ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, IL 60439

ANL/MCS-TM-199

Users Guide for the ANL IBM SP_x

by

William Gropp and Ewing Lusk

Mathematics and Computer Science Division

Technical Memorandum No. 199

<http://www.mcs.anl.gov/sp1/guide-r2/guide-r2.html>

December 1994

MASTER

This work was supported by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DC

DISCLAIMER

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, make 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.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Contents

Abstract	v
1 Introduction	1
1.1 Getting Started	1
1.2 Usage Rules	1
1.3 Comments on This Manual	1
2 Machine Configuration	2
2.1 Hardware	2
2.1.1 Details on the High-Performance Switch	2
2.2 Software	2
2.3 What Is Not Provided	3
3 Programming	4
3.1 Transport Layers	4
3.1.1 Ethernet/IP	4
3.1.2 Switch/IP	4
3.1.3 Switch/us	4
3.1.4 Identifying the Transport Layer in Use	4
3.2 Parallel Programming Libraries	5
3.2.1 MPL, POE, and PE	5
3.2.2 Chameleon	5
3.2.3 Fortran M	6
3.2.4 MPI	6
3.2.5 p4	7
4 Using the System	9
4.1 Compiling and Linking Applications	9
4.2 Running Applications	10
4.3 Parallel Unix Tools	10
4.4 Displaying System Information	10
4.5 Scheduling Use of the SPx	10
4.6 Additional Information	10
5 The File Systems	12
6 The Fiber Channel	13
7 Correctness and Performance Debugging	14
7.1 pdbx and xpdbx	14
7.2 dbx	14
7.3 Chameleon Options	14

7.4 upshot	15
7.5 vt	15
8 Nonportable Programming	16
8.1 Using MPL	16
9 Benchmarking	17
10 Known Problems	18
10.1 IBM Documentation	18
11 In Case of Difficulty	20
12 Reporting Problems	25
13 Miscellaneous	26
13.1 Selecting Interrupt-driven MPL	26
13.2 Interrupt-driven Receives in MPL	26
13.3 Threads	26
Acknowledgments	27
Bibliography	28

Users Guide for the ANL IBM SPx

by

William Gropp and Ewing Lusk

Abstract

This guide presents the features of the IBM SPx installed in the Mathematics and Computer Science Division at Argonne National Laboratory. The guide describes the available hardware and software, access policies, and hints for using the system productively.

Chapter 1

Introduction

This guide presents the features of the IBM SPx installed in the Mathematics and Computer Science Division at Argonne National Laboratory. The guide describes the available hardware and software, access policies, and hints for using the system productively.

This document is available online through any World Wide Web (WWW) reader, such as `xmosaic`, `tkwww`, or `www`. Many of the tools mentioned in this document also have on-line documentation. In particular, `p4` and Chameleon have info documentation; Chameleon also has an extensive set of man pages.

1.1 Getting Started

This section describes how to get an account and get started on the ANL SPx.

To get an account, send e-mail to `spaccount@mcs.anl.gov` and request an SPx account. Fill out the form that is returned to you. You will be notified when your application is approved or denied.

Publications resulting from your work on the Argonne SP1 should use the following acknowledgment.

The author(s) gratefully acknowledge use of the Argonne High-Performance Computing Research Facility. The HPCRF is funded principally by the U.S. Department of Energy Office of Scientific Computing.

To get started once you have an account, log into either `bonnie.mcs.anl.gov` or `clyde.mcs.anl.gov`. These are IBM RS/6000s that serve as compile servers and file system servers for the IBM SPx. If you edit your `.login` or `.cshrc` files, make sure that you understand the consequences of any changes that you make to your `PATH` or `MANPATH`.

To change your default shell, send mail to `spsupport@mcs.anl.gov`. The only shells supported are `csh`, `ksh`, `sh`, and `tcsh`.

The nodes of the SPx are named `spnode1` through `spnode128`. Names of the form `spnode001` also work. Access to these is through a locally developed scheduler that is described below. If you have received interactive access to some of the nodes, you may log into them by using `rlogin`, `telnet`, or `rsh`.

1.2 Usage Rules

Usage of the SPx is now managed by the `spsubmit` program. Information may be found on the Web at <http://www.mcs.anl.gov/home/rayl/scheduler/scheduler.html>.

1.3 Comments on This Manual

Please send any comments on this manual to `gropp@mcs.anl.gov`. We are particularly interested in new items for Chapters 10 and 11 [Known Problems] and [In Case of Difficulty].

Chapter 2

Machine Configuration

This chapter discusses the hardware and software configuration of the Argonne SPx, as well as what is not provided on this system.

2.1 Hardware

The Argonne SPx consists of 128 nodes and two compile servers. Each node is essentially an RS/6000 model 370. This model has a 62.5 MHz clock, a 32 KB data cache, and a 32 KB instruction cache. Key features of this system are as follows.

- 128 MB of memory per node
- 1 GB local disk on each node (400 MB are available to users; the rest is for paging and the operating system)
- Full Unix on each node (IBM AIX 3.2.5)
- Accessibility of each node by Ethernet from the Internet
- High-performance Omega switch (63 μ sec latency, 35 MB/sec bandwidth when using "user space" option)

In addition, the ANL SPx has a large high-performance file system (220 GB of RAID disk and a 6 TB automated tape library).

The peak performance of each node is 125 MFlops (1 64-bit floating-point add and 1 floating-point multiply in each clock cycle). In practice, each node can achieve between 15 and 70 MFlops on Fortran code. Higher performance can be reached by using the BLAS or ESSL routines.

