

NOV 21 1988

STP/STP/STP

ornl

ORNL
MASTER COPY

ORNL/TM-11375

OAK RIDGE
NATIONAL
LABORATORY

MARTIN MARIETTA

**PVM: A Framework for
Parallel Distributed Computing**

V. S. Sunderam

OPERATED BY
MARTIN MARIETTA ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831; prices available from (615) 576-8401, FTS 626-8401.

Available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161.

NTIS price codes—Printed Copy: A03 Microfiche A01

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ORNL/TM-11375

**Engineering Physics and Mathematics Division
Mathematical Sciences Section**

PVM: A Framework for Parallel Distributed Computing

**V. S. Sunderam
Department of Math and Computer Science
Emory University
Atlanta, GA 30322**

Date Published:

**Research performed at the Mathematical Sciences Section of
Oak Ridge National Laboratory under the auspices of the Faculty Research
Participation Program of Oak Ridge Associated Universities, and supported
by the Applied Mathematical Sciences subprogram of the Office of Energy
Research, U.S. Department of Energy.**

**Prepared by the
Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831
operated by
MARTIN MARIETTA ENERGY SYSTEMS, INC.
for the
U.S. DEPARTMENT OF ENERGY
under Contract No. DE-ACO5-84OR21400**

CONTENTS

Abstract	1
1. Introduction	1
2. The User Interface	3
2.1 Processes and Process Initiation	4
2.2 Data Transfer and Barrier Synchronization	6
2.3 Shared Memory and Mututal Exclusion	7
2.4 Miscellaneous Facilities	9
3. PVM System Design and Implementation	9
3.1 Basic Facilities	10
3.2 Process Control	12
4. Results and Experiences	13
5. Conclusions & Future Work	16
References	17

PVM : A Framework for Parallel Distributed Computing

V. S. Sunderam

Department of Math and Computer Science
Emory University, Atlanta, GA 30322

ABSTRACT

The PVM system is a programming environment for the development and execution of large concurrent or parallel applications that consist of many interacting, but relatively independent, components. It is intended to operate on a collection of heterogeneous computing elements interconnected by one or more networks. The participating processors may be scalar machines, multiprocessors, or special-purpose computers, enabling application components to execute on the architecture most appropriate to the algorithm. PVM provides a straightforward and general interface that permits the description of various types of algorithms (and their interactions), while the underlying infrastructure permits the execution of applications on a virtual computing environment that supports multiple parallel computation models. PVM contains facilities for concurrent, sequential, or conditional execution of application components, is portable to a variety of architectures, and supports failure detection (and certain forms of recovery) at the process and processor levels.

1. Introduction

In recent years, parallel and distributed processing have been conjectured to be the most promising solution to the computing requirements of the future. Significant advances in parallel algorithms and architectures have demonstrated the potential for applying concurrent computation techniques to a wide variety of problems. However, most of the research efforts have concentrated either upon *computational models* [1] such as pipelining, shared variables, dataflow computing, and message passing, or upon machine *architectures*; relatively little attention has been given to software development environments or program construction techniques that are required in order to translate algorithms into operational programs. This aspect is becoming more important as parallel processing progresses from the solution of stand-alone, mathematically

precise, problems to larger and more complex software systems. Such systems often consist of many interacting components, each with its unique requirements. Unfortunately, a coherent and consistent framework for the specification and development of such systems does not exist.

The primary reason for this situation is the high degree of architecture dependency exhibited by efficient parallel solutions to typical problems. Another reason is the fundamental incompatibility between the different parallel programming paradigms. A third factor is the nature of the problems themselves — typical applications have thus far been confined to isolated, well-defined problems — usually programmed in the most suitable, machine specific language. As applications grow larger and more complex, however, these languages (and indeed, specific machines) will not be uniformly appropriate or efficient for all the components of a parallel system. Certain components may be ideally suited for solution on a hypercube, for example, while others may require extensive vector processing. Some algorithms may be best implemented using the shared-memory paradigm, while the user-interface aspects of the system might require a graphics engine. A real example of the above is the Global Environment Simulation project [2], a large simulation effort to study contaminant concentrations and dispersal characteristics as a function of various environmental factors. The computational requirements of this simulation are vector processing (for fluid flow analysis), distributed multiprocessing (modeling contaminant transport), high-speed scalar computation (simulation of temperature effects), and real-time graphics for user interaction.

It should be noted that most typical computing environments already possess the hardware base required to solve such large, parallel applications. High speed local networks with graphics workstations, high-performance scalar engines, an occasional multiprocessor, and perhaps a vector computer are the norm rather than the exception, and will continue to be over the next few years. However, to harness this collection of capabilities and to utilize it productively requires considerable efforts in coordination and reconciliation between different computation models and architectures — all of which has to be done manually. The PVM (Parallel Virtual Machine) project is an attempt to provide a unified framework within which large parallel systems can be developed in a straightforward and efficient manner. The overall objective of this project is to permit a collection of heterogeneous machines on a network to be viewed as a general purpose concurrent computation resource. Application algorithms are expressed using the most suitable paradigm; the PVM system executes them on the most appropriate hardware available, either directly or by emulating the particular computation model. Furthermore, it is frequently desired to incorporate existing software (preferably with little or no modifications) into a larger system; the PVM system is designed to enable this in a convenient and natural manner.

The PVM system provides a set of user interface primitives that may be incorporated into existing procedural languages. Primitives exist for the invocation of processes, message transmission and reception, broadcasting, synchronization via barriers, mutual exclusion, and shared memory. Processes may be initiated synchronously or asynchronously, and may be conditioned upon the initiation or termination of another process, or upon the availability of data values. Message transmission as well as file output may be preceded by invocations of specially provided primitives to ensure that data is transmitted or stored in a machine independent form. Application systems may be programmed using these primitives in the language of choice; different components may even be programmed in different languages. The PVM constructs therefore permit

the most appropriate programming paradigm and language to be used for each individual component of a parallel system while retaining the ability for components to interact.

The PVM system consists of support software that executes on participating hosts on a network; the network may be local, wide-area or a combination, and the host pool may be varied dynamically. Hosts may be scalar machines, workstations, or parallel processors — the latter being considered an atomic computational resource by PVM. This support software interprets requests generated by the user-level constructs and carries out the necessary actions in a machine independent manner. In addition to implementing reliable and sequenced data transfer, distributed consensus, and mutual exclusion, the PVM software is also responsible for the detection of process and processor failures (including certain forms of deadlock) and executing user-defined corrective actions. It should be mentioned that the PVM user-interface primitives have been partly derived from and are a superset of the portable programming constructs described in [3]; an application written using these primitives may therefore also execute directly on a specific multiprocessor when necessary.

Several projects similar to PVM have been undertaken in the past, and some are ongoing. A few representative examples are listed below, with comparisons to PVM. The DPUP library [4] emulates a loosely coupled multiprocessor on a local network, as does the dsim [5] system and the Cosmic environment [6]. The two latter systems require the preconfiguration of a virtual machine on which applications execute and support only basic message passing mechanisms. The Amber project [15] is somewhat different in that the targeted environment is a collection of homogeneous *multi*-processors. One of the operating modes within DPUP, as well as projects such as Marionette [7] and MSPCM [8], uses the master-slave approach, where a central controlling process is responsible for or is involved in every system event. In addition to affecting performance and being an unnatural model for certain classes of problems, this central process is critical, and its failure leads to a complete collapse of the entire system. Another shortcoming common to all the above is the use of virtual circuits for network communication; in addition to overheads that may not be justifiable, practical limits on the number of connections affect the scalability of applications. In addition, failure resiliency and debugging support is minimal. The PVM system is completely distributed, supports a dynamic host pool, and assumes only that an unreliable, unsequenced datagram delivery mechanism is available. From the application's point of view, PVM constructs are substantially more general in nature and encompass both the message passing and shared memory paradigms; yet, by substituting alternative libraries, unmodified programs may execute on specific multiprocessors.

