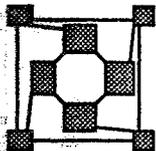


Express Fortran

Reference Guide

Version 3.0



© ParaSoft Corporation, 1988, 1989, 1990

Table of Contents

Chapter 1: System Commands 2

Tools providing services in support of *Express* applications

1	Executing <i>Express</i> commands in “non-windowing” operating systems.	3
2	Executing <i>Express</i> commands in “windowing” systems	3
3	Specifying numeric data in switches	3
4	Manual Page Layout.	4

Chapter 2: FORTRAN runtime library 52

Library routines available to *Express* programs written in FORTRAN

1	High Level Communication System.	53
2	Hardware Dependent Communication System.	54
3	Synchronization	55
4	Decomposition Tools	56
5	Multitasking Support	56
6	Processor Allocation and Control.	57
7	I/O Services	58
8	Graphics	59
9	Standard variables and the <i>/XPRESS/</i> common block.	59
10	Manual Page Layout.	60

Appendix A: Classification of routines. 238

A listing of the *Express* routines, broken down by functionality

Appendix B: Library Availability. 246

The correspondence between C and FORTRAN libraries and the synchronization properties of *Express* functions

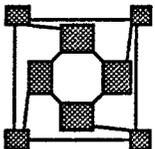
1	Correspondence between C and FORTRAN	247
2	Synchronization Rules	247
3	Libraries and Programming Models	247
4	NOTES	248

Appendix C: Index of Routines. 254

An alphabetical listing of routines, variables, commands and macros

Appendix D: Index	258
General index to <i>Express</i> and the examples from the text	

© 1988, 1989, 1990
ParaSoft Corporation
All rights reserved.



© ParaSoft Corporation, 1988, 1989, 1990

1 Executing *Express* commands in “non-windowing” operating systems

When running *Express* under operating systems with conventional line-oriented user interfaces such as UNIX, MS-DOS, VMS and similar, commands are executed by typing their names at the command line prompt.

Usage generally follows the conventional UNIX style with options being indicated by the ‘-’ character, e.g.,

```
cubix -n 2 -t 120 toyland
```

The particular command line string ‘-?’ provides a brief on-line summary of the options and purpose of a command. While this may help in jogging the memory this manual should be consulted for complete details.

On most machines you will need to add a new directory to the set which is searched when looking for executable programs. The exact mechanism for doing this varies from one machine to another and the details for individual operating systems can be found in the introductory guide to *Express* for that machine. If you find that messages such as

```
Command not found
```

or

```
No such file or directory
```

appear whenever you try to execute one of the *Express* commands then you should check that the appropriate directory really has been added to your search path. If this seems to be correct you should next try running the `excustom` program which will ensure that the *Express* installation is internally consistent.

In keeping with the conventional style all commands exit with status 0 upon successful termination and with non-zero values if errors occur.

2 Executing *Express* commands in “windowing” systems

In windowing systems such as MicroSoft Windows and the Macintosh, *Express* programs are usually executed by selecting icons from the screen. In most cases a dialog box will then be presented allowing the entry of parameters. In most cases the entries to be made have a one-to-one correspondence to the switches used in the line-oriented interfaces. Usually some mechanism is also provided to “Abort” or “Cancel” the program without executing any commands.

Note that only the line-oriented interfaces are completely documented in this reference. In most cases this causes few problems since the switches and “boxes” are obviously related to one another. In cases where confusion may arise the introductory guide to *Express* on your system should be consulted for more help.

3 Specifying numeric data in switches

Many of the parameters necessary to the commands listed in this section have numerical values - the number of processors to use, the number of bytes to display, the position at which to load the *Express* kernel, etc. In *most* cases these values can be entered with the usual C-style notation as either decimal, octal or hexadecimal values.

Consider the `exdump` command, for example. One of its arguments specifies the address from which data should be extracted - the '-B' switch. Typically one knows this value as some "hex" constant and would therefore use a command of the form

```
exdump -B 0x79000.....
```

Alternatively you could use either octal or decimal notation replacing this by

```
exdump -B 01710000.....
```

or

```
exdump -B 495616.....
```

to achieve the same effect. Similar remarks apply to most of the other *Express* commands - you can execute a *Cubix* program on 16 nodes with any of the command switches

```
cubix -n 16.....
```

```
cubix -n 020.....
```

```
cubix -n 0x10.....
```

4 Manual Page Layout

The manual pages are, for better or worse, modeled after those often found in UNIX documentation which means that each manual page has several well-defined sections. The overall structure is shown below.

```

acctool
-----
NAME
    acctool - Analyze accounting data

SYNOPSIS
    acctool [-p] [-a dir] [-f logfile]

DOMAIN
    Available on SUN host machines only

DESCRIPTION
    acctool is used to analyze the us.....

OPTIONS
    -p          Suppress graphics
    -a dir      Name of directory containing accounting data
    -f logfile  Write output to logfile.

EXAMPLES
    acctool -a /home/kastor/accounting
    Analyze data from the directory /home/kastor/...

NOTES/WARNINGS/BUGS
    None

SEE ALSO
    Excustom

```

Header contains the name of the manual page which is usually the same as the command described. ←

The various sections and their contents are:

- NAME** Repeats the name associated with the manual page and a brief one-line description of the purpose of the associated routines
- SYNOPSIS** Summarizes the arguments used by the indicated command. Arguments enclosed in '[', ']' pairs are optional. If more than one command is described on a particular page then all are listed in this section
- DOMAIN** Describes the machines on which the command is available and any restrictions on when it may be used.
- DESCRIPTION** Describes the purpose of each command and lists the actions caused by its

most important arguments. This section is the most important reference material for each command.

OPTIONS

This section lists all the supported arguments for each command and the actions caused by specifying them.

EXAMPLES

Usually several examples are presented of the use of each function showing the most important arguments and switches.

NOTES

If present this section contains useful information about oddities in the implementations of a particular command. It may also repeat important information from the **DESCRIPTION** section.

WARNINGS

If the command has peculiar side effects or is "dangerous" in some way it will be noted in this (optional) section. Any non-intuitive behavior is also noted here.

BUGS

Currently known bugs and/or unimplemented switches are noted in this (optional) section.

SEE ALSO

Related commands and/or routines from the various *Express* libraries are noted in this section. Using this information is usually the quickest way to build a complete picture of the interaction between the various utilities.

of 0000

10/1/84

10/1/84

10/1/84

10/1/84

NAME

acctool - Analyze parallel computer usage under *Express*

SYNOPSIS

```
acctool [-p] [-a account_dir] [-f logfile]
```

DOMAIN

This command is available on SUN host computers only.

DESCRIPTION

acctool is used to analyze accounting data previously obtained from *Express* programs.

If the accounting system has been enabled on a particular host every *Express* program writes an entry into a system data file whenever it allocates or deallocates nodes. Special entries are also assigned whenever the system crashes or is reinitialized. acctool analyzes this data in an interactive fashion displaying the usage of resources on a machine-by-machine basis.

Results are reported for all users, in hours, or on a single job basis for individual users, in seconds. Statistics are managed on a monthly basis with options to restrict attention to particular months or ranges of months.

The operation of the accounting system is controlled by the *excustom* command. One of this system's options is whether or not to enable the accounting system. If enabled a place must be indicated for the accounting information to be maintained.

OPTIONS

- a *account_dir*
By default acctool looks in the current directory for the data files describing the system configuration and accounting data. This switch allows an alternative directory to be specified.
- f *logfile*
All information provided by acctool appears on the display device. If this switch is given a "log file" will also be kept containing the identical information. (In the Sunview version of the program this effect is obtained by entering a name in the log file field of the control panel.)
- p
By default acctool operates in the Sunview environment providing a simplified user interface. If Sunview is not supported on your system this switch enables a line-oriented interface in which the user is prompted to enter various options from the keyboard.
- ?
Print usage message.

EXAMPLES

The following command executes the profiling tool in a windowing environment and searches the directory /home/kastor/accounting for the necessary databases.

```
acctool -a /home/kastor/accounting
```

factool

SEE ALSO

excustom.

st
sub
ed
d
y

sub

NAME

cnftool - Configure Transputer systems.

SYNOPSIS

cnftool [-p] [-d]

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command is used to configure or reconfigure a transputer network for use with *Express*. Two interfaces are available; with the '-p' switch a simple line oriented interface leads the user through the configuration process. Without this switch a menu driven utility allows the user to specify the physical transputer interconnect and also to add additional hosts to an existing system.

One of the features of the system is a "worm" program which can be used to detect the initial hardware configuration on statically wired systems. This program has a simple searching algorithm which examines the links on each node and attempts to find a node connected to each. As each link is examined and another node detected the program recursively examines other nodes which may be attached. Note that this can only be achieved if the system has "physical" rather than electrical connections. Hardware which has INMOS' link switch cannot be examined by this method since the links are initially disconnected on hardware reset.

Details of the use of this system can be found in the accompanying documentation, "Configuring Transputer Systems: cnftool".

OPTIONS

- p By default cnftool supports a menu driven graphical interface. This switch enables a simpler, but more tedious, line interface for system configuration.
- d Run silently - the system is configured in much the same way as with the '-p' switch except that a "general" network topology is also selected automatically. No user interaction is required unless the "worm" program fails to operate successfully.
- ? Print usage message.

SEE ALSO

"CnfTool: Configuring Transputer Systems"

NAME

ctool - Analyze Communication Profile

SYNOPSIS

```
ctool [-b nbins] [-p] [log_file_name]
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command is used to examine and analyze the log file created with the communication profiler commands. The only argument is the name of the file containing the profile data which may be omitted if it has the default value "cprof.out".

If the "-p" switch is given this command presents a separate table on `stdout` from each node. The information contained in each table is:

- An identifier showing which node the following data is from.
- A summary of the calculation, communication and I/O times in the processor. In making this classification all inter-node and basic host-node communication comes under the heading "Communication" while genuine I/O requests such as calls to read, write, printf, fopen, etc. are counted as I/O.
- A summary of the time spent in, number of calls to and errors incurred in each communication function called by the processor. This information is use to give a quick breakdown of the total communication pattern. The "error" count is also a good place to look for obscure bugs. Each function makes some consistency checks on the supplied arguments and returns an error if they are inconsistent.
- A breakdown of the values returned by the communication functions. The return values are binned logarithmically - the column headed "8" indicates the frequency of return values in the inclusive range 8 thru 15. The exact interpretation of this data depends on the particular function being invoked but is usually related to the message length involved in the call. By default data from ten logarithmic bins is included in the output although the '-b' switch is provided to override this default.

One very important use of this system is the detection of programs which are sending too much data in their messages. These will show up very clearly in the histogram output.

This data appears on `stdout`.

If the `ctool` command is invoked without the "-p" switch then a graphical interface allows data to be presented in graphical form. The package is menu-driven and (hopefully) quite straightforward to use. A full list of the available options is presented in Subsection 3.4 of the *PM* manual.

OPTIONS

-p Suppress graphical output. The analysis results are presented in tabular form

on stream stdout.

-b nbins Specifies an alternate number of logarithmic bins to display when used in conjunction with the '-p' switch. (Default 10).

EXAMPLES

To examine the profile data in a file called "phase3.prof" execute the command

```
ctool phase3.prof
```

SEE ALSO

cubix in the *Express* documentation.

NAME

cubix - Host slave process for node programs.

SYNOPSIS

