

October 1990

Order Number: 311532-006



**iPSC[®]/2 and iPSC[®]/860
USER'S GUIDE**



intel[®] Corporation

Copyright ©1990 by Intel Scientific Computers, Beaverton, Oregon. All rights reserved. No part of this work may be reproduced or copied in any form or by any means...graphic, electronic, or mechanical including photocopying, taping, or information storage and retrieval systems...without the express written consent of Intel Corporation. The information in this document is subject to change without notice.

Intel Corporation make no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Intel Corporation assumes no responsibility for any errors that may appear in this document. Intel Corporation makes no commitment to update or to keep current the information contained in this document.

Intel Corporation assumes no responsibility for the use of any circuitry other than circuitry embodied in an Intel product. No other circuit patent licenses are implied.

Intel software products are copyrighted by and shall remain the property of Intel Corporation. Use, duplication or disclosure is subject to restrictions stated in Intel's software license, or as define in ASPR-7-104.9(a)(9).

The following are trademarks of Intel Corporation and its affiliates and may be used only to identify Intel products:

286	iDBP	iPSC	Plug-A-Bubble
386	iDIS	iRMX	PROMPT
4-SITE	iLBX	iSBC	Promware
Above	im	iSBX	QueX
BITBUS	Im	iSDM	QUEST Programming
COMMPuter	iMDDX	iSXM	Quick-Pulse
Concurrent File System	iMMX	KEPROM	Ripplemode
Concurrent Workbench	Insite	Library Manager	RMX/80
CREDIT	int l _e	MAP-NET	RUPI
Data Pipeline	int IBOS _e	MCS	Seamless
Direct-Connect Module	Intelevison	Megachassis	SLD
FASTPATH	Intellec	MICROMAINFRAME	SugarCube
GENIUS	int ligen _e t Identifier	MULTIBUS	UPI
i	int ligen _e t Programming	MULTICHANNEL	VLSiCEL
i ² ICE	Intellink	MULTIMODULE	
i860	iOSP	ONCE	
ICE	iPDS	OpenNET	
iCEL		OTP	
iCS		PC BUBBLE	

Ada is a registered trademark of the U.S. Government, Ada Joint Program Office

APSO is a service mark of Verdix Corporation

Ethernet is a registered trademark of XEROX Corporation

Excelan is a trademark of Excelan Corporation

EXOS is a trademark or equipment designator of Excelan Corporation

FORGE is a trademark of Pacific-Sierra Research Corporation

Green Hills Software, C-386, and FORTRAN-386 are trademarks of Green Hills Software, Inc.

GVAS is a trademark of Verdix Corporation

Lucid and Lucid Common Lisp are trademarks of Lucid, Inc.

NFS is a trademark of Sun Microsystems

Sun Microsystems and the combination of Sun and a numeric suffix are trademarks of Sun Microsystems

The X Window System is a trademark of Massachusetts Institute of Technology

UNIX is a trademark of AT&T

VADS and Verdix are registered trademarks of Verdix Corporation

VAST2 is a registered trademark of Pacific-Sierra Research Corporation

VMS and VAX are trademarks of Digital Equipment Corporation

VP/ix is a trademark of INTERACTIVE Systems Corporation and Phoenix Technologies, Ltd.

XENIX is a trademark of Microsoft Corporation

REV.	REVISION HISTORY	DATE
-001	Original Issue	12/87
-002	Revision	03/88
-003	Revision	03/89
-004	Revision	10/89
-005	Preliminary	12/89
-006	Revision	06/90
—	Revised by Change Notice 312003-001	10/90

RESTRICTED RIGHTS

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the rights in Technical Data and Computer Software clause at 52.227-7013. Intel Corporation, 3065 Bowers Avenue, Santa Clara, California 95051.

PRELIMINARY

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION

INTRODUCTION	1-1
iPSC® SYSTEM ARCHITECTURE	1-2
Compute Nodes	1-4
I/O Nodes	1-5
HOST SYSTEM SOFTWARE	1-5
NODE SYSTEM SOFTWARE	1-6
SOFTWARE DEVELOPMENT ENVIRONMENT	1-6
RUNTIME ENVIRONMENT	1-8
The User Interface	1-8
System Calls	1-9

CHAPTER 2

USING iPSC® SYSTEM COMMANDS

INTRODUCTION	2-1
A QUICK EXAMPLE	2-2
ALLOCATING AND RELEASING CUBES	2-4
Choosing the Node Memory Size (only with CX nodes)	2-6
Choosing VX Nodes (only with CX nodes)	2-7
Choosing SX Nodes (only with CX nodes)	2-7
Choosing CX Node Coprocessors (only with CX nodes)	2-8
Choosing Hybrid Cubes (only with CX nodes)	2-8
Choosing a Particular Set of Contiguous Nodes	2-9
Allocating Multiple Cubes	2-9
The Current Cube	2-10
MANAGING PROCESSES	2-12
INPUT/OUTPUT WITH THE HOST	2-16
COMPILING AND LINKING A C APPLICATION	2-17
Compiling a Host Program	2-17
Compiling a Node Program	2-17
Compiling on the Remote Host	2-20.a
COMPILING AND LINKING A FORTRAN APPLICATION	2-21
Compiling a Host Program	2-21
Compiling a Node Program	2-22
Compiling on the Remote Host	2-25
RUNTIME PROFILING	2-26
REBOOTING THE CUBE	2-27

