Early Applications in the Message-Passing Interface (MPI)*

Anthony Skjellum‡
Department of Computer Science &
NSF Engineering Research Center for Computational Field Simulation
Mississippi State University
PO Box 6176
Mississippi State, MS 39762

Ewing Lusk, William Gropp§
Argonne National Laboratory
Mathematics and Computer Science Division
9700 South Cass Avenue
Argonne, IL 60439

June 15, 1994
Revised: December 14, 1994

*Accepted to Special Issue of *International Journal of Supercomputing Applications*. Revision 1.1; December 14, 1994.

†Work supported in part by the NSF Engineering Research Center for Computational Field Simulation, Mississippi State University.

‡Corresponding author; e-mail: *tony@cs.msstate.edu*; phone: 601-325-8435.

§Work supported by U.S. Department of Energy under contract W-31-109-Eng-38.

Abstract

We describe a number of early efforts to make use of the Message Passing Interface (MPI) standard in applications, based on an informal survey conducted in May-June, 1994. Rather than a definitive statement of all MPI development work, this paper addresses initial successes, progress, and impressions that application developers have with MPI, according to the responses received. We summarize the important aspects of each survey response, and draw conclusions about the spread of MPI into applications.

An understanding of message-passing, and access to the MPI standard are prerequisites for appreciating this paper. Some background material is provided to ease this requirement.

1 Introduction

In this paper, we describe a number of early efforts to make use of the Message Passing Interface (MPI) standard in real applications (Forum 1994a; Forum 1994b). An informal survey of efforts is reported here, together with our commentary. We summarize the responses, highlighting important facets of each contributed application survey, and describe an overall impression of the state of MPI applications.

1.1 A Brief History of MPI

MPI was developed as part of a multi-national effort (supported partially by NSF, ARPA, and ESPRIT) to replace the many, marginally compatible vendor and portability systems currently in use with a coherent system that exploits the experience of these previous message-passing systems while providing greater functionality, unified design, and an emphasis on user access to high performance protocols and hardware. Participants included a core of 40 or so forum members, including developers of all the major then existing message-passing systems (*e.g.*, P4, PVM, NX), as well as computer vendors, and application developers. MPI encompasses and extrapolates the best ideas from existing practice, and adds features designed to secure high performance for applications on the hardware that exists today, and for hardware expected in the next few years.

The initial round of that MPI standardization effort began in April 1992 and culminated at Supercomputing '93 (November, 1993), where the standard was presented to the public[1]. By the time of its public introduction, several implementations were already underway, notably the Argonne/Mississippi State "model implementation" (also known as MPICH) (Doss, Gropp, Lusk, and Skjellum 1993), which was also demonstrated at Supercomputing '93; other portable implementations have followed (*e.g.*, LAM (Burns, Daoud, and Vaigl 1994), Chimp MPI (Edinburgh Parallel Computer Centre)). During the first half of 1994, IBM

[1]However, the standard document was finalized only in May, 1994, with on-going errata, so the existence of several applications at the writing of this paper in June, 1994 is quite significant.

launched a commercial MPI effort (*e.g.*, (Franke, Hochschild, Pattnaik, and Snir 1993)); IBM demonstrated its experimental "MPI-F" at Supercomputing '93 as well. Cray Research, through an alliance with Edinburgh Parallel Computing Centre, are working to support MPI on the Cray T3D. nCUBE has made MPI the native message passing system for its next-generation systems. Intel, Thinking Machines, Meiko, and others also have implementation efforts or plans; the MPICH and other portable implementations run on virtually all the machines for which commercial versions remain to be created.

1.2 Some Features of MPI Relevant to Applications

In order to make what follows more accessible, it is helpful to describe certain features of MPI. The first of these is the dual concept of process groups and communicators. Here we draw on (Forum 1994a; Forum 1994b).

Groups and Communicators MPI designates communicating processes with communicators (intra-communicators and inter-communicators). An intra-communicator consists of a group of processes, described first, and "context" described thereafter. An MPI program currently gets "`MPI_COMM_WORLD`" when it starts executing; this communicator defines the initial group and communication space for processes to use. Other communication relationships can then be created during the course of program execution.

A group is an *ordered* collection of processes. Processes in a group therefore have a unique integer rank (starting from zero). Once created, a group never changes size (hence, we sometimes call MPI groups "static groups.") A process group provides the rank property for send-and-receive communication and indicates which processes participate for collective communication (synchronization, global sum, etc). MPI's concept of designating groups as boundaries for communication is extremely important to its ability to scope messages. Scoping of messages allows any parallel code to insulate itself from user code or other libraries in the same application.

Intra-communicators include a notion of context of messages. A context keeps messages sent by a process

from arriving in the wrong scope of a receiving process. One can conceive of planes of message-passing, with processes belonging to as many hypothetical planes as makes sense for application requirements, but any message is restricted to a single plane. This naming convention leads to simple scoping rules for parallel libraries with simple-to-use techniques for assuring that a library has safe communication space (lexical scoping of messages). Finally, communicators guarantee that pending sends and receives cannot foul up a collective communication operation. This further simplifies parallel programming. Intra-communicators are the dominant programming idiom at present, so the word "communicator" is often used loosely to refer to them, though inter-communicators are just as interesting and useful.

Specifically, an inter-communicator involves a "local" and a "remote" group, and appropriate context information, and is consequently a simple generalization of the intra-communicator. The intra-communicator allows a local set of processes to describe point-to-point communication to a remote group, and vice versa. MPI emphasizes point-to-point communication for inter-communicators, though extensions are possible (Skjellum, Doss, and Viswanathan 1994).