```
cubix [-?] [-n nodes] [-d doc] [-P] [-t time] [-S]
[-T plot_option] [-E custom_file] [-f file] [-fp]
[-mc|x|e] [-D] [-x] program [arg1] [arg2]...
```

DESCRIPTION

This command provides an interface between node applications and the host filesystem and operating system utilities. It is also responsible for node allocation and the communication of command line arguments to a node program.

This command, with the '-S' option, starts up a *Cubix* server process. Instead of loading a user application as is usually the case the server merely waits for I/O requests from any node in the parallel computer system.

While all standard I/O and graphics requests are available the stateless nature of the server may make its operation slightly strange in a multi-user environment. Routines which affect the state of the system such as *chdir* will have ramifications beyond those normally expected. In this case, for instance, a call which changes the active directory of the server for one user may invalidate requests to open files for another user.

OPTIONS

- | | |
|----------------|---|
| -n nodes | Allocate nodes processor for this process. Default 1. |
| -d doc | Alternative to -n switch. Specify size of processor group logarithmically in manner suitable for hypercubes (i.e., doc = 0 for 1 node, doc = 1 for two, doc = 2 for four, etc.) Default 0. |
| -P | Load the program into its processors but do not start it running. This option is useful in connection with the node debugger <i>ndb</i> . |
| -t time | Time out the process after the given number of seconds. This can be useful in detecting 'hung' programs. The default is no time out. |
| -T option | Specify a graphics option for programs that use <i>Plotix</i> . |
| -fp | Execute the program on attached "vector" nodes, if available. |
| -f file | Read the programs to be run and their arguments from the specified file. This option is used whenever different node programs are required or different arguments should be passed to different nodes. The file format is basically single lines containing a range of nodes, an executable program and an argument list. See the examples below. |
| -E custom_file | Directs <i>cubix</i> to use an alternative system customization file rather than the system default. |
| -m[x c e] | Enable the performance monitoring tools. The 'x', 'c' and 'e' characters |

refer to the execution, communication and event driven profiling systems respectively and may be combined. For more details refer to the *PM* manual.

- o start_node
Specifies which nodes are to be allocated to the program. An attempt will be made to allocate consecutively numbered nodes starting at the indicated number. If this cannot be done the `cubix` command will terminate.
- S
Enter server mode. (Used on multi-host systems.)
- arg1 arg2... These arguments are passed to the node main program as the conventional runtime parameters `argc`, `argv`.
- D
Enable system debugging. With this option set `cubix` prints a huge amount of data about the system as it runs. Should be of little interest to most users.
- x
In windowing versions of *Express* such as the Macintosh or MicroSoft Windows this switch forces the *Express* kernel to be re-loaded before beginning the user program. It is essentially equivalent to running the `exinit` program from the shell.
- ?
Print usage message

EXAMPLES

```
cubix -n 4 noddy
```

Loads the program `noddy` into four processors. No arguments will be passed to `main()` other than the program name in `argv[0]`.

```
cubix -d 1 -t 120 -mce longjob 3.14 2.72
```

Loads the executable `longjob` into two nodes with a total execution time limit of two hours. Also passes two extra arguments to the node program. Finally enables both communication and event driven performance monitoring tools.

```
cubix -n4 -Tega plotter
```

Run the program `plotter` in four nodes and enable graphics output on an IBM Enhanced Graphics Adapter.

```
cubix -n 1 -P buggy
```

This sequence loads a single node with the user code `buggy` but halts execution before the users main routine. The job is run in background mode so that debugging can be carried out.

```
cubix -n 4 -f loadfile
```

This sequence allocates four nodes and then loads programs according to the instructions found in the file `loadfile`. Basically the format is single lines containing either a node number or a range of nodes followed by a program name and argument list. Blank lines are ignored and `#` introduces comments. Continuation lines, backslashes and quotes are processed in the conventional manner. As an example consider the following sample `loadfile`

cubix

```
# This is a command file specifying how node programs
# should be loaded into the cube.
  0-1 proga foo bar
  3 progb horse dog cow
  2 progc really\ one\ argument
```

Note that a range of nodes is indicated for `proga` and that the backslash symbol is used to concatenate tokens into a single argument - in the above case `progc` would have only two arguments the name `progc` and the string `really one` argument.

EXIT CODE

The `cubix` process exits to the shell with the same exit code as used in the call to `exit()` in the node program.

DIAGNOSTICS

Among the errors detected by `cubix` are requests for more nodes than are available and missing program files. After validating that the specified program is indeed an executable image it is loaded into the machine using the `exload` system call. This produces messages about the size of file to be loaded and a single 'b' character for each 1024 byte block loaded. A common situation is that in which the previous job crashed the node operating system in which case the loader will say loading some number of bytes but no 'b's appear, or many 'b's appear and the final 'E' but the program does nothing after the `Starting` message. This is usually a good time to run `exinit`.

Upon exit `cubix` reports the elapsed time divided between `user` and `system`. The latter is time spent performing system functions such as program loading and is always rather small. It is provided simply for compatibility with other systems running `cubix` applications.

NAME

etool - Analyze Event Profile

SYNOPSIS

```
etool [-p] [-t] [log_file_name]
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command is used to examine and analyze the event log created with the `eprof` commands. The only argument is the name of the file containing the profile data which may be omitted if it has the default value "eprof.out".

This package is exclusively graphical and menu-driven. The most complete source of reference is the discussion in Section 4.4 of the *PM* manual.

OPTIONS

- p Suppress graphical output. The analysis results are presented in tabular form on stream `stdout`.
- t Display only the data from the "toggle" events.

EXAMPLES

To examine the profile data in a file called "phase3.prof" execute the command

```
etool phase3.prof
```

SEE ALSO

cubix in the *Express* documentation.

excustom

NAME

excustom - Reconfigure *Express*.

SYNOPSIS

```
excustom [-r] [-?] custom_file
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

excustom is used to modify the system parameters which describe a particular implementation or version of the *Express* system. All system variables are maintained in a particular file called the "*Express* customization file" which is located in an operating system dependent location. By default **excustom** will modify this file although another may be indicated by the optional *custom_file* argument.

In non-windowing environments **excustom** prompts, in turn, for values of all important system parameters, offering defaults based on the values found in the current customization file. If you do not enter any input on a particular line the original value will be taken. In windowing systems a pop-up display is typically used to offer the current values of all parameters for modification. You can then change individual entries at will. An additional option restores all entries to "sensible" defaults which guarantee that *Express* will operate correctly. (This option is obtained in the non-windowing environment by specifying the '-r' switch when invoking **excustom**.)

The **excustom** tool typically asks only about top level information from which it derives all other related data using the "macro" mechanism discussed below. In some cases you may need to modify individual system parameters at a finer level of detail. This is achieved by simply editing the customization file with a suitable text editor or word processor. (You can find the name of the default customization file with the '-?' command line option.) The exact meaning of all system variables is shown in the accompanying "Excustom" manual.

In order for the customization file to take effect the system must be reloaded with the **exinit** command.

OPTIONS

- r** By default **excustom** prompts you with the current system parameters as obtained from the customization file. With this switch "sensible" defaults are used instead of the current values.
- custom_file** This argument requests that the modification process be applied to the named file rather than the default system configuration file. This allows **excustom** to be used by applications which maintain their own customization systems.
- ?** Print the name of the default system configuration file.

IMPLEMENTATION

The customization file is a line oriented ASCII file which contains definitions of important system variables, one to a line. Lines beginning with either ';' or '#' characters are treated as comments. Other lines take either of the symbolic forms

```
NAME=text  
MACRO:=text
```

As is suggested by the above notation the former type are merely assignments to *Express* system parameters while the second define macros that may be further used in the customization file to simplify definitions of multiple related objects. A good example might be the default start-up information required by the debugger, *ndb*. As part of its configuration information it needs to know the location of the on-line help facility and also the system start-up file which contains the definitions of system commands. Since these are often in the same or related directories one might imagine two entries in the customization file of the type

```
NDB_HELPDIR=c:\parasoft\help  
NDB_STARTUP=c:\parasoft\lib
```

These entries could, however, be replaced using the macro replacement facility with the lines

```
PARASOFT:=c:\parasoft  
NDB_HELPDIR={PARASOFT}\help  
NDB_STARTUP={PARASOFT}\lib
```

Notice that the value of the PARASOFT macro is indicated with the '{' and '}' characters.

While three lines may seem more complex than the original two the use of the PARASOFT macro means that the *Express* system can be moved from one directory to another by simply changing the macro rather than each line of the customization file.

SEE ALSO

excustom (subroutine), "Customizing *Express*".

exdump

NAME

exdump - Retrieve data from node RAM files

SYNOPSIS

```
exdump [-?] [-B base] [-d doc] [-l length] [-n nodes]
[-N node] [-N node-node] [-o file] [-t threshold]
[-p pid] [-s start] [device]
```

DOMAIN

This command is available at the command line prompt on the host processor.

DESCRIPTION

This command is used to retrieve the debugging information stored in the internal RAM file under *Cubix*. It can be used either as a post-mortem dump or while a process is still running. If set up correctly data can be retrieved after machine initialization with `exinit`. The `device` argument specifies which array the dump is to be taken from - in the current implementation this should be left to its machine dependent default.

By default the dump is assumed to contain ASCII data and continues until several consecutive unprintable characters are seen. An alternative is to dump in "binary mode" in which case data is just read from the node file and sent to `stdout`. In this case options are available to both control the amount of data printed and also redirect the output to a file - printing binary data to a terminal has rather detrimental effects on its behavior.

The detailed use of the RAM file and its manipulation are fully described in the accompanying *Cubix* documentation - "Programming Parallel Computers Without Programming Hosts".

OPTIONS

- B Dump binary data instead of ASCII. By default 16 Kbytes will be taken from each node.
- b base Defines the base address of the RAM file. Decimal, octal and hex constants are valid base values. Note that this option potentially interacts with the linker/locator and also the parameters used in the `ramfopen` call. Consult the Appendix discussing RAM files in the *Cubix* manual.
- d doc Dump data from 2^{doc} nodes. This is an alternative to the `-n` option designed for hypercube users.
- l length Specify amount of data to be dumped from each node. In the default ASCII mode less will be read if the data ends early.
- n nodes Specify number of nodes from which to dump data.
- t threshold As currently implemented `exdump` is most useful for retrieving printable ASCII information. It continues reading data until `threshold` successive unprintable characters are seen and then moves onto the next node. The

-
- default threshold is five.
- N node Read the RAM file from processor node.
 - N node1-node2
 Read RAM files from the inclusive range of nodes node1-node2.
 - o file Redirect output to the named file. Default output is to stdout.
 - p pid If the process whose file is to be examined is still active then its process ID should be specified and its RAM file will be read.
 - s start Specify the physical node number from which the dump is to start. This is useful in cases where the program ran in high numbered nodes and you are dumping data after the program has stopped. Since the default allocation strategy is to allocate the lowest numbered nodes with the required size it is occasionally necessary to use this switch to “grab” the higher numbered processors.
 - ? Print usage message

EXAMPLES

```
exdump -d 1 -s 2
```

This command reads the RAM file from the default address in two nodes. The two processors will, if possible, be those starting at node two in the array. This form of the command is often used either after the node has “hung” in communication (and the nodes had to be reset with `exinit`) or when the process finished but with some error. Note that `exinit` normally initializes the contents of memory while loading *Express* so it is necessary to use the `excustom` facility to prevent this if we wish to preserve RAM file data.

```
exdump -n 1 -B 0x1000
```

Retrieves the data from a single node starting at address 1000 (hex). This form is used in conjunction with the `ramfopen` call in *Cubix*.

```
exdump -p 376 -n 4 -N2-3
```

This option retrieves the information currently contained in the RAM file of the process whose *process ID* is 376. Data will only be dumped from nodes 2 and 3 in the group allocated by the process.

```
exdump -b -B0x80001000 -o xdump.out -l 4096 -n2
```

Dump 4 Kbytes of data in binary mode from two nodes. Write the output to the file `xdump.out`.

NOTES

Numeric parameters may be specified in decimal, octal or hex using the usual C style notation: 123 is decimal, 0123 is octal and 0x123 is hex. Switch values may follow immediately after their switches or there may be intervening spaces: ‘-B0x1000’ and ‘-B 0x1000’ are both valid.

exdump

SEE ALSO

"Cubix: Programming Parallel Computers Without Programming Hosts."

exinit (command)

NAME

exinit - Reboot and reload *Express* kernel.

SYNOPSIS

```
exinit [-K] [-m] [custom_file]
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

exinit must be executed before any routines may access the parallel machine - failure to do so will result in the failure of all attempts to allocate processors. It loads the *Express* kernel and starts it running in the node processors. It also performs any necessary hardware configuration on systems which support such abilities.

exinit attempts to check that no node processes are actually executing before resetting the hardware. If node programs are detected exinit will report and suggest the use of the '-K' switch. If this switch is supplied any node programs will be killed before the reset operation is performed.