The following section describes the user interface and the important design aspects of PVM. An overview of the PVM support software, with an emphasis on the protocol algorithms and key implementation features follows. Preliminary results and performance figures are then presented, and the concluding section reports on continuing and future work.

2. The User Interface

The application views the PVM system as a very general and flexible parallel computation resource that supports common parallel programming paradigms. Application programs access

these resources by invoking function calls from within common procedural languages such as C or Fortran. Such an interface was selected primarily for portability reasons — most multiprocessor applications are currently written in procedural languages with embedded, machine-specific, function calls that perform process spawning, message reception and transmission, and shared memory operations. The PVM primitives have been made the same as or very similar to the union of these functions, thereby enabling previously written applications to be ported readily to the PVM environment and also permitting PVM to execute applications or components thereof on specific machines when possible. The PVM user interface syntax and semantics are presented in this section with illustrative examples using the C language interface.

2.1. Processes and Process Initiation

In the PVM system, an application is considered to consist of *components*. For example, a simulation application might consist of a partial differential equation component, a matrix solution component, and a user interface component. It should be pointed out that this definition of a component is perhaps unconventional; usually, the term implies a phase or portion of an application that is embodied in a subroutine — such as "the forward-substitution component of a matrix solver". However, the PVM system is a large-granularity environment, primarily targeted at applications that are collections of relatively independent programs. In view of this, a PVM component corresponds not to a phase in the traditional sense, but rather to a larger unit of an application. From the system point of view, a component corresponds to an object file that is capable of being executed as a user-level process. A compiled C program that performs LU factorization is an example of a component. It is the responsibility of the user to compile component programs to all target architectures on which that component may execute. Depending upon the target machine, the compiled version of a component may either link against the PVM primitives, or machine specific libraries, or both. A component is therefore a static entity and is identified by a *name*; associations between component names and executable versions are set up as discussed in the following paragraphs.

A complete description of application components, i.e. the component name and all corresponding executables (each with an architecture tag), is obtained by the PVM support software. This information is gathered either from a file or from a startup process, as will be explained below. An example of a component description file is shown in Figure 1. This table illustrates that a component, identified by a name, may be manifested as several different executable files; and conversely, that multiple component names may map onto the same executable. The first feature permits the PVM system to execute components at the most suitable location, while the second allows the user to force a specific location as will be explained below.

A *process* is an executing instance of a component and is identified by the component name and a positive instance number. Processes may be initiated from within components or from a "startup" process that may be manually executed on any participating host. A process is initiated by invoking the *initiate* primitive with the component name as an argument; the instance number of the initiated process is returned to the user. Prior to executing any PVM construct, however, a process must invoke the *enroll* function; this establishes a (machine dependent) mechanism by which a user process may communicate with the PVM system. A typical section of code executed

Name	Location	Object file	Architecture
factor	iPSC	/u0/host/factor1	ipsc
factor	msrsun	/usr/alg/math/factor	sun3
factor	msrsun	/usr/alg4/math/factor	sun4
chol	csvax2	/usr/matrix/chol	vax
chol	vmsvax	JOE:CHOL.OBJ	vms
tool	msrsun	/usr/view/graph/obj	sun3
factor2	iPSC	/u0/host/factor1	ipsc

Figure 1 : Example Component Description File

by a startup process is shown in Figure 2. It should be noted in the example shown that the physical location of the initiated processes is transparent to the invoking process; the PVM system determines the best machine on which to execute a process based upon the current host pool, the alternative architectures on which a component may execute, and the load factor on those machines. However, a specific location may be forced by declaring a new component name (as in the last line of the component description file above) and initiating that component.

```
...
enroll("startup");
for (i=0;i<10;i++)
    instance[i] = initiate("factor");
...
```

Figure 2 : Initiation of multiple component instances

The `initiate` mechanism is, by default, asynchronous. Control is returned to the invoking process as soon as the instance number of the process is available. However, under certain circumstances, it may be necessary to initiate a component only after another process has terminated. The `initiateP` variant allows this by permitting the user to defer initiation of a component until after another has terminated. For example,

```
initiateP("factor", "matmul", 3);
```

will initiate an instance of "factor" only after instance number 3 of "matmul" has terminated. A third-argument value of 0 will cause "factor" to be initiated only when *all* instances of "matmul" terminate. In an analogous fashion, `initiateD` is used to execute components conditional upon the occurrence of a user-signaled event, normally the availability of data. Thus,

```
initiated("chol", "dataset7");
```

will delay the execution of "chol" until some other process signals the occurrence of the "dataset7" event, by invoking the `ready("dataset7")` primitive. All variants of `initiate` return a negative result if a process could not be initiated, thereby enabling the invoker to take appropriate action. The global component dependencies of the application may therefore be specified within

a startup process by the use of appropriate **initiate** primitives or variants, embedded within common selection and iteration control flow constructs available in the host language. Of course, a component itself may be composed of several subcomponents — whose dependencies and execution order are indicated in an analogous manner within that component. Two other constructs termed **terminate** and **waitprocess** are also provided. Both take a component name and an instance number (or 0 to mean all instances) as arguments; the first aborts the process while the second blocks the caller until the process completes.

2.2. Data Transfer and Barrier Synchronization

Inter-process communication via message passing is one of the basic facilities supported by PVM. In the interest of portability and wide applicability, the primitives to accomplish message transfer have been derived from existing implementations (e.g [9]), including those described in [3]. Certain aspects, however, are necessarily different; primary among them is addressing. Since the physical location of processes is deliberately transparent to user programs, message destinations are identified by a {component name, instance number} pair. Furthermore, owing to the heterogeneous nature of the underlying hardware that PVM executes upon, it is necessary for user programs to send and receive typed data in a machine independent form. To enable this, a set of conversion routines has been provided — user programs invoke these routines to construct message buffers and to retrieve data values from incoming messages.

In keeping with popular message passing mechanisms, the PVM **send** and **receive** constructs incorporate a "type" argument. This is the only argument to **receive**, while **send** requires a destination component name and instance number as additional arguments. The type parameter permits the selective reception of messages and has been found to be extremely useful in practical applications. It should be noted that neither the data buffer itself nor its length appear explicitly as arguments — owing to data representation and size differences on different machines, user programs should only access messages using the conversion routines. Shown in Figure 3 is an example of data transfer between two component processes.

In order for a receiving process to obtain additional information about the most recently received message, the **recvinfo** construct is provided; this returns the name and instance number of the sending process and the message length. In addition, two variants of the **recv** construct are provided. The first, **recv1**, permits the user to specify the maximum number of messages of other types that may arrive in the interim (i.e. while waiting for a message of the specified type). If a message of the anticipated type does not arrive within this window, an error value is returned to the program, thus enabling the detection of and possible recovery from incorrect program behavior or unacceptable levels of asynchrony. The second variant, **recv2**, allows the specification of a timeout value and is valuable in preventing certain forms of deadlock as well as in user-level detection of failed components. Also provided is the **broadcast** primitive that sends a message to all instances of a specified component.