Virtual Topologies Virtual topologies allow the application programmer to attach numerical names other than group ranks to processes, so that the domain decomposition or other application-relevant mapping of data to processes is intuitive and convenient. MPI provides such mappings for the convenience of the application programmer; both cartesian grids and general graph mappings are supported. As examples, cartesian grids of 1-, 2-, and 3-dimensions are quite useful for linear algebra codes and solution of PDEs. Graph topologies might be useful for solution of problems with sparse connectivity or for combinatorial problems (see (Gropp, Lusk, and Skjellum 1994) for more details).

Datatypes MPI provides basic datatypes as part of communication calls. In the simplest case, the first-class types of each language are available as datatypes to send, receive, and collective communication calls. MPI programmers send arrays of types and do not normally send arrays of bytes, unless they choose to do

so. These datatypes encapsulate any heterogeneous conversion needed. Hence, an MPI programmer does not worry about XDR conversion, as this is part of the user program (in fact, something faster than conventional XDR is likely in real implementations).

In addition to pre-defined datatypes, MPI supports aggregate or derived datatypes (ideally defined once and used repeatedly). These allow arbitrary gather-scatter patterns to be created. In this mode, the application programmer does not explicitly "buffer up" data into a contiguous area of memory prior to calling MPI send, or must the programmer scatter a contiguous buffer on the receive end, unless they prefer to work in this pattern. Such a strategy can reduce the number of copies needed to effect data transfer. As such, it both simplifies and possibly speeds up message passing.

1.3 Applications Represented

Application programmers were informed of the opportunity to contribute to this paper, through an informal survey mechanism published on popular mailing lists, and on USENet groups. The responses described here represent all but one of the respondents during the period alloted for the survey (one response was too preliminary to contain any reasonable information about the work undertaken. Responses came in the following categories:

- Applications:
  - Many-body Methods for Light Nuclei,
  - Solution of Unsteady Incompressible Viscous Flows,
  - NAMMU — Groundwater Modeling for Continuous Media,
  - PARMICS-MP — Microscopic Traffic Simulation,
  - Volume Visualization,
  - Reservoir Simulation (not included).

- Application-Enabling Libraries:
  - A Fast Particle Summation Library,
  - PRISM Eigensolver Library,
  - CVL — C Vector Library,
  - "ParX" (Parallel X environment).

1.4 Organization

This paper is organized as follows. First, we outline the survey itself. Then, each of the surveyed applications is described; for clarity, we have not followed the format of the survey in reporting the application responses. Subsequent to this discussion is an interpretation of the responses to the survey, followed by the summary, conclusions, and future directions section.

2 The Survey

In order to collect data from the widest possible source of application programmers, but consistent with a short deadline, we posted requests for application work on major USENet newsgroups (*comp.parallel*, *comp.parallel.pvm*), as well as to the MPI "reflectors" maintained by University of Tennessee (*mpi-comm@cs.utk.edu*) and Argonne National Laboratory (*mpi-users@mcs.anl.gov*). Responses were solicited for applications underway, completed, or even just getting started.

The following information was requested for contributions to the survey: a description of the MPI-based application, researchers actively involved, and experiences thus far. In addition, we requested a description of the application area, goals of application project (high speed, cheap cycles, new physics, and so on), and proper citations to the project, and techniques upon which it is based.

We asked respondents to provide anectodal information about process of converting from other message-passing systems (if applicable), including their level of effort, motivation for conversion, and degree of satisfaction. We asked them to determine which features of MPI have proven most helpful in applications development and/or conversion, such as

- intra-communicators (single-process-group communication),
- inter-communicators (two-process-group communication),
- virtual topologies (naming conventions for processes),
- pack/unpack capabilities (for PVM compatibility),
- datatypes (gather-scatter specifications).

We also asked respondents to describe which parts of MPI are most understandable and/or most confusing. We asked them to show early indications of performance results, while indicating which MPI implementation or implementations they have used in conjunction with the applications. Finally, respondents were invited to provide miscellaneous feedback from application programmers and users, if relevant. For instance, we asked them to discuss whether revised and/or additional features were needed in MPI.

3 Applications

In this section, we describe five responses to the survey, each of which is a pure application project.

3.1 Many-body Methods for Light Nuclei

Researchers (Steven Pieper and Robert Wiringa[2]) in Argonne's Physics Division, in collaboration with a colleague at the University of Illinois at Urbana-Champaign, are computing the properties of light (up to 40 neutrons and protons) nuclei using realistic two- and three-nucleon interactions. This research involves developing many-body methods for reliably computing the properties of a nucleus for complicated forces that are strongly dependent on the spins and charge states of the nucleons.

Unlike the Coulomb force used in atomic or condensed-matter calculations, no useful fundamental theory tells us what this force is. The two-body force can be partially constrained by fitting nucleon-nucleon scattering data, but many-body calculations are required to test other properties of this force as well as the three-body interaction. Thus the researchers are refining their knowledge of the forces and at the same time using this knowledge to make predictions about nuclei.

3.1.1 Goals

The goal of the application project is to reveal new physics: that is, to compute binding energies more accurately than ever before. New calculations being done on the IBM SP1 parallel computer at Argonne are using a new nuclear interaction and are obtaining much better (when compared with experiment) results for the binding energy and density profile of oxygen than had previously been obtained. In particular, physicists are now able to find parameters of the minimal energy solution for sixteen nucleons much more precisely than ever before. Before this computation, theory did not clearly predict that the nucleus of an oxygen atom would remain stable; the previous computational results would have allowed it to disintegrate into four helium atoms. Work on calcium (forty nucleons), which could not be contemplated before, is now also in progress.