CHAPTER 3

USING iPSC® SYSTEM CALLS

INTRODUCTION	3-1
CUBE CONTROL	3-1
Allocating, Loading, and Releasing a Cube	3-2
Controlling Processes	3-4
Redirecting Input and Output	3-6
Handling Errors and Exceptions	3-8
MESSAGE PASSING	3-9
Message Characteristics	3-9
Synchronous and Asynchronous Message Passing Calls	3-10
SYNCHRONOUS SENDS AND RECEIVES	3-10
ASYNCHRONOUS SENDS AND RECEIVES	3-11
Pending Messages	3-13
Getting Information About Pending or Received Messages	3-14
Flushing and Canceling Messages	3-16
Treating a Message as an Interrupt	3-18
GLOBAL OPERATIONS	3-20
MESSAGE PASSING WITH FORTRAN COMMONS	3-23
BYTE SWAPPING	3-24
USING GRAY CODES, BLINKING THE LED, AND TIMING EXECUTION	3-28
PERFORMING EXTENDED ARITHMETIC	3-31

CHAPTER 4 DESIGNING A CONCURRENT APPLICATION

INTRODUCTION	4-1
Separating the User Interface from the Computation	4-2
Load Balancing	4-2
DOMAIN DECOMPOSITION	4-2
CONTROL DECOMPOSITION	4-4
Designing a Communication Strategy	4-5
Generalizing the Number of Nodes	4-5
EXAMPLE APPLICATION: CALCULATING PI	4-6
EXAMPLE APPLICATION: MATRIX*VECTOR MULTIPLICATION	4-11
EXAMPLE APPLICATION: THE N-QUEENS PROBLEM	4-13

CHAPTER 5 DEBUGGING

INTRODUCTION	5-1
A QUICK DECON EXAMPLE	5-1
DECON REQUIREMENTS	5-7
USING DECON	5-8
Loading Programs and Setting the Context	5-8
The Debug Context	5-9
Aliases	5-9
The DECON Configuration File	5-10
Setting and Displaying Breakpoints	5-10
Controlling Execution	5-11
Displaying and Modifying Program Variables	5-13
Displaying the Message and Receive Queues	5-14
MORE DECON FEATURES	5-17

CHAPTER 6 PROGRAMMING TECHNIQUES

INTRODUCTION	6-1
MEMORY ACCESSES	6-1
MESSAGE PASSING	6-2
Application Buffer Alignment (CX and RX nodes, running C and Fortran only)	6-2
Avoid Message Buffering (CX and RX nodes, running C and Fortran only)	6-2
Use Static Buffers (CX nodes, running C only)	6-3
ERROR CHECKING (CX and RX nodes, running C only)	6-4
MISCELLANEOUS	6-4

CHAPTER 7 THE CONCURRENT FILE SYSTEM

CONCURRENT FILE I/O	7-1
CFS COMMANDS	7-2
CFS SYSTEM CALLS	7-5
Concurrent File	7-6
I/O Modes	7-7
Reading and Writing Files	7-12
USING THE CFS TO STORE NODE PROGRAMS	7-14
USING THE NODE SHELL	7-14
REMOTE HOST ACCESS TO NSH	7-17
CFS ROUTINE SYNCHRONIZATION	7-18

CHAPTER 8

TCP/IP ON THE NODES

INTRODUCTION	8-1
PERFORMANCE CONSIDERATIONS	8-2
THE IPHOST	8-2
NODE TCP/IP SPECIFICS	8-3
Addressing Scheme and Socket Type	8-3
Include Files	8-3
Server/Client Relationship	8-4
COMPILING A NODE PROGRAM WITH TCP/IP	8-4
The Node TCP/IP System Calls	8-5
How to Create a Socket, Connect to It, and Transfer Data	8-8
AN EXAMPLE OF A NODE SERVER (pi Example)	8-9
AN EXAMPLE OF A HOST CLIENT (pi Example)	8-12

CHAPTER 9

THE X WINDOW SYSTEM ON THE iPSC®/860

INTRODUCTION	9-1
COMPILING AND LINKING X WINDOW SYSTEM APPLICATIONS	9-2
RESOURCES	9-4
NODE CONNECTION TO THE SERVER	9-4
SETTING THE DISPLAY ENVIRONMENT VARIABLE	9-5