Synchronization via barriers is a common construct in many applications. Under PVM, barrier synchronization is accomplished using the **barrier** construct. An instance of a component invoking this construct will block until all instances of the component also arrive at the barrier.

```
/* Sending Process */
/*-----*/
init send(); /* Initialize send buffer */
putstring("The square root of "); /* Store values in */
putint(2); /* machine independent */
putstring("is "); /* form */
putfloat(1.414);
send("receiver", 4, 99); /* Instance 4; type 99 */

/* Receiving Process */
/*-----*/
char msg1[32], msg2[4];
int num; float sqnum;
recv(99); /* Receive msg of type 99 */
getstring(msg1); /* Extract values in */
getint(&num); /* a machine specific */
getstring(msg2); /* manner */
getfloat(&sqnum);
```

Figure 3 : User process data transfer

The PVM system attempts to detect and correct barrier deadlocks by notifying invoking processes if some instances of a component terminate before they reach a barrier — live processes return from a barrier call with a negative result value in such situations. In addition to barriers, or as an alternative, the `waituntil` construct is also provided as a means of synchronization. This construct (suggested in [10]) takes an event name as an argument and blocks until another process indicates the occurrence of that event by using the `ready` primitive mentioned earlier.

2.3. Shared Memory and Mutual Exclusion

The use of shared memory to synchronize and communicate between processes is a convenient and well understood paradigm, and the PVM system provides such an interface for algorithms that are best expressed in these terms. It should be noted, however, that in most cases this facility is emulated on distributed memory machines and some performance degradation should be anticipated. The primitives provided are modeled once again after popular, existing implementations. A shared memory segment is first allocated by invoking the `shmget` construct that takes a string valued identifier and a segment size in bytes. To acquire a shared memory segment for use, a user process invokes the `shmat` construct, specifying the segment identifier, the address within the process at which the segment is to be mapped, a flag indicating whether the segment is to be mapped read-only or read-write, and a timeout value. This construct implicitly incorporates a lock operation; if mutually exclusive access to the segment cannot be provided, the invoking process is suspended — for a period not to exceed the specified timeout value.

The attach operation described above maps a contiguous, untyped block of bytes at the specified address. In most situations however, shared memory segments will be used to store and manipulate typed data. In order to permit this among dissimilar machines, typed variants of the attach construct are provided. For example, the `shmatint` construct takes an integer pointer as its second argument, while the `shmatfloat` variant is used for shared memory regions that hold floating point values. (It should be noted that typed data transfer between dissimilar architectures could lead to loss of precision or to truncation owing to wordsize differences. Both message passing and shared memory mechanisms are subject to this drawback. The PVM system attempts to minimize this by utilizing the largest size possible for typed data values.) When a process no longer needs exclusive access to a region, it invokes the `shmdt` construct (or a typed variant) whereupon the lock is released and the region unmapped. Finally, the `shmfree` construct is used to deallocate a segment of shared memory when it is no longer required. Shown in Figure 4 is an example of the use of these constructs to pass an array of real numbers between two processes.

```
/* Process A */
/*-----*/
if (shmget("matrx",1024) error());/* Allocation failure */
while                               /* Try to lock & map seg */
    (shmatfloat("matrx",fp,"RW",5));
for (i=0;i<256;i++) *fp++ = a[i]; /* Fill in shmem segment */
shmdtfloat("matrx");             /* Unlock & unmap region */

/* Process B */
/*-----*/
while                               /* Lock & map; note:reader*/
    (shmatfloat("matrx",fp,"R",5));/* may lock before writer */
for (i=0;i<256;i++) a[i] = *fp++; /* Read out values */
shmdtfloat("matrx");             /* Unlock & unmap region */
shmfree("matrx");               /* Deallocate mem segment */
```

Figure 4 : Use of shared memory for IPC

While shared memory is perhaps the most common resource that processes require mutually exclusive access to, it is possible that the PVM environment contains other resources that processes must access in a similar manner. To accommodate such requirements, a generalized locking facility is also provided. The `lock` construct permits the logical locking of an entity that is named by a string argument; the PVM system blocks other processes wishing to lock this entity until the possessor invokes the `unlock` construct. For example, different components of a large application may wish to output results periodically to a user terminal. To avoid interference and to distinguish the source of the output, components may adopt a convention that requires locking "terminal" before printing messages or results. Another situation where such a facility could be useful may be found in the shared memory example in Figure 4. In that example, it is easy to see that the processes may access the shared memory segment in an *incorrect order* even though each will have exclusive access to it. A possible rectification of this situation is to use the `lock` construct as shown in Figure 5; however, in practice it is more likely that a transmitted message or the `waituntil` facility will be used to resolve such situations.

```
/* Process A */
/*-----*/
lock("fillmatrix",5);
/* Allocate, attach, fill, and detach shared mem segment. */
unlock("fillmatrix");

/* Process B */
/*-----*/
loop:
    lock("fillmatrix",5);
    if (shmatfloat(...) == SEGMENT_NONEXISTENT) {
        unlock("fillmatrix");
        sleep(1);
        goto loop;
    }
/* Read values out of shared mem segment, detach, & free */
```

Figure 5 : Use of the lock construct

2.4. Miscellaneous Facilities

In addition to the primary constructs described in the preceding sections, a few miscellaneous constructs are also provided. The `status` construct takes a component name and instance number as arguments and returns status and location information regarding that component. The `entercomp` construct permits dynamic additions to the component description table. The `shmstat` construct is used to obtain information about active shared memory regions, while the `lockstat` primitive reports the status of active locks. A complete list of all the user interface constructs along with their argument lists and a one-line description is given in the appendix.

3. PVM System Design and Implementation

The PVM support software executes as a user-level process on each host in the participant pool. An initial set of participating hosts is statically identified; additions or deletions are possible during operation by means of an administration interface. The PVM system is designed to be implemented in a manner that requires no operating system changes or modifications, and porting efforts to varied operating system environments are minimal. The PVM support process (termed *pvmd*) on a host is responsible for all application component processes executing on that host; however, there is no central or master instance of *pvmd*. Control is completely distributed (by virtue of all *pvmd* processes possessing global knowledge) in the interest of avoiding performance bottlenecks and increasing fault tolerance. The *pvmd* processes are initiated on each participating host either manually, through the administration interface, or via a machine/OS dependent mechanism such as *inetd* in the Unix environment. In this section, the key design aspects of the *pvmd* software are discussed with an emphasis on the protocol algorithms used.

3.1. Basic Facilities

In terms of network capabilities, the PVM system assumes only that unreliable, unsequenced, point-to-point data transfer (but with data integrity) facilities are supported by the hardware platform on which it executes. The required reliability and sequencing, as well as other necessary operations such as broadcast, are built into the PVM system in the interest of efficiency and portability. While it is true that most operating systems in existence already support reliable and sequenced data delivery, in most cases this is via the use of virtual circuits ---- for the projected use of PVM the overheads and scalability limitations of using such a service directly did not warrant its adoption. In the test implementations of PVM, the UDP [11] protocol was used; this deliberate choice of a simple datagram protocol also permits relatively simple porting or protocol conversion when PVM is to be installed under a different operating system environment.

Across the network, *pvmd* processes communicate using UDP datagrams. The "well known port" approach is used for addressing; all incoming messages are received by *pvmd* processes on a predetermined port number. For user-process to user-process communication, the following scheme is employed. The first communication instance between any two entities is routed through the *pvmd* processes on the source and destination machines. Location and port number information is appended to this exchange; the PVM routines (linked to the user process) that implement `send` and `recv` cache this information, thus enabling direct communication for subsequent exchanges. Local user processes communicate with *pvmd* using the most efficient machine dependent mechanism available and the development of this mechanism is deemed part of the installation procedure. However, the generic version of *pvmd* may be adopted; this utilizes UDP datagrams once again, via the loopback interface if one is available. The *pvmd* process uses a different, predetermined port number for incoming messages from all local user processes.