2.1.1 Details on the High-Performance Switch

References for Omega networks include [8] and [7].

2.2 Software

Since each SPx node is running a full Unix, most of the usual Unix tools are available. Users may log directly into any SPx node by using `telnet`, `rlogin`, or `rsh`, as long as they have been granted access by the `spsubmit` program. The tools include

- multiple parallel programming environments (see Section 3 [Programming], page 4),
- IBM's ESSL library, and
- performance debugging tools.

2.3 What Is Not Provided

While the nodes support a full AIX, several services are not supported. These include mail, printing, and access to other file systems (with the exception of `/sphone`; see Section 5 [The File Systems], page 12).

We have made no effort to provide PVM or PVM/e on our system.

The only mathematics software libraries on the SPx are the standard C library `'libm'` (which contains some special functions, such as Bessel functions), IBM's ESSL (includes the BLAS and LAPACK), and various research libraries. Use `-lessl` on your link line to include ESSL.

We do not support the usual MCS environment on the nodes. In particular, if you use a `.cshrc` or `.login` file that needs things like `'/mcs/etc/path'`, it will fail, and you will need to modify it (if you are using `'/mcs/etc/path'`, you should look at the man page on `software`. This is the supported method of getting the basic environment, though it may cause problems for `sh` users).

We also provide no user-support services. We run a "dark" machine room; this means that there is no "operator" to respond to user requests, particularly outside of normal working hours. There are no consultants to help with writing or debugging your programs. We do respond to mail about problems with the machine, and the various research groups often express interest in helping people use their tools and in suggesting new features. See Section 12 [Reporting Problems], page 25 for more information.

Please see Section 10 [Known Problems], page 18 for a list of known problems. See Section 11 [In Case of Difficulty], page 20, for help in identifying problems; please check it before sending mail about a problem with the machine.

Chapter 3

Programming

3.1 Transport Layers

There are a number of *transport layers*, or ways to communicate between nodes, on the SPx. For most uses, programmers will not use these directly; rather, they will use one of the portable programming libraries. However, since the programming libraries use these transport layers to actually accomplish the communication, it is important to understand them so that the proper transport layer can be chosen.

The available transport layers are Ethernet/IP, Switch/IP, and Switch/us. Only the first two support multiple parallel jobs on the same node. Switch/us provides better performance (particularly reduced latency) but can run only as a single process-per-node. Switch/us jobs may coexist with other processes on the same node as long as only one process is using the Switch/us mode.

3.1.1 Ethernet/IP

All nodes in the SPx are connected by Ethernet, with node names `spnode001` through `spnode128` (leading zeros may be ignored; thus `spnode001` and `spnode1` are the same node). With this choice of transport layer, the SPx looks just like a collection of workstations. This method does suffer from the same drawbacks as any Ethernet-connected system: high latency (about 1 msec), low bandwidth (1 MB/sec), and low scalability (the 1 MB/sec is shared among all processors).

3.1.2 Switch/IP

All of the nodes in the SPx are also connected by the high-performance switch (HPS). One easy way to use the HPS is to use the same code but with different node names. Using `swnode...` in place of `spnode...` causes the HPS to be used. This may be used with operations such as `rsh` and `rcp`, and by parallel programming libraries such as MPL and MPI.

3.1.3 Switch/us

The same high-performance switch that provides the Switch/IP interface also has a higher-performance version that uses a proprietary interface for accessing the HPS. This version is called the "user space" version because it bypasses the AIX operating system (a source of the high latency in the Switch/IP method) and accesses the HPS directly. A drawback to this approach is that only one process per node (or processor) may use this mode to access the HPS.

3.1.4 Identifying the Transport Layer in Use

The transport layer used depends on the parallel environment chosen and possibly the hostnames. The relationships are shown in Table 3.1.

Table 3.1: Relationship between node names and transport layers

Nodename	Environment	Transport Layer
spnodex	Unix	Ethernet
swnodex	Unix	Switch/IP
fcnodex	Unix	Fiber Channel
any node	MP_EUILIB=us	Switch/us
any node	MP_EUILIB=ip	Switch/IP

3.2 Parallel Programming Libraries

We strongly encourage all programmers to use one or more of the systems described in this section when developing parallel programs for the SPx. These have been used on other systems and have made it easy to port significant applications to the SPx very quickly. These systems also offer a spectrum of tradeoffs between functionality and performance; the Chameleon system in particular has no overhead when used in "production" mode.

3.2.1 MPL, POE, and PE

MPL (formerly EUI) is IBM's message-passing interface to the high-performance switch. POE is the parallel operating environment and controls the running of MPL programs. PE is the parallel environment and is IBM's term for all of the parallel tools (such as MPL, POE, pdbx, etc.) PE supports a parallel symbolic debugger (pdbx; xpdbx is not yet available) and a performance visualization tool (vt). MPL is documented in [6]. Any environment that makes use of Switch/us uses POE; all options described in the man page on poe (man poe) may be used.

3.2.2 Chameleon

Chameleon is a lightweight, portable message-passing system. It provides access to a wide range of transport layers, including MPL, MPI, PVM (not supported on the ANL SPx), and p4. Chameleon provides a common startup model that simplifies choosing a transport layer.

Examples: '/home/gropp/tools.n/comm/examples' contains C and Fortran examples as well as makefiles.

Documentation: '/home/gropp/tools.n/docs/tutorial/parallel.tex'. The script '/home/gropp/tools.n/bin/toolman' will start an xman for the manual pages for the Chameleon routines (and others). See also [4].