APPENDIX A
IPSC® SYSTEM COMMANDS,
SYSTEM CALLS, and ROUTINES

APPENDIX B
IPSC® SYSTEM FEATURES

APPENDIX C
IPSC® SYSTEM SPECIFICATIONS

SYSTEM SPECIFICATIONS C-1

IPSC®/860 RX COMPUTE NODE C-2

IPSC®/2 CX COMPUTE NODE C-2

IPSC® SYSTEM I/O NODE C-3

SYSTEM RESOURCE MANAGER C-3

ELECTRICAL AND ENVIRONMENTAL C-4

 Electrical C-4

 Safety/rfi/emi Standards System Is Designed To Meet C-4

 Environmental C-5

 Physical C-5

APPENDIX D DECON EXAMPLE SOURCE FILES

MAKEFILE	D-1
D_PROMPT.C	D-2
D_HOST.C	D-3
D_NODE.C	D-5

APPENDIX E TCP/IP SYSTEM CALLS

PRELIMINARY

LIST OF ILLUSTRATIONS

Figure 1-1. Two iPSC® Systems: One with One Compact Cabinet and One Standard Cabinet, and the Other with One Compact Cabinet	1-3
Figure 1-2. Compiling and Linking an Application for an iPSC® System	1-7
Figure 3-1. The Relationship Between the Application and System Buffers	3-13
Figure 3-2. The Operation of flushmsg() and msgcancel()	3-17
Figure 3-3. Operation of a gdsum()	3-22
Figure 4-1. Using Domain Decomposition to Achieve Load Balancing	4-3
Figure 4-2. The Decomposition Used for the Pi Example	4-7
Figure 4-3. Calculating Pi: Fortran Code for a Sequential Version	4-8
Figure 4-4. Calculating Pi: Fortran Host Code for a Parallel Version	4-9
Figure 4-5. Calculating Pi: Fortran Node Code for a Parallel Version	4-10
Figure 4-6. Matrix Vector Multiplication: Fortran Code Fragment	4-12
Figure 4-7. The N-Queens Solution Tree for a 4 x 4 Board	4-14
Figure 5-1. Control Flow for the DECON Demonstration Program (C version)	5-3

LIST OF TABLES

Table 2-1. The cc Options for Node Programs	2-18.a
Table 2-2. The pgcc Options for Node Programs	2-18.a
Table 2-3. The f77 Options for Node Programs	2-23
Table 2-4. The pgf77 Options for Node Programs	2-23.a
Table 8-1. Node TCP/IP System Calls	8-5
Table 9-1. X Window Libraries	9-3
Table 9-2. X Window Libraries for Advanced Applications	9-3
Table A-1. SRM Cube Command Summary	A-1
Table A-2. Summary of System Calls for Cube Control (C Version)	A-3
Table A-3. Summary of Routines for Cube Control (Fortran Version)	A-5
Table A-4. Summary of System Calls for Message Passing (C Version)	A-7
Table A-5. Summary of Routines for Message Passing (Fortran Version)	A-10
Table A-6. Global Operations (C Version)	A-13
Table A-7. Global Operations (Fortran Version)	A-16
Table A-8. Byte-Swapping System Calls for Remote Host/Cube Communication (C Version)	A-19
Table A-9. Byte-Swapping Routines for Remote Host/Cube Communication (Fortran Version)	A-20
Table A-10. Miscellaneous System Calls (C Version)	A-21
Table A-11. Miscellaneous Routines (Fortran Version)	A-22
Table A-12. Summary of System Calls for Concurrent File I/O (C Version)	A-23
Table A-13. Summary of Routines for Concurrent File I/O (Fortran Version)	A-25

Table A-14. Summary of Mathematical System Calls (C Version) A-27

Table A-15. Summary of Mathematical Routines (Fortran Version) A-28

Table A-16. C-Shell Built-Ins and UNIX Utilities That Run Under the Node's C-Shell A-29

Table A-17. UNIX-Compatible System Calls (C Version) A-30

Table A-18. UNIX Compatible I/O Library Calls (C Only Version) A-31

Table A-19. Summary of SRM UNIX Extensions A-32

Table A-20. Node's C-Shell Cube Command Summary A-33

Table A-21. Unique Node TCP/IP System Calls A-34

Table B-1. iPSC® System Features B-1

Table F-1. Node TCP/IP System Calls E-1

PRELIMINARY

PRELIMINARY

COMPILING AND LINKING A C APPLICATION

This section explains how to compile and link a C application program to run on the host and the nodes.

First, be sure to include the file *cube.h* in both host and node programs. This file resides in */usr/include* on the SRM. Include it as `#include <cube.h>`.

Compiling a Host Program

When compiling a host program on the SRM, include the `-host` option on the `cc` invocation line. For example, consider the pi example. It has three source files: *host.c*, *prompt.c*, and *node.c*. The files *host.c* and *prompt.c* are compiled and linked to make the executable host program called *host* as follows:

```
% cc -o host host.c prompt.c -host
```

Compiling a Node Program

When compiling a node program on the SRM, include the `-node` option on the `cc` invocation line. Once again consider the pi example. The file *node.c* is compiled and linked to make the executable node program as follows:

```
% cc -o node node.c -node
```

If the node program is intended for an RX node, be sure to include the `cc` compiler's `-i860` option.

```
% cc -o node node.c -i860 -node
```

Or, you can use the `pgcc` compiler:

```
% pgcc -o node node.c -node
```

If the node program is intended for a CX node and makes use of an SX processor, include the `-sx` option. Do not include the `-sx` option if you do not have an SX processor. That's because if you don't have the SX option, you do have an Intel 80387 numeric coprocessor. The SX processor and the Intel 80387 numeric coprocessor use different registers to return their results, and an executable program should contain compiled modules for only one of the two. For example, if your nodes had SX processors and you were doing floating-point calculations, the `cc` invocation line would look as follows:

```
% cc -o node node.c -sx -node
```

If the node program uses a VX processor, also include the `-vx` and either the `-vec` or `-vecdb` option. The `-vx` and `-vecdb` switches are inappropriate for host programs.