To achieve reliable and sequenced point-to-point communications, the *pvmd* processes use a positive acknowledgment scheme and an additional header that contains sequence numbers as well as fragmentation and reassembly information. Unacknowledged transmissions are retried a parameterized number of times after which the recipient process or processor is presumed to be inoperative. The sequence numbers are destination specific and are used by the message recipient for sequencing as well as for duplicate detection. The header is placed at the end of a UDP datagram to reduce copying overheads, and single datagram sizes are restricted to the smallest MTU (maximum transmission unit) of all participating hosts. When first initiated, *pvmd* processes determine the protocol specific addresses of all participating hosts and proceed to service incoming requests from the network or user processes in an infinite loop.

Each *pvmd* process maintains information concerning the location and status of all application component processes. A user `send` is addressed to a component name and instance number; the local *pvmd* determines the physical location of that process and forwards the message to the remote *pvmd*. As described, the user process library performs this translation for the second and subsequent messages. The source component name and instance number are appended to the message, enabling message delivery with an indication of the sender's identity. As mentioned, executing the `enroll` construct is a precondition to user process participation --- in the UDP implementation, this supplies to the local *pvmd* the receiving port number of the user process.

Broadcast is a commonly performed operation in the PVM system, both because applications desire such a facility and since it is inherent to the completely distributed nature of the PVM support software. All *pvmd* processes maintain information regarding all processes, shared memory segments, and locks, to guard against loss of context and state in the event of failures. User process broadcasts are first delivered sequentially to local recipient user processes after which the local *pvmd* process broadcasts over the network to all other *pvmd* processes that in turn, deliver the message to their local user process recipients. Although most computing environments support a network broadcast facility, *pvmd* broadcast is implemented in PVM using point-to-point messages with recursive doubling. This decision was made in the interest of portability and efficiency; given that network broadcast is unreliable, acknowledgments are necessary from each recipient, resulting in $O(p)$ time ($p + 1$ sequential steps are to be performed by the originator), while recursive doubling broadcast is accomplished in $O(\log_2 p)$ time, where p is the number of processors. The participating pool of hosts is logically numbered from 0 to $p-1$, and the originator (or root) of the broadcast is part of the broadcast message. There is one *pvmd* process per host, which represents that processor. Broadcast proceeds in "rounds", with the number of processors contributing to the broadcast effort doubling in each round. In any round, processor i transmits to $i + 2^r \bmod p$ and receives an acknowledgment. A processor j joins the broadcast effort at round r_j , where $r_j = \text{no. of significant bits in } (j - \text{root}) \bmod p$. In the event of processor failure, the *pvmd* process that first detects the failure assumes the broadcast duties of the failed processor. If the quantity $2(2^r)$ is less than $2 \log_2 p$, failure notification is piggybacked on the broadcast, at the end of which the remaining processors are individually informed. Otherwise, the detecting processor initiates another broadcast with failure information, at completion of the current broadcast.

The *pvmd* processes execute a finite state machine which gives precedence to messages (requests) incoming while another activity is in progress. Such a scheme is adopted to avoid deadlock; two processes transmitting to each other may both wait indefinitely for each others acknowledgment if this precedence rule were not followed. It should also be pointed out that in the case of PVM hardware platforms where wide area networks are involved, the choice of an appropriate timeout value can significantly affect the performance of the data transfer mechanisms and the broadcast process. Further, the present implementation does not perform any optimizations in the broadcast scheme when a geographically distant host is at a non-leaf position in the broadcast spanning tree.

Mutual exclusion is another primitive required both in response to user requests as well as for *pvmd* coordination. Examples are exclusive access to emulated shared memory, general resource locking, and assignment of unique instance numbers for application component processes. Distributed mutual exclusion is normally achieved by unanimous or majority consensus; a requesting process that receives permission from a certain number of processes is deemed to have acquired the lock. Different strategies, varying in their approach, efficiency, and level of failure resiliency have been proposed and representative methods are described in [12, 13, 14]. The strategy adopted in PVM is somewhat different from these approaches, but the algorithm is efficient and, more importantly, is integrated with the required distribution to all *pvmds* of lock location information.

A *pvmd* process, either for its own purposes or on behalf of a local user process, attempts to obtain a lock by broadcasting a "claim" for the lock. Since all *pvmd* processes possess knowledge regarding the use (and location) of all locks, such an attempt will, of course, only be made when a lock is known to be free at the start of the claim. In the absence of conflicting claims (a situation most likely to be encountered in practice), the requester, after the broadcast has been completed, can assume that the lock has been successfully obtained. In the process, all other *pvmd*s update their lock table information, and (implicitly) grant the requester permission for exclusive access to the particular resource.

It is of course possible that two processes may initiate claims on the same resource before either has received the other's request. In PVM, such situations are resolved using a heuristic that assumes that communication between any pair of processors takes the same amount of time. In particular, consider two processors (*pvmd* processes) *A* and *B* that wish to acquire the same lock, and another processor *C*. Note that $\{B,C\}$ and $\{A,C\}$ are in the broadcast spanning trees of *A* and *B* respectively, possibly at different depths. It may be assumed without loss of generality that *A*'s processor number is less than that of *B*. Under the constant time assumption, both *A* and *B* will receive each others claims before their broadcast is completed. When *B* receives *A*'s claim, it computes the number of rounds that *A*'s broadcast has proceeded; if this number is greater than the number of rounds that its own broadcast has proceeded, *B* surrenders its claim to the lock. An identical (first-claim, first served) policy is followed by *A*. If the broadcast progress metrics are the same, the lower numbered processor is given priority and is considered to have obtained the lock. The passive processor *C* also makes the same decision since it has the capability of computing the number of rounds of broadcast progress that each claimant has made when the second such broadcast arrives.

In practice however, communication times between arbitrary pairs of processors may not be constant; further, intervening messages of other types may skew the propagation time of a broadcast claim. For practical safety therefore, the originators of conflicting claims exchange a confirmatory message — with the claim being abandoned if their respective notions of the successful claimant are not in agreement. In such a situation, the lower numbered processor broadcasts a "reset lock" message, and the entire process is started afresh, but without competition from the "losing" processor. In case of process or processor failures, the strategy of all *pvmd* processes possessing all information is used to avoid undesirable situations. If a process or processor holding a lock terminates without releasing it, the particular resource is marked as "defunct"; further requests to the resource are denied until an explicit reset is performed.

3.2. Process control

The initiation order and process dependencies of application components are described by the use of appropriate **initiate** constructs embedded within host language control flow statements as described in the preceding section. This implies that it is not possible to determine statically the application process flow graph as component initiations may be conditional or repeated based upon parameters known only at execution time. The PVM system therefore performs process initiations in response to requests based upon the resources available and load conditions at the moment of the request — rather than by constructing a predetermined static schedule and process

to processor assignments.

When an application component process makes an *initiate* request, the local *pvmd* process first determines a candidate pool of target hosts based upon the information in the component description file. One host is then selected from this pool based upon the following algorithm:

- (1) Select next host from pool in round-robin manner, based upon all initiations that originated here.
- (2) Obtain load metric (decayed average of number of processes in run queue) from this potential target host.
- (3) If this quantity is less than a prespecified threshold, select this host.
- (4) Otherwise, repeat the process. If no host has a load factor below the threshold, the host with the lowest load is the selected target.