Supported Transport Layers: MPL, MPI, p4

Special Comments: The Chameleon makefiles provide portability by using names defined on the make command line. You should set ARCH=rs6000 and BOPT=g (for debugging) or BOPT=O (for production) and COMM=p4, COMM=eui, or COMM=mpl for the interfaces p4, MPL, and MPI, respectively. If you choose not to use the Chameleon makefiles, be sure that you get all of the required options (e.g., C programs require -D_POSIX_SOURCE).

Contact: gropp@mcs.anl.gov for more information.

How to compile and link: Examples may be found in '/usr/local/tools.core/comm/examples'. Chameleon uses fairly complicated makefiles to achieve portability to a wide variety of systems; you should look at the Chameleon manual [4] for more details. The value of ARCH for the SPx is rs6000. If you are using the usual Chameleon makefile, an appropriate make line for MPL is

make ARCH=rs6000 COMM=eui BOPT=0

How to run: Chameleon provides a nearly consistent interface for all transport layers. The special cases are detailed here by transport layer:

COMM=eui Run as an MPL program.

COMM=p4 To use p4 with Ethernet, you need to do

setenv TOOLSHOSTS /sphome/gropp/hosts

To use p4 with the high-performance switch, you need to use

setenv TOOLSHOSTS /sphome/gropp/hosts.hps

and the command-line argument **-p4 altnet**.

COMM=pvm Not yet supported.

COMM=pvm3 Not yet supported.

Note that when using **COMM=eui**, you must use the MPL syntax to run a parallel program. In particular, you need to set the environment variables **MP_PROCS**, **MP_EUILIB**, etc. You can not use the **-np** command-line argument to choose the number of processes. See the section on running MPL jobs in the **spsubmit** documentation (<http://www.mcs.anl.gov/home/rayl/scheduler/scheduler.html>).

3.2.3 Fortran M

Fortran M is a small set of extensions to Fortran that supports a *modular* approach to the construction of sequential and parallel programs. Fortran M programs use *channels* to plug together *processes* that may be written in Fortran M or Fortran 77. Processes communicate by sending and receiving messages on channels. Channels and processes can be created dynamically, but programs remain deterministic unless specialized nondeterministic constructs are used.

Examples: `/usr/local/fm/examples` contains Fortran M examples as well as makefiles.

Documentation: Bound hardcopies (the ANL Technical Report) may be obtained from the contact listed below.

Supported Transport Layers: Ethernet/IP and Switch/IP

Special Comments: See the "Network Specifics" section of the manual for details on running Fortran M on the SPx. The hostnames of the switch interface (i.e., `swnode1`) should be used. The latest version of the Fortran M compiler is installed in `/usr/local/fm`.

Contact: `fortran-m@mcs.anl.gov` for more information.

3.2.4 MPI

MPI (Message Passing Interface) is a new message-passing system "standard" that has recently been defined by a broadly based group of parallel computing vendors, library writers (including us), and users. The current draft is now in the public-comment stage [3]. It was completed in the spring of 1994. The standard draft is available by anonymous ftp from `info.mcs.anl.gov` in `pub/mpi/mpi-report.ps.Z` or on the WWW in <http://www.mcs.anl.gov/mpi/mpi-report/mpi-report.html>.

There are two implementations: one being done by us and a group from Mississippi State University, and the other being done by IBM. Both are relatively complete and run now on the SPx.

3.2.4.1 The IBM Version

An experimental IBM implementation of MPI, from IBM Yorktown, is available. This version works with POE/MPL system. It is located in `/usr/lpp/mpif`.

Examples: The directory `/usr/lpp/mpif/samples` contains examples in both Fortran and C as well as a makefile.

Documentation: There is documentation in `/usr/lpp/mpif/docu/mpif.ps` and `/usr/lpp/mpif/docu/errata.tex`. Documentation on MPI is available in the MPI WWW page. (<http://www.mcs.anl.gov/mpi>)

Supported Transport Layers: Switch/IP and Switch/us.

Contact: gropp@mcs.anl.gov or lusk@mcs.anl.gov for more information.

How to compile and link: No special options are needed when compiling. The include file `'mpi.h'` is located in `/usr/lpp/mpif/include` and is needed by an C program. The include file `'mpif.h'` should be included by any Fortran routine that uses an MPI call. However, you need to use `mpicc` instead of `mpcc` and `mpixlf` instead of `mpxlf`.

How to run: MPI-F programs are run using much the same approach as MPL programs. To run a program

```
poe a.out ...
```

use the same environment variables (e.g., `MP_EUILIB`) to control the transport layer, choice of nodes, and level of messages.

3.2.4.2 The MPICH Version

Examples: `/usr/local mpi/examples` contains examples in Fortran and C, and a makefile.

Documentation: No implementation-specific documentation is yet available. Our implementation uses Chameleon, so the usual Chameleon methods of preparing programs and starting jobs apply. There are man pages for all of the MPI routines in `/usr/local/mpi/man`.

Supported Transport Layers: Switch/IP, Switch/US, Ethernet/IP.

Contact: gropp@mcs.anl.gov or lusk@mcs.anl.gov for more information.

3.2.5 p4

p4 [1, 2] is a portable message-passing system that runs on a very wide variety of parallel systems and workstations. It is in use at approximately 200 sites around the world. Existing p4 programs will run unchanged on the SPx. The current version is version 1.4, but experimental versions (such as 1.4a) and later releases (1.5) may appear from time to time.