---

[2]Theoretical Physics Group Physics Division, Argonne National Laboratory, Argonne IL 60439. e-mail: *pieper@phy.anl.gov, wiringa@phy.anl.gov.*

3.1.2 Details

This program was originally written for portable parallel execution using p4; it evidently took one afternoon to convert it to MPI. The authors' motivation was to learn about MPI and prepare the code for continued portability on high-performance machines as further MPI implementations become available. The parallel algorithm used was straightforward as no advanced features of MPI were needed. Understandably, the most helpful aspect of MPI was that it is similar to existing message-passing systems, rendering porting straightforward as well.

This application used the portable model MPI implementation developed at Argonne National Laboratory and Mississippi State University. One feature that the p4 implementation provided that was unavailable in MPI is a form of broadcast that is received by an ordinary receive instead of a corresponding broadcast operation; this was missed by the developers. Table 1 contains information about the size and performance of this application on the 128-node SP-1 at Argonne.

3.2 Solution of Unsteady Incompressible Viscous Flows

The computation of flow regimes over complex configurations involves the numerical solution of a system of coupled non-linear partial differential equations, the Navier-Stokes equations. Researchers in the Computational Fluid Dynamics Laboratory at the Mississippi State University-NSF Engineering Research Center for Computational Field Simulation have developed an implicit, finite-volume code (known as UNCLE) for solving the unsteady three-dimensional incompressible Euler and Navier-Stokes equations using dynamic multiblock grids (Taylor, Busby, Jiang, Arabshahi, Sreenivas, and Whitfield 1993; Taylor and Whitfield 1991; Whitfield and Taylor 1991). The flow solver can be used in a variety of applications ranging from maneuvering underwater vehicles to centrifugal compressors. Ramesh Pankajakshan[3] and W. Roger Briley,

[3]Contact address: MSU ERC, PO Box 6176, Mississippi State, MS 39762. e-mail: *ramesh@erc.msstate.edu*.

also of the MSU ERC, have undertaken the scalable parallelization research and development of this solver.

Long runtimes and large memory requirements restrict the size and complexity of the problems that can be handled using the sequential version of the UNCLE code. Therefore, a scalable portable parallel version is being developed that can take advantage of existing as well as emerging parallel platforms. The goal is to use parallel computing to enable solution of large-scale field simulation problems on both parallel supercomputers and idle and/or dedicated workstation clusters. The message-passing interface required for the parallel implementation has to support collective operations within user-defined groups as well as provide safe communication contexts for overlapping sets of collective operations. Motivations for using MPI include these features, as well as portability, ease of software development, high performance expectations, and broad base of participants in MPI effort to guaranteed continued access and evolution as hardware rogresses.

3.2.1 Details

The parallel formulation (Pankajakshan and Briley 1994) employs spatial decomposition of the overall grid into sub-blocks assigned to separate processes, and exploits coarse-grained parallelism and message passing within sub-iterations of an implicit time-dependent solution algorithm. The solution at points shared by neighboring processes is updated between each sub-iteration by means of a message exchange.

Key areas of the parallel implementation include

- initialization of the flow field,
- duplication of stored data for points near block interfaces,
- exchange of data during sub-iterations, for points having duplicated storage, and
- treatment of line searches and integrals along coordinates emanating from solid boundaries, which arise from the particular algebraic turbulence model used.

In the parallel implementation of this code (Pankajakshan and Briley 1994), the domain is partitioned into a number of nearly equally sized sub-domains, each of which is assigned to a node of the multicomputer. The local data dependencies at the boundary of each block are taken into account by a two-cell-deep layer of buffer cells whose values are updated from the appropriate neighboring block. The data duplication and updating at the block boundaries is implemented using the `MPI_Sendrecv` routine. The connectivity of the processes gives rise to a Cartesian virtual topology having empty nodes in it. Each process is tagged using an ordered triplet that represents its coordinate on the virtual cartesian grid (topology). These coordinates are then used to define communicators for processes that are oriented along any of the three axes of the grid. This involves repeated use of the `MPI_Comm_split` routine. The line searches and integrals of the turbulence model are implemented using the `MPI_Bcast` routine and user-defined `MPI_Allreduce` operations over subsets of processes.

Work has progressed to the point of a working (but as yet preliminary) MPI application. Constructing this MPI code from an existing serial FORTRAN code took about three weeks by one of the researchers (Pankajakshan), who was familiar with the existing serial code and who also had a working knowledge of MPI and C programming at the outset. From their experience so far, the researchers find that MPI appears well suited for this application. MPI was especially useful in performing collective operations on subsets of processes, a need which arises in connection with line searches and integrals in the algebraic turbulence model. Communicators and datatypes are features of MPI which have been or will be helpful in developing this application. However, datatypes were seen as the most confusing aspect of MPI.

The parallel code has used both the IBM's MPI-F and Argonne/Mississippi State implementations of MPI. The code has run on the IBM-SP1 at the High Performance Computing Research Facility, Mathematics and Computer Science Division, Argonne National Laboratory, and on an MSU Sun workstation/ethernet cluster.

3.3 NAMMU — Groundwater Modeling

NAMMU is a package for predicting groundwater flow, heat, and mass transport through continuous porous media. The application is based on an efficient implementation of the finite-element method, which in turns makes heavy use of the Harwell Frontal Direct matrix solver routine 'MA42'. The work described here is a parallelization of the solver through domain decomposition of the mesh. NAMMU was developed by AEA Technology, modification of the MA42 package was done by Harwell, and the message passing (in MPI) was done by the Edinburgh Parallel Computing Centre. The authors of NAMMU are as follows: Hon W. Yau (Edinburgh Parallel Computing Centre)[4], K. Andrew Cliffe (AEA Technology, Harwell), Jennifer A. Scott (Atlas Centre, Harwell), David Brear (AEA Technology, Harwell). Hon Yau is the designer (and one of the implementors) of the parallelized version of NAMMU.

