by HPC Standards) was selected as organiser of the European section of the HUG, with responsibility for organising the next HUG workshop, which is planned to be held in Europe in the spring of 1998. A number of HUG-related activities have been initiated at the VCPC as a result of this role. E.g. a HUG

Web site is being created to disseminate information on HPF compilers, tools, applications, tutorials, etc, and a HUG mailing list has been established.

I wrote a detailed report about the HPFF meeting and future plans for the HUG that can be found on the WWW at **http://www.vcpc.univie.ac.at/news/santa-fe.html**

A3.2.3 EPCC (Alistair Ewing), First Annual HPF User Group Meeting, Santa Fe, 24-26 Feb 1997

The first HPF User Group Meeting was held in Santa Fe to provide an opportunity for users of HPF to meet each other, share ideas and experience and get up-to-date information about implementations and future directions. In each of these areas, I found the meeting extremely useful and stimulating.

My presentation at the meeting was "Optimising HPF Codes: A User's Experience" and aimed at outlining the types of techniques available to users to improve the performance of their HPF codes. These techniques range from single processor optimisations to making best use of locality, to exploiting tuned third party libraries to including explicit message-passing calls in the HPF framework. My conclusion was that there is much scope for optimising HPF codes and I outlined a strategy for such optimisation and other further work of interest. The full abstract is at

http://www.lanl.gov/HPF/submitted_Abstracts.html,

with the talk itself available from

http://www.epcc.ed.uk/~ake/HPF

From the feedback I received I feel that the talk was well-received, by both users and developers alike, and although not providing general solutions, demonstrated what is possible given a standard set of optimisation techniques by a user of HPF and which is widely applicable. Moreover, it demonstrated that leading edge work in this field is being undertaken in Europe.

The funding from the Standards Project was invaluable for making attendance of this meeting possible, and I was very pleased to mention this at the start of my talk. Not only was I able to promote European participation and input into Standards meetings, but I was also able to demonstrate the research contributing to the European HPF effort. Ironically, it took a trip to the US to meet many other European researchers involved in HPF, however, I feel that these contacts will significantly increase the strength of the European involvement in future HPF developments. In addition to this, the experience I gained from other HPF developments very naturally feeds into EPCC's HPF training manuals and will enhance the technical application support for the users of our facility.

A3.3 MPI-2 Meetings in 1996

A3.3.1 PALLAS (Hans-Christian Hoppe), MPI Forum meetings, June 5-7 1996, July 17-19 1996, September 3-6 1996, October 8-11 1996, November 17-22 1996.

PALLAS activities in HPC-Standards

This is a summary of the MPI-2 activities by PALLAS in the HPC-Standards project.

Meetings attended

Up to the end of November 1996, Hans-Christian Hoppe has attended the following meetings:

- MPI Forum meeting June 5-7
- MPI Forum meeting July 17-19
- MPI Forum meeting September 3-6
- MPI Forum meeting October 8-11
- MPI BOF meeting, November 17-22

PALLAS could and did participate in all binding votes taken during these meetings, as well as in all general discussions.

Input to the MPI-2 Standardization process

Significant input was given to the following subcommittees, with the main focus on Fortran-90, language interoperability and external interfaces:

- Dynamic Process Creation
- External interfaces
- Fortran 90 Bindings
- Miscellaneous (Language interoperability)

European benefit

PALLAS was able to represent the interests of a number of Esprit-Projects: PHAROS, PARASOL and CISPAR, in particular with respect to Fortran-90, interoperation of C and Fortran, and the coupling of several MPI applications (CISPAR). Working in the “External interfaces” subcommittee (parallel trace generation, VAMPIR) could protect investment into message-passing programming tools from PPPE, and the further development of those tools was expedited significantly.

Dissemination of MPI-2 issues was done in the above-mentioned projects, as well as on several other occasions. First-hand knowledge from the discussions in the MPI Forum proved to be a crucial asset.

A3.3.2 Dolphin Interconnect Solutions (James Cownie), MPI Forum meetings, July 17-19 1996, September 3-6 1996.

I attended two MPI-2 meetings in Chicago funded by the HPC Standards project. The MPI-2 process is beginning to converge, and has presented a draft standard for public comment at SuperComputing '96. European input to MPI-2 has had some effects (for instance the MPI_Name_put and MPI_Name_get functions for communicator naming which will make debugging and profiling of MPI-2 programs more intuitive, and benefit European products such as Pallas' "Vampir"). On a wider canvas my attendance at the MPI-2 meetings meant that I was able to provide informal feedback on the status of the process to many interested European parallel programmers at the European Parallel Tools meeting in Paris.

A3.3.3 EPCC (Lyndon Clarke) MPI Forum meetings, September 3-6 1996, October 8-11 1996.

I have attended four MPI-2 Forum meetings during 1996, two of which were funded by the ESPRIT "HPC Standards" project: January 24-26, Vienna; June 5-7, Chicago; September 3-6, Chicago; October 8-11, Chicago. In these meetings my primary areas of activity have been: Dynamic processes; Single sided communication; Parallel I/O. I have also closely followed the languages bindings discussions for Fortran 90 and C++.