Examples: `/usr/local/p4-1.4a/SP2/examples` on `bonnie` and `clyde` contains C and Fortran examples, example makefiles, and a makefile that can be used to construct your makefiles with the appropriate options for the SPx's various transport layers.

Documentation: `/home/lusk/ibm/p4-1.4b/doc` contains the latexinfo source for the manual [2], an ASCII version, and the Postscript for a reference card. `/home/lusk/p4.manual` contains the Postscript for the manual itself in (`'p4.ps'`). The manual is also online via `info` and any of the inside-Emacs or stand-alone `info` readers (e.g., `gnu-info`).

Supported Transport Layers: Ethernet/IP, Switch/IP, Switch/us.

Special Comments: See the section of the p4 manual entitled "Running on Specific Machines" for how to specify each of the transport layers.

Contact: lusk@mcs.anl.gov for more information.

How to compile and link: For MPL, you must use a p4 made with P4ARCH=SP1_EUIH. At Argonne, these libraries are in '/usr/local/p4-1.4a/SP2' where x.y are currently 1.3c.

How to run: To use p4 over MPL on the SPx, invoke your program from a node with

```
<programe> -procs <numprocs> <user args>
```

Specific nodes can be chosen by the same method as for other POE jobs. The p4 procgroup file should contain the single line

```
local <numprocs-1> <pathname of program>
```

To use Switch/IP, you need a procgroup file of the form

```
swnode1 0 executable  
swnode2 1 executable  
swnode3 1 executable
```

The line swnode 0 executable replaces the usual local 0.

Chapter 4

Using the System

This section describes how to run parallel programs and how to monitor the state of the system. (IBM provides a powerful system monitoring utility, but it is only available to `root`. The monitoring tools described here were locally developed.)

4.1 Compiling and Linking Applications

All compilation and linking should be done on the compile servers `bonnie.mcs.anl.gov` and `clyde.mcs.anl.gov`. The compile servers should not be used to run any other programs.

The compilers are named `xlc` (for C) and `xlf` (for Fortran). (For MPL programs you must use `mpcc` and `mpxlf` instead.) These support most of the usual options but sometimes with unusual (for Unix systems) names. Useful options include

- `-c` Compile to a `.o` file
- `-O3` Produce optimized code
- `-qstrict` Ensure that optimizations do not alter program semantics
- `-qsource -qxref` Produce a source listing with cross-references
- `-qlist` Produce an object listing. This is equivalent to `-S` on other Unix systems
- `-qwait=-1` Wait indefinitely for a network license token. You may need this if too many people are using the compilers at once, where "too many" is defined by the license server.

The GNU C++ compiler is available in `/usr/local/gcc/bin/c++`. It is not supported for the SPx but should work.

We also have the IBM C++ compiler `xlc`; however, we do not have any information on it (no man pages and no info information). We have been able to glean that to link Fortran subroutines to a program whose main program is in C++, add the library `-lxlf` to your link line.

The following arguments are available only to Fortran programmers:

- `-qautodbl=dblpad` Promote `REAL` declarations to `DOUBLE PRECISION`. This can be useful with Cray codes that require more than 32-bit precision (the default precision of `REAL` on RS/6000's). However, any double-precision or `REAL*8` statements will be promoted to `REAL*16` (128 bits), which is significantly slower and probably not what is wanted.
- `-qdpce` Treat floating-point constants as double precision.

-qextname Add an underscore to all external names. This makes Fortran routines available to C programs that expect Fortran to add an underscore suffix; this is common (but not standard) behavior of many Unix Fortran compilers. Note that some software packages do not expect Fortran users to use this option; their Fortran interfaces adjust to the default behavior of the Fortran compiler. Using **-qextname** with these libraries can lead to unresolved references.

When linking, the following can be useful:

-qloadmap:filename Produce a load map. This is handy for identifying from which library an object file was loaded, or for seeing the total memory requirements of your module.

-o name Provide a name for the module. The usual Unix default of `a.out` is used if no name is specified.

4.2 Running Applications

Each parallel programming system has its own requirements for running an application. In this section, we describe some issues that are common to all systems.

All access is through the scheduler <http://www.mcs.anl.gov/home/rayl/scheduler/scheduler.html>. That document contains information on running programs on the SPx.

4.3 Parallel Unix Tools

We have developed parallel analogs of the common Unix tools `cp`, `ls`, `ps`, and others. These are scalable and relatively fast. They simplify the task of executing commands on all or part of the SPx. However, they are currently unavailable on the SPx. We expect to provide them again in the near future.

4.4 Displaying System Information

Currently, there are four X Windows tools for displaying the state of the system. These are `wish` scripts that require `tk-3.3`. This code is available in the MCS Division. Since `tk` is in the public domain, other sites should have no problem acquiring it. In addition, there is a program provided by IBM that gives some information on POE.

Currently, the only supported tool is `xspusage`. This shows which nodes are in use and which are available.

The program `'spusage'` will display the users of the the SPx nodes in ASCII format. This is useful on a text terminal or if running X over a slow connection.

4.5 Scheduling Use of the SPx

All access to the SPx is made through the use of the `spsubmit` program. This program maintains queues of pending jobs. The program `spq` lists the contents of the queues. `sphelp` lists the programs that are available for submitting jobs and managing them. On-line documentation is provided through WWW.

4.6 Additional Information

There are three IBM manuals on the Fortran (and C) compilers:

The **Language Reference** explains, in addition to describing the Fortran language, a large number of IBM-provided system-interface type of subroutines.

The **User's Guide** describes compiling, linking, use of the preprocessors, interaction of Fortran with files, etc.