Once the target host is identified, the local *pvmd* sends the *initiate* request to the *pvmd* process on the remote host, where the application component is initiated. The remote *pvmd* then broadcasts notification of this event to all processors, simultaneously claiming an instance number for this initiation (by simply incrementing the last previous instance number for the component). Conflicting claims for the same instance number are again resolved as in the case of multiple claims to a lock, with a "losing" processor using a higher value. Once again, consistent conflict resolution is confirmed by an exchange of messages between all claimants and reset actions are performed in the case of disagreement. Application process termination information is also broadcast to all *pvmd* processes. Conditional variants of *initiate* are saved by the local *pvmd*, and this queue is inspected and appropriate action taken when the particular event occurs.

In the PVM system, shared memory is emulated by first creating an image of a memory segment on secondary storage. A file of the requested size is created; for efficiency and failure resiliency reasons, the local *pvmd* (the processor at which the creation request originated) attempts to locate the file on a device that is accessible to other processors via a network file system. Mutual exclusion, both for creation as well as for access, is achieved as described earlier. A *pvmd* process that has acquired a lock (on behalf of a local application process) copies the file into the requested address space; this is done directly if the file is accessible directly, and with the assistance of the remote *pvmd* if it is not. A user release request results in the specified memory area being copied back to the file unless the lock request was for read-only access. It should be noted that creation, locking, unlocking, and deallocation (resulting in file removal) events are broadcast to all *pvmds*; given the conflict resolution rules and highest priority to incoming requests, undesirable inconsistencies are avoided.

4. Results and Experiences

To facilitate its use and to determine its effectiveness, the PVM system has been implemented on a variety of machines including Sun 3/50 workstations, Vax 11/785 and Sun 4/280 servers, a 64 node Intel iPSC/2 hypercube, and a 12 processor Sequent shared-memory multiprocessor. The minimal assumptions made regarding the underlying facilities available greatly simplified the implementation efforts; the software could be ported (from a base Sun 3

implementation) to all the environments with changes necessary only in data representation and conversion utilities. In the implementation, the *pvmd* processes run independently, while the user level routines are supplied as a set of libraries to be linked in with application components. Initial experiences with the use of PVM are reported in this section.

The efficiency of user-level data transfer is perhaps the most critical aspect of any distributed computing environment. In the PVM system, low latency data transfer has been provided without sacrificing location transparency, and a datagram protocol is used so that overheads of more heavyweight protocols are avoided. It is anticipated that a large proportion of the use of PVM will be constrained to local networks, with only a few applications wishing to execute components on geographically distant hosts. The protocols used by PVM therefore incur the overheads of retransmission and sequencing only when the underlying network quality is poor; more than 95% of local network communications typically succeed on the first attempt. Table 1 shows the message delivery times for varying message sizes under PVM between two Sun 3/50 systems on a 10 MB/s Ethernet. It should be noted that these figures represent elapsed time from the start of message transmission to the receipt of positive acknowledgment and are averages over several runs performed under varying host and network loads.

Message size (bytes)	8	128	256	512	1024
First instance	15 ms	18 ms	22 ms	25 ms	30 ms
Subsequent instances	6 ms	8 ms	10 ms	12 ms	15 ms

Table 1 : User process data transfer times in PVM

Broadcast, since it is used heavily within PVM, is another important factor in the performance of the system. It was observed that for broadcast among *pvmd* processes, the calculated performance of the recursive doubling algorithm is consistent with actual behavior. Acknowledged message transmission on a single branch of the broadcast spanning tree required between 4 and 9 milliseconds for a (typical) 100-byte message, depending upon the speed and load on the processors involved. This translated to measured figures of 15, 28, 35, and 50 milliseconds for typical broadcasts to 3, 7, 15, and 31 hosts respectively. For user process broadcasts, the figures vary widely, owing to the fact that *pvmd* processes deliver broadcasts sequentially to local recipients. Thus, if a large percentage of the user process broadcast group were physically executing on one host, the sequential delivery time for that host would dominate the total broadcast time. Table 2 shows typical time requirements for user broadcast, under the assumption that the broadcast groups are evenly distributed among participating hosts.

Owing to the manner in which process initiation and mutual exclusion are implemented, the time taken for these operations are almost identical to that for 100-byte broadcasts between *pvmd* processes. In the current implementation, barrier synchronization is also performed using broadcasts. When a user process executes a barrier call, the local *pvmd* process first waits until all participating user components on this host also arrive at the barrier. This aggregate notification is then broadcast to the other *pvmd* processes. When all local and remote user processes are known to have arrived at the barrier, each *pvmd* signals the processes local to its processor to proceed. The

No. of processors	No. of processes	Message size (bytes)				
		8	128	256	512	1024
4	4	18 ms	22 ms	26 ms	30 ms	35 ms
4	8	21 ms	26 ms	30 ms	35 ms	41 ms
8	8	39 ms	48 ms	60 ms	75 ms	90 ms
8	16	47 ms	57 ms	70 ms	86 ms	105 ms
16	16	65 ms	76 ms	91 ms	110 ms	130 ms
16	32	95 ms	110 ms	125 ms	145 ms	180 ms

Table 2 : User process broadcast timings in PVM

performance of the barrier PVM primitive was tested using a sample application that invoked the barrier construct twice consecutively, with the second invocation being timed. Shown in Table 3 are the observed measurements. When the number of participants is large, the burst of simultaneous broadcasts resulted in lost messages, accounting for the disproportionate increase in time.

No. of processors	No. of processes	Time (msecs)
4	4	40
4	8	100
8	8	150
8	16	250
16	16	1200
16	32	1550
32	32	4000
32	64	7100

Table 3 : Barrier synchronization timings in PVM

Although the PVM system has not yet been used for a heterogeneous application (one in which different components have different requirements), two existing multiprocessor codes have been ported to run on the system. The first is numerical integration using the rectangle rule. The results from this experiment were uninterestingly predictable; on PVM, scaling in the number of processors or the number of rectangles resulted in a linear performance increase. Furthermore, the performance ratio between PVM and the iPSC/2 multiprocessor for this problem was constant — and consistent with the inherent processor speed differences. The second application is Cholesky matrix factorization [16] — an application that has a relatively high communication to computation ratio. Table 4 shows the elapsed times for this problem run on a network of Sun 3/50 machines for varying problem sizes. Shown in parentheses adjacent to each timing are the corresponding times for running the same program on an Intel iPSC/2 hypercube. In this experiment, no attempt was made to place more than one component process on a host; however, it should be noted that each participating host was also being used simultaneously for general purpose editing, compilation, etc.

No. of processors	Problem size (Order of Matrix)							
	100		200		500		1000	
2	6	(2)	35	(30)	260	(245)	1950	(1923)
4	7	(2)	22	(16)	130	(120)	990	(970)
8	9	(1.5)	17	(9)	75	(64)	610	(490)
16	14	(1)	12	(6)	46	(34)	342	(255)

Table 4 : Times (in seconds) for Cholesky factorization

The anomaly apparent in the first column of Table 4 was traced down to the nature of the Cholesky algorithm — for a small matrix, the simultaneous broadcasts of every process's matrix column to a (relatively) large number of processors resulted in a high percentage of dropped packets, leading to retransmissions and elapsed timeouts. For larger problems however, it can be seen that the performance of PVM is acceptable at the least, considering that the application is a substantially communication oriented one and that general purpose machines on a local network were used. It should also be pointed out that these figures are 2 to 4 times better than those for other distributed multiprocessor simulators such as dsim[5]. Furthermore, the factorization program was built for performance measurement purposes and therefore internally generated the matrix elements and did not output the factorized results. Given the usual difficulty and inefficiencies in I/O from within the nodes in a distributed memory machine, it is expected that PVM will compare much more favorably against hardware multiprocessors when significant amounts of I/O are performed.