The **-vx** option brings in the interface library for the vector board. The **-vec** option brings in the library of vector routines that execute on the VX processor, the SX processor, or the 80387 numeric coprocessor.

PREPARED BY

PRELIMINARY

The `-vecdb` switch is like `-vec`, except that it brings in a debug version of the vector library routines. This debug version performs some additional testing such as checking for proper data types and alignment.

You can combine `-sx` and `-vx` switches. For example, the following Fortran compilation command makes an executable for a node that has VX and SX processors:

```
% cc -o node_tst node_tst.c -sx -vx -vec -node
```

If you are compiling node program `node.c` for use on an RX node with TCP/IP libraries, issue one of the following command lines on the SRM:

```
% cc -o node node.c -lsocknode -node -i860
```

```
% pgcc -o node node.c -lsocknode -node
```

Refer to the Chapter 8 for more information on the version of TCP/IP that runs on RX nodes.

Table 2-1 summarizes the `cc` compile options for node programs. Table 2-2 summarizes the `pgcc` compile options for node programs.

Table 2-1. The cc Options for Node Programs

cc Options	Description
-node	386-based node that has an 80387 numeric coprocessor (CX node).
-sx -node	386-based node that has an SX processor (SX node).
-vec -node	386-based node that has an 80387 numeric coprocessor and uses the 80387 version of the vector routines (CX node).
-sx -vec -node	386-based node that has SX processor and uses the vector routines that run on the SX processor (SX node).
-vx -vec -node	386-based node that has an 80387 numeric coprocessor and a VX processor (VX node).
-sx -vx -vec -node	386-based node that has both SX and VX processors (SXVX node)
-vx -vecdb -node	386-based node that has an 80387 numeric coprocessor and a VX processor and uses debug vector routines (VX node).
-sx -vx -vecdb -node	386-based node that has both SX and VX processors and uses debug vector routines (SXVX node).
-i860 -node	i860-based node (RX node).
-i860 -lvec -node	i860-based node that uses the i860 version of the vector routines (RX node). There are no name clashes between the i860 and the 386 vector libraries (VecLib).
-lsocknode -node	Links your program to all the socket libraries used by node TCP/IP.

Table 2-2. The pgcc Options for Node Programs

pgcc Options	Description
-node	i860-based node (RX node).
-lvec -node	i860-based node that uses the i860 version of the vector routines (RX node). There are no name clashes between the i860 and the 386 vector libraries (VecLib).
-lsocknode -node	Links your program to all the socket libraries used by node TCP/IP.

DRAFT

When they are required, the `-i860`, `-sx`, and `-vx` switches must appear in both the compile and link steps. The `-vec`, `-lvec`, `-vecdb`, `-node`, `-host`, and `-lsocknode` switches are needed only on the link step. The `-i860`, `-sx`, `-vx`, and `-node` flags result in the definitions of preprocessor symbols `__i860`, `__SX`, `__VX`, and `__NODE`, respectively. These can be used for conditional compilation in programs that run in multiple environments. The symbol `__i386` is defined by default when `-i860` is not specified (cc only).

For example, consider the following makefile. It is set up to enable you to compile the pi example used in this manual for CX, SX, or RX nodes, simply by invoking `make` with the appropriate argument.

```
#
# This file is used to compile and link the host.c, prompt.c,
# and node.c files for the pi numerical integration example.
#
help:
    @echo
    @echo "You must specify the type of node you wish to build a node"
    @echo "executable for, choose one of the following:"
    @echo
    @echo "    make cx      (for 386 nodes with 387 coprocessors)"
    @echo "    make sx      (for 386 nodes with SX coprocessors)"
    @echo "    make rx      (for i860 nodes)"
    @echo

cx:    host node          #Use default compile and link flags

sx:
    make host
    make "CFLAGS=-O -sx" "LDFLAGS=-sx" node

rx:
    make host
    make "CFLAGS=-O -i860" "LDFLAGS=-i860" node

host:  host.o prompt.o
    cc -o host host.o prompt.o -host

node:  node.o fx.o
    cc -o node node.o fx.o $(LDFLAGS) -node

clean:
    rm host node host.o node.o prompt.o
```

Assume that you want to compile the pi example for use on SX nodes. You must ensure that the `-sx` appears both in the predefined macro `CFLAGS` and on the `cc` line. `make`'s implicit rules bring in `CFLAGS` on the compile step. The `makefile` macro `LDFLAGS` explicitly places the `-sx` on the `cc` line for the link step.