It is important that all node processors be reset before loading the *Express* kernel since otherwise parts of the network may be inaccessible. Most hardware systems have intelligent reset lines so that several boards may be reset one by another. If you are connecting several hosts together the exreset command is available to reset a particular subset of the network. It does not, however, reload the *Express* kernel. This must be done with the exinit command from some other console.

The optional custom_file argument is used to specify an alternative customization file when downloading *Express*. This allows temporary modifications to the system parameters for testing and also allows custom applications to maintain their own customization information.

A very important point to note about exinit is that, by default, it destroys the contents of the node memory while loading *Express*. This behavior is normally quite reasonable with the exception that one may wish to preserve the contents of some RAM file for use with exdump or ndb. In this case the '-m' switch can be used to prevent memory initialization. Alternatively the excustom tool has an option which forces the kernel to be loaded without destroying memory by providing an explicit "start address". A good way to proceed, therefore, is to make a system customization file which contains the load address and then to re-load *Express* by telling exinit to use this file rather than the system default.

OPTIONS

- K By default exinit aborts if any processes are still running in the parallel machine. This switch causes an attempt to be made to kill all such programs before resetting and reloading *Express*.
- m address Load *Express* into the nodes at the indicated address without destroying

exinit

the contents of memory. This is useful in conjunction with the RAM file system for debugging after system crashes. The address used will depend on your hardware configuration.

`custom_file` Indicates that a system customization file other than the default should be used to load *Express*.

`-?` Print usage message.

EXAMPLES

```
exinit -m 0x80069000 -K
```

Reinitialize the machine by killing all currently executing processes and loading *Express* at the indicated address. The current contents of node memory will remain intact, except for the region near 0x80069000 which will be overwritten by the kernel.

SEE ALSO

`exstat`, “CnfTool: Configuring *Express*”, “Using *Express* with Multiple Hosts”, “Excustom: Customizing *Express*”.

NAME

exreset - Reset a group of nodes.

SYNOPSIS

exreset

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command is used to reset a set of boards without loading the *Express* kernel. If your hardware configuration is capable of supporting a tree-like reset path in which all nodes can be reset from a central “master” console this command will be totally unnecessary since `exinit` will be able to reset and load *Express* into all nodes.

If your hardware does not support this chaining of reset signals then you will need to partition the nodes into groups, each of which can be reset from a particular host. The `exreset` command will then perform this operation allowing a subsequent `exinit` to load *Express* into the entire network.

SEE ALSO

`exinit`, “`CnfTool: Configuring Express`”.

exstat

NAME

exstat - Display node usage information.

SYNOPSIS

```
exstat [-l] [device]
```

DESCRIPTION

exstat is used to present statistics about the current node usage on the multiprocessor device. The default value will be either `/dev/transputer` or `/dev/ncube` depending on the hardware installed.

Without the `-l` switch only the number of allocated nodes and the total number of nodes are presented. The inclusion of this switch also provides information about which process is allocated which nodes, and which processes share which nodes.

Until the system has been initialized with a call to `exinit` this utility will return the rather disconcerting result that there are no nodes available.

OPTIONS

- `-l` Produce an extended (long and informative) listing which includes the process I. D. and physical node origin of all active processes.
- `-?` Print usage message.

SEE ALSO

`exinit`

NAME

ndb - Symbolic, source and assembly level debugger for parallel computers.

SYNOPSIS

```
ndb [-?] [-I incdir] [-p procid] [-d doc] [-n nodes] file
```

DESCRIPTION

ndb is an interactive symbolic debugger for use on parallel computers. Two styles of interface are available depending upon the particular hardware/compiler combination available.

The simplest interface is a source level debugger patterned after the UNIX utility dbx. At this level the user is able to examine code, set breakpoints and examine variables at the level of the original C or FORTRAN source code.

The lower level interface is designed for machine level debugging and is based on the conventional assembly level debugger adb. It allows for the examination of both data and assembly level code as well as the setting of run-time breakpoints. ndb incorporates a superset of adb commands which should facilitate its use.

ndb is always available at this second level and the commands associated with its use are described completely in this document. The availability of the source level interface is subject to compiler/hardware restrictions. The associated commands are described in this manual but may be unavailable in some implementations.

In order to effectively debug parallel programs a simple extension to the conventional debugger syntax has been made. This is the concept of a "set" of processors. Each command is executed on some group of nodes which can be defined and altered by the user. Several common groups of processors are predefined and user defined sets are also available.

COMMAND LINE OPTIONS

- | | |
|-----------|---|
| -d doc | Specify dimension of subcube to be debugged. Default is 0 (1 node). |
| -n nodes | Alternative to -d switch - specify the number of nodes rather than its logarithm. |
| -p procid | Debug a background process. This option is useful in conjunction with the -P switch to cubix which loads a program and stops it at its starting point. It is also used to perform post-mortem debugging on processes which are "hung". |
| -I dir | Specify a directory to be searched when looking for source code. By default only the current directory is searched. This switch may be repeated multiple times and the associated directories are searched in the order they are specified. |
| file | Specify the program whose symbol table is to be read. Unless this name is specified no symbol table entries will be available which significantly reduces the capabilities of ndb. |

-? Display information about supported run-time switches.

USAGE

An introductory guide to the debugger is available elsewhere so the following sections merely explain the syntax of the various commands. In nearly all cases the low level syntax is exactly as in the regular UNIX command `adb` while the source level command shares the same syntax as `dbx`.

1 Help

A certain amount of on-line help is available by entering the `help` command. Various topics may be selected for further perusal by entering

```
help topic
```

where `topic` is the name of the required subject. The syntax for a particular command can be found by using

```
help keyword
```

where *keyword* is the identifier whose usage is required.

2 Sets

Each command is executed on a collection of nodes called a "set". A 'current set' is maintained by `ndb` which is used unless overridden by prefixing a command with the `on` keyword. There are three verbs used in manipulating sets.

<code>on</code>	The following set specification is used for the command that follows it and then the current set reverts to its previous value.
<code>pick</code>	The current set is changed to the nodes listed.
<code>setdef</code>	Define a new set containing the specified nodes. The resulting set is assigned an index number which is returned and can be used in future set specifications.

The set specifications are built up from (comma separated) elements of the following types. (In the following the symbol `#` denotes a decimal, integer, constant).

<code>#</code>	A single node.
<code>node #</code>	A single node.
<code># to #</code>	An inclusive range of nodes.
<code># thru #</code>	An inclusive range of nodes.
<code># - #</code>	An inclusive range of nodes.
<code>all</code>	All nodes in the subcube.
<code>even, odd</code>	Either the even or odd parity nodes defined by the number of bits set in the processor number.
<code>set #</code>	The set with index <code>#</code> as given by a previous <code>setdef</code> instruction. (Note that set numbers are indicated when using the <code>setdef</code>

command or with the “show sets” command.

`nof #` The “hypercube” neighbors of node #.

`neighbors` The “hypercube” neighbors of all the nodes listed so far.

`board n0;n1`
Specifies node `n1` on attached peripheral board `n0`. Ranges of nodes may also be given as `n1-n2`.

3 Displaying Source Code

The simplest command for displaying source code is `list`. With no parameters this prints out ten lines of program starting at the “current line”. This latter is set implicitly during program execution by the `show state`, where and single stepping commands. It may be explicitly altered with

`file name` Set the “current line” to the first line of the named file. If no name is given display the current file.

`func name` Set the “current line” to the first line of the named function. With no parameters display the current function.

The `list` command has parameters itself which are either one or two integers separated by a comma. The various combinations of positive and negative values are used to indicate ranges of lines to display. A few examples should clarify the details.

`list 20` Display line 20 only.

`list 20, 50` Display lines 20 through 50 inclusive.

`list -5` Display ten lines starting 5 before the current line. This option provides a “window” facility.

`list -5, 20`
Display 20 lines starting 5 before the current line.

`list function`
Display the first ten lines of the named function.

When source files are named, either explicitly in `file` commands or implicitly during program execution an internal directory search path is used to look for appropriate source files. Two methods are available for altering this path. When starting an `ndb` session the “-I” command line argument names a directory which should be appended to the current search list. Alternately the `use` command can be used - it is followed by a list of directories which replaces the current list. Thus

```
use . ../src ../lib ../tst
```

might be given to name several directories in which source files are to be found. The order of the entries in this list is important - directories are searched from left to right so possible name clashes may have to be considered. Of course, if `ndb` fails to find the correct version of a source file with its automatic search then the `file` command can be used with a full pathname to override `ndb`’s choice.

The use command with no parameters displays the current search path.

4 Stack Operations

The most useful command for finding the current status of the program is `where`. With no arguments this command displays the top 8 levels of stack activity naming subroutines called and displaying their arguments. If less than 8 subroutines have been called the list terminates with the system initialization routine. If more, or less than 8 levels are required then a numeric parameter may be given; `on all where 3` displays the top three function calls in all nodes.

It is important to note that these stack tracing operations require that several probing messages be sent to each node. As a result they may work incorrectly when the node program is actually executing since the stack may be significantly different each time leading to inconsistent results. If this is a problem then one of the single step commands can be used to effectively “stop” the nodes allowing `where` to succeed. The program can be later continued with the `cont` command.

Another useful function in connection with tracing subroutine calls is `isin`. One can say

```
on all isin main
```

to obtain a listing of the activity of all nodes which currently have the named routine in their stacktrace. This is a useful form of data reduction since it allows the user to immediately discover which nodes are in the wrong place.

`dump routine` displays the calling sequence and local variables of the named routine if it occurs in the stack backtrace. If no function is specified then that containing the current program counter is dumped.

Note that these commands may interact with the language specification flag discussed in context of the `ndbenv` command. Often different language compilers use different calling sequences that cannot be dynamically inferred from the actual code. In this case a stack trace may be invalid unless the language switch is set correctly. To change from C (the default) to Fortran, for example, one issues the command

```
ndbenv Fortran
```

Commands such as this are often best placed in the user’s initialization file `.ndbinit`.

5 Displaying data

The simplest command for showing data values is `print` which takes a list of expressions and prints their values according to the variable types indicated by the program. One can, for example, type

```
print 1+2.5, my_struct
```

to which the response might be

```
1+2.5 = 3.5
my_struct = {
  i = 12
  f = 1.44
```

}

Some compilers do not preserve information about symbol types in which case you have to explicitly indicate in what manner you wish to see the data displayed. Occasionally you may also wish to override ndb's choice of formats for a printed variable.

Data display requests take the symbolic form

```
address, count / format
```

or

```
address, count ? format
```

The first form reads data from the nodes themselves while the second accesses the actual executable on disk. For this reason the second form is to be preferred when looking at assembly code while the first is most common for actual data.

Essentially the `format` field of the command is an instruction which explains how to print data. This command is repeated `count` times starting at `address`. The `address` and `count` fields may contain any valid expression (as explained below) while the `format` field contains any number of modifier characters which denote how a particular datum is to be formatted. The particular characters and their interpretation is as follows

D	32 bit integer
d	16 bit integer
O	32 bit octal
o	16 bit octal
X	32 bit hex
x	16 bit hex
U	32 bit unsigned
u	16 bit unsigned
b	8 bit octal
B	8 bit binary
c	8 bit ASCII
C	8 bit ASCII with interpretation of control characters
s	Null terminated string
S	Null terminated string with control character interpretation
i	Disassemble instruction
I	Source module and line number (No '.' increment)
a	Symbolic address (No '.' increment)
p	Symbolic address
Y	Date and time
t	Tab
r	Space
n	Carriage return
+	Increment '.' by current amount
-	Decrement '.' by current amount
^	Backup '.' by current amount

Each of these characters can be preceded by a repeat count. ndb maintains an idea of the current address on each node which is referred to by the special symbol '.'. Each of the formatting instructions (except those specifically mentioned) increments this quantity by the size of the object to be displayed.