5. Conclusions & Future Work

The primary motivation for the PVM project is derived from the existing and anticipated need for a general, flexible, and inexpensive concurrent programming environment. The successful implementation of the system has demonstrated that such a framework can be provided and can execute on existing hardware bases, with the benefits of a procedural programming interface and straightforward constructs for access to various resources. The most noteworthy features of PVM are the support of multiple models of computation and a heterogeneous collection of machines; the provided framework enables interaction between application components and machine architectures that are normally incompatible. Anticipating that large and complex parallel systems will require error indication and failure detection capabilities, such features have been built into the PVM system and are likely to be valuable. From the performance point of view, PVM has proven to be acceptable even for applications with a high communication to computation ratio — although its primary intent is to support applications with much larger grainsize and less interaction. Perhaps of more importance in certain situations is the ability of PVM to utilize resources that already exist and would be wasted otherwise, not to mention its value as a prototyping tool for new algorithms or applications. The simplicity of porting the PVM system as well as application software also enhances its appeal and will contribute to its increased use.

There are, however, several aspects to PVM that require further work; some efforts are ongoing while others are planned for the future. It is evident that the data transfer, broadcast, and mutual exclusion protocols are the most crucial parts of the system, and work is in progress to improve these. Barrier synchronization should be implemented using a more effective strategy such as dimensional exchange to reduce message losses when a large number of processes arrive at a barrier simultaneously. Conflict resolution in the locking algorithm presently uses a heuristic method — a provably correct formalization of this will be undertaken soon. From the application point of view, certain additional features might be desirable such as (1) the ability to coalesce emulated and real shared memory, and (2) to dynamically optimize message passing, locking etc., depending on the architecture on which a component instance is executing. Also planned for the future are a graphical interface for the specification of component execution order and interactions, as well as debugging and execution history trace facilities.

Acknowledgements

The author would like to thank M. T. Heath, G. A. Geist, T. H. Dunigan, and D. A. Poplawski for helpful discussions during the course of this work and for their comments on earlier versions of this paper.

References

- [1] K. Hwang, F. A. Briggs, *Computer Architecture and Parallel Processing*, McGraw-Hill, New York, 1984.
- [2] H. Narang, R. Flanery, J. Drake, *Design of a Simulation Interface for a Parallel Computing Environment*, Oak Ridge National Laboratories Report, preprint.
- [3] G. A. Geist, M. T. Heath, B. W. Peyton, P. H. Worley, *A Machine Independent Communication Library*, Proceedings of the Hypercube Concurrent Computers Conference 1989, J. Gustafson ed., to appear.
- [4] T. J. Gardner, I. M. Gerard, C. R. Mowers, E. Nemeth, R. B. Schnabel, *DPUP : A Distributed Processing Utilities Package*, Computer Science Technical Report, University of Colorado, Boulder, 1986.
- [5] T. H. Dunigan, *Hypercube Simulation on a Local Area Network*, Oak Ridge National Laboratory Report ORNL/TM-10685, November 1988.
- [6] C. Seitz, J. Seizovic, W. K. Su, *The C programmer's Abbreviated Guide to Multicomputer Programming*, Caltech Computer Science Report, CS-TR-88-1, January 1988.
- [7] M. Sullivan, D. Anderson, *Marionette: A System for Parallel Distributed Programming Using a Master/Slave Model*, Proceedings of the 9th International Conference on Distributed Computing Systems, June 1989, pp. 181-188.
- [8] G. Riccardi, B. Traversat, U. Chandra, *A Master-Slaves Parallel Computation Model*, Supercomputer Computations Research Institute Report, Florida State University, June 1989.

- [9] Intel iPSC/2 Programmer's Reference Manual, Intel Corporation, Beaverton, OR, March 1988.
- [10] A. Karp, *Programming for Parallelism*, IEEE Computer, May 1987, pp. 43-57.
- [11] J. B. Postel, *User Datagram Protocol*, Internet Request for Comments RFC-768, August 1980.
- [12] N. Maekawa, *A \sqrt{n} Algorithm for Mutual Exclusion in Decentralized Systems*, ACM Transactions on Computer Systems, May 1985, pp. 145-159.
- [13] K. Raymond, *A Tree Based Algorithm for Distributed Mutual Exclusion*, ACM Transactions on Computer Systems, February 1989, pp. 61-77.
- [14] D. Agarwal, A. E. Abbadi, *An Efficient Solution to the Distributed Mutual Exclusion Problem*, Proceedings of the Principles of Distributed Computing Conference, Edmonton, August 1989 (to appear).
- [15] J. S. Chase *et. al.*, *The Amber System: Parallel Programming on a Network of Multiprocessors*, to appear in 12th SOSP, Litchfield Park, November 1989.
- [16] G. A. Geist, M. T. Heath, *Matrix Factorization on a Hypercube Multiprocessors*, in Hypercube Multiprocessors 1986, SIAM, Philadelphia, 1986, pp. 161-180.

Appendix — PVM User Interface Constructs

enroll(<name>)

Enroll calling process as <name>. Returns instance number.

initiate(<component name>)

Start new process specified by <component name>. Returns instance number.

initiateP(<name1>,<name2>,<inum>)

Start new instance of <name1> when <name2>/<inum> terminates.

initiated(<name>,<event>)

Start new instance of <name> when <event> occurs.

ready(<event>)

Inform PVM system of occurrence of <event>.

terminate(<name>,<inum>)

Terminate instance <inum> of component <name>.

waitprocess(<name>,<inum>)

Block caller until instance <inum> of <name> terminates.

send(<name>,<inum>,<type>)

Send message of specified type to specified destination process. Negative return value on failure.

recv(<type>)

Receive message of specified type.

recv1(<type>,<other_limit>)

Receive message of specified type; <other_limit> msgs of other types allowed. Negative return value on failure.

recv2(<type1>,<timeout>)

Receive message of specified type within <timeout> seconds. Negative return value on failure.

putstring(<string>)

Store <string> in send buffer in machine independent form.

putint(<num>)

Store integer in send buffer in machine independent form.

putfloat(<fnum>)

Store real number in send buffer in machine independent form.

getstring(<string_ptr>)

Retrieve string from receive buffer in machine dependent form.

getint(<integer_ptr>)

Retrieve integer from receive buffer in machine dependent form.

getfloat(<float_ptr>)

Retrieve real number from receive buffer in machine dependent form.

recvinfo(<string_ptr>,<inum_ptr>,<len_ptr>)

Return source name, instance number, and length of last received message.

broadcast(<name>)

Broadcast send buffer to all instances of <name>. Returns actual number of recipients.

barrier()

Blocks caller until all instances arrive at barrier. Negative return value if some instances have terminated.

waituntil(<event>)

Blocks caller until specified event occurs.

shmget(<key>,<size>)

Allocates shared memory segment of specified size identified by <key>. A negative return value indicates that the key value is already in use.

shmat(<key>,<ptr>,<flag>,<timeout>)

Locks shared memory segment identified by <key> and maps segment at caller's address space starting at <ptr>. "R" or "RW" are possible flag values; a negative value is returned if the attach does not succeed within <timeout> seconds. A 0 timeout value causes the caller to block until successful.