Consider, also, the following makefile, which does the same thing using the `pgcc` compiler.

```
#
# This file is used to compile and link the host.c, prompt.c,
# and node.c files for the pi numerical integration example.
#
help:
    @echo
    @echo "You must specify the type of node you wish to build a node"
    @echo "executable for, choose one of the following:"
    @echo
    @echo "    make cx      (for 386 nodes with 387 coprocessors)"
    @echo "    make sx      (for 386 nodes with SX coprocessors)"
    @echo "    make rx      (for i860 nodes)"
    @echo

cx:
    make "COMPILER=cc" host
    make "COMPILER=cc" node

sx:
    make "COMPILER=cc" host
    make "COMPILER=cc" "CFLAGS=-O -sx" node

rx:
    make "COMPILER=cc" host
    make "COMPILER=pgcc" "CFLAGS=-O" node

host:
    host.o prompt.o
    cc -o host host.o prompt.o -host

node:
    node.o fx.o
    $(COMPILER) $(CFLAGS) -o node node.o fx.o -node

clean:
    rm host node host.o node.o prompt.o

.c.o:
    $(COMPILER) $(CFLAGS) -c $<
```

Compiling on the Remote Host

When running on a remote host, use the remote host's C compiler to compile the host code. Be sure to use the `-lhost` switch and not the `-host` switch.

Use `rcc` to compile the node program. The remote C compiler `rcc` copies the node source files to the SRM, compiles and links this file on the SRM, and sends the resulting executable file back to the remote host. If you have a network of IPSC systems, use the `-h` switch to choose a particular SRM. Otherwise `rcc` chooses the first available SRM listed in the `srms` file. For example, to choose the SRM called *jaxom*, issue the compile command:

```
% rcc -h jaxom -o node node.c -node
```

Including the `-cpp` option with `rcc` causes the C preprocessor to run on the remote host before the node source is copied to the SRM. This affects how include files are resolved.

- Without `-cpp` all include files are resolved on the SRM.
- With `-cpp`, files delimited with quotation marks are resolved on the remote host. For example,

```
#include "rhost.inc"
```

looks in the current directory for the file *rhost.inc* while

```
#include "/usr/ipsc/include/rhost.inc"
```

looks in the directory */usr/ipsc/include* on the remote host.

- The `-I` switch defines a prefix for include files. For example, the compilation line,

```
% rcc -cpp -I/usr/ipsc/include -o node node.c -node
```

sets the include prefix to */usr/ipsc/include*. Then,

```
#include "rhost.inc"
```

in the source file resolves as */usr/ipsc/include/rhost.inc* on the remote host. If an include file is not found on the remote host, the prefix is used to search for the file on the SRM.

- Even with `-cpp`, all *top-level* include files delimited with `<>` are resolved on the SRM. However, if an include file delimited with `<>` is *in* a file resolved on the remote host (that is, it is not a top-level include file), then `rcc` looks for the file in */usr/include* on the remote host.

PRELIMINARY

For example, consider the following situation. The node source code contains the line `#include "rhost.inc"` and the file *rhost.inc* contains the line `#include <suntool.h>`. Because *rhost.inc* is in quotation marks, `rcc -cpp` resolves it on the remote host, and because the line `#include <suntool.h>` is in *rhost.inc*, *suntool.h* is not a top-level include file. The result is that `rcc` looks for *suntool.h* in */usr/include* on the remote host. If the line `#include <suntool.h>` were in the node source file, *suntool.h* would be a top-level include file and its delimiting `<>` would cause `rcc` to look for it on the SRM.

COMPILING AND LINKING A FORTRAN APPLICATION

This section explains how to compile and link a Fortran application program to run on the host and the nodes.

First be sure to include the file *fcube.h* in both host and node programs. This file resides in */usr/include* on the SRM. Include it as `include '/usr/include/fcube.h'`.

Delimit include files with single quotes (' '). The compiler `f77` uses the quoted string as a pathname for the include file. For example:

```
include 'myfile.inc'
```

This example tells `f77` to look in the current directory for *myfile.inc*.

If the file to be compiled is called *node.F* (with an uppercase *F* extension, rather than a lowercase *f*), `f77` will run the C preprocessor on the file. This enables you to use lines like the following in a Fortran program:

```
#include <fcube.h>

#define MAX 87
```

Compiling a Host Program

When compiling a host program on the SRM, include the `-host` option on the `f77` invocation line. For example, consider the *pi* example. It has four source files: *host.f*, *prompt.f*, *fx.f*, and *node.f*. The files *host.f* and *prompt.f* are compiled and linked to make the executable host program called *host* as follows:

```
% f77 -o host host.f prompt.f -host
```

By including `-vec` you can use the vector routines that run on the SRM's 80387 numeric coprocessor.

Compiling a Node Program

When compiling a node program on the SRM, include the `-node` option on the `f77` invocation line. Once again consider the pi example. The files `node.f` and `fx.f` are compiled and linked to make an executable node program for a CX node as follows:

```
% f77 -o node node.f fx.f -node
```

If the node program is intended for an RX node, be sure to include the `f77` compiler's `-i860` option.

```
% f77 -o node node.f fx.f -i860 -node
```

Or, you can use the `pgf77` compiler:

```
% pgf77 -o node node.f fx.f -node
```

If the node program is intended for a CX node and makes use of an SX processor, include the `-sx` option. Do not include the `-sx` option if you do not have an SX processor. That's because if you don't have the `SX` option, you do have an Intel 80387 numeric coprocessor. The `SX` processor and the Intel 80387 numeric coprocessor use different registers to return their results, and an executable program should contain compiled modules for only one of the two.