The goal of this work is to produce a portable, parallelized version of NAMMU that will best exploit parallel architectures. Together with two other AEA Technology safety assessment codes, this will form a suite of parallelized groundwater simulation codes.

3.3.1 Details

The conversion effort was found to be considerably eased by the availability of a high quality in-house implementation of the MPI specifications. Moreover, a 'toy' MPI implementation of the solver using MA42 had already been successfully demonstrated prior to work on the NAMMU package. From the design of the NAMMU implementation to the successful execution of the available test cases, took approximately three months. This included roughly 20 person-days for the conversion of NAMMU to use domain decomposition, and one person-month to write and debug the MPI message-passing constructs needed.

The authors were motivated to convert to MPI because they wanted a portable implementation that

[4]Contact address: Edinburgh Parallel Computing Centre (EPCC), James Clerk Maxwell Building, King's Buildings, Mayfield Road, Edinburgh EH9–3JZ, U.K. e-mail: *hwyau@epcc.ed.ac.uk.*

would demonstrate good scalability within the limitations of the parallel algorithm. The authors reported a high degree of satisfaction with the current MPI specifications, in particular with the richness of features. They noted that it seemed a thoroughly thought-out product. In this regard, only a duplicate of the global group of processes (represented in `MPI_COMM_WORLD`) was used in the application, since the intension was to use one or more of the collective communications functions. Derived datatypes were heavily used, as an efficient alternative to multiple messages and packing/unpacking; this feature had much to be commended.

The most obvious source of problem is with the Specifications Document (Forum 1994a). As the only currently available introduction to MPI it is seen by the developers as daunting — in particular to those with no prior message passing experience. It is, however, also seen as reasonably complete, though in need of more Fortran examples and a glossary of terms.

An in-house EPCC implementation of MPI, based on CHIMP, was used. This version of MPI runs on a variety of machines, including Sun and Silicon Graphics workstations, and the Meiko CS-1/i860.

3.4 PARMICS-MP — Microscopic Traffic Simulation

PARAMICS-MP builds on earlier work on the PARAMICS project, which developed a PARAllel MICroscopic traffic Simulator. The PARAMICS system, in its current state, can simulate over 250,000 vehicles at speeds faster than real time, on a road network based on real data for the Scottish Trunk Road Network. As far as the researchers know, this is the largest microscopic traffic simulation ever implemented. PARMICS-MP is being developed by Gordon Cameron, Gordon Duncan[5], David McArthur, of the Edinburgh Parallel Computing Centre.

The primary goal for PARAMICS-MP is to port the existing PARAMICS simulator, originally developed as a data-parallel implementation for the Connection Machine CM-200, to a Message-Passing version running on the recently purchased 256-processor Cray T3D located in Edinburgh. The "MP" version will use the

[5]Contact address: e-mail: *gd@epcc.ed.ac.uk*; phone: +44 (0)31 650 5818.

MPI library currently under development at EPCC.

It is expected that the T3D's superior computational power will enable the integration to the simulator of a significantly more detailed vehicle dynamics model than is currently possible. In addition, more advanced simulation features will be added, such as the effects of dynamic routing as proposed by reactive Road Traffic Information systems.

The current status of the project is that a prototype version of the MP simulator is running on a distributed system, using EPCC's own message passing system CHIMP. It is intended that as soon as the MPI implementation being developed locally for the T3D becomes available the application will be converted to use it. *This move is part of an EPCC organization-wide policy of migration to the MPI standard in all forthcoming applications.*

3.5 Volume Visualization

Volume visualization techniques were created to allow researchers to examine the data and spatial relationships contained in volume data sets, but, in general, the algorithms are relatively slow on even moderately sized data sets. This problem is compounded as the size of the data increases, and additional problems with memory size may be encountered in generating a visualization for large data sets on a conventional machine. It is thought that some of these problems can be addressed by distributed memory parallel computers, by allowing the parallel machine to perform the calculation of the scene, and employing conventional machines to display that scene. Parallel computers have added potential over conventional machines for scene rendering in that a large data volume can be distributed over the processors in the machine and operated on concurrently, reducing at once the computational time and and memory requirement per processor (scalability) (Lorensen and Cline 1987; Westover 1990; West, Stephens, and Turcotte 1994b; West, Stephens, and Turcotte 1994a).

John E. West[6], Michael Stephens, and Louis Turcotte, of the Waterways Experiment Station (WES) are investigating the application of MPI to the development of algorithms for performing volume visualization of large data sets on parallel machines. The goal is to create preliminary versions of volume visualization algorithms for parallel machines both to investigate the use of a specific machine, the nCUBE/2 at WES, for this type of application, and to gain insight into the issues surrounding parallel visualization in general. The algorithms under examination represent the two general classes of algorithms available for extracting information out of 3-D data sets: marching cubes, which is prototypical of surface-extraction algorithms, and splatting, a volume-rendering technique.

It is expected that after the preliminary study is completed, some or all of the algorithms will be ported to other parallel machines and evaluated for hardware-specific performance features. Using MPI is expected to reduce the cost of this software port dramatically as well as simplify code development.

3.5.1 Details

Conversion issues do not apply to this project, since the code is being created from scratch using MPI. According to the researchers, this represents a distinct advantage for them, in that, not only is the initial development time expected to be decreased, but the time for subsequent ports of the code to different platforms will also be reduced.