shmatint(<key>,<integer_ptr>,<flag>,<timeout>)

Variant of **shmat** that maps segment in typed form as integer array.

shmatfloat(<key>,<float_ptr>,<flag>,<timeout>)

Variant of **shmat** that maps segment in typed form as float array.

shmdt(<key>,<ptr>)

Unlocks and unmaps specified shared memory segment from process' address space indicated by <ptr>.

shmdtint(<key>,<integer_ptr>)

Unlocks and unmaps specified shared memory segment in a typed form from process' address space indicated by <integer_ptr>.

shmdtfloat(<key>,<float_ptr>)

Unlocks and unmaps specified shared memory segment in a typed form from process' address space indicated by <float_ptr>.

shmfree(<key>)

Deallocates specified shared memory segment.

lock(<resource_name>,<timeout>)

Permits exclusive access to abstract resource identified by string-valued <resource_name>. Negative return value indicates resource could not be acquired within specified timeout period.

unlock(<resource_name>)

Releases lock on previously acquired resource.

status(<component_name>,<inum>,<stat_ptr>,<loc_ptr>)

Takes string valued component name and an instance number and returns the status (0=nonexistent, 1=active) and location (processor number in the range 0 — p) of that component instance.

entercomp(<name>,<loc_machine>,<obj_file>,<arch>)

Permits a component description to be added; specifying the component name, the object file name and the machine on which it is located, and the type of architecture that the object will execute on.

shmstat(<key_ptr_ptr>,<stat_ptr>)

Obtains information about shared memory regions. Array of strings holds the key values; array of integers holds status (0=free, 1=locked) information. Return value specifies total number of active regions.

lockstat(<key_ptr_ptr>)

Returns total number of active locks with array of strings holding the key values.

Notes:

Among the PVM constructs described above, those concerned with machine dependent data representation are to be implemented as part of the installation procedure for architectures not represented in the generic distribution of the software. Also deemed part of this procedure are constructs to handle other data types such as double precision, boolean, enumerated types, etc.

INTERNAL DISTRIBUTION

- | | | | |
|--------|-------------------|--------|---|
| 1. | B. R. Appleton | 22. | C. H. Romine |
| 2. | J. J. Dongarra | 23. | R. C. Ward |
| 3. | J. B. Drake | 24. | P. H. Worley |
| 4. | G. A. Geist | 25. | Central Research Library |
| 5-6. | R. F. Harbison | 26. | ORNL Patent Office |
| 7. | M. T. Heath | 27. | K-25 Plant Library |
| 8. | M. R. Leuze | 28. | Y-12 Technical Library/
Document Reference Station |
| 9-13. | F. C. Maienschein | 29. | Laboratory Records - RC |
| 14. | E. G. Ng | 30-31. | Laboratory Records Department |
| 15. | C. E. Oliver | | |
| 16. | G. Ostrouchov | | |
| 17-21. | S. A. Raby | | |

EXTERNAL DISTRIBUTION

32. Dr. Donald M. Austin, Office of Scientific Computing, Office of Energy Research, ER-7, Germantown Building, U.S. Department of Energy, Washington, DC 20545
33. Dr. Robert G. Babb, Department of Computer Science and Engineering, Oregon Graduate Center, 19600 N.W. Walker Road, Beaverton, OR 97006
34. Lawrence J. Baker, Exxon Production Research Company, P.O. Box 2189, Houston, TX 77252-2189
35. Dr. Jesse L. Barlow, Department of Computer Science, Pennsylvania State University, University Park, PA 16802
36. Dr. Edward H. Barsis, Computer Science and Mathematics, P.O. Box Box 5800, Sandia National Laboratory, Albuquerque, NM 87185
37. Dr. Chris Bischof, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
38. Prof. Ake Bjorck, Department of Mathematics, Linkoping University, Linkoping 58183, Sweden
39. Dr. James C. Browne, Department of Computer Sciences, University of Texas, Austin, TX 78712
40. Dr. Bill L. Buzbee, Scientific Computing Division, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307
41. Dr. Donald A. Calahan, Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109

42. Dr. Tony Chan, Department of Mathematics, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90024
43. Dr. Jagdish Chandra, Army Research Office, P.O. Box 12211, Research Office, P.O. Box 12211, Research Triangle Park, NC 27709
44. Dr. Eleanor Chu, Department of Computer Science, University of Waterloo, Ontario, Canada N2L 3G1
45. Prof. Tom Coleman, Department of Computer Science, Cornell University, Ithaca, NY 14853
46. Dr. Paul Concus, Mathematics and Computing, Lawrence Berkeley Laboratory, Berkeley, CA 94720
47. Prof. Andy Conn, Department of Combinatorics, and Optimization, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
48. Dr. Jane K. Cullum, IBM T.J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598
49. Dr. George Cybenko, Computer Science Department, University of Illinois, Urbana, IL 61801
50. Dr. George J. Davis, Department of Mathematics, Georgia State University, Atlanta, GA 30303
51. Prof. John J. Dorning, Department of Nuclear Engineering and Physics, Thornton Hall, McCormick Road, University of Virginia, Charlottesville, VA 22901
52. Dr. Iain Duff, CSS Division, Harwell Laboratory, Didcot, Oxon OX11 0RA, England
53. Prof. Pat Eberlein, Department of Computer Science, SUNY/Buffalo, Buffalo, NY 14260
54. Dr. Stanley Eisenstat, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
55. Dr. Lars Elden, Department of Mathematics, Linkoping University, 58183 Linkoping, Sweden
56. Dr. Howard C. Elman, Computer Science Department, University of Maryland, College Park, MD 20742
57. Dr. Albert M. Erisman, Boeing Computer Services, 565 Andover Park West, Tukwila, WA 98188
58. Dr. Peter Fenyves, General Motors Research Laboratory, Department 15, GM Technical Center, Warren, MI 48090
59. Prof. David Fisher, Department of Mathematics, Harvey Mudd College, Claremont, CA 91711
60. Dr. Geoffrey C. Fox, Booth Computing Center 158-79, California Institute of Technology, Pasadena, CA 91125

61. Dr. Paul O. Frederickson, Computing Division, Los Alamos National Laboratory, Los Alamos, NM 87545
62. Dr. Fred N. Fritsch, L-300, Mathematics and Statistics Division, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550
63. Dr. Robert E. Funderlic, Department of Computer Science, North Carolina State University, Raleigh, NC 27650
64. Dr. Dennis B. Gannon, Computer Science Department, Indiana University, Bloomington, IN 47405
65. Dr. David M. Gay, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
66. Dr. C. William Gear, Computer Science Department, University of Illinois, Urbana, IL 61801
67. Dr. W. Morven Gentleman, Division of Electrical Engineering, National Research Council, Building M-50, Room 344, Montreal Road, Ottawa, Ontario, Canada K1A 0R8
68. Dr. Alan George, Vice President, Academic and Provost, Needles Hall, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
69. Dr. John Gilbert, Xerox Palo Alto Research Center, 3333 Coyote Hill Road Palo Alto, CA 94304
70. Dr. Jacob D. Goldstein, The Analytic Sciences Corporation, 55 Walkers Brook Drive, Reading, MA 01867
71. Prof. Gene H. Golub, Department of Computer Science, Stanford University, Stanford, CA 94305
72. Dr. Joseph F. Grear, Division 8331, Sandia National Laboratories, Livermore, CA 94550
73. Dr. Per Christian Hansen, Copenhagen University Observatory, @Ø@ster Voldgade 3, DK-1350 Copenhagen K, Denmark
74. Prof. Robert M. Haralick, Boeing Clairmont Egtvedt Prof., Department of Electrical Engineering, Director, Intelligent Systems Laboratory, University of Washington, 402 Electrical Engineering, Building, FT-10, Seattle, Washington 98195
75. Dr. Don E. Heller, Physics and Computer Science Department, Shell Development Co., P.O. Box 481, Houston, TX 77001
76. Dr. F. J. Helton, GA Technologies, P.O. Box 81608, San Diego, CA 92188
77. Dr. Charles J. Holland, Air Force Office of Scientific Research, Building 410, Bolling Air Force Base, Washington, DC 20332
78. Dr. Robert E. Huddleston, Computation Department, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94550
79. Dr. Ilse Ipsen, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520