For example, if your nodes had `SX` processors and you were doing floating-point calculations, the `f77` invocation line would look as follows:

```
% f77 -o node node.f -sx -node
```

If the node program uses a `VX` processor, also include the `-vx` and either the `-vec` or `-vecdb` option. The `-vx` and `-vecdb` switches are inappropriate for host programs.

The `-vx` option brings in the interface library for the vector board. The `-vec` option brings in the library of vector routines that execute on the `VX` processor, the `SX` processor, or the 80387 numeric coprocessor.

The `-vecdb` switch is like `-vec`, except that it brings in a debug version of the vector library routines. This debug version performs some additional testing such as checking for proper data types and alignment.

You can combine `-sx` and `-vx` switches. For example, the following Fortran compilation command makes an executable for a node that has `VX` and `SX` processors:

```
% f77 -o node_tst node_tst.f -sx -vx -vec -node
```

Table 2-3 summarizes the `f77` compile options for node programs. Table 2-4 summarizes the `pgf77` compile options for node programs. When they are required, the `-i860`, `-sx`, and `-vx` switches must appear in both the compile and link steps. The `-vec`, `-lvec`, `-vecdb`, `-node`, `-host`, and `-lsocknode` switches are needed only on the link step. The `-i860`, `-sx`, `-vx`, and `-node` flags result in the definitions of preprocessor symbols `__i860`, `__SX`, `__VX`, and `__NODE`, RESPECTIVELY. These can be used for conditional compilation in programs that run in multiple environments. The symbol `__i386` is defined by default when `-i860` is not specified (`f77` only).

Table 2-3. The `f77` Options for Node Programs

f77 Options	Description
-node	386-based node that has an 80387 numeric coprocessor (CX node).
-sx -node	386-based node that has an SX processor (SX node).
-vec -node	386-based node that has an 80387 numeric coprocessor and uses the 80387 version of the vector routines (CX node).
-sx -vec -node	386-based node that has SX processor and uses the vector routines that run on the SX processor (SX node).
-vx -vec -node	386-based node that has an 80387 numeric coprocessor and a VX processor (VX node).
-sx -vx -vec -node	386-based node that has both SX and VX processors (SXVX node)
-vx -vecdb -node	386-based node that has an 80387 numeric coprocessor and a VX processor and uses debug vector routines (VX node).
-sx -vx -vecdb -node	386-based node that has both SX and VX processors and uses debug vector routines (SXVX node).
-i860 -node	i860-based node (RX node).
-i860 -lvec -node	i860-based node that uses the i860 version of the vector routines (RX node).

Table 2-4. The pgf77 Options for Node Programs

pgf77 Options	Description
-node	i860-based node (RX node).
-ivec -node	i860-based node that uses the i860 version of the vector routines (RX node).

PRELIMINARY

```

# This file is used to compile and link the host.f, prompt.f,
# fx.f, and node.f files for the pi numerical integration example.
#
help:
    @echo
    @echo "You must specify the type of node you wish to build a node"
    @echo "executable for, choose one of the following:"
    @echo
    @echo "    make cx      (for 386 nodes with 387 coprocessors)"
    @echo "    make sx      (for 386 nodes with SX coprocessors)"
    @echo "    make rx      (for i860 nodes)"
    @echo

cx:   host node      #Use default compile and link flags

sx:
    make host
    make "F77FLAGS=-sx" "FLDFLAGS=-sx" node

rx:
    make host
    make "F77FLAGS=-i860" "FLDFLAGS=-i860" node

host: host.o prompt.o
      f77 -o host host.o prompt.o -host

node: node.o fx.o
      f77 -o node node.o fx.o $(FLDFLAGS) -node

clean:
      rm host node host.o node.o prompt.o fx.o

```

Assume that you want to compile the pi example for use on SX nodes. You must ensure that the `-sx` appears both in the predefined macro `F77FLAGS` and on the `cc` line. `make`'s implicit rules bring in `F77FLAGS` on the compile step. The `makefile` macro `FLDFLAGS` explicitly places the `-sx` on the `cc` line for the link step.

Consider, also, the following makefile, which does the same thing using the `pgf77` compiler.

```
#
# This file is used to compile and link the host.f, prompt.f,
# and node.f files for the pi numerical integration example.
#
help:
    @echo
    @echo "You must specify the type of node you wish to build a node"
    @echo "executable for, choose one of the following:"
    @echo
    @echo "    make cx      (for 386 nodes with 387 coprocessors)"
    @echo "    make sx      (for 386 nodes with SX coprocessors)"
    @echo "    make rx      (for i860 nodes)"
    @echo

cx:
    make "COMPILER=f77" host
    make "COMPILER=f77" node

sx:
    make "COMPILER=f77" host
    make "COMPILER=f77" "FFLAGS=-O -sx" node

rx:
    make "COMPILER=f77" host
    make "COMPILER=pgf77" "FFLAGS=-O" node

host:
    host.o prompt.o
    f77 -o host host.o prompt.o -host

node:
    node.o fx.o
    $(COMPILER) $(FFLAGS) -o node node.o fx.o -node

clean:
    rm host node host.o node.o prompt.o

.f.o:
    $(COMPILER) $(FFLAGS) -c $<
```

PRELIMINARY

The example presented later in this chapter has the node program accept the name of its iphost as a command line argument, provided when you loaded it with the load command. For example, assume that you load the program *node* as

```
% load node iphostname
```

Then, the node program can attach to *iphostname* by executing the call

```
setiphost(argv[1]);
```

NODE TCP/IP SPECIFICS

This section describes the particular version of TCP/IP that is supported on the nodes.