Through the various HPC and MPI activities at EPCC we are fortunate to have contact with a number of European users of MPI-1 and potential users of MPI-2. These contacts have identified that dynamic processes, single sided communication, parallel input/output, and language bindings for Fortran 90 and, to a lesser extent, C++ are key result areas for the European MPI user community.

During the two meetings at which attendance was sponsored by "HPC Standards" the MPI-2 Forum has made considerable technical and procedural progress in the key result areas identified. The MPI-2 chapters on Dynamic Processes and Single Sided Communications have received second formal votes and are therefore more or less finalised. The MPI-2 chapter on Parallel I/O, which was added to the scope of MPI-2 during 1996, has received a first formal reading and is just on schedule for completion during 1997Q1.

(A very detailed report of the 8-11 October workshop, by Lyndon Clarke, can be found at the HPC Standards Web site.)

A3.3.4 GMD (Klaus Wolf), MPI Forum meetings June 5-7 1996, July 17-19 1996, September 3-6 1996, October 8-11 1996.

MPI-2 and the dedicated Needs in European Software-Projects

In the last few years lots of industrial simulation applications were ported onto parallel systems to gain more computational power (e.g. in EUROPORT 1+2). Besides other parallelisation techniques message passing seems to be the most powerful and accepted one. And since MPI is the upcoming standard definition for message-passing libraries, portability and adaptation among different hardware- and software-environments becomes more and more realistic.

However some problems are still unsolved, if two or more independent parallel MPI-applications shall be coupled and work together. Previously standalone address-spaces (unique MPI_WORLDs) have to cooperate now - across their borders - with address-spaces in other worlds.

In three CEC funded projects (SOFTPPAR 8451, CISPAR 20161 and WINPAR 23516) MPI was and will be the implementation base for parallel software tools. Some concrete problems arose when using only MPI-1 as message-passing library: the management of separate address spaces, the integration of dynamic process sets or the possibility of asymmetric communication was always an open issue and needed lots of 'work-arounds'.

MPI-2 definition has a deep interest in 'Dynamic Process Management', 'Language Bindings for C++ and F90' and 'One-sided Communication'. This gives the implementor of a parallel library more chances to realize his tool according to the requirements coming from the applicationers.

To conclude, for the European software-projects mentioned above, MPI-2 is not only a future's standard, but a requirement.

So far, we have attended the following MPI meetings:

MPI, 5-7 June, 1996, Chicago,
MPI, 17-19 July, 1996, Chicago,
MPI, 4-7 September, 1996, Chicago,
MPI, 8-11 October, 1996, Chicago.

A3.4 MPI-2 Meetings in 1997

Most contributions concerning MPI are either lengthy reports or minutes. These are available at the project Web site. One example is included here.

A3.4.1 University of Southampton (Panagiotis Melas), MPI-2 Workshop, University of Edinburgh on 13th and 14th of February 1997

Abstract

This is a report on the **MPI-2 workshop** held at The University of Edinburgh on 13th and 14th of February 1997. This workshop was hosted by the Edinburgh Parallel Computing Centre and supported by the ESPRIT Project "HPC Standards".

1 Summary

MPI (Message Passing Interface) is rapidly gaining wide acceptance as the standard for the message-passing programming. All major HPC hardware vendors support MPI or have announced that they will do so. Mature public domain implementations of MPI are now available and are in wide spread use.

The MPI Forum reconvened during 1995 to address new functionality and to develop an MPI-2 standard. MPI-2 is a superset of MPI and is expected to be finalised in the first half of 1997. The new functionality in MPI-2 includes: dynamic process creation; single sided communication; parallel input/output; and real time extensions.

MPI-2 was described during a successful BOF at Supercomputing 1996. This workshop was organized by EPCC, and sponsored by the ESPRIT project "HPC Standards".

The workshop was organised in three main sessions;

MPI-2 overview: included presentations by a number of key members of the MPI Forum: Rusty Lusk, Steve Huss-Lederman, Bill Saphir, Bill Gropp and Bill Nitzberg.

Developer Session: European members of the MPI Forum presented MPI implementations and MPI related products (SGI/Cray, Digital, NEC, Hitachi, Dolphin, Pallas, etc.).

User Session: In the last session there were presentations of MPI applications for end users, experiences and perspectives of MPI (GMD/SCAI, University of Stuttgart, Basel, INRIA, etc).

The workshop provided an excellent opportunity for users and developers to learn more about MPI-2 and met with members of the MPI-2 Forum and the MPI developer community. The workshop was also an excellent and timely opportunity to provide feedback to the MPI-2 Forum.

2 First Day

On the first day 09:00-09:30 in the morning registration took place. Richard Kenway, director of EPCC, welcomed all the participants and opened the workshop.