Thus far, the features of MPI that have proven most helpful are the communicators and topologies. Having developed codes for varying platforms in the past using native message passing constructs (NX, Vertex, etc), these researchers attest to the fact that the presence of the communicators and utility functions that are available for their manipulation greatly simplify the development of a communication protocol within an application. The topology functions have also been useful in performing the data distribution. They support the creation of a process topology that matches the original data topology in only a few simple steps. This

[6]This work is being performed at the DOD High Performance Computing Center at the Army Corps of Engineers Waterways Experiment Station. e-mail: *hatter@athena.wes.army.mil*; phone: 601-634-4610.

further simplifies the data distribution and communication problems typically encountered in other message passing protocols. It is expected that datatypes, and in particular strided datatypes, will also prove especially useful in the future.

4 Application-Enabling Libraries

In this section, we describe application-related responses to the survey. Each of these responses represents a project that is either part of an application, or enables a class of applications.

4.1 A Fast Particle Summation Library

A general-purpose "Fast Particle Summation" library that is, in turn used, in several applications[7], is under development by John K. Salmon[8] of Caltech, Michael S. Warren[9] of Los Alamos, and David J. Edelsohn[10] of Syracuse University. The code itself and it's original application in gravitational N-body simulations has been described in (Salmon, Winckelmans, and Warren 1994; Warren and Salmon 1993; Warren and Salmon 1994).

4.1.1 Details

The application is currently written to the researchers' own set of message-passing primitives, (*e.g.*, asynchronous send, asynchronous receive, exchange, etc). These primitives are by design extremely simple. They

[7] A description of an application solving potential flow problems and another involving generation of Gaussian random fields is in preparation. All the relevant papers, including some with an emphasis on astrophysics, rather than computation, are available by anonymous ftp in *ftp://ftp.ccsf.caltech.edu/nbody*, or on Mosaic (HTML), as *http://www.ccsf.caltech.edu/johns/papers.html.* Also, papers describing the evolution and interaction of disk galaxies by David Edelsohn (NPAC, Syracuse University) and Bruce Elmegreen (IBM Research) are in preparation.

[8] Physics Department, California Institute of Technology, *johns@ccsf.caltech.edu,* phone: +1 818 395 2907.

[9] Theoretical Astrophysics, Los Alamos National Laboratory, *msw@tdo-serv.lanl.gov*, phone: +1 505 665 5023.

[10] Northeast Parallel Architectures Center, Syracuse University, *edelsohn@npac.syr.edu*, phone: +1 914 945 3011.

are implemented using the message-passing programming interface of each system. Because the primitives previously had been ported to the IBM MPL (EUI), converting to MPI was greatly simplified, in view of the similarities between the two message-passing definitions. The most difficult element is adapting to system peculiarities for querying about the status of outstanding operations and statistics about received messages. The researchers ported to MPI because it was simple to accomplish, expanded their target environments, and improved flexibility and maintenance for future systems on which this application will run. At the moment the authors only use the lowest level of MPI functionality. Using inter-communicators for interactive, real-time control and display of simulation results is a long-term goal for them.

Of the features that the authors wanted to use, the multiple completion functions are much less flexible than they desired. They would have preferred functions that act upon dynamic sets or collections of communication objects which can be updated efficiently and then used by a generalized POSIX "select()" function to allow rapid discovery of completion of one of many outstanding communication operations, because they may have hundreds of such outstanding operations. Their conclusion is that the current MPI array semantics does not allow for an efficient implementation of the functionality they desire for their applications and limits usefulness to applications with a relatively short and mostly static list of communication handles.[11]

The IBM MPI-F prototype implementation running on an IBM SP1 with the High Performance Switch achieves a latency of 30 microseconds and a bandwidth of 8.7 MBytes per second. This implementation is an early research prototype developed within IBM Research that conforms to and implements the entire, final MPI document (Forum 1994a). Comparing the application running with MPI-F and MPL/p (EUI-H) on the IBM SP1, they observed no significant change in performance (Franke, Hochschild, Pattnaik, and Snir 1993).

[11] In our experience, many applications do not have more than a handful of outstanding communication operations pending, unlike the algorithms considered by these researchers. Alternative functionality proposed by these researchers near the end of the MPI standardization process, was consequently rejected as an alternative. Had it been proposed as a supplement to the easy-to-use arrays of status objects, rather than a replacement, its inclusion would have been much more likely. Undoubtedly, we will revisit this issue for MPI2.

4.2 CVL - C Vector Library

CVL (C Vector Library) is a library of low-level vector routines callable from C. It is used as a basis for implementing NESL, a portable nested data-parallel language (Blelloch, Chatterjee, Hardwick, Sipelstein, and Zagha 1994) and also by the Proteus project at UNC to implement the parallel prototyping language Proteus. The library includes a wide variety of vector operations such as elementwise function applications, scans, reductions and permutations. Most CVL routines are defined for both segmented and unsegmented vectors, since segmentation is critical for implementing nested data parallelism. Previously, CVL has been ported to each new parallel machine architecture (Cray, CM-2, CM-5, MasPar, etc), using the manufacturer's own low-level communication routines. The main goal of this project is the development of a portable, interactive environment for programming a wide range of supercomputers, using a very high-level data-parallel language that supports nested parallelism. We are especially interested in the implementation of algorithms with irregular or dynamic structures.