80. Ms. Elizabeth Jessup, Department of Computer Science, Yale University, P.O. Box 2158, Yale Station, New Haven, CT 06520
81. Prof. Barry Joe, Department of Computer Science, University of Alberta, Edmonton, Alberta, Canada T6G 2H1
82. Dr. Harry Jordan, Department of Electrical and Computer Engineering, University of Colorado, Boulder, CO 80309
83. Dr. Bo Kagstrom, Institute of Information Processing, University of Umea, 5-901 87 Umea, Sweden
84. Dr. Hans Kaper, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
85. Dr. Linda Kaufman, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
86. Dr. Robert J. Kee, Applied Mathematics Division 8331, Sandia National Laboratories, Livermore, CA 94550
87. Dr. Kenneth Kennedy, Department of Computer Science, Rice University, P.O. Box 1892, Houston, TX 77001
88. Ms. Virginia Klema, Statistics Center, E40-131, MIT, Cambridge, MA 02139
89. Dr. Richard Lau, Office of Naval Research, 1030 E.Green Street, Pasadena, CA 91101
90. Dr. Alan J. Laub, Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106
91. Dr. Robert L. Launer, Army Research Office, P.O. Box 12211, Research Triangle Park, NC 27709
92. Dr. Charles Lawson, MS 301-490, Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109
93. Prof. Peter D. Lax, Director, Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012
94. Dr. James E. Leiss, 13013 Chestnut Oak Drive, Gaithersburg, MD 20878
95. Dr. John G. Lewis, Boeing Computer Services, P.O. Box 24346, M/S 7L-21, Seattle, WA 98124-0346
96. Dr. Heather M. Liddell, Director, Center for Parallel Computing, Department of Computer Science and Statistics, Queen Mary College, University of London, Mile End Road, London E1 4NS, England
97. Dr. Joseph Liu, Department of Computer Science, York University, 4700 Keele Street, Downsview, Ontario, Canada M3J 1P3
98. Dr. Franklin Luk, Electrical Engineering Department, Cornell University, Ithaca, NY 14853
99. James G. Malone, General Motors Research Laboratories, Warren, MI 48090-9055
100. Dr. Thomas A. Manteuffel, Computing Division, Los Alamos National Laboratory, Los Alamos, NM 87545

101. Dr. Bernard McDonald, National Science Foundation, 1800 G Street, NW, Washington, DC 20550
102. Dr. James McGraw, Lawrence Livermore National Laboratory, L-306, P.O. Box 808, Livermore, CA 94550
103. Dr. Paul C. Messina, California Institute of Technology, Mail Code 158-79, Pasadena, CA 91125
104. Dr. Cleve Moler, Ardent Computers, 550 Del Ray Avenue, Sunnyvale, CA 94086
105. Prof. Neville Moray, Department of Mechanical and Industrial Engineering, University of Illinois, 1206 West Green Street, Urbana, IL 61801
106. Dr. Brent Morris, National Security Agency, Ft. George G. Meade, MD 20755
107. Dr. Dianne P. O'Leary, Computer Science Department, University of Maryland, College Park, MD 20742
108. Dr. James M. Ortega, Department of Applied Mathematics, University of Virginia, Charlottesville, VA 22903
109. Prof. Chris Paige, Department of Computer Science, McGill University, 805 Sherbrooke Street W., Montreal, Quebec, Canada H3A 2K6
110. Dr. John F. Palmer, NCUBE Corporation, 915 E. LaVie Lane, Tempe, AZ 85284
111. Prof. Roy P. Pargas, Department of Computer Science, Clemson University, Clemson, SC 29634-1906
112. Prof. Beresford N. Parlett, Department of Mathematics, University of California, Berkeley, CA 94720
113. Prof. Merrell Patrick, Department of Computer Science, Duke University, Durham, NC 27706
114. Dr. Robert J. Plemmons, Departments of Mathematics and Computer Science, North Carolina State University, Raleigh, NC 27650
115. Dr. Alex Pothan, Department of Computer Science, Pennsylvania State University, University Park, PA 16802
116. Dr. John K. Reid, CSS Division, Building 8.9, AERE Harwell, Didcot, Oxon, England OX11 0RA
117. Dr. John R. Rice, Computer Science Department, Purdue University, West Lafayette, IN 47907
118. Dr. Garry Rodrigue, Numerical Mathematics Group, Lawrence Livermore Laboratory, Livermore, CA 94550
119. Dr. Donald J. Rose, Department of Computer Science, Duke University, Durham, NC 27706
120. Dr. Ahmed H. Sameh, Computer Science Department, University of Illinois, Urbana, IL 61801

121. Dr. Michael Saunders, Systems Optimization Laboratory, Operations Research Department, Stanford University, Stanford, CA 94305
122. Dr. Robert Schreiber, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180
123. Dr. Martin H. Schultz, Department of Computer Science, Yale University, P.O. Box 2158 Yale Station, New Haven, CT 06520
124. Dr. David S. Scott, Intel Scientific Computers, 15201 N.W.Greenbrier Parkway, Beaverton, OR 97006
125. Dr. Lawrence F. Shampine, Mathematics Department, Southern Methodist University, Dallas, TX 75275
126. Dr. Kermit Sigmon, Department of Mathematics, University of Florida, Gainesville, FL 32611
127. Dr. Danny C. Sorensen, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
128. Prof. G. W. Stewart, Computer Science Department, University of Maryland, College Park, MD 20742
- 129-133. Dr. V. S. Sunderam, Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322
134. Dr. Kosmo D. Tatalias, Atlantic Aerospace Electronics Corporation, 6404 Ivy Lane, Suite 300, Breenbelt, MD 20770-1406
135. Prof. Charles Van Loan, Department of Computer Science, Cornell University, Ithaca, NY 14853
136. Dr. Robert G. Voigt, ICASE, MS 132-C, NASA Langley Research Center, Hampton, VA 23665
137. Prof. Mary F. Wheeler, Mathematics Department, University of Houston, 4800 Calhoun, Houston, TX 77204-3476
138. Dr. Andrew B. White, Computing Division, Los Alamos National Laboratory, Los Alamos, NM 87545
139. Dr. Arthur Wouk, Army Research Office, P.O. Box 12211, Research Triangle Park, NC 27709
140. Dr. Margaret Wright, Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974
141. Dr. A. Yerebin, Department of Numerical Mathematics of the USSR Academy of Sciences, Gorki Street 11, Moscow, 103905, USSR
142. Office of Assistant Manager for Energy Research and Development, U.S. Department of Energy, Oak Ridge Operations Office, P.O. Box 2001 Oak Ridge, TN 37831-8600
- 143-152. Office of Scientific & Technical Information, P.O. Box 62, Oak Ridge, TN 37831