2.1 MPI-2 Developers Overview Session:

Chair: Lyndon Clarke, EPCC

2.1.1 MPI-2 Overview, Edwin Lusk

Edwin Lusk gave a brief overview about message passing concepts, how it started from MPI-1 and which are the motivations for MPI-2, then he described the basic structure of MPI-2 and some other MPI issues.

An elementary concept of MPI as a model of parallel computation is a process. Each process controls access to its own space in memory, a parallel computation consists of a number of processes, data can be moved from one process's address space to another's by a matching pair of operations (e.g. send/receive).

By the time MPI-1 came out, the message-passing paradigm was well understood with many applications and efficient match to hardware. Unfortunately most of the systems that time were not portable, while portable ones were not efficient or complete.

Then Lusk gave a brief history of MPI which started in April 1992, up to the MPI-1.1 release in May 1995. MPI-1 was very successful and it was quickly adapted by many MPP vendors (IBM, Intel, Meiko, HP-Convex, Cray, etc).

The success of MPI led developers to push its limits, and request extensions for parallel I/O, process management, C++ bindings, threads, etc. These requests made MPI-2 necessary. The new standard, MPI-2, should include C++ and Fortran bindings, external interfaces, extended collective operations, and language interoperability. Extensions to the message-passing modes should include dynamic process management and one-sided operations. New features of parallel I/O and real-time protocols are included as well.

Likely structure of MPI-2:

MPI-1.2 Corrections, clarifications, version number.

MPI-2.0 Rounding out MPI-1 with useful extensions, dynamic process management module, one-sided operations module and Parallel I/O module.

MPI Journal of Development, will discuss Real time MPI extensions, threads, full F90 interface.

2.1.2 Dynamic Processes, William Saphir

MPI-1 does not provide any mechanism for process management, all processes are initialised during the start-up. There is a need to have a variable number of processes, and use as many processes you can but if there is a need, then you can return some processes back.

The requirements of process management should make sense in every parallel environment (MPPs to workstation clusters). MPI should not take over any OS responsibilities, it must continue guarantee determinism, and it should not compromise over performance and should be compatible with MPI-1.

The MPI solution should not make assumptions about the runtime environment (actually ADI - Abstruct Device Interface) expected to interact when it exists), when MPI must know something it should get information from an info-object.

What is a communicator? a group of processes plus a communication channel, an intercommunicator is a set of groups of processes plus a logical set on communication channels between all processes.

MPI_SPAWN is a collective operation over a parent communicator to an intercommunicator. It is not a replacement of PVM_SPAWN function, MPI programs should initialised on a single process and later on it can spawn to as many processes as the program requires or the system has.

Info-Object is a set of (key info) pairs, both key and value are strings (keys: hoses, arch, wdir, path, file, soft) i.e. soft is a reserved key that holds the max number of processes for MPI_SPAWN, if it failed to start that max number it monitors the soft which should have a number of processes less that max.

MPI-2 allows processes that started independently to establish communication, so an application 1 should be a kind or ‘‘server” while application-2 could be a ‘‘client” (four major routings spawn, accept, listen, connect) i.e. client/server applications can use an interface for two parts of a cooperating application.

Safety delete processes provision mechanism.

2.1.3 MPI-2 Single Sided, Steve Huss-Laderman

In some applications with irregular distributed data it could happen one and only process to have all the knowledge for a call, therefore that process should be able to send or receive a message on its own. This approach gets closer to shared Memory (SM) approach, it can be found available on some current systems, or it is closer to hardware on some other systems.

Remote Memory Access (RMA) extends the communication mechanisms of MPI by allowing one process to specify all communication parameters, both for the sending side and for the receiving side. Each process can compute what data it needs to access or update at other processes workspace, however processes may not know which data in their own memory needed to be accessed or updated by remote process or they do not know the identity of these processes.

Message-passing communication achieves two effects: communication of data from sender to receiver; and synchronization of sender with receiver. The RMA design separates these two functions. Three RMA communication calls are provided MPI_PUT, MPI_GET and MPI_ACCUMULATE.

Window creation, the initialisation operation allows each process to specify in a collective operation a memory window that is made accessible to accesses by remote processes. The information is attached to an opaque object returned by the call, which can be used to perform RMA operations.

MPI provides three synchronisation calls: MPI_WIN_BARRIER a collective synchronisation call, MPI_WIN_START, COMPLETE, POST, WAIT, and a pair of shared and exclusive locks MPI_WIN_LOCK and MPI_WIN_UNLOCK.

The design goals of single sided communication are similar to MPI-1 for send/recv functions:

* protection (communicators mechanism)
* multiple options (buffered, ready, synchr.)
* efficient on machines with hardware semantics
* possible on all machines
* general memory access
* support common practice.

2.1.4 MPI-2 Parallel I/O, Bill Nitzberg

UNIX provides one model for a portable file system, MPI I/O is based on that interface with enhanced portability and optimisation for open/close/read/write functions with extensions of:

* collective data access
* arbitrary flat to partition
* scatter/gather memory operation

Plus

* Basic file interoperability between systems
* user directed optimisation via portable hint
* non-blocking data access

I/O should be layerable on top of the MPI-2 external interface.