Several other commands of this format are allowed and are denoted by the modifiers listed below

v	Dump data as both hex and ASCII. The count field denotes the number of 16 byte lines of data to show. A repeat count before the v character requests hex data values of that length (in bytes), i.e., 4v requests a dump in 32 bit words.
l value mask	This command searches from the given address through count bytes looking for a value which, when "AND"ed with mask is equal to value. The default search length is 4096 bytes. Warning: this option is VERY slow
L value mask	Searches for a 32 bit match. See previous modifier.
w value	Write the specified 16 bit value at address
W value	Write the given 32 bit value at address.

6 Expressions

ndb recognizes most of the usual arithmetic operators in expressions. Symbolic constants are also recognized with or without the preceding '_' added by the C compiler or the conversion to upper case performed in Fortran. The scope rules for simple variable names is to look in the current function (as denoted by the register PC or the most recent func command) and then the external variable table. References to local variables in other than the present function can be made by specifying a full symbol name of the form

function`variable

NOTE that certain keywords are reserved for the use of ndb and thus cannot be used as variable names. Since none of these words begin with an '_' character the variable with the same name can always be referenced by including the underbar.

The various legal expression elements are

.	The value of the current address.
<name	The value of the named register.
(exp)	The value of the enclosed expression
name	Address of the named symbol using the scope rule that the function denoted by the current program counter is searched first followed by the external variable table. (Can be overridden with the func command.)
routine`name	The value of the variable name in the given subroutine which

must be in the current stack backtrace.

The following are allowed operators in expressions

+	Binary addition
-	Either binary subtraction or unary minus
*	Either a pointer dereference or binary multiplication
%	Binary integer division
&	Binary AND operator
	Binary OR operator
^	Binary XOR operator
~	Unary NOT operator
#	Round first operand to next multiple of second
<<	Left shift.
>>	Right shift
@	Pointer dereference

In addition to using expressions to specify addresses it is also possible to use ndb as a regular integer calculator. The values of expressions are printed by following the expression with '=' and a format specifier as indicated in the previous section on displaying data. Thus

```
0x1234 + 16*(1<<3^{ }2 - 3) = X
```

prints an uninteresting 32 bit hex value.

7 The "show" command

Various special commands have been added to the usual syntax to take advantage of some special features of the parallel machine. These commands are all of the form

```
show something
```

where the something is chosen from the list (Other options may be available on your system, type "help show" for details.)

breakpoints	List active breakpoints.
pregs	Internal processor communication registers.
queues	Unread messages for this node.
regs	General processor registers, current instruction and source file location.
sets	User defined sets.
state	Process state, current instruction and source file location.
times	Idle and active times in this processor.

Note that only an initial substring of the names listed above is necessary to pick options so that `show st` is equivalent to `show state`.

8 Arrays

The simplest way to print out array elements is with the `print` command described previously. If you have to resort to the lower level formatting commands for some reason then array indices are indicated in a different way.

This syntax follows the C and Fortran notation with the addition that the user must specify the declared dimensions of the array as well as the indices required. In C, therefore, the syntax to print out the element `lattice[2][4][5]` from a 10x 10x 10 array as a 32 bit integer is

```
lattice[2;10][4;10][5;10]/D
```

where the values after the semi-colons are the declared dimensions of the array. In order to perform offset calculations `ndb` has to know the size of an individual array element. By default the value is taken to be 4, correct for integer and 32 bit floating point data types. If the data item is actually of a different size - e.g., a byte or a structure then this can be specified in braces after the array name. Thus to print out values from an array of sixteen byte structures one might use

```
complicated{16}[3;8][4;12]/ddfff
```

The notation for Fortran style arrays is similar with the array indices being paired up with dimensions via semi-colons. An example might be

```
array(3;4, 5;120)/f
```

`ndb` understands the difference in ordering between multi dimensional arrays in C and Fortran as well as the fact that Fortran array indices start at 1. It also nests array dimensions arbitrarily deeply.

9 High level job control

Several high level commands are available for running and controlling user programs. The first set are used to start up either the debugger or the user application

```
run string
```

The command indicated by the "string" parameter is executed and the debugger attempts to attach to the resulting process as though started with an appropriate "-p" option. If no `string` is given then the previous `run` command is re-executed. I/O redirection is allowed with the usual '<', '>', '>&', '>>' syntax.

```
debug program process
```

This command can be used to name the program that is to be debugged after `ndb` has started. This is useful if more than one executable is loaded into the machine since it allows switching between symbol tables. The `process` argument is optional and specifies the process-ID number of the program that is to be debugged. The `program` argument can be the single character '-' in which case no symbol table will be loaded but a process-ID may still be given.

`io` This command is used to redirect terminal input to the user program. By default `ndb` reads all characters typed and interprets them as debugger commands. After this command all terminal input is sent to the user program. To issue further debugger commands use the keyboard interrupt sequence (usually CTRL-C) to return control to `ndb`.

`kill` Kill the program being debugged. Confirmation is requested.

As well as these functions commands are also available to control the execution of the user code at a finer level through breakpoints, single stepping etc. The commands are

`stop in name` Insert a breakpoint at the first line of the named function.

`stop at number` Insert a breakpoint at the indicated line in the current source file. Note that this command interacts with the `file` and `func` commands discussed earlier.

`stop variable` Continue execution in single step mode and halt the program when the named variable changes value. This command executes rather slowly due to the interpretive nature of the processing required.

`cont` Continue execution from a breakpoint, or single step. This command interacts with the "wait" flag of the `ndbenv` command - by default the `ndb` prompt appears immediately allowing further commands to be entered. Alternately one can specify that `ndb` should continuously poll the nodes until another breakpoint is found or the application terminates before prompting again. This latter behavior is most common in sequential debuggers but slows down `ndb` somewhat as polling is an inherently slow process.

`step n` Single step the program over "n" lines of source code (default 1). If function calls are detected then the single stepping process enters each subroutine. If the current node "set" contains a single processor then this command will display the source lines as they are stepped past.

`next n` Single step over "n" lines of source code without entering any new functions. (Default 1 line). This option is similar to `step` but avoids the problems of having to single step through system functions etc. Source lines are displayed as processed if the current "set" contains a single node.

`status` Display the list of active breakpoints indicating which nodes they are present in, where they are placed and an index number used for deletions.

`delete n` Delete the breakpoint with index number "n" as determined from the status command.

Note that any of these commands may be prefixed by a "set" specification to allow different actions to be performed n distinct nodes. Thus to insert breakpoints in only the first three nodes one might use

```
on 0-2 stop at 23
```

10 Miscellaneous commands

Several miscellaneous commands are available to make debugging easier.

`sh string`
`! string` Any command line that begins with '!' or sh is executed by the shell.

`pwd` Show current directory.

`cd directory` Change to an new directory. This is occasionally useful for finding source files since the default search path starts with the "current directory".

`source file` Read ndb commands from the named source file. This is useful for performing repetitive tasks or for making data dumps. Consider also the `$>` command which redirects the output from the debugger. By default ndb attempts to find a file named `.ndbinit` in either your home directory or the current directory whenever started and reads initialization commands from it.

`alias s1 s2` Define a new command. Henceforth the command `s1` will be treated exactly like the command `s2`. The command

```
alias l list
```

for example, allows one to use the single character 'l' instead of the `list` command. It is also possible to set up aliases with arguments and defaults using the UNIX C-shell syntax. The command

```
alias myuse use !:{1-.} !:{2-../src} !:3 !:4 !:5
```

defines a new command for setting the source code search path. 5 arguments are specified and the first two have defaults "." and "../src" so that the simple command `myuse` can be issued without any arguments to set the search path to '. ../src' or arguments can be specified to set the path to other things.

`quit` Exit ndb. If the user program started within ndb a "kill" command will be given and you will be asked whether to terminate the program or not. If the program started outside of ndb it will be left alone.

11 The `ndbenv` command

This command defines the specific “environment” in which `ndb` is working. The currently implemented settings are the high level language being debugged, the “wait” state, the “repeat” mode and the “symbol match length”. To see the options currently in effect type

```
ndbenv
```

which might yield

```
Language:           C
Wait mode:          FALSE
Autorepeat:         OFF
Symbol match length: 8192
```

Each of these options is explained below.

Language	Certain features of <code>ndb</code> depend implicitly upon the high level language being debugged - for example array indexing and stack tracing. By default <code>ndb</code> is in the C mode suitable for the “C” language but may be switched over to Fortran with the command <code>ndbenv F</code> .
Wait mode	This parameter controls the behavior of <code>ndb</code> upon receiving a <code>cont</code> or <code>run</code> command. By default the prompt immediately reappears and the user is able to enter further debugging commands while the node program continues to execute. If the wait state is set to <code>TRUE</code> with the command <code>ndbenv wait</code> then <code>ndb</code> continuously polls the nodes and only returns control to the user when all nodes have stopped at breakpoints or with some error. This mode can be turned off again with <code>ndbenv nowait</code> .
Autorepeat	By default <code>ndb</code> repeats the last command entered whenever the user command is a single carriage return. This feature can be disabled with the command <code>ndbenv norep</code> .
Symbol Match Length	When translating memory address into variable names <code>ndb</code> uses a cutoff to avoid translating system memory addresses into user variable names - i.e., addresses further than this length above a known symbol will be translated into hex values rather than “name+offset”. By default this cutoff is 32768 bytes. On occasions it may be necessary to increase this number so that large functions appear by name in stack traces rather than as hex numbers.

12 Assembly Level Debugging

In addition to the `adb` implementation effectively described in the next few sections the following commands are available for debugging at the machine code level.

<code>listi address</code>	Display ten machine instructions from the given address. If none is given continue from the last address specified.
<code>stepi n</code>	Similar to the <code>step</code> instruction but considers only the machine code. Encountered subroutines are entered and the

	machine registers are displayed if the current node "set" contains only a single node.
<code>nexti n</code>	Similar to <code>stepi</code> but passes over subroutine calls.
<code>stopi address</code>	Place a breakpoint at the named address.

13 Assembly Level Job Control

Various commands are available which allow one to control the execution of a node program. They are all of the general form

`arg1, arg2 : modifier string`

in which `arg1` and `arg2` may be any general expression and the various modifiers are listed below. (Note that some cases do not require arguments in which case `arg1` and `arg2` can be omitted)

- b Set a breakpoint at address `arg1`. Note that only 8 breakpoints may be set in any node at one time so an attempt to set more will result in a request from `ndb` to delete an entry.
- d Delete the breakpoint at address `arg1`.
- s Step processor over a single machine instruction.
- c Continue as from a breakpoint.
- C Continue from breakpoint but instead of returning control to `ndb` immediately wait for the node specified as `arg1` to hit a breakpoint. If `arg1` is omitted wait for node 0.
- k Kill the process inside the machine.
- K Compare `arg2` bytes of memory starting at address `arg1` on all nodes in the active set.
- r Run the command specified by the `string` argument under the control of the debugger. I/O redirection is available with the usual constructions '>', '>>', '<', '<&'. Note the comment above on terminal input to the running process.
- R Run a command, as with the `r` specifier, above, but wait for the process specified in `arg1` to hit a breakpoint before returning control to `ndb`. If `arg1` is omitted, wait for node 0.

`ndb` leaves the debugger in control of the terminal even when continuing from breakpoints. This is contrary to conventional sequential debuggers which normally switch over to sending input to the debugged process whenever it is running. This distinction is made because of the distributed nature of parallel applications where it is not unusual to have some nodes running while looking at the state of others. If the running process requires terminal input the single command `io` switches control from `ndb` to the user process sending all further keyboard input to that process. To return control to the debugger use the interrupt sequence (usually CTRL-C).