The MPI port of CVL took approximately two person months. It benefited from using algorithms for a distributed memory architecture initially developed for the CM-5 port, although these had to be adapated to use coarse-grain rather than fine-grain communication. The overwhelming motivation for writing an MPI port of CVL was the prospect of not having to do further ports in the foreseeable future. That is, the effort required to port to future machines would be reduced to tuning the MPI port to best fit a particular manufacturer's implementation of MPI. The developer has been satisfied with MPI's portability and ease of use; of all the CVL ports to MPP architectures, he reports that this was the easiest.

The developer notes that the most helpful features of MPI for the CVL implementation are as follows:

- Built-in support for scans and reductions, although this is not as complete as it could have been,
- Non-blocking sends and receives to overlap communication and computation,
- Choice of send semantics with respect to blocking and buffering; this makes it easy to get a quick-

and-dirty version running, which can then be refined to increase performance.

Complete performance results are given in (Hardwick 1994), MPI CVL comes close to the performance of a machine-specific CVL implementation on the CM-5, the only platform on which the developers can currently perform a comparison. Furthermore, the developer noted that better support in MPI for fine-grained communication (e.g., put/get, active messages, or dynamic all-to-all communication) would simplify the code for MPI CVL and could also improve its performance. He also noted that it would be helpful if MPI could support exclusive scans as well as the inclusive variety: it is trivial to convert an exclusive scan into an inclusive scan, but the reverse involves extra communication for operations with no inverse (*e.g.*, max, min). Segmented scans and reductions, as supplied by High Performance Fortran, would also be an advantage, although they can be implemented together using communicators or user-defined scan operators in MPI.

4.3 PRISM — Scalable Eigensolvers

The PRISM (Bischof, Huss-Lederman, Sun, Tsao, and Turnbull 1994) project is investigating scalable eigensolvers for parallel distributed memory computers. The goal is to produce a public domain library that runs on a variety of platforms. To achieve this goal the project has looked at various issues including linear algebra kernels such as matrix multiplication and matrix factorization, data layout and communication optimization. Christian Bischof[12], Steven Huss-Lederman[13], Xiaobai Sun[14], Anna Tsao[15], Thomas Turnbull are working on PRISM. The PRISM project is aimed at producing a library to solve eigenproblems on a variety of distributed memory computers. This library will be used by a variety of groups, such as computational chemists, to supply more computational power to the solution of their problems.

[12] Argonne National Laboratory
[13] Supercomputing Research Center, Lanham, MD.
[14] Duke University, Department of Computer Science
[15] Supercomputing Research Center, Lanham, MD.

4.3.1 Details

The developers found the effort of converting to MPI to be modest. A simple port was found to be reasonably easy. A slightly harder task has been taking advantage of features available in MPI to which they did not previously have access. The PRISM project wants to produce a package that runs efficiently on a large number of parallel systems. They did not want to be tied to a specific vendor or national laboratory for message passing software.

The developers have used communicators to allow our library to work in a safe communication space. They also use topologies to define a logical 2D grid/torus that they work on as well as use row and column subgroups for collective communications. They have only made limited use of derived datatypes, because proper use of this feature requires modest rewriting of these sections of the code. However, they anticipate using derived datatypes to a greater extent in the future. The developers also plan to utilize persistent communication handles since they repetitively do operations such as shifting to the next processor.

The developers have found the performance of MPI to be satisfactory. They are particularily interested in collective operations. They have found that the performance of MPI has been improving with time. Details can be found in (Bischof, Huss-Lederman, Sun, Tsao, and Turnbull 1995). The PRISM developers have used the MPI implementation from Argonne National Laboratory/Mississippi State (MPICH) on Sun clusters, Intel Delta and Paragon, TMC CM5, IBM SP1 and Meiko CS2. They have also used the MPI implementation from IBM (MPIF) on the SP1.

4.4 ParX

Steve Ball, of Australian National University[16], has developed ParX. ParX is a port of the X Window System to a distributed-memory multicomputer (in the first instance, the Fujitsu AP1000). Both the Xlib library and an X server are to be ported. This will allow an application running on the multicomputer to

[16] Contact address: *Steve.Ball@cs.anu.edu.au.*

display graphics on a local framebuffer or a remote display using the same application programmer interface. There is no need to off-load data to a host processor subsequently to generate X calls. There will be an implementation of the Xlib library that allows an application program to connect to an X server (local or remote) as well as an implementation of an X server for a distributed framebuffer. This server will itself be distributed across the multicomputer.

The developer was just starting a re-write of the Xlib implementation, when X11R6 was released as well as an implementation of MPI for the AP1000. Thus no extra effort has been necessary to convert to MPI, apart from learning about MPI itself. In the not-too-distant future, the developer will be porting ParX's Xlib to the CM5 (CMMD), as CM5 users require MIMD X (as opposed to CMX11, which is SIMD). MPI allows the code to be written portably to the CM5, as well as other distributed systems. MPI has satisfied all of the developer's requirements, apart from the semantics of broadcasting. A non-blocking, non-synchronizing broadcast should overcome this. The developer noted that intra- and inter-communicators are vital for this project, as it is a library that needs to insulate itself from the user's application. Pack/unpack may be useful later for optimizing memory access.

The point-to-point and collective operations are fairly straight-forward (the author skipped over the datatype features). Communicators took a little while to work through, but they also turned out to be straight-forward, according to the author. MPI release 0.1 for the Fujitsu AP1000 (port done by ANU), evidently based in part on Argonne/Mississippi State University portable implementation.

5 Interpretation of the Survey

We interpret responses here in three categories: Conversion to MPI, Target Architectures, and Needed Improvements to MPI.