Data portioning definition: each file is regarded as a linear byte stream, supporting random access to any set of its bytes.

Displacement-data position, is an absolute byte position relative to the beginning of a file. File type describes data distribution, which process has the right to access which etype.

View defines the current set of data visible and accessible from an open file as an order set of displacement, etype, and filetype.

Data access routines include:

* transfer (read/write calls)
* positioning access routines (explicit offsets/filepointers)
* synchronism routines (blocking/non blocking)
* coordination (independent/collective)

File interoperability is the ability to read the information previously written to a file, it is specified in MPI_OPEN function. There are three aspects to file interoperability outside of a single environment, transferring the bits, converting between different file structures, or different machine representations.

File consistency, defines the outcome of multiple accesses to a single file. MPI provides three levels of consistency: a sequential among all accesses with a single file handle, a sequential among all accesses using file handles created from a single collective open, and a weak consistency among accesses. Overlapping accesses are not by definition consistent.

2.1.5 MPI-2 Language Bindings, William Saphir

Fortran 90 is the current international Fortran standard containing numerous features such as modules, derived data types, strict type checking, function overloading, optional arguments, array sections, etc.

MPI implementations with Fortran binding have some problems e.g. data scoping, strong typing problems, sequence association and naming of basic datatypes.

MPI defines three levels of support for Fortran:

* Basic Fortran support (original Fortran 77 bindings on MPI-1)
* Extended Fortran support (first level plus some additional features of Fortran 90)
* Full Fortran support (not yet specified).

MPI design was based around the notion of objects, C++ bindings are included as part of MPI-2 function specifications. The C++ binding interface to MPI consists of a small lightweight set of classes. There is a one-to-one mapping between MPI functions and their C++ bindings, existing C bindings can still be used in C++ programs.

2.1.6 MPI-2 Extended Collective, William Gropp

MPI-2 provides an extended set of collective communication not only within communicators but also within intracommunicators as well. Collective communication for MPI-2 covers:

* Extension of MPI collective operation to intercommunicators
* Non-Blocking collective communication, separation for data delivery/computation

* Persistent collective communication
* New or extended collective operations

* Point-to-point channels offering early binding of a send/receive pair, or an analogue of a real time channel.

2.1.7 MPI-2 External Interfaces, Steve Huss-Lederman

MPI users are able now to create functionality similar to what is present in MPI, i.e. I/O functionality could be layered on top of MPI. Users can add error codes and new classes on the top of MPI. There are few constraints, with new operations integrated to MPI that should not slow down other MPI functions.

Layerability can help users to create requests providing additional operations on MPI or to handle in a general fashion info requests returned from MPI functions or error codes.

Thread support is not required in all MPI implementations. Initialisation could define either there should be use of threads or not.

2.2 Developer Session I

In this session developers from MPP companies presented their MPI implementations and plans for MPI. The main characteristic from these presentations is that they area all very interested in MPI-2 and ready to support MPI-2 implementations.

2.2.1 Future SGI/Cray Research Plans for MPI, Gordon Miller, SGI/Cray

SGI/Cray, currently support three different MPI implementations for their Cray T3x series computers.

2.2.2 Implementation of MPI for Digital SMP Clusters, Tom Stones, Digital

Major aspects of Digital's MPI implementation are high bandwidth, low latency interconnections channels, their MPI library is optimised for SMP and clusters of MP. The current MPI implementation is based on MPICH.

2.2.3 MPI implementation for the NEC SX-4 supercomputer, Rolf Hempel, NEC

At the moment NEC has two implementations for MPI-1 known as MPI/SX which supports either the TCP/IP based interconnection model or SM for their SX-4 supercomputer. There is under development an independent implementation of MPICH as well.

3 Friday 14th

3.0.1 ''Left to my own devices" Mark Fallon, Hitachi

Hitachi's MPI implementation is based on MPICH. Their main product (SR2201) is an MPP computer scalable to 2048 nodes, using a crossbar for the interconnection among nodes with a peak performance of 600 GFlop. The ADI implementation used three protocols according to the length of the message (short, medium, and long) in order to minimise latency.

3.0.2 Debugging MPI programs with TotalView, Jim Cownie, Dolphin

TotalView is a debugging tool for MPI programs from Dolphin. It is a symbolic debugger for C, C++ and Fortran, and covers a wide range of platforms (Sun, DEC, IBM, Meiko, Win/NT, NEC, SGI, HP, etc).

3.0.3 Tools for Optimizing MPI Applications, Hans-Christian Hoppe, Pallas

Pallas provides a wide range of software tools for developing, tuning, and maintaining parallel applications. Their products include VAMPIRTrace MPI profiling library, a visualisation and analysis tool (VAMPIR), and DIMEDAS a performance prediction tool for message-passing programs.

3.1 User Session I

The last two sessions included presentation from the application-user point of view. The presentations were either from institutes/organisations or academic research groups.