14 Assembly Level System control

Various commands are available to control the way ndb interprets and outputs its results and to access some of the more machine specific requests. They all take the general form

```
arg1 $ modifier string
```

where `arg1` is any legal ndb expression and the modifiers are as follows

- b List all active breakpoints. The notation for the nodes on which the breakpoint is active is essentially a bit mask with each bit (reading from left to right) denoting a single processor.
- c Traceback of all active C procedures together with their arguments interpreted as 32 bit hexadecimal constants.
- C As in the 'c' option above but prints out the values of all known local variables. Note that the appropriate compiler options must be used to compile information about local variables.
- d Set default base for numbers to 10.
- e Print out all external variables and their values interpreted as 32 bit hex constants.
- f Traceback of all active Fortran subroutine calls. No argument information is supplied by the compiler so the first few elements off the stack are interpreted as 32 bit hex constants.
- F Fortran traceback showing all local variables are 32 bit hex constants. Note that the appropriate compiler switch must be used to include information about local variables.
- m Print memory map of current program showing sizes of various data areas.
- n Show internal processor communication registers.
- o Set default base for input to octal
- q Quit from ndb. If you entered ndb via the `-p` command line option the node process is left alone. Otherwise it is killed.
- r Print general processor registers together with an interpretation of the current instruction and the source line/module information.
- s Set the maximum offset from the public symbol to `arg1` for which ndb still interprets an address as being within that function.
- t Show a one line status summary for each processor showing the current state, program location and source file/line number information
- w Set the output page width to `arg1`.
- > Redirect output to the file named in `string`

WARNINGS

Error checking in ndb is rather primitive. Furthermore if an error is actually detected it will

quite probably be misdiagnosed. Certain words are reserved for use in commands and cannot, therefore, be used as variable names. The full list of reserved words is as follows: on, setdef, pick, thru, to, set, node, even, odd, all, show, help, quit, io, neighbors, nof.

Programs which put the nodes into strange states may also affect the debugger in odd ways.

DIAGNOSTICS

The prompt issued by ndb attempts to indicate the current set to which commands will be applied. Most variations are self-explanatory except the mysterious word `array` which indicates a node combination too complicated to figure out.

Syntactical errors on input generate many splendid messages, some of which might even complain about errors.

If no program is given on the command line a warning is issued about the lack of a symbol table.

Various out of memory errors produce both fatal and non-fatal diagnostics. Error recovery from these cases may or may not work.

Attempting to load a non-standard executable program will fail and produce a message suggesting corrective action.

BUGS

Printing non-floating point values with the `f` or `F` formats occasionally leads to core dumps. This sometimes happens even with legal floating point values under XENIX due to deficiencies in the run-time support.

The exact abilities of ndb depend a lot on the underlying operating system and hardware characteristics. As a result it is not possible to implement all features of ndb in all *Express* versions.

SEE ALSO

“NDB: A Guide to Parallel Debugging under *Express*.”

NAME

tcc - Compile and link *Express* C and C⁺⁺ programs for Transputers.

SYNOPSIS

```
tcc [-B address] [-c] [-o outfile] [-Dname[=value]]
    [Idirname] [-Uname] [-E] [-g] [-dryrun] [-K] [-llibname]
    [-r] [-P] [-S] [-T0] [-T4] [-T8] [-e name] [-N] [-x] [-n]
    [@filelist] files...
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command provides an interface to the Logical Systems C compiler useful for compiling programs to be run under *Express*. Filenames ending with the '.c' suffix are taken to be C source code and are compiled while those ending with '.tal' are treated as assembly code source and assembled. In both cases the resulting output files have the '.trl' suffix. Note that the preprocessor is run on assembler files by default allowing some of the advanced features of the Logical Systems assembler to be used.

After compiling all source files tcc proceeds, by default, to link the resulting object files into an executable program. If no '-o' switch is provided this will have the name trans.tld. By default libraries are searched which allow access to the *Express* communication routines only. The *Cubix* and *Plotix* subsystems are included with the -lcubix and -lplotix switches respectively. (It is important to note that programs linked with either of these switches *CANNOT* be executed with normal host programs in the "host-node" mode. Similarly programs compiled without one of these switches will not run with the cubix program.)

In addition to producing the executable image two additional files are (optionally) produced which have suffices '.sym' and '.map'. The former is used by the source level debugger ndb while the latter is of general interest - it contains information about the memory locations of program variables and which libraries and object files were searched.

By default all compilation/linking is performed for T800 transputers. Note that object files and libraries are not necessarily interchangeable between the two CPU types since instructions may be generated that are not supported on both types of hardware. The '-T4' switch is provided to force the generation of programs suitable for execution on T400 series systems. The '-T0' switch attempts to generate code independent of the transputer type by omitting instructions specific to only one model. Note that this switch does not support floating point operations.

OPTIONS

-B address Specify alternate load address for program. By default loading is performed at the beginning of "off-chip" memory. (See "Excustom" in the *Express* users guide" for more information.)

tcc

- c Compile only - do not proceed to link resulting object files.
- dryrun Print the commands to be executed without actually performing any of them. This options implies both '-x' and '-N'.
- Dname
- Dname=value Define preprocessor symbol and optionally assign a value.
- e name Specify an alternate entry point. By default the user program is entered through initialization routines required by *Express*.
- E Run preprocessor only. Output is left in a file with the suffix '.pp'.
- f2 Use 32 bit floating point arithmetic for all "double" variables. (Default is 64 bit.) Other options are also available - see the Logical Systems documentation for more details.
- g Include additional symbol table information for source level debugging. This switch adds additional code at entry and exit of *ALL* subroutine calls to enable stack-tracing which can significantly slow down execution.
- K Disable stack tracing. Used to suppress stack tracing, even when '-g' switch is given.
- Idirectory Add a directory to the path searched when looking for '#include' files.
- lcubix Search the *Cubix* library for unresolved symbols in addition to that required by *Express*.
- lplotix Search both *Cubix* and *Plotix* libraries for unresolved symbols.
- n Execute the link phase of compilation on transputers rather than the host system. (Only available on some systems.)
- N Keep all intermediate files. (Default is to delete them after use.)
- o name Specify an alternate name for the executable program produced by the linker. Default is 'trans.tld'.
- P Run preprocessor only. Output remains in a file with the extension '.pp'.
- r Generate "position-independent" code which can be relocated at runtime.
- S Produce assembly code listing of C source program.
- T4 Compile for T400 series transputers.
- T8 Compile for T800 series transputers.
- Uname Undefine a preprocessor symbol. Reverses the effect of '-D' switches or '#define' statements.
- x Display all commands before executing them.
- @filename Take "filename" to be a file containing a list of source or object files to be compiled or linked, one name to a line.

-? Print usage message.

EXAMPLES

```
tcc -c hello.c
```

Compile, but do not link the C source file `hello.c`. The resulting object file will be called `hello.tr1` and will be for the T800 series transputers.

```
tcc -o prog f1.tr1 f2.c f3.tr1 -lcubix
```

Compile file `f2.c` and proceed to link it with `f1.tr1`, `f2.tr1` and the *Cubix* libraries to make an executable program called `prog`. This executable will run on T800 transputers.

```
tcc -T4 -o prog4 g1.tr1 g2.c g3.tr1 -lcubix
```

This example is the same as the previous one but the resulting executable file, called `prog4` will run only on T400 series transputers. Notice that one cannot mix transputer types so the object files `g1.tr1` and `g2.tr1` must have also been produced with the `-T4` switch.

```
tcc -S -T4 foo.c
```

Generate an assembly code listing of the C source file `foo.c` suitable for a T400 transputer system.

```
tcc -Imyinc -DCUBIX -c noddy.c
```

Compile but do not link the C source code in the file `noddy.c` for a T800 series transputer system. Additionally define the `CUBIX` symbol and search the directory `myinc` when attempting to satisfy `#include` statements.

```
tcc -g -c noddy.c
```

Compile, but do not link, the file `noddy.c` for a T800 series transputer. Include both source line numbering information and also additional entry/exit subroutine calls to enable stack tracing. Note that the code resulting from this file will execute rather more slowly than would be the result if the `'-g'` switch were omitted.

```
tcc -o prog -g prog.bin subs.bin -lcubix
```

In this case the two named object files are linked together to produce an executable program called `prog`. In addition a symbol table called `prog.sym` will be created for use with the source level debugger, `ndb`.

MULTIPLE INPUT FILES

Some operating systems impose constraints on the length of a command line which preclude the linking of large programs with many input files using the standard `tcc` syntax. In this case `tcc` allows the list of filenames to be provided in a file and passed to the compiler using the `'@'` syntax. Consider, for example, a program made up of ten object files with names `"object0.tr1"`, `"object1.tr1"` and so on up to `"object9.tr1"`. In this case we would create a file containing the ten lines

```
object0.tr1  
object1.tr1
```

tcc

```
object2.tr1
```

```
....
```

```
object9.tr1
```

and save it with a name such as "link.lst". We could then invoke tcc with a command such as

```
tcc -o prog -g @link.lst -lcubix
```

to link the program with the *Cubix* libraries, build a symbol table for debugging and name the output file prog. Note that the suffix '.lnk' should not be used since tcc uses that name internally.

NAME

tcc3l - Compile and link *Express* C Transputer node programs.

SYNOPSIS

```
tcc3l [-B address] [-c] [-o outfile] [-Dname[=value]]
      [-Idirname] [-Uname] [-dryrun] [-i] [-g] [-llibname]
      [-T4] [-T8] [-x] [-N] [@filelist] files...
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command provides an interface to the 3L C compiler useful for compiling programs to be run under *Express*. Filenames ending with the '.c' suffix are taken to be C source code and are compiled. In both cases the resulting output files have the '.bin' suffix.

After compiling all source files `tcc` proceeds, by default, to link the resulting object files into an executable program. If no '-o' switch is provided this will have the name `trans.tld`. By default libraries are searched which allow access to the *Express* communication routines only. The *Cubix* and *Plotix* subsystems are included with the `-lcubix` and `-lplotix` switches respectively. (It is important to note that programs linked with either of these switches *CANNOT* be executed with normal host programs in the "host-node" mode. Similarly programs compiled without one of these switches will not run with the `cubix` program.)

In addition to producing the executable image two additional files are (optionally) produced which have suffices '.sym' and '.map'. The former is used by the source level debugger `ndb` while the latter is of general interest - it contains information about the memory locations of program variables and which libraries and object files were searched. The map file also contains the error messages, if any, from the linker. If the program aborts with a message such as

```
Failed to find .b4 file
```

this usually indicates that the link process failed with some error which can be located by searching for the string "ERROR" in the ".map" file.

By default all compilation/linking is performed for T800 transputers. Note that object files and libraries are not necessarily interchangeable between the two CPU types since instructions may be generated that are not supported on both types of hardware. The `-T4` switch is provided to force the generation of programs suitable for execution on T400 series systems.

It is important to note that the 3L compilers/linkers execute directly on the first transputer in the attached network and destroy and *Express* programs executing there, including the *Express* kernel itself. It is necessary to re-load the system with the `exinit` command before running any program.

OPTIONS

- B Specify alternate load address for program. By default loading is performed at the beginning of "off-chip" memory.
- c Compile only - do not proceed to link resulting object files.
- dryrun Print the commands to be executed without actually performing any of them. This options implies both '-x' and '-N'.
- Dname
- Dname=value Define preprocessor symbol and optionally assign a value.
- g Include additional symbol table information for source level debugging. Used at link time to force the generation of a symbol table for use with the debugger, ndb.
- i Do not search any of the built-in default directories for include files. Rely solely on the definition of the 3LCC_INC environment variable.
- Idirectory Add a directory to the path searched when looking for '#include' files.
- lcubix Search the *Cubix* library for unresolved symbols in addition to that required by *Express*.
- lplotix Search both *Cubix* and *Plotix* libraries for unresolved symbols.
- N Keep all intermediate files instead of deleting them.
- o name Specify an alternate name for the executable program produced by the linker. Default is 'trans.tld'.
- T4 Compile for T400 series transputers.
- T8 Compile for T800 series transputers.
- Uname Undefine a preprocessor symbol. Reverses the effect of '-D' switches or '#define' statements.
- x Generate a listing of all command lines before they are executed. This option is useful if certain commands need to be run by hand.
- @filename Take "filename" to be a file containing a list of source or object files to be compiled or linked, one name to a line.
- ? Print usage message.