The Optimization and Tuning Guide for the XL Fortran and XL CCompilers give a very useful (and readable) description of optimization considerations for the RS/6000 architecture.

If you are using an X Windows system to access the compile server, you may view these manuals on line by using IBM's `info` system. Before invoking `info`, define your display to the compile server:

```
setenv DISPLAY myid.mcs.anl.gov:0.0    (csh)
export DISPLAY=myid.mcs.anl.gov:0.0    (ksh)
info
```

where `myid.mcs.anl.gov` is the TCP/IP address of your computer or X-station. When `info` finally comes up (it can take several minutes), select the "List of Books" panel, find the desired manual, and double-click on it.

When connecting to Sun workstations, `info` (and most IBM X Windows tools) will generate several warning messages. These may be ignored.

Chapter 5

The File Systems

There are currently two file systems accessible to the nodes: `/sphome` (shared) and `/tmp` (local). The third file system (the 220 GB RAID arrays) is not yet available (as of 9/1/94) and will not be discussed in this document.

The file system `/sphome` is shared between the nodes and the compile servers. This file system is NFS mounted.

Warning:

There is a problem where files that are copied over an existing file do not always seem to be noticed. You should always `rm` a file before you copy over it.

Be frugal with your use of `/sphome`. There is only 1 GB of space in `/sphome`, and this must be shared by all users, including those with large input and output files. Try not to keep any large files in `/sphome` except when you need them for a run (don't move files back and forth during the day but also don't leave a large output file on `/sphome` once it is generated). The program `xspou` may be used to display the usage of `/sphome`.

Each local disk holds 400 KBytes and may also be used for temporary data files and for executables, but again, you should not leave large files on them. The program `xspidf` (see Section 4.3 [Parallel Unix Tools], page 10) can be used to display the amount of space available on each node's `/tmp` partition. (The rest of the local disk holds the local `/var` partition and the swap area.)

The SPx nodes do *not* mount the MCS file systems. Thus, the only way to move files is either by first moving them to `/sphome` or by using `rcp`.

For additional information on other file systems being considered or developed for the ANL SPx, see `/home/nickless/Documents/io.ps`. Contact `nickless@mcs.anl.gov` for more information.

Important Note: Neither `/sphome` nor `/tmp` is backed up. In addition, the `/tmp` areas will be cleaned of old files on a regular basis; large files may be removed at any time.

Chapter 6

The Fiber Channel

NOTE: THE FIBER CHANNEL IS NOT AVAILABLE YET; ANY INFORMATION IN THIS SECTION IS SUBJECT TO CHANGE WITHOUT NOTICE

Nodes $2 + 4i$ have Fiber Channel connections for $i = 0, \dots, 31$.

Fiber Channel supports both TCP/IP and direct (`ioctl`) interfaces. The node names for Fiber Channel are `fcnodei`.

Performance is about 2.5 MB/sec for IP and 17 MB/sec for direct.

p4 provides access to Fiber Channel. Here is a sample p4 procgroup file:

```
fcnode2 1 /sphome/lusk/p4test/systest
fcnode6 1 /sphome/lusk/p4test/systest
```

Chapter 7

Correctness and Performance Debugging

Several tools may be used to do correctness and performance debugging of parallel programs. This section describes these tools and how to use them.

7.1 pdbx and xpdbx

IBM's POE provides parallel versions of `dbx` and `xdbx` called `pdbx` and `xpdbx`, respectively. These work with POE/MPL programs and allow you to run and debug parallel applications.

7.2 dbx

It is also possible to use `dbx` and `xdbx` with parallel programs by logging into each node and using the command

```
dbx process_id
```

after the parallel programs have started. You can get the `process_id` with

```
ps -ef | grep $LOGNAME
```

7.3 Chameleon Options

Chameleon provides a number of features that support both correctness and performance debugging. These apply only to programs compiled with `BOPT=g` or `BOPT=OpG`. These features may be selected by command-line switch when the Chameleon application is started:

- event** Generate an event log that may be viewed with `upshot` (see Section 7.4 [upshot], page 15).
- summary** Generate a summary of message passing activity, giving the amount of time spent in each process sending and receiving messages
- trace** Produce a line of output for each message-passing operation (including the beginning and ending of receives). This is useful for finding the cause of deadlock in parallel applications.

Figure 7.1: Sample upshot output

7.4 upshot

Upshot [5] is a portable X Windows program for visualizing the behavior of a parallel program. Both Chameleon and p4 can generate the event logs that upshot reads. To generate an event log from Chameleon, use the `-event` flag

```
a.out -np 4 -event
```

This generates a file 'b1'. To view this file with upshot, use

```
upshot -l b1
```

In order for a Chameleon program to be able to generate this event file, the `BOPT=g` or `BOPT=Opg` version must be used.

Currently, you should have '/home/gropp/bin/sun4' in your path to use upshot on the Sun workstations. The version in '/usr/local/bin' will be updated in the near future.

Sample output from upshot is shown in Figure 7.1.

Chameleon can generate event logs with any transport layer.

7.5 vt

The POE (Parallel Operating Environment) provided by IBM includes a performance visualization tool called vt. This tool, like upshot, can analyze a tracefile that is generated by a parallel program.

Chapter 8

Nonportable Programming

It is possible to use MPL without using the portability tools. This section describes how to run these programs. There are man pages for the individual routines (see `'/usr/man/cat1'` on `bonnie` or `clyde`).