3.1.1 Coupling Fluids and structures Codes on MPI: A User Report on MPI1/MPI2, Klaus Wolf, GMD/SCAI

Germany's national Research Center for Information Technology, SCAI (Institute of Algorithms and Scientific Computing) is interesting in developing a general communications library for coupled problems (COCOLIB) and the interfaces to it. COCOLIB will enable the coupling of message-passing versions of existing codes.

3.1.2 Usage of MPI at the High Performance Computing Centre Stuttgart, Steps forwards Metacomputing, Michael Resch, University of Stuttgart.

University of Stuttgart are interested in MPI-2 as they have a wide range of hardware and software platforms and applications as well. In particular they are interested in aspects such as interoperability and metacomputing.

3.1.3 Communication harnesses- an applications perspective, Robert Allan, Daresbury Laboratory.

Robert Allan gave a brief history of HPC emphasizing portability, efficiency and the message-passing paradigm provided on MPPs.

3.1.4 MPI-DDL: A distributed Data Library for MPI, Niandong Fang, Univeritaet Basel

This presentation discussed a solution for data parallelism and distributed data (which is not supported directly by MPI). Adding an extra layer on top of MPI could overcome this problem?

3.2 User Session II

3.2.1 Parallel finite element solution of the wave equation using MPI, Michel Kern, INRIA Rocquencourt

MPI is used to solve problems with finite element solution of the wave equation.

3.2.2 Direct astrophysical N-body-simulations of star clusters and their implementation on massively parallels computers, Univesitaet Heidelberg Rainer Spurzem

This was a presentation of the Astrophysical department of Heidelberg's University, a complex scientific problem that uses the message passing programming model and MPI.

3.2.3 Comparing the Performance of MPI on the Cray Research T3E and IBM SP2, Glenn Luecke, Iowa State University

This was a presentation from Iowa State University on the performance of the T3E and SP-2 on a collection of communication benchmarks using MPI for message passing.

4 Panel Session

The workshop closed with an hourly panel session (Rusty Lusk, Steve Huss-Lederman, William Gropp, Bill Nitzberg, William Saphir) and an open discussion. Users and developers or MPI had the chance to ask and exchange ideas about the new standards and any other aspects not covered in the workshop.

A3.5 PARKBENCH Meetings in 1996

A3.5.1 University of Southampton (Tony Hey), Supercomputing ’96, Pittsburgh

Report on PARKBENCH "Birds of a Feather" (BOF) session, Supercomputing '96, Pittsburgh, PA, U.S.A

I was chairman of a well-attended PARKBENCH BOF session at Supercomputing '96 in Pittsburgh. After a brief introduction from me, Erich Strohmaier gave an update on the status of the present release of the PARKBENCH benchmarks, and discussed plans for future versions. These include the HPF versions of the codes, and the possibility of shared memory versions of the benchmarks. I gave a presentation on the future development of GBIS, the Graphical Benchmark Information System. This is being rewritten in Java, with a Web interface, and it is proposed that the benchmark results should be stored in a Relational or Object Relational database.

Subhash Saini from NASA then gave a presentation of the latest results on the NAS Parallel Benchmark subset of PARKBENCH, including the new Origin 2000.

The final presentation was by me, on the launch of an electronic journal of "Performance Evaluation and Modelling for Computer Systems" at the beginning of 1997. The venture is part sponsored by the HPCnet Network of Excellence, and is backed by the PARKBENCH consortium.

A3.5.2 University of Portsmouth (Mark Baker), PARKBENCH meeting, 31st Oct 1996, Knoxville, Tennessee

Report on PARKBENCH meeting, 31st Oct 1996, Knoxville, Tennessee, USA

Future work

I proposed that PARKBENCH should initiate the design, development and implementation of a suite of HPF and Shared-Memory benchmarks that mirror the functionality of the current suite of PARKBENCH message passing codes (low-level, kernels and compact applications). It was proposed by the PARKBENCH committee that I should initiate a group to discuss, design, implement and test a suite of HPF PARKBENCH codes.

I discussed my interaction with Oxford computing laboratory, who are interested in producing a set of benchmarks similar to PARKBENCH. It was proposed that I should ask Oxford to go ahead and produce a set of low-level and kernel codes which would, in the first instance, be kept in the Netlib repository.

I also proposed that PARKBENCH should initiate a suite of (much needed) I/O benchmarks. The committee proposed that I should put together a proposal for a suite of I/O benchmark(s) and submit them at the next PARKBENCH meeting.

I discussed with the committee the need to update the mechanism of inputting, storing and presenting the PARKBENCH performance data. I proposed that the data should be held in a object-relational database and be displayed and manipulated via a Java-based GUI. The committee asked that I, on behalf of the Universities of Portsmouth and Southampton, should produce a detailed proposal for acceptance at the next PARKBENCH meeting.