Addressing Scheme and Socket Type

The TCP/IP that is available to node programs follows the AF_INET addressing scheme. This addressing scheme uses Internet addresses to identify the machine. Port numbers allow more than one socket on a machine. Also, the node TCP/IP uses only the SOCK_STREAM type of sockets.

The consequence is that when you create a socket with the socket() call, you must specify AF_INET and SOCK_STREAM. Also, be sure to use the Internet socket structure, struct sockaddr_in. This structure is defined in */usr/include/CMC/netinet/in.h*.

Finally, by default sockets created for use with node TCP/IP are blocking sockets. When you issue an accept() call, your program blocks until someone makes a connection. You can make sockets nonblocking with fcntl().

Include Files

Include files reside by default on the SRM under */usr/include/ipsc/CMC*. The directory structure under *CMC* is as you would expect for TCP/IP. For example, if your makefile specifies the -I option as *-I/usr/include/ipsc*, then to include *socket.h*, your program should contain the line,

```
#include <CMC/sys/socket.h>
```

There is one exception, however. When you use the byte order calls (such as htons()), you must also include */usr/include/ipsc/CMC/ntoh.h*. Under the UNIX operating system, these calls require only the header files, */usr/include/ipsc/sys/types.h* and */usr/include/ipsc/sys/in.h*. Under the NX operating system, they also require */usr/include/ipsc/CMC/ntoh.h*.

Server/Client Relationship

TCP/IP convention is to have the server do the `accept()`. The server doesn't have previous knowledge of its caller. It gets that information from the `accept()`. A client on the other hand must know from whom to request service. It must put that information in the `connect()` call.

A node can be a server or a client. However, in an iPSC system, the nodes are usually considered to be a compute server, and the host is the client. When you use TCP/IP to send data between a node and the host, the node is usually the acceptor and the host is the connector.

TCP/IP convention is to have the server `fork()` a child to handle a client request after a successful `accept()`. If you do this in an RX node, the NX operating system loads the child on another node in the cube. For this to be successful, you must have at least one node in your allocated cube that does not have a process running on it. You also might consider writing your application so that it does not `fork()` processes to handle requests.

If a name server is configured for the SRM, the iPSC system will use it. The calls `sethostent()` and `endhostent()` are available to open and close the `/etc/hosts` file, but `gethostent()` is not provided. You have to use either `gethostbyname()` or `gethostbyaddr()` to access the file.

Also, when you use `gethostbyname()`, note that the argument is the name of the iphost, not the name of the SRM.

COMPILING A NODE PROGRAM WITH TCP/IP

The TCP/IP calls are in `libsocknode.a`. Include the switch `-lsocknode` on the linking step of your compile. For example, to compile the program `node.c` for use on an RX node with the TCP/IP libraries, issue one of the following commands on the SRM:

```
% cc -o node node.c -lsocknode -i860 -node
```

```
% pgcc -o node node.c -lsocknode -node
```

If you are developing on the remote host, use `rcc` instead of `cc`. If you use a makefile, the pertinent lines are as follows:

- For the `cc` compiler:

```
CFLAGS= -i860
node: node.o
    cc -o node node.o -lsocknode -i860 -node
host: host.c
    cc -o host host.c -host
```

- For the pgcc compiler:

```
node: node.o
      pgcc -o node node.o -lsocknode -node
host: host.c
      cc -o host host.c -host
```

Note that you may want to put **-node** as part of CFLAGS. The built-in rule for generating a C object from a C source uses CFLAGS. Putting **-node** into CFLAGS would include **-node** on the compilation step. This results in the definition of the preprocessor symbol `__NODE`, which your application may find useful.

PRELIMINARY

INDEX

A

accept() 8-8
adminproc 1-6
adminproc.srm 1-5
AF_INET 8-3
allocating a cube 3-2
allocating and releasing cubes 2-4
application buffer 3-13
asynchronous receive 3-11
asynchronous send 3-11
attachcube 2-10
attachcube() 3-2

B

bind() 8-12
BLAS 4-11
blocking sockets 8-3
byte ordering convention 3-24
byte swapping 3-24