INCLUDE FILE PROCESSING

The rules regarding the searching for include files in the 3L compiler are quite tricky. On UNIX systems some attempt is made to locate system include files according to the customization information supplied when installing the system. While this method is usually effective it can lead to extremely long command lines which cannot be processed by the 3L compiler. To avoid this situation the '-i' switch should be given, which suppresses the default search completely. In this case only those directories specified in the 3LCC_INC environment variable will be searched when looking for include files.

Under MS-DOS no attempt is made to locate default include file directories since the resulting command lines are nearly always too long for processing. In this case `tcc3l` will not execute unless the `3LCC_INC` variable is defined. If no such variable is found a suggestion will be made as to the correct assignment.

EXAMPLES

```
tcc3L -c hello.c
```

Compile, but do not link the C source file `hello.c`. The resulting object file will be called `hello.bin` and will be for the T800 series transputers.

```
tcc3L -o prog f1.bin f2.c f3.bin -lcubix
```

Compile file `f2.c` and proceed to link it with `f1.bin`, `f2.bin` and the *Cubix* libraries to make an executable program called `prog`. This executable will run on T800 transputers.

```
tcc -T4 -o prog4 g1.bin g2.c g3.bin -lcubix
```

This example is the same as the previous one but the resulting executable file, called `prog4` will run only on T400 series transputers. Notice that one cannot mix transputer types so the object files `g1.bin` and `g2.bin` must have also been produced with the `-T4` switch.

```
tcc3L -S -T4 foo.c
```

Generate an assembly code listing of the C source file `foo.c` suitable for a T400 transputer system.

```
tcc3L -Imyinc -DCUBIX -c noddy.c
```

Compile but do not link the C source code in the file `noddy.c` for a T800 series transputer system. Additionally define the `CUBIX` symbol and search the directory `myinc` when attempting to satisfy `#include` statements.

```
tcc3L -o prog -g prog.bin subs.bin -lcubix
```

In this case the two named object files are linked together to produce an executable program called `prog`. In addition a symbol table called `prog.sym` will be created for use with the source level debugger, `ndb`.

MULTIPLE INPUT FILES

Some operating systems impose constraints on the length of a command line which preclude the linking of large programs with many input files using the standard `tcc3l` syntax. In this case `tcc3l` allows the list of filenames to be provided in a file and passed to the compiler using the '@' syntax. Consider, for example, a program made up of ten object files with names "object0.bin", "object1.bin" and so on up to "object9.bin". In this case we would create a file containing the ten lines

```
object0.bin
object1.bin
object2.bin
...
object9.bin
```

tcc31

and save it with a name such as "link.lst". We could then invoke tcc31 with a command such as

```
tcc31 -o prog -g @link.lst -lcubix
```

to link the program with the *Cubix* libraries, build a symbol table for debugging and name the output file prog. Note that the suffix '.lnk' should not be used since tcc31 uses that name internally.

DIAGNOSTICS

If the linking procedure fails for some reason a rather uninformative message similar to

```
Failed to find .b4 file
```

is often generated. In this case the ".map" file should be consulted for error messages. (A good way to do this is to search for the string "ERROR" with a text editor or similar.)

NAME

tfc - Compile and link *Express* FORTRAN Transputer node programs

SYNOPSIS

```
tfc [-c] [-o outfile] [-g] [-llibname] [-T4] [-T8]
    [-dryrun] [-x] [-N] [@filelist] files...
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command provides an interface to the 3L FORTRAN compiler useful for compiling programs to be run under *Express*. Filenames ending with the '.f' suffix are taken to be FORTRAN source code and are compiled. The resulting output files have the '.bin' suffix.

After compiling all source files `tfc` proceeds, by default, to link the resulting object files into an executable program. If no '-o' switch is provided this will have the name `trans.tld`. By default libraries are searched which allow access to the *Express* communication routines only. The *Cubix* and *Plotix* subsystems are included with the `-lcubix` and `-lplotix` switches respectively. (It is important to note that programs linked with either of these switches *CANNOT* be executed with normal host programs in the "host-node" mode. Similarly programs compiled without one of these switches will not run with the `cubix` program.)

In addition to producing the executable image two additional files are (optionally) produced which have suffices '.sym' and '.map'. The former is used by the source level debugger `ndb` while the latter is of general interest - it contains information about the memory locations of program variables and which libraries and object files were searched.

By default all compilation/linking is performed for T800 transputers. Note that object files and libraries are not necessarily interchangeable between the two CPU types since instructions may be generated that are not supported on both types of hardware. The `-T4` switch is provided to force the generation of programs suitable for execution on T400 series systems.

OPTIONS

- | | |
|-----------------------|--|
| <code>-c</code> | Compile only - do not proceed to link resulting object files. |
| <code>-dryrun</code> | Print the commands to be executed without actually performing any of them. This options implies both '-x' and '-N'. |
| <code>-g</code> | Include additional symbol table information for source level debugging. If specified at link time force the generation of the '.sym' file for debugging. |
| <code>-lcubix</code> | Search the <i>Cubix</i> library for unresolved symbols in addition to that required by <i>Express</i> . |
| <code>-lplotix</code> | Search both <i>Cubix</i> and <i>Plotix</i> libraries for unresolved symbols. |

tfc

- N Keep all intermediate files instead of deleting them.
- o name Specify an alternate name for the executable program produced by the linker. Default is 'trans.tld'.
- T4 Compile for T400 series transputers.
- T8 Compile for T800 series transputers.
- x Print each command before executing it.
- @filename Take "filename" to be a file containing a list of source or object files to be compiled or linked, one name to a line.
- ? Print usage message.

EXAMPLES

```
tfc -c hello.f
```

Compile, but do not link the Fortran source file `hello.f`. The resulting object file will be called `hello.bin` and will be for the T800 series transputers.

```
tfc -o prog f1.bin f2.f f3.bin -lcubix
```

Compile file `f2.f` and proceed to link it with `f1.bin`, `f2.bin` and the *Cubix* libraries to make an executable program called `prog`. This executable will run on T800 transputers.

```
tfc -T4 -o prog4 g1.bin g2.f g3.bin -lcubix
```

This example is the same as the previous one but the resulting executable file, called `prog4` will run only on T400 series transputers. Notice that one cannot mix transputer types so the object files `g1.bin` and `g2.bin` must have also been produced with the `-T4` switch.

```
tfc -o prog -g prog.bin subs.bin -lcubix
```

In this case the two named object files are linked together to produce an executable program called `prog`. In addition a symbol table called `prog.sym` will be created for use with the source level debugger, `ndb`.

MULTIPLE INPUT FILES

Some operating systems impose constraints on the length of a command line which preclude the linking of large programs with many input files using the standard `tcc` syntax. In this case `tfc` allows the list of filenames to be provided in a file and passed to the compiler using the '@' syntax. Consider, for example, a program made up of ten object files with names "object0.bin", "object1.bin" and so on up to "object9.bin". In this case we would create a file containing the ten lines

```
object0.bin  
object1.bin  
object2.bin  
....  
object9.bin
```

and save it with a name such as "link.lst". We could then invoke `tfc` with a command such as

```
tfc -o prog -g @link.lst -lcubix
```

to link the program with the *Cubix* libraries, build a symbol table for debugging and name the output file `prog`. Note that the suffix `.lnk` should not be used since `tfc` uses that name internally.

DIAGNOSTICS

If the linking procedure fails for some reason a rather uninformative message similar to

```
Failed to find .b4 file
```

is often generated. In this case the `.map` file should be consulted for error messages. (A good way to do this is to search for the string "ERROR" with a text editor or similar.)

xtool

NAME

`xtool` - Analyze Execution Profile

SYNOPSIS

```
xtool program_name [log_file_name]
```

DOMAIN

This command is available at the system prompt on the host processor.

DESCRIPTION

This command is used to examine and analyze the log file created with the execution profiler, `xprof`, commands. The first argument is the name of the executable program to be profiled and the second is the name of the file containing the profile data. This may be omitted if it has the default value "`xprof.out`". Note that the execution profiler relies on data contained in a symbol table for correct functioning. This can usually be made by specifying the `-g` switch when linking the program - the same procedure as used for debugging with "`ndb`".

This command presents a separate table on `stdout` from each node. The information contained in each table is:

- An identifier showing which node the following data is from.
- A summary of the busy and idle time in each processor. In this regard we measure CPU time so that the only "idle" time is when the CPU is not actively executing the process such as when waiting for a message to arrive. All other classes of activity are counted as "busy". Note that this interpretation is different from that of `ctool` which distinguishes between calculation and communication time.
- A count of the number of profiling "misses". Since the buffer supplied to the profiling function `profil` may not be large enough to encapsulate the entire program it is possible that the execution profiler will "miss" occasionally - i.e., the program will be executing at an address which lies outside the region mapped by the `profil` call when it tries to log the profile event. In this case the "miss" counter is incremented. The ratio of hits to misses is presented to give a guide to the effectiveness of the profile obtained - a lot of misses means that the routines in the profile list may not, in fact, be the most heavily used.
- A profiling list containing the most heavily used 20 functions in the program. Each shows the fraction of the total profiling events that it corresponds to.

This data appears on `stdout`.

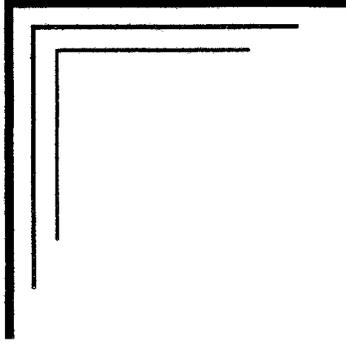
EXAMPLES

To examine the profile data in a file called "`phase3.prof`" created by the program `master` execute the command

```
xtool master phase3.prof
```

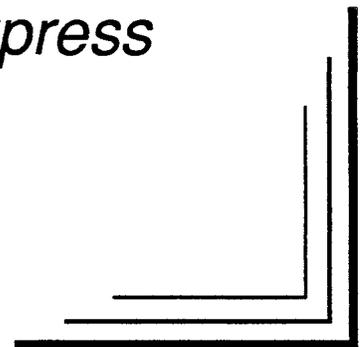
SEE ALSO

`cubix` in the *Express* documentation.



FORTRAN runtime library

Library routines available to *Express*
programs written in FORTRAN



This (large) section of the manual is devoted to a listing of the contents of the subroutine library which is invoked by *Express* programs.

Since parallel processing is an inherently complex activity the capabilities of *Express* are correspondingly broad. This, in turn, leads to a very extensive set of functions which may appear daunting to those familiar with other parallel processing systems or totally unfamiliar with parallel computing. These users should not, however, be put off by the long list of routines given in this section since we have found that practical applications use only a small fraction of the available number. Unfortunately different programs tend to use different small subsets of the total list which makes predictions difficult.

As a help in selecting the appropriate functions we have tried to indicate routines with similar or related functionality in the "SEE ALSO" section at the end of each manual page. In conjunction with the full manual and the numerous "EXAMPLES" this should give a reasonable guide.

One of the most important pieces of information contained on each manual page is in the "DOMAIN" section. This paragraph tells you whether the routine in question is available to programs running on the "host" computer or to those running on the nodes of the parallel computer system. In the latter case there is also an indication of which library switch is required to gain access to the routine. Note that this information must be used in conjunction with that concerning the "Host-Node" and/or *Cubix* programming models.

If you are using the former style of computation then only routines shown as appearing in all node libraries may be called from your "node" programs. Routines shown as appearing in the *Cubix* or *Plotix* libraries cannot be called from such programs.

On the other hand, if you are using the *Cubix* model of computation you may freely call *nearly all* of the routines described in this manual being careful only to specify the *Plotix* libraries for the graphical routines. The exceptions in this case are those routines which specifically interface to similar routines in the host processor - since you will not be writing a program to run on the "host" you cannot call the corresponding routine there! Typical examples are the CP/ELT combinations such as KCPCP and KCPELT. To achieve the effect of these routines in *Cubix* programs one would instead use KCPEND.

The information regarding which routines are available in which libraries and to which type of programming models they belong is summarized in a later section of this manual where we also show the correspondence between the various language variants of *Express*.

The various routines contained in *Express* can be classified according to their functionality in several broad categories. The following sections attempt to indicate some of the important features of each and also to supply, in a condensed form, some information about important *Express* parameters and the header files necessary to use them.

1 High Level Communication System

This section describes the communication system available to application programs under *Express*. Several levels of functionality are provided although some features are common to all. While one may use the system to send messages to specific destination processors by specifying their processor numbers one can also use the primitives in a "topology-independent" manner.

The KXGRID system allows problems to be specified in the domain of the user data structures and

can be used to determine processor numbers automatically for use in the communication primitives. Using this system it is possible to design applications that have absolutely no knowledge of the underlying hardware topology and which will, in fact, execute transparently on any hardware that supports *Express*. Similar routines are available to dynamically configure an application to the available processing resources at runtime.

Several concepts underpin the entire communication system and can be summarized as follows;

- All messages have “destinations”. This merely specifies the node to which the message will be sent.
- All messages are “typed”. As well as the information concerning what data is to be sent and to whom every message has a `type` field which allows receivers to distinguish between various messages.
- The message reception mechanism has an “acceptance” criterion. All read routines may specify source and type information which constrains the range of messages which may be read. This information may either limit attention to specific node/type combinations or various degrees of “dontcare” behavior may be specified using the wildcard value `NO CARE`.
- Both blocking and non-blocking read functions are supplied.
- Messages are “atomic”. A single read operation corresponds to a single write operation. If the sender transmits more data than is read then the excess are discarded and may NOT be read with another read request. If less data are sent than were requested then the message is read and a return code indicates the discrepancy - another read request may not make up the difference unless another write request is also made.

On a more functional level the following generalities may also be observed:

- The “node” and “type” information associated with a message are always returned to the calling routine. In read requests, for example, a wildcard specification will be overwritten with the actual parameter value.
- The general ordering of arguments is: what, how much, where and type - i.e., the first arguments specify what data is to be transmitted, the second how much, the third indicates to whom the data should be sent and the last argument denotes the type of the associated message. This standard leads to an obvious calling sequence for the simplest “read” and “write” operations