8.1 Using MPL

MPL programs must be compiled with `mpx1f` and `mpx1c` (in fact, only the main program must be compiled with this routines, but it is safer to use them for all routines). The switch `-us` selects Switch/us; the switch `-ip` selects Switch/IP.

Chapter 9

Benchmarking

Benchmarking on the SPx is more difficult than on most other MPPs because the SPx provides a friendly multiprocessing environment on each node. Avoiding the effects of contention with other users requires that the machine be scheduled as single user. To schedule single-user time, see Section 4.5 [Scheduling Use of the SPx], page 10.

It is important to distinguish between elapsed (also called wall-clock) and CPU time when running a parallel program. Depending on the transport layer that is used, CPU time may not include the time in which processes were idle while waiting for messages to arrive. It is best when benchmarking to compute both the elapsed and CPU time; for most purposes, elapsed time is the better indicator of true performance.

The Chameleon package (see Section 3.2.2 [Chameleon], page 5) provides portable routines for measuring both elapsed and CPU time.

```
#include "system/system.h"
double start_cpu, cpu_time, start_elapsed, elapsed_time;
...
start_cpu      = SYGetCPUTime();
start_elapsed = SYGetElapsedTime();
... work to time ...
elapsed_time   = SYGetElapsedTime() - start_elapsed;
cpu_time       = SYGetCPUTime() - start_cpu;
```

Fortran versions of these are also available.

Chapter 10

Known Problems

1. NFS file updates. We are seeing a problem where changing a file on `/sphome` does not seem to actually change the file. Until this is fixed, the best work-around is always to `rm` the file rather than `cp`'ing over it. If you change a program, rebuild it, and then run it and don't see the effect of your change, you are probably seeing this problem.

2. Linking takes *forever*. To speed linking, you can (sometimes) do the following.

- (a) If you are relinking an application, use your old executable as *input* to the linker, along with *only* the modules that you have changed. For example, if the executable is `a.out` and the file `changed.f` contains the *only* changes, then

```
    xlf changed.f a.out
```

will generate a new `a.out` executable. You will also get an error message about "xlf: 1501-218 file a.out contains an incorrect file suffix": you can ignore this message. If you use an executable name with no suffix, you will not get this message.

This also works with `xlc` (including the error message).

- (b) If you are using one or more object libraries, consider moving them to a local disk. On *bonnie*, this means to `/sphome`.
- (c) Consider prebinding any libraries. For example,

```
    ld -r foo.a bar.a -o foobar.o
```

will produce a file, `foobar.o`, that is "prebound"; that is, all of the references between functions and symbols within these files are resolved. This can speed the time to link, particularly if you make sure that this file is local (on *bonnie* or *clyde*, on `/sphome`).

- (d) (**Experts only**): Consider building a *shared* library. Information on this may be found by using `info` and looking for "How to Create a Shared Library."
- (e) Check that only four `biod` processes are running on the compile server. In experiments, using 16 `biod` processes caused a significant increase in linking time compared to 4 `biod` processes.

3. `info` generates "X Error of failed request: BadValue" when it starts. You can ignore this message (it has to do with the fonts used by the X11 server).

10.1 IBM Documentation

This section lists various SP-related documents. Those documents followed by "[== some date, 1993]" are orderable, bound IBM documents. There is at least one copy of these documents in the lab and Tim Lehmann has unbound equivalent copies (with the noted date). All documents not followed by "[== some

date, 1993]" are available unbound (from Tim Lehmann; some cannot be ordered, and some came with the software). There's a day or two turnaround to have copies made for the big documents.

The following SP documents exist for Release 1.0

- IBM 9076 Scalable POWERparallel Systems: Administration Guide, SH26-7221-00, September, 1993 [== -, 1993]
- IBM 9076 Scalable POWERparallel Systems: Maintenance Information, SY66-0299-00, September, 1993 [== -, 1993]
- IBM 9076 Scalable POWERparallel Systems: Planning and Installation Guide, SA23-2481-0, September, 1993 [== -, 1993]
- IBM 9076 System Support Programs, AIX 3.2.4, Release Notes, Release 01, Mod level 00, no date

The following PFS document exists

- Parallel File System External User Interface, Proposal, Version 1.0, September 17, 1993

The following Vesta Documents exist

- Overview of the Vesta Parallel File System, no date
- Satisfying the I/O Requirements of Massively Parallel Supercomputers, no date
- Interfacing Vesta to an External File System, Version 0.3, August 23, 1993
- Vesta File System Programmer's Reference, Version 0.82, August 10, 1993

Chapter 11

In Case of Difficulty

This chapter describes some common problems and possible solutions or work-arounds. You can also send mail to spsupport@mcs.anl.gov if you have questions.

1. Q: I ran my program and it seems to be running the "wrong" program (e.g., print statements that you added do not seem to work).

A: You may be running afoul of the NFS bug (see Section 10 [Known Problems], page 18). Remove the executable and object files and rebuild the program.

2. Q: When running my MPL program, the application never seems to start.

A: Unknown at present.

3. Q: I specify Switch/us but get error messages like this:

```
setenv MP_EUILIB us
setenv MP_HOSTFILE /sphome/<myhostlist>
setenv MP_INFOLEVEL 0
setenv MP_PROCS 4
poe a.out -procs 4
ERROR: 0031-619 mp_euilib specifies us, adapter not us
ERROR: 0031-619 mp_euilib specifies us, adapter not us
ERROR: 0031-619 mp_euilib specifies us, adapter not us
ERROR: 0031-619 mp_euilib specifies us, adapter not us
```

A: This may mean that you linked the executable with `-bnso`. This specifies no shared libraries and caused `mpxlf` or `mpcc` to link in only the Switch/ip libraries.