Together with Vladimir Getov, I proposed that the Electronic Benchmarking journal, first mooted by Tony Hey at the PARKBENCH meeting in March 1995, should be set up and the first edition published on the WWW by Spring 1997. It was proposed that Jack Dongarra and Tony Hey should be Editors, with Vladimir Getov, Erich Strohmaier and myself as Associate Editors. It was proposed that a call for papers should be prepared for distribution at SC96 in Pittsburgh, USA. It was also proposed that the Associate Editors would prepare the journal for publication.

A3.5.3 Genias Software (Erik Riedel), PARKBENCH Committee Meeting, October 31st, 1996

Why GENIAS participates

In early 1996, GENIAS Software GmbH was selected by the HPC-Standards Consortium to act as one of the European Representatives in the PARKBENCH Committee. In October (October 31st, 1996) Erik Riedel from GENIAS attended the PARKBENCH Meeting in Knoxville TN.

GENIAS has long-term experience in benchmarking. For example:

1979: first kernel benchmark for the DLR (German Research Establishment for Aerospace and Aeronautics, Goettingen), and test of Cray-1 and CDC Star-100
1980: enhancement of DLR Benchmark by several algorithm modules and CFD kernels, test of Hitachi VP and Denelcor HEP
1981 - 1985: test of all existing vector and parallel supercomputers
1985 - 1989: test of nearly all vector minisupercomputers
1989: cofounder of EuroBen, the European Benchmark Initiative
1990 - 1993: Project Leader of the PARANUSS Project for developing a general parallel benchmark including basic kernels, algorithms and applications
1993 - today: tests and evaluations of nearly all existing parallel computers and workstation clusters

Because most parallel computers are built in the US without any participation or contribution from Europe, it is undoubtedly of the utmost importance to participate in US and international activities in this field. Therefore we see our (Genias') involvement as a means to transfer information and knowledge about benchmarking technology to the European HPCN community.

An important example is the PARKBENCH Initiative. PARKBENCH not only develops a standard set of parallel programs to test and evaluate parallel supercomputers, which can be used by everybody, vendors, users, and researchers, it also tests and evaluates supercomputers to optimize and calibrate the benchmark programs to provide a stable and fair benchmark suite to the HPCN community.

We recommend regular visits to the Web Pages provided by the Universities of Tennessee and Southampton for further information. In addition, GENIAS will include a pointer to these PARKBENCH Web Pages in its own Web pages, soon.

Report on PARKBENCH meeting, October 31st, 1996, Hilton Hotel, Knoxville, Tennessee, USA

After a brief introduction to the participants (Jack Dongarra, UT, Ron Sercely, HP/CONVEX Christian Almer, UT, Edgar T. Kolns, IBM, Mark Baker, Uni of Portsmouth, Sabhashi Saini, NASA, Adolfy Hoisi, CORNEL, Erich Strohmaier, UT, Erik Riedel, GENIAS) Erich Strohmaier gave a comprehensive overview of the situation of the current release. One of his main focuses was low level benchmarking. During a discussion about how to improve some algorithms Vladimir Gretov interjected that, in his mind, a Benchmark is not necessarily the most efficient algorithm; the worst case could also be a good test. It comes up that the Makefile and the directory structure are too complicated. This will be improved by moving the libraries to the corresponding directories.

The next discussion was about the setup of the benchmark on a computer: the number of processes are not in input files and the programs always have to be recompiled. The common meaning is to not change the input files.

After this, Erich Strohmaier presented some benchmark diagrams. In a first diagram about the RINF1 Kernel 4 low level benchmark the SGI PC90 was superior to all other machines in this test case, because of the 4 MB Cache and the benchmark size with only 4 MB. As a solution among others it was proposed to enhance the test cases from 17 to 30 cases or to change the stride from 8 to a higher value.

Similar problems exist in the low-level benchmark Poly2. Here also the SGI PC 90 is the outclassing winner because the tests are so small that the whole test is placed in its cache. A possible solution could be to indicate the memory sizes and the maximum cache space by the vendors.

Another problem is with the timers on the different machines. Some participants are concerned about potential influence of some vendors on the benchmark through different timers. Subhash Saini proposed that all vendors should use the MPI timer MPI_wtime. If no MPI timer is available the vendors have to rebuild one. This timer has to count micro seconds at least.

More tests and comparisons with HPF for a low-level HPF benchmark were considered. This low-level HPF benchmark should provide the possibility to the user to ascertain that his application runs well on a specific machine.

A further discussion about handling MPI 2.0 as soon as it is available and if MPI2.0 will be implemented, was adjourned to the Supercomputing Conference in Pittsburgh.

Also the usage of other languages in the benchmark was discussed. Mark Baker mentioned BSP (block system language, Oxford), which already has a NAS implementation. Sabhashi Saini proposed to place this bid on the Internet to see the reaction. Ron Sercely mentioned that P-Threads, are quite different from MPI or PVM and that it could be convenient to make benchmarks for P-Threads. Mark Baker presented the new database design which contains an object database plus a Java script, which generates a PostScript output. Sabhashi Saini showed the results of a comparison of automatic parallelism and a hand-optimised parallelism with MPI.