```
ISTAT = KXREAD (BUFFER, LENGTH, ISRC, ITYPE)
ISTAT = KXWRIT (BUFFER, LENGTH, IDST, ITYPE)
```

Some calls which both read and write data have the above sequence duplicated for both operations so the `exchange` function, for example, has the calling sequence

```
ISTAT = KXCHAN (INBUF, INCNT, ISRC, ITYPE,
               OUTBUF, OUTCNT, IDST, OTYPE)
```

2 Hardware Dependent Communication System

Express has been carefully designed to allow programs to be written which will execute

transparently on a wide variety of different parallel architectures. As a result, none of the basic primitives make any reference to the underlying hardware configuration. On occasions, however, portability may be a less important goal than absolute performance on a particular piece of hardware. To support those users who have this type of constraint an *Express* subsystem has been provided with a “raw” interface to the communication hardware. Typically the use of these routines disables most of the higher level processing of which *Express* is capable such as the debugging and performance analysis tools. For this reason we suggest that its use be adopted as the final stage in the development of any parallel processing project after whatever bugs and/or performance questions have been resolved by using the full *Express* system.

3 Synchronization

One of the key concepts which underlies all of *Express* concerns interprocessor synchronization. In some sense this issue is the key to all of parallel processing - different algorithms can most often be classified not by the particular scientific or other field from which they arise but by the way in which they necessitate interprocessor synchronization. In *Express* we classify two types of behavior:

- **Asynchronous**
“Asynchronous” system calls can be made in any node at any time regardless of the activities currently occurring in other nodes. One can consider that the node making the call is operating totally in isolation.
- **Loosely synchronous**
A “loosely synchronous” system call can be perceived as a barrier to the further progress of the program. When one node makes a loosely synchronous call it waits for all other nodes to make the *same* system call (albeit with possibly different arguments). When all nodes have made the call every node proceeds. This concept might be classed “synchronous” but this is too restrictive - it is quite permissible for one node to make the “loosely synchronous” call far ahead of the other nodes. All nodes will, however, be synchronized *after* the call completes.

Note that these behaviors are not (usually) *states* of the system but are applied individually to different function calls. The function `KXWRITE`, for example, which sends an *Express* message may always be made asynchronously - i.e., any node may send a message at any time. Similarly any node may call `KABORT` to terminate a program at any time. On the other hand, `KMULTI`, the system call which switches between file I/O modes must *always* be made “loosely synchronously”.

Because the synchronization properties of a parallel program are often the key to its construction and optimization, the situation is actually more complex than just discussed.

The default state of *Express* is that every system call has an associated synchronization property. These states are listed in section 3 of this manual. Also available (in the *Cubix* library) is a global override function, `KCBXSY`, which switches all system calls to asynchronous mode.

At a slightly more useful level, each open file has its own synchronization property. This allows, for example, a program to have a global input stream for basic parameters, individually (and asynchronously) accessed data files for operational data and error reporting, distributed (synchronized) files for output data, etc. In each of these modes different requirements are made by *Express* on what can and cannot be done to the files.

Even within the “asynchronous” functions there are different levels of behavior. The KXWRIT function mentioned earlier, for example, may be called at any time in a user program, but it does not return to its caller until a message has been transmitted to the receiving node. The analogous KXSEND system call also sends a message to another processor but returns immediately to its caller without waiting for the data to be transmitted. While both of these calls are “asynchronous” in the sense that the start of the operation may occur in any node at any time KXSEND is clearly “more” asynchronous than KXWRIT because the point at which the buffer containing the data which has been sent can be re-used is not known when the call returns.

This discussion may have convinced you that the topic of interprocessor synchronization is too complex to ever be understood. This is not, however, the case. While it is true that many of the elementary bugs in *Express* are caused by violations of some synchronization constraint they are remarkably easy to find and eliminate using tools like the *ParaSoft* debugger, *ndb*. Furthermore, the existence of these synchronization constraints tends to help rather than hinder the development process. Much care has gone into the I/O system, for example, to make the synchronization as natural as possible. Typically we find that the message

```
abort (-1)
```

(which is the response of the *cubix* program to a violation of a synchronization rule) is indicative of an error in the user application which might otherwise have gone unnoticed or else caused other problems to occur later on.

4 Decomposition Tools

This section describes the utilities used to automatically distribute problems among parallel processors. The *Express* manuals describe a set of communication primitives designed to allow “topology independent communication”. Problems can be specified in their own natural domain - two dimensions for image processing applications and three for aircraft simulation, for example. The utilities in this section are then provided to assign the “processor numbers” used by the communication routines described in the previous section.

Also available is another utility which allows applications access to certain important runtime parameters. In conjunction with the other utilities this allows programs to be dynamically configured, at runtime, to the system on which they execute. This allows, for example, a program developed on a four processor system to be run on a 1000 node production machine by merely changing a single command line parameter.

5 Multitasking Support

Express supports a powerful *remote* multitasking facility which allows programs running on any processor in the system to initiate a “task” on another node of the parallel computer. This system is built around the KXHAND function which associates a program segment with a particular message type. Upon the arrival of a registered message type the indicated program section is triggered as a separate task which is then free to pursue its own independent execution path.

In support of this multitasking facility is a set of semaphore operations designed to allow two or more processes on a node to cooperatively update shared data.

6 Processor Allocation and Control

This section describes in detail the control functions at the lowest levels of *Express*. They are used in "Host" programs to allocate groups of processors, load programs and start execution. Note that this section will not concern you if you intend to use the *Cubix* programming model since the *cubix* program takes care of the necessary steps automatically.

The unifying concept of this section is that of the *processor group*. This is the fundamental unit of processor allocation - processors are allocated to processor groups which are then treated as a unit. When programs are to be loaded into processor groups the *processor group index* must be specified.

```
PROGRAM FIRST
INTEGER PGIND

C
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Set up Express and its common block.
C
CALL KXINIT
C
C-- Allocate four transputer nodes anywhere in array.
C
PGIND = KXOPEN('/dev/transputer', 4, NOCARE)
IF (PGIND .LT. 0) THEN
    WRITE(6,*) 'Failed to allocate 4 processors'
    STOP
ENDIF
C
C-- Load application, "noddy" into all processors.
C
ISTAT = KXLOAD(PGIND, 'noddy')
IF (ISTAT .LT. 0) THEN
    WRITE(6,*) 'Failed to load program'
    STOP
ENDIF
```

As well as allowing a single host process to allocate and manipulate more than one group of processors it is also possible for two or more users to simultaneously allocate and work with groups of processors. (Provided, of course, that the host operating system allows multitasking. This feature is not, for example, supported under MS-DOS.) It is even possible for multiple host processes to share access to the same group of processors. This mechanism allows multiple, disjoint, front end processes such as a file serving utility and a complex graphical user interface to both have access to the same group of nodes. Routines are available to ensure that the processes do not interfere with each other.

All the routines in section return -1 to indicate errors. Some also write diagnostic messages and some cause immediate termination of the calling process. In any case the parallel machine should remain intact and available for use by other applications and users.

7 I/O Services

The *Cubix* I/O library is available to programs using the *Cubix* programming model and associated server process. In principle, It provides a full set of utilities as well as many extensions designed explicitly for parallel processing. Unfortunately, due to the abysmal lack of any standardized form of FORTRAN I/O the current FORTRAN implementations of *Cubix* fall somewhat short of perfection. For this reason a highly non-standard but useful set of SUBROUTINES and FUNCTIONS have been made available to supplement the normal FORTRAN language statements. Using these (parallel) extensions it is possible to make use of the full power of the *Cubix* system albeit at the expense of losing portability.

The standard FORTRAN I/O interface is not documented here. You should consult whatever documentation is supplied with your compiler for details of its capabilities. The following table summarizes the current status of *Cubix* as regards compatibility with normal FORTRAN I/O

language statements:

I/O mode	Terminal	Formatted	Unformatted
singl, input	OK	OK	OK
multi, input	Not available	Not available	Note 1
async, input	OK	Note 2	Note 3
singl, output	OK	OK	OK
multi, output	OK	Note 4	Note 1
async, output	OK	Note 2	Notes 2,3

NOTES:

1. The simplest interface is provided by the KMREAD, KMWRITE, KMRD2D and KMWT2D functions.
2. Asynchronous mode is most effective when each node access a different file.
3. The simplest interface is currently provided by the KREAD and KWRITE functions.
4. The behavior of formatted file output is somewhat erratic as various compilers interpret FORTRAN files in unusual ways.

The following is the list of extensions to the normal FORTRAN I/O system which are supported in all versions of *Express*, together with the list of restrictions imposed on each

by the parallel nature of the underlying system.

KREAD ¹	KWRITE ¹
KMREAD ²	KMWRT ²
KMRD2D ²	KMWT2D ²

NOTES:

1. These routines must be called “loosely synchronously” and with identical arguments in each node unless the UNIT argument is in `async` mode.
2. Must be called “loosely synchronously” in all processors unless the stream argument is in `async` mode. If the stream argument is in `multi` mode the arguments may differ from node to node but the function must still be called “loosely synchronously”.

8 Graphics

The *Plotix* library is supplied to allow both parallel programs running in the *Cubix* programming model and “host” programs access to device independent graphics in a portable manner. The library contains about twenty routines which are sufficient to cover the majority of graphical tasks while not being an implementation of any particular standard.

9 Standard variables and the */XPRESS/* common block.

Central to the use of *Express* is the labelled common block “*/XPRESS/*” which should be included whenever *Express* functions are being used. This defines a number of important parameters which have wide usage in the system. Before any use of such values can be made, however, they must be initialized by calling the start-up routine, *KXINIT*. This routine should be called at the top of every *Express* FORTRAN program which should, therefore, have a structure similar to the following