4. Q: Whenever I run a Switch/us job, I get these error messages:

```
spnode002.mcs.anl.gov:/sphome/gropp(4) poe ring
WARNING: 0031-402 Using css0 as euidevice for User Space job
WARNING: 0031-403 Forcing dedicated adapter for User Space job
WARNING: 0031-403 Forcing dedicated adapter for User Space job
...
WARNING: 0031-403 Forcing dedicated adapter for User Space job
```

with $p+1$ lines of output if I am running on p processors.

A: To disable this feature, you can do

```
setenv MP_INFOLEVEL 0
```

but this will disable all (or almost all) informational and warning messages. There is no known way to select specific messages or to reduce the output to a single line for all processors. The default value of MP_INFOLEVEL is 1; there appears to be no way to change the default.

5. Q: My job fails with these messages:

```
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 023-005 Failed to connect() spnode001.mcs.anl.gov, errno = 79.
jm: 0023-161 connect jmd exceeded max times.
jm: 0023-023 Quitting...
ERROR: 0031-117 Unable to contact Resource Manager
ERROR: 0031-635 Non-zero status -1 returned from pm_mgr_init
```

A: This is caused when the MPL job manager dies (for reasons unknown at this time).

6. Q: I get error messages like this when I run an MPL job (using Switch/ip):

```
ERROR: 0031-161 EOF on socket connection with task 28
ERROR: 0031-160 I/O error on socket connection with task 29
ERROR: 0031-619 : Bad file number
ERROR: 0031-161 EOF on socket connection with task 29
ERROR: 0031-160 I/O error on socket connection with task 31
ERROR: 0031-619 : Bad file number
ERROR: 0031-161 EOF on socket connection with task 31
...
```

A: Unknown at this time.

7. Q: I get error messages like this when I run an MPL job:

```
ERROR: 0031-630 pm_contact: read timeout occurred; nprocs = 32
ERROR: 0031-618 The following nodes were not contacted:
ERROR: 0031-623 task 5: hostname spnode009.mcs.anl.gov
ERROR: 0031-635 Non-zero status -1 returned from pm_mgr_init
```

A: Unknown at this time.

8. Q: I try to compile my program and I get an error message:

```
xlfc -c utility.f
NETLS: StartNetLS: communications failure (network computing system/RPC runtime)
The error code from the last failed command is -1.
```

A: Welcome to the world of licensed software. Our compilers (and possibly other products) are controlled by a network license server. This error message means that the license server could not be contacted; please notify spsupport. Another error that you might encounter is some message about there being no licenses left; you can ask for the compiler to wait for a license with the option -qwait=-1.

9. Q: My program is in '/sphome/\$LOGNAME/myname', but when I try to execute it, I get

```
myprog: Command not found.
```

A: This means that your PATH does not contain the current directory (usually a single "." in the PATH. Either modify your PATH or prefix the program name with ./:

```
./myprog ...
```

Also see Section 1.1 [Getting Started], page 1.

10. Q: My program runs correctly when I compile it -g but not when I compile it with optimizations (-O).

A: The compiler may be generating incorrect code. Try using bisection on your source files: compile half with -g and half with -O, and test again. Continue replacing files compiled with -O with ones compiled with -g until the code works. You may be able to simply use a few modules compiled with -g in your application.

If you can produce a simple example (a few dozen lines of code) that you are sure is incorrectly compiled, send it to spsupport.

11. Q: I tried to login but got messages like

```
cat: cannot open /mcs/etc/path
arch: Command not found.
hostname: Command not found.
biff: Command not found.
tset: Command not found.
hostname: Command not found.
```

A: Your '.login' or '.cshrc' file contains references to the MCS filesystem (e.g., '/mcs/etc/path'). These are unavailable on the nodes. Modify your files accordingly. See '/sphome/INIT/initcsh' for a script to create initial '.login' and '.cshrc' files.

12. Q: I tried to use xspusage but I got the error message

```
bonnie % Xlib: connection to "neptoon:0.0" refused by server
Xlib: Client is not authorized to connect to Server
couldn't connect to display "neptoon:0"
```

A: You need to do xhost + bonnie on the X11 server you are using (in this example, neptoon). Some other tools will require that you have allowed connections from the individual nodes; you may want to add xhost + spnode\$i for i from 1 to 128.

13. Q: My parallel program runs on other parallel machines but seems to deadlock on the SPx when using MPL, MPI, or Chameleon.

A: The following parallel program can deadlock on *any* system when the size of the message being sent is large enough:

```
send( to=partner, data, len, tag )
recv( from=partner, data, maxlen, tag )
```

where these are blocking sends and receives (mp_bsend in MPL and PIbsend in Chameleon). For many systems, deadlock does not occur until the message is very long (often 128 KBytes or more). The limit for MPL is smaller.

To fix this, you have several choices:

- Reorder your send and receive calls so that they are paired up. For example, if there are always an even number of processors, you could use

```

if (myid is even) {
    send( to=partner, data, len, tag )
    recv( from=partner, data, maxlen, tag )
}
else {
    recv( from=partner, data, maxlen, tag )
    send( to=partner, data, len, tag )
}

```

Another possible source of problems is trying to receive messages in a different order than they were sent. For example, if processor 1 does

```

send( to=0, data, len, tag1 )
send( to=0, data, len, tag2 )

```

and processor 0 does

```

recv( from=1, data, len, tag2 )
recv( from=1, data, len, tag1 )

```

and these are blocking receives, then when using MPL, this program may deadlock. To fix this, reorder either the sends or receives.