In the administrative part of the meeting it was decided that there will be only one mailing list (PARKBENCH-common) instead of all existing PARKBENCH mailing lists. Further, new benchmark tests will be considered in the area of I/O and C language tests, some low-level sequences of which will be coded in C++.

Mark Baker reported that several people asked him about Web benchmarks. A proposal will be there in the next meeting. A date for the next meeting should be in spring 1997 on a Friday. There is an agreement to have the next meeting after April 5th but not during the HPCN and IPPS.

A3.6 PARKBENCH Meetings in 1997

A3.6.1 University of Westminster (Vladimir Getov), PARKBENCH meeting, 9 May 97, Knoxville, Tennessee

According to the tentative agenda for the last PARKBENCH Meeting, which was distributed at the beginning of March, I was in charge of two items; the low-level communication benchmarks (with Erich Strochmaier – UTK and Ron Sercely - HP/CXTC) and the new Java low-level benchmarks.

The low-level benchmarks have always been a strong European area within the International PARKBENCH Committee. In fact, the origins of the current low-level PARKBENCHmarks come from two European initiatives - Euroben and Genesis.

At this meeting my task was to propose further developments of the low-level communication benchmarks, as a result of the recent advances in high-performance computer architectures and in particular the communication subsystems. A number of pilot improvements to the COMMS1 code had already been completed by Roger Hockney after e-mail discussions with Ron Sercely, Charles Grassl and myself. We had also taken time measurements on 8 different kinds of state-of-the-art platforms. After my report and a discussion at the beginning of the meeting, the revised version of the COMMS1 benchmark was accepted in principal by the committee (see the enclosed minutes). The discussion also resulted in a decision for further changes in order to improve the internal organisation of the low-level benchmarks (split the measurements code from the performance analysis code) and the reporting of results. Erich Strohmaier and I will be coordinating the implementation of these changes.

During the last academic year we have been developing preliminary Java versions of some of the PARKBENCH benchmarks in the Performance Engineering group at the University of Westminster. In particular, a Java-to-C interface generator has been developed in order to ease the access of existing libraries from Java. At the meeting I presented the first experiments with our MPI wrapper in Java including performance results on the 16-node IBM SP-2 platform in Southampton. The reaction of the committee members was very useful and important for the proposal of Java benchmarks to be presented at one of the next PARKBENCH meetings.

In addition, I was also actively involved in the discussions on

- PEMCS electronics journal (as an associate editor)
- Future PARKBENCH publications and reports
- Status of funding proposals

A3.6.2 University of Westminster (Vladimir Getov), Supercomputing’97 Conference, San Jose, 15-21 November, 1997, PARKBENCH BOF Session

As usual, the last PARKBENCH Meeting for 1997 was organised during the Supercomputing'97 Conference in San Jose in the form of a BOF session on 19 November 1997 (see enclosed agenda). The session was chaired by Professor Tony Hey. I reported shortly about the current status and availability of the PARKBENCH Curve Fitting Tool as part of the Low-Level Performance Evaluation Toolset). This tool developed at the University of Westminster by Roger Hockney is now publicly available on the Web and is being used for analysis and interpretation of supercomputer performance results.

I was also actively involved in the discussions on

- PEMCS electronics journal (as an associate editor)
- Current and Future PARKBENCH releases, as well as the status of funding proposals

In addition to the PARKBENCH meeting, I presented our research results in High-Performance Parallel Programming in Java at the research posters exhibition at SC'97. This piece of research introduces a way of successfully tackling the difficulties in binding both scientific and message-passing libraries to Java. We have created a tool (JCI) for automating the creation of portable interfaces to native existing libraries in Fortran-77 and C. Evaluation results on the IBM SP2 at Cornell were presented using the IBM HPCJava compiler, which generates a native code for the RS6000 architecture.

Agenda of the PARKBENCH BOF

- Introduction, background, WWW-Server
- Current Release of PARKBENCH
- Low Level Performance Evaluation Tools
- LinAlg Kernel Benchmarks
- NAS Parallel Benchmarks, including latest results
- Plans for the next Release
- Electronic Journal of Performance Evaluation and Modelling for Computer Systems
- Questions from the floor / discussion

A3.6.3 University of Portsmouth (Mark Baker), Supercomputing'97 Conference, San Jose, 15-21 November, 1997, PARKBENCH BOF Session

Summary of Objectives for Attending SC'97 in San Jose, Ca.:

My general objective was to attend a number of technical and panel sessions, look around the vendor exhibition, as well as meet up with EU and US colleagues to talk about current and potential future collaborations.