5.1 Conversion to MPI

Surveyed applications reported a variety of effort levels in conversion. Conversion periods ranging from a few hours, to person weeks, to person months were reported. Surveyed applications used either C or FORTRAN, with no particular bias.

In addition to using simple features (send, receive), some of the applications exploited advanced MPI features, notably virtual topologies. Some applications used topologies, but only simple communicator strategies (copies of `MPI_COMM_WORLD`). Others exploited both topologies and multiple communicators. In many cases, datatypes were not used but those surveyed intended to use them in the future. We attribute this to the early nature of implementations available when this report was written; in several cases datatypes were not fully implemented by the time the responses were received. However, respondents who tried datatypes found them almost universally confusing.

Although the performance results described in the foregoing are limited, it is clear that researchers are already getting good results on MPI-based applications. The Light Nuclei results is the only one whose performance is quantified here, but it is evident that others have also made timings, particularly the work on the Fast Summation Library, where the researchers mention comparable performance to the fastest non-MPI message-passing system available on the IBM SP-1. For other researchers, portability and development have preceeded application tuning, so their detailed performance results will come only later. Optimization of MPI implementations will make performance results obtained in 1995 much higher than those that might have been quoted here, in any event.

Portability was universally a theme in all the responses received. Furthermore, there were only positive comments about the process of converting to MPI, with the exception that some of the researchers had specific needs for collective operations not addressed in the initial version of MPI. Furthermore, the complexity of the datatypes feature of MPI was shown repeatedly. The interface was seen to be confusing.

5.2 Target Architectures

The variety of machines mentioned by the researchers was extremely interesting. Reseachers target the IBM SP-1, CM-5, and nCUBE/2, newer machines like the Cray T3D, and less well known machines, like the Fujitsu AP-1000, were also mentioned. All are currently able to run a portable version of MPI, some will have commercial versions of MPI in 1994, and many more in 1995. This indicates significant penetration of the MPI standard after only a few months. Furthermore, the variety of implementations of MPI in use were interesting. The Argonne/Mississippi State implementation, IBM implementation, and Edinburgh Parallel Computing Centre implementations were mentioned. It is clear that at least two commercial MPI versions will emerge in 1994, as indicated in the body of this paper. Thus, application development is not at all premature in MPI.

It is interesting to note that several surveys mentioned using MPI on clusters of workstations, but there were no high-performance networks mentioned as yet. Although we know of a working version of MPI for Fiberchannel, availability of MPI for high-speed networks is probably not significant as yet, as gauged by the applications surveyed. We expect that, in the future, MPI will have a wide audience on high-performance clusters, nonetheless.

5.3 Needed Improvements to MPI

Respondents spoke to the issues of needed improvements to MPI. In the area of communication primitives, a common request was for a broadcast primitive to be matched up with a simple receive. Furthermore, non-blocking (or asynchronous) collective communication operations were requested. Currently, all such operations block the participating processes. The interface to derived datatypes was found unclear and more documentation and examples of such datatypes was requested. The desire to have more documentation and examples (of varying difficulty, in both C and FORTRAN) was also raised widely. Finally, one set of authors noted that MPI has limited scalability when a huge number of request objects are employed. They

suggest that a more efficient mechanism for request objects be added to MPI. Finally, when respondents commented on the standards document, they indicated that it needs to be made more understandable, have more examples (in Fortran as well as C), and that other simpler explanations are needed. We noted that several such projects (such as an annotated reference manual) are fortunately underway.

6 Summary, Conclusions, Future Work

We surveyed application programmers concerning early applications in the Message-Passing Interface (MPI). We received ten responses, nine of which are described in this paper. Five of the responses were purely application projects, a sixth aimed at a class of applications, and the latter four responses were generic, application-enabling libraries.

We were not able to get extensive performance feedback at this early stage. However, the variety of application efforts, and target architectures is extremely positive. We expect than *many* projects will be reporting performance results during 1995 and beyond, so that this paper will serve as a supplement to such detailed case studies.

This paper will is necessarily the first step of an on-going study of the progress of these (and hopefully additional) key applications under MPI. As the number of applications continue to grow, we expect that a paper of this kind would continue to evolve, showing application classes that MPI impacts, and describing MPI at work with other major software systems on such applications.

7 Acknowledgements

We would like to thank the following researchers for contributing to the survey and to this paper:

- S. Pieper and R. Wiringa, Physics Division, Argonne National Laboratory,

- R. Pankajakshan and W. R. Briley, and J. Zhu, Mississippi State University, NSF Engineering Research Center for Computational Field Simulation.
- J. West, M. Stevens, L. Turcotte, Army Corps of Engineerings, Vicksburg,
- H. W. Yau, Edinburgh Parallel Computing Centre,
- K. A. Cliffe and D. Brear, AEA Technology, Harwell,
- S. Lederman, A. Tsao, and T. Turnbull of SRC,
- C. Bischof of Argonne National Laboratory,
- X. Sun, of Duke University,
- J. A. Scott, Atlas Centre, Harwell,
- J. Salmon of Caltech,
- M. S. Warren of Los Alamos National Laboratory,
- D. J. Edelson of Northeast Parallel Architectures Center, Syracuse University,
- J. C. Hardwick of Carnegie Mellon University, and
- S. Ball, of Australian National University.

We wish to thank the members of the MPI Forum as a whole, for their efforts in creating MPI.

References

Bischof, C., S. Huss-Lederman, X. Sun, A. Tsao, and T. Turnbull (1994, March). Parallel performance of a symmetric eigensolver based on the invariant subspace decomposition approach. In *Proc. of Scalable High Performance Computing Conference 1994*.