- Use nonblocking sends and receives instead.
- If you are using Chameleon, you can use the command line argument `-eui nsend=0` to cause Chameleon to provide buffering from user-space for all blocking sends. This works with the `COMM=eui`.

14. Q: My C program include `fcntl.h` but does not find symbols like `O_RDONLY` (needed for `open`).

A: You need to add `-D_POSIX_SOURCE` to your compile flags. This is usually required if you use `xlc` instead of `cc`.

15. Q: I can't seem to run MPL jobs in the background or with `nohup` job; they just seems to die when I log off. (This may happen only to `ksh` users.)

A: (quick answer): Don't log off; just leave an `xterm` running the program.

(long answer; not recently tested): If you are using `ksh`, and you start a normal job or a `p4-socket` job on one of the nodes, using

```
nohup job.....
```

then when you log off, the job continues, providing you used `telnet` to get to the node (`rlogin` probably won't work).

If you make a script for the job and put

```
#!/bin/csh
```

at the beginning of the script and then run

```
nohup script.name &
```

it still dies.

However, if you enter `csh` first,

```

csh
nohup script.name &

```

then all is ok.

16. Q: My Fortran program uses **NAMelist** and no longer works.

A: There are reports that the behavior of **NAMelist** has changed. Now **xlf** conforms to Fortran 90, and the Fortran 90 standard specifies the form and interpretation of **NAMelist**. The earlier version of the Fortran standard, Fortran 77, did *not* include **NAMelist**. However, many Fortran compilers implemented some form of **NAMelist**. The problem comes from incompatibilities between these previous nonstandard definitions of **NAMelist** and the Fortran 90 specification.

Chapter 12

Reporting Problems

If you believe that the SPx has a hardware or software problem, please send mail to `spsupport`. Please be as specific as you can; include samples of code or output, if relevant.

The mail alias `spusers` should be used for sending mail to fellow users of the SPx. This list is appropriate for mail about new tools, user-group meetings, pleas for reducing disk space usage, and the like.

If the `/sphone` file system fills up, you should do the following:

1. Use the command `df /sphone` to verify that the file system really is full.
2. Remove whatever you can. The command

```
du -k /sphone/$LOGNAME
```

will show (in kilobytes) the disk usage of all of your subdirectories.

3. Send a message to the other SPx users alerting them to the fact that `/sphone` is full. This can be done by sending a message to `spusers@mcs.anl.gov`. You might want to include (some) of the output of

```
du -k -s /sphone/* | sort -r -n | head
```

which will show the disk usage of each user (note that `/sphone/pass` is a mounted file system and hence does not take space away from the `/sphone` file system).

Chapter 13

Miscellaneous

13.1 Selecting Interrupt-driven MPL

By default, MPL uses a polling method to check for incoming messages. This has the advantage of reducing the number of interrupts that must be handled, but can in some cases introduce long delays as processes wait for messages to be delivered. Programs containing large amounts of nonblocking communications can benefit from setting the environment variable `MP_CSS_INTERRUPT` to `yes` before running. This works with Chameleon and p4 as well as for MPL jobs started with `poe`.

13.2 Interrupt-driven Receives in MPL

As an *experimental* feature, our version of MPL includes a way to specify that a function should be called when a message of a particular type is received. Please contact Mike Minkoff (minkoff@mcs.anl.gov) for more information.

13.3 Threads

DCE threads are available on all nodes and compile servers.

Acknowledgments

We thank all the users who have contributed to this guide. Particular thanks go to Steve Pieper of the ANL Physics Division, who has been an aggressive, pioneering, and always helpful user and has uncovered many problems for the first time. His initial efforts to document his findings on the use of the SP provided the initial text for the first edition of this guide, and the authors of this edition gratefully acknowledge his contributions.

The work described in this report has benefited from conversations with and use by a large number of people. Steven Tuecke provided the descriptions for Fortran M and a number of valuable comments. Thanks to the many people who provided comments on the early draft of this document. In particular, Bill Nickless and Paul Plassmann made a number of useful suggestions. Tim Lehmann read the document with particular care, identifying a number of places where files or commands referred to had moved.

Bibliography

- [1] Ralph Butler and Ewing Lusk. Monitors, messages, and clusters: The p4 parallel programming system. *Journal of Parallel Computing*. To appear (also Argonne National Laboratory, Mathematics and Computer Science Division preprint MCS-P362-0493).
- [2] Ralph Butler and Ewing Lusk. User's guide to the p4 parallel programming system. Technical Report ANL-92/17, Argonne National Laboratory, October 1992.
- [3] Message Passing Interface Forum. Document for a standard message-passing interface. Technical Report CS-93-214, University of Tennessee, November 1993.
- [4] William D. Gropp and Barry F. Smith. Chameleon parallel programming tools users manual. Technical Report ANL-93/23, Argonne National Laboratory, March 1993.
- [5] Virginia Herrarte and Ewing Lusk. Studying parallel program behavior with Upshot. Technical Report ANL-91/15, Argonne National Laboratory, August 1991.
- [6] IBM. *IBM AIX Parallel Environment Parallel Programming Subroutine Reference Release 2.0*, June 1994.
- [7] D. H. Lawrie. Access and alignment of data in an array processor. *IEEE Transactions on Computers*, C-24(12):1145-1155, December 1975.
- [8] H. J. Siegel. *Interconnection Networks for Large Scale Parallel Processing*. Lexington Books, 1985.