C

cbackup 7-3
cc
 -cpp option 2-20.a
 -i860 option 2-17
 -lsocknode option 8-4
 -node option 2-17
cc options 2-18.a
CFLAGS 2-20, 2-24
commons
 different views 3-23
commser 1-5
compact cabinet 1-4
compile and link steps 2-19, 2-23
compiling a Fortran host program 2-21
compiling a Fortran node program 2-22
compiling a Fortran program on a remote host 2-25
compiling and linking a C application 2-17
compiling and linking a Fortran application 2-21
compiling on a remote host 2-20.a
completion code 3-5
Concurrent File System 1-4, 7-2

connect() 8-4
cprobe() 3-13
createstruc() 3-25
crecv() 3-10
crestore 7-4
critical code 3-19
csend() 3-10
csendrecv() 3-10
cube.h 2-17
cubeinfo 2-10
 -a switch 2-11
 -h switch 2-11
 -n switch 2-11
 -s switch 2-11
cubeinfo() 3-2
current cube 2-10
CX node
 features 1-2
CX nodes 1-2

D

dclock() 3-29
DCM 3-21
dimension 2-5
Direct-Connect™ Module 1-2

E

e..() 3-31
ecmp() 3-32
e-cube routing algorithm 3-21
ediv() 3-32

endhostent() 8-4
ermo 3-8
esize_t 3-31
exception numbers 3-8
extended arithmetic 3-31

F

F extension on Fortran files 2-21
f77 2-21
f77 options 2-23
flick() 3-4
flushmsg() 3-5, 3-16, 3-17
fork() 8-4
fserver 1-5

G

gdsum() 3-21, 8-9
getcube 2-4, 3-6
 -c switch 2-9
 n specifier 2-9
 redirecting output of 2-5
 -t switch 2-5
gethostbyaddr() 8-12
gethostbyname() 8-9, 8-12
getiphosts() 8-2
getpeername() 8-12
ginv() 3-28
global operations 3-20
gray codes 3-28
gray() 3-28

H

h...() 3-18
 handler() 3-8
 host option 2-17
 host system software 1-5
 hrecv() 3-4, 3-18
 hsend() 3-19
 hsendrecv() 3-19
 HTOCL() 3-24
 htonl() 8-9
 hton() 8-9
 hwclock() 3-28
 hybrid cube 2-8

I

I/O nodes 1-5
 include files 2-20.a–2-21, 2-25
 info...() 3-14
 infocount() 3-14
 iocube 2-11
 iphost 8-1
 iprobe() 3-13
 irecv() 3-11
 isend() 3-11
 isendrecv() 3-11

K

killcube 2-15
 node shell 7-15
 killproc() 3-5
 killsyslog() 3-6

L

LED 3-28
 lhost switch 2-20.a
 libsocknode.a 8-4
 lifeline 1-5
 linking TCP/IP programs 8-4
 load 2-12
 -H switch 2-14
 -p switch 2-13
 program arguments 2-14
 loader 1-5
 loader.srm 1-5

M

managing processes 2-12
 masktrap() 3-19
 mclock() 3-29
 memory size of node 2-6
 memset() 8-12
 message buffer 3-16
 message ID 3-9
 message length 3-9
 message passing 1-2, 3-8, 3-9
 message type 3-9, 3-16
 message type selection 3-9
 messages as interrupts 3-18
 msgcancel() 3-11, 3-17
 msgdone() 3-11
 msgwait() 3-11
 myhost() 3-3
 mynode() 3-3
 mypid() 3-3

N

newserver 2-16
newserver() 3-6
node shell 7-14
node shell prompt 7-14
node system software 1-6
nodedim() 3-3
nsh 7-14
ntohl() 8-9
ntohs() 8-9
numnodes() 3-3
NX/2 pid 3-3, 3-14

O

oad() 3-2

P

PBX I/O node 8-1
pending message 3-13
pending messages 3-16
perror() 3-8
pgcc options 2-18.a
pgf77 options 2-23.a
pi example 2-17
port number 8-3, 8-14

R

rcam 1-5
rcc 2-20.a
read() 8-9
relcube 2-4
-c switch 2-10

relcube() 3-2
relstruc() 3-25
remote host software 1-6
restrictvol() 7-2
rhost.inc 2-20.a
ripscd 1-6
runtime profiling 2-26
RX node 2-17
RX nodes 1-2, 3-8, 3-10
RX process 3-4

S

scopy() 4-11
SCSI controller 1-5
sdot() 4-11
server/client 8-4
sethostent() 8-4
setiphost() 8-2
setpid() 3-3
setsyslog() 3-6
showvol 7-2
signal() 3-8
sigset() 3-8
SOCK_STREAM 8-3
socket node 8-1
standard cabinet 1-4
startcube 2-14
stoe() 3-32
SX node 2-7
SX processor 2-17
synchronous receive 3-10
synchronous send 3-10

syslog 2-16

syslog process 3-6

system buffer 3-13

system calls 1-9

System Resource Manager 1-1, 1-2, 1-6

T

TCP/IP host program 8-12

TCP/IP include files 8-3

TCP/IP node program 8-9

TCP/IP system calls 8-5

TCP/IP system calls supported 8-5, 8-6, E-1

TCP/IP with an iPSC/2 system 8-1

TTYs field 2-11

U

underscore calls 3-8

UNIX system calls 1-10

user interface commands 1-8

V

VAST2 1-4

VecLib 4-11

volume number of disk drive 7-2

VX node 2-7

VX processor 1-4, 2-17, 2-22

W

waitall() 3-4

waitcube 2-15
node shell 7-15

waitone() 3-4

write() 8-9

PRELIMINARY