Bischof, C., S. Huss-Lederman, X. Sun, A. Tsao, and T. Turnbull (1995, February). A case study of mpi: Portable and efficient libraries. In *Proceedings, Seventh SIAM Conference on Parallel Processing for Scientific Computing*.

Blelloch, G. E., S. Chatterjee, J. C. Hardwick, J. Sipelstein, and M. Zagha (1994, April). Implementation of a portable nested data-parallel language. *Journal of Parallel and Distributed Computing 21*(1), 4–14.

Burns, G., R. Daoud, and J. Vaigl (1994). Lam: An open cluster environment for mpi. URL ftp://tbag.osc.edu/pub/lam/lam-papers.tar.Z.

Cameron, G, Smith, M, White, M, Wylie, B.J.N, and McArthur, D (1994). The PARAMICS Parallel Microscopic Traffic Simulator. In *UK IT Forum Conference Digest*.

Cameron, G, Wylie, B.J.N, and McArthur, D (1994). PARAMICS: Moving Vehicles on the Connection Machine. In *Submitted to Supercomputing '94*.

Doss, N., W. Gropp, E. Lusk, and A. Skjellum (1993). A model implementation of MPI. Technical report, Argonne National Laboratory.

Forum, M. P. I. (1994a, April). Document for a standard message-passing interface. Technical Report Technical Report No. CS-93-214 (revised), University of Tennessee. Available on **netlib**.

Forum, M. P. I. (1994b, April). MPI: A message-passing interface standard. Computer Science Dept. Technical Report CS-94-230, University of Tennessee, Knoxville, TN. (To appear in the International

Journal of Supercomputer Applications, Volume 8, Number 3/4, 1994).

Franke, H., P. Hochschild, P. Pattnaik, and M. Snir (1993). An efficient implementation of MPI. (contact: Hubertus Franke, frankeh@watson.ibm.com).

Gropp, W., E. Lusk, and A. Skjellum (1994, November). *Using MPI: Portable Parallel Programming with the Message-Passing Interface*. MIT Press.

Hardwick, J. C. (1994, November). Porting a vector library: a comparison of mpi, paris, cmmd and pvm. Technical Report CMU-CS-94-200, School of Computer Science, Carnegie Mellon University.

Lorensen, W. and H. Cline (1987, July). Marching cubes: A high resolution 3d surface construction algorithm. *Computer Graphics 21*(4), 163–169. Proceedings of SIGGRAPH 1987.

McArthur, D (1994). The PARAMICS Model: Present and Future Directions. In *current working paper, to be submitted to the Transportation Research Journal.*

Pankajakshan, R. and W. R. Briley (1994). Parallel solution of unsteady incompressible viscous flows. (in preparation).

Salmon, J. K., G. S. Winckelmans, and M. S. Warren (1994). Fast parallel treecodes for gravitational and fluid dynamical N-body problems. *J. Supercomputer Appl. 8.* (to appear) available as ftp/nbody/ijsa.ps.Z@sampson.ccsf.caltech.edu.

Skjellum, A., N. E. Doss, and K. Viswanathan (1994, July). Inter-communicator Extensions to MPI in the MPIX (MPI eXtension) Library. Submitted to ICAE Journal special issue on Distributed Computing.

Taylor, L. K., J. A. Busby, M. Y. Jiang, A. Arabshahi, K. Sreenivas, and D. L. Whitfield (1993, August 2–5). Time accurate incompressible Navier-Stokes simulation of the flapping foil experiment. Presented at the Sixth International Conference on Numerical Ship Hydrodynamics, Iowa City, Iowa.

Taylor, L. K. and D. L. Whitfield (1991, June 24–26). Unsteady three-dimensional incompressible Euler and Navier-Stokes solver for stationary and dynamic grids. AIAA 91-1650. Presented at the AIAA

22nd Fluid Dynamics, Plasma Dynamics and Lasers Conference, Honolulu, Hawaii.

Warren, M. S. and J. K. Salmon (1993). A parallel hashed oct-tree N-body algorithm. In *Supercomputing '93*, Los Alamitos, pp. 12–21. IEEE Comp. Soc.

Warren, M. S. and J. K. Salmon (1994). A portable parallel particle program. *Computer Physics Communications*. (submitted).

West, J. E., M. M. Stephens, and L. H. Turcotte (1994a, October). Adaptation of volume visualization techniques to mimd architectures using mpi. In *Proceedings of the IEEE Scalable Parallel Libraries Conference*. Starkville, MS.

West, J. E., M. M. Stephens, and L. H. Turcotte (1994b). An investigation of the parallel implementation of volume visualization techniques on the ncube 2. Technical Report Submitted for publication, DoD High Performance Computing Center, USAE Waterways Experiment Station.

Westover, L. (1990, August). Footprint evaluation for volume rendering. *Computer Graphics 24*(4), 367–376. Proceedings of SIGGRAPH '90.

Whitfield, D. L. and L. K. Taylor (1991, June 24–26). Discretized Newton-relaxation solution of high resolution flux-difference split schemes. AIAA-91-1539. Presented at the AIAA 10th Computational Fluid Dynamics Conference, Honolulu, Hawaii.

Wylie, B.J.N, Cameron, G, White, M, Smith, M, and McArthur, D (1993). PARAMICS: Parallel Microscopic Traffic Simulator. In *Proceedings of Second European Connection Machine Users Conference.*

Lines of Fortran code and comments	23,000
Lines in time-critical subroutines	1000
Total program single-processor speed (megaflops)	48
Total 128-node speed (gigaflops)	5.9
Speed in the critical subroutines (megaflops)	36-56

Table 1: Summary of the Argonne nuclear structure code.