I attended the PARKBENCH Birds of a Feather (BOF) meeting (http://www.supercomp.org/sc97/bofs.html#W, Wednesday, November 19 - 5:30 PM) and presented a PEMCS status talk and gave a call for papers - see - http://www.sis.port.ac.uk/~mab/Talks/SC97/PARKBENCH/

I attended the Seamless Computing BOF meeting, (http://www.supercomp.org/sc97/bofs.html#Tu

Tuesday, November 18 - 5:30 PM) and presented a talk, and gave a call for papers - see -

http://www.sis.port.ac.uk/~mab/Talks/SC97/Seamless/CFP/

http://www.sis.port.ac.uk/~mab/Talks/SC97/Seamless/Corba-Framework/

I attended the High Performance Computing on Windows NT BOF meeting (http://www.supercomp.org/sc97/bofs.html#W - Wednesday, November 19 - 5:30 PM)

- see - http://www.sis.port.ac.uk/~mab/Papers/PC-NOW/

I met up with former CRPC colleagues (http://www.crpc.rice.edu) and discussed a number of on-going projects and potential collaborations (NHSE + PARKBENCH), for example see -

http://www.sis.port.ac.uk/~mab/Talks/NHSEReview97/

I met up with former NPAC colleagues (http://www.npac.syr.edu) and discussed a number of on-going projects and potential collaborations, see http://www.sis.port.ac.uk/~mab/Papers/PC-NOW/ for example.

Appendix 4: The PARKBENCH COMMS1 Benchmark

The COMMS1, or pingpong, benchmark measures the basic communication properties of a message-passing MIMD computer. A message of variable length, n, is sent from a master node to a slave node. The slave node receives the message into a Fortran data array, and immediately returns it to the master. Half the time for this message pingpong is recorded as the time, T, to send a message of length, N. The time as a function of message length is fitted by least squares using the parameters RINF and NHALF to the following linear timing model:

T = (N+NHALF)/RINF

when the communication rate (bandwidth) is given by,

R=RINF/(1+NHALF/N) = PI0*N/(1+N/NHALF)

and the startup time (latency) is

T0=NHALF/RINF=1/PIO

PI0 is known as the specific performance. In general, we may say that RINF characterizes the long-message performance and PI0 the short-message performance. The COMMS1 benchmark computes all four of the above parameters, RINF, NHALF, T0, and PI0, because each emphasizes a different aspect of performance. However only two of them are independent. In the case that there are different modes of transmission for messages shorter or longer than a certain length, the benchmark can read in this breakpoint and perform a separate least-squares fit for the two regions. An example is the Intel iPSC/860 which has a different message protocol for messages shorter than and longer than 100 byte.

Appendix 5: Features of the PICT Tool and the New 3-Parameter COMMS Model

A5.1 Features of the PICT Tool

The PICT tool provides the following features:

1. Automatic plotting of Low-Level PARKBENCH output files from a URL anywhere on the Web (At present limited to New COMMS1 and Raw data, but easily extended to original COMMS1 and RINF1). This is useful for a quick comparison of raw data.

2. Automatic plotting of both 2 and 3-parameter (see below) curve-fits which are produced by the benchmarks. Good for checking the quality of the fits.

3. Allows manual rescaling of the graph range to suit the data, either by typing in the required range values or by dragging out a range box with the mouse.

4. Allows the 2-parameter and 3-parameter performance curves to be manually moved about the graph in order to fine tune the fits. The curve follows the mouse and the RMS and MAX percentage errors are shown as the curve moves. Alternatively parameter values can be typed in and the Manual button pressed when the curve for these values will be plotted.

5. The data file being plotted can be VIEWed and a HELP button provides a description of the action of each button in a separate windows.

A5.2 Three-Parameter Performance Model

The two-parameter (rinf,nhalf) characterisation of performance is satisfactory in any situation where the time is approximately linear with respect to some other variable such as message or vector length. A least-squares 2-parameter fit is calculated by the COMMS1 benchmark whilst the benchmark measurements are being made, and the results are read by the PICT tool from the COMMS1 output file. The PICT tool can then be used to view the quality of the fit.

However some recent experimental COMMS1 data from communication subsystems using libraries of different message-passing protocols shows a variation of bandwidth with message length that does not fit this model well.

A generalisation of the linear model is therefore required to cope with this new situation. The *variable-power model* introduces a third parameter to add some extra flexibility with the least increase in complexity. It also reduces to the existing linear form when the new third parameter is unity.

The variable-power model represents performance with the three parameters (rivp, navp, gamvp) and the equation

r = rivp / [1+ (navp / n)^gamvp] ^(1 / gamvp)

This model becomes the same as the linear two-parameter model when gamvp=1 in which case rivp becomes rinf and navp becomes nhalf.

A 3-Point variable-power fit is calculated by the new COMMS1 benchmark, and the results are read by the PICT tool from the new COMMS1 output file. The PICT tool can then be used to view the quality of the fit.

Details of the new model were presented at the HPC-Standards funded PARKBENCH workshop. The following paper is available at the PEMCS Journal Web site:

"PARKBENCH Low-Level Performance Profiles, Computational Similarity, and a New 3-Parameter Performance Model", Roger W. Hockney, School of Computer Science, Westminster University, UK, November 1997:

http://hpc-journals.ecs.soton.ac.uk/Workshops/PEMCS/fall-97/talks/Roger-Hockney/perfprof1.html