```
PROGRAM MYPROG
C
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Start up Express by initializing its COMMON block.
C
CALL KXINIT
C
C-- Can now proceed to use Express .....
C
```

The values in the common block are used for a variety of purposes as explained below.

NO CARE	This value is used to indicate that the source or type of an incoming message are of no consequence. Note that it is illegal to send a message with type
---------	--

	NOCARE even if you really don't care!
NORDER	This value is used to indicate that the parallel I/O system should operate in a mode in which data from nodes is sent to and/or received from the host in order of increasing processor number. It is used in conjunction with the KORDER, KMREAD and KMWRT function calls.
NONODE	This macro value is used by the KXGRID functions to indicate that no node is attached to the user decomposition in the indicated direction. Such a case might arise, for example, when solving partial differential equations on a finite space - some nodes have no neighbors in some directions since they lie on the edge of the domain.
IHOST	An integer variable which contains the "processor number" used by node programs to send/receive data to the host. By default this is the machine that loaded the node program although it is possible to override this.
IALNOD	A value used in the "global" communication routines (KXCOMB, KXBROD, KXCONC, etc.) to indicate that a particular operation should be applied to all nodes in a parallel processing system. Never includes the host processor(s).
IALPRC	A value similar in use to IALNOD except that it implicitly includes all host processors attached to the system.

10 Manual Page Layout

The manual pages are, for better or worse, modeled after those often found in UNIX documentation which means that each manual page has several well-defined sections. The overall structure is

shown below.

Header contains the name of the manual page which is usually the same as the routine described.

KABORT
NAME KABORT - Immediately terminate program
SYNOPSIS SUBROUTINE KABORT (STATUS) INTEGER STATUS
DOMAIN Available to node programs compiled with the <i>Cubix</i> or <i>Plotix</i> libraries only
DESCRIPTION The KABORT mechanism.....
EXAMPLES This function is most useful in..... <pre>CALL KASYNC (6) IF (ERROR .GT. 99) THEN WRITE (6,*) 'Death!! ..' CALL KABORT (15) ENDIF</pre>
BUGS/WARNINGS None.
SEE ALSO " <i>Cubix</i> : Programming parallel computers without programming hosts"

The various sections and their contents are:

- | | |
|--------------------|---|
| NAME | Repeats the name associated with the manual page and a brief one-line description of the purpose of the associated routines |
| SYNOPSIS | Summarizes the arguments used by the indicated routines. If more than one routines is described on a particular page then all are listed in this section |
| DOMAIN | Describes the libraries in which the routine is to be found and any restrictions on when it may be used. |
| DESCRIPTION | Describes the purpose of each routine and lists the actions caused by its most important arguments. This section is the most important reference material |

for each command.

EXAMPLES

One or more examples of the use of each routine are shown on each manual page. This section probably represents the best information on how the various arguments are put together in “real” examples. This section is also useful for demonstrating the order in which function calls should be made and which ones are necessary at which points in the execution of *Express* programs.

WARNINGS

If the routine has peculiar side effects or is “dangerous” in some way it will be noted in this (optional) section. Any non-intuitive behavior is also noted here.

BUGS

Currently known bugs and/or unimplemented routines are noted in this (optional) section.

SEE ALSO

Related commands and/or routines from the various *Express* libraries are noted in this section. Using this information is usually the quickest way to build a complete picture of the interaction between the various utilities.

NAME

KABORT - Immediately abort program.

SYNOPSIS

```
SUBROUTINE KABORT (STATUS)
INTEGER STATUS
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

The KABORT routines causes the immediate termination of the parallel program. A message is displayed on the host processor showing the processor number of the aborting node and also the STATUS value supplied to the call. If examined with the debugger, ncb, the aborting nodes will appear to be at breakpoints.

EXAMPLE

This function is most useful for dealing with asynchronous problems which require termination of the program but which might not occur in every node. If the termination condition is known to occur in each node the regular STOP should probably be preferred. A problem that may merit KABORT is the failure of the KOPENP graphics function for some reason.

```
C
C-- Assume all arrays, parameters defined elsewhere..
C
      ISTAT = KOPENP (GBUFFR, IGSIZE)
      IF (ISTAT .LT. 0) THEN
          CALL KABORT (12)
      ENDIF
```

KASPEC

NAM

KASPEC - Inquire device aspect ratio.

SYNOPSIS

```
SUBROUTINE KASPEC (DEVX, DEY)  
  REAL DEVX, DEY
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

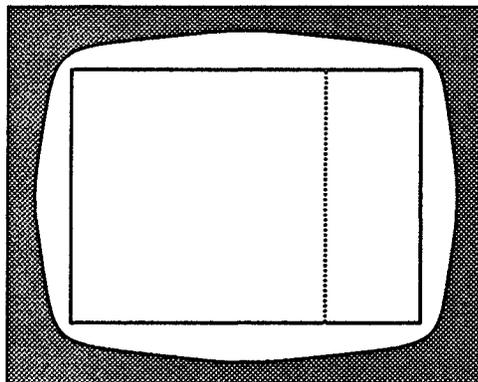
DESCRIPTION

This routine returns the "size" of the display surface. This concept is less than well-defined; it can either mean the number of pixels in each direction or the physical size of the display device. We choose to return the latter values since they seem to be more useful. In particular the sizes returned are the width and height of the display surface, in inches.

EXAMPLE

The following code shows the interaction between the KASPEC and KVPOR calls which allows the user image to always appear in the largest square region of the display surface independent of the actual shape of the device screen. We assume that the user has supplied a routine, DRAWSQ, that draws a square in the user coordinate space. We show the output both on the default viewport and also after using KASPEC to make a square window.

```
      REAL XFAC, YFAC, RATIO  
;  
;-- Draw a square in user coordinate space.  
;  
      CALL DRAWSQ  
;  
      CALL KASPEC (XFAC, YFAC)  
      RATIO = XFAC/YFAC  
      IF (RATIO .GT. 1.0) THEN  
          CALL KVPOR (0., 0., 1./RATIO, 1.)  
      ELSE  
          CALL KVPOR (0., 0., 1., 1./RATIO)  
      ENDIF  
      CALL KLINEM(1)  
;  
;-- Redraw square - should now look square!!  
;  
      CALL DRAWSQ
```



SEE ALSO

KSPACE, KVPOR

KBOX

NAME

KBOX - Draw and fill rectangles.

SYNOPSIS

```
SUBROUTINE KBOX(X0, Y0, X1, Y1, COLOR, EDGE)
REAL X0, Y0, X1, Y1
INTEGER COLOR, EDGE
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

This routine provides a simple interface to the polygon drawing primitives for the common case of rectangular regions. A rectangle will be drawn whose bottom left corner has position (X0, Y0) and whose top right corner is at (X1, Y1). The COLOR argument indicates the manner in which the region should be filled. Positive values of COLOR translate into solid colors in the same manner as the arguments to the line color primitive, KCOLOR. Negative values yield device dependent shading patterns. If the EDGE argument is non-zero then the boundary of the region will be drawn in the color most recently specified in a call to the KCOLOR function.

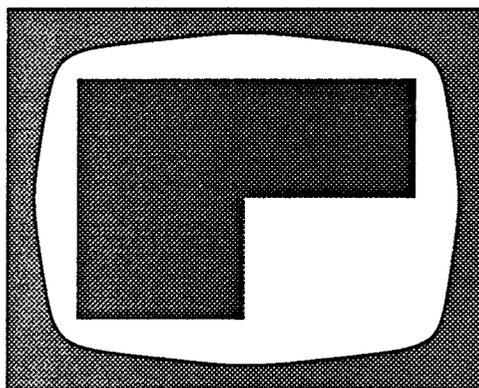
All coordinates are expressed relative to the most recent call to KSPACE.

Note that filling with COLOR = 0 and EDGE = 0 results in a "selective erase" - specific areas of the screen can be erased.

EXAMPLE

The following code draws a simple box in the foreground color and then takes a "bite" out of it by drawing and filling in the background color.

```
C
C-- Define coordinate system.
C
      CALL KSPACE(0.,0.,10., 10.)
C
      CALL KBOX(1.0, 1.0, 9.0, 9.0, 1, 1)
      CALL KBOX(5.0, 1.0, 9.0, 5.0, 0, 0)
C
      CALL KSENDP
```



Note that this code achieves the same effect as that shown on the KPANEL manual page but is much simpler. Also note that filling rectangular regions can often be done by hardware even in cases where no general hardware polygon fill is available. In these situations KBOX will be significantly faster than the equivalent calls to either KPOLGN or the KPANEL routines.

SEE ALSO

KPANEL, KCOLOR

NAME

KCALHO - Interface to user host routines from *Cubix* program

SYNOPSIS

```
INTEGER FUNCTION KCALHO (FUNC, ODAT, OCNT,  
                        IDAT, ICNT, SENT, STAT)
```

```
INTEGER FUNC, ODAT (*), OCNT, IDAT (*), ICNT, SENT, STAT
```

```
INTEGER FUNCTION (FUNC, ODAT, OCNT)
```

```
INTEGER FUNC, ODAT (*), OCNT
```

```
INTEGER FUNCTION KRETHO (IDAT, ICNT, SENT, STAT)
```

```
INTEGER IDATA (*), ICNT, SENT, STAT
```

DOMAIN

These routines may only be called in programs compiled with the *Cubix* or *Plotix* libraries. Furthermore special software is necessary to re-link the part of the *Cubix* server which executes on the system host.

DESCRIPTION

These routines provide an interface between normal *Cubix* programs and user written routines which must run on the host computer. The simplest routine, *KCALHO*, causes a user written routine (denoted by the integer *FUNC*, explained later) to begin execution. This routine can be passed up to 512 bytes of data from the buffer *ODAT* which it will receive as an argument. The exact number of bytes to be sent to the host routine is specified in the *OCNT* argument.

After processing the host routine is also allowed to send up to 512 bytes of information back to the node program which invoked it. This data will be stored at address *IDAT* in the node program. The *ICNT* argument indicates the maximum number of bytes which should be copied to the node's *IDAT* buffer. If more are sent from the host they are ignored. In any case the *SENT* variable argument will be initialized to the number of bytes which the host attempted to transfer, whether larger or smaller than *ICNT*.

Finally the variable *STAT* argument will be set to the value returned by the host routine.

The 512 byte restriction is imposed by the fact that for reasons of speed the data transferred to/from the host routine is not copied to a "safe" user buffer but remains in the system buffer to which it was sent. If this poses too strict a restriction on the abilities of the node program then an alternative interface is provided by the *KSTRHO* and *KRETHO* functions.

KSTRHO is responsible for starting up the host routine and passing it up to 512 bytes of data in the same manner as indicated by the first three arguments to *KCALHO*. *KRETHO* performs the operation of the last four arguments to *KCALHO* which it interprets in an identical fashion. The advantage of this interface is that between the calls to *KSTRHO* and *KRETHO* in the node program the host and node codes are running in a mode identical to the "host-node" programming model and can, therefore, communicate data at will using the regular

Express systems calls (KXREAD, KXWRIT, KXBROD, etc.)

Note, however, that this is a rather “double-edged” advantage. On the one hand it allows the host routine and the node program to communicate data at will avoiding the 512 byte restriction for data transactions in both directions. On the other hand the advantages of the *Cubix* programming model are suspended until the host routine terminates. As a result the node program *cannot* use any *Cubix* I/O or *Plotix* graphical commands until the call to KRETHO completes. Similarly the user will have to resolve potential byte swapping and alignment problems due to incompatible host and node CPU’s which might otherwise have been covered up by the *Cubix* programming model. For these reasons, therefore, the interface through the KCALHO routine is to be preferred.

EXAMPLE

The following code segment is used to execute the host routine with index 3 in the host. A simple integer is sent to this routine and a small array of integers is returned.

```

        INTEGER TOT, I, STAT, HSTAT, NGOT
        INTEGER*4 INDAT(32)
        INTEGER*4 MYVAL
C
C-- We might have to swap the stuff sent to the host....
C
        CALL KXSWAW(MYVAL, MYVAL, 4)
        ISTAT = KCALHO(3, MYVAL, 4, INDAT, 4*32, NGOT, HSTAT)
C
C-- If either the host or nodes reported an error give up now.
C
        IF(ISTAT .LT. 0 .OR. HSTAT .LT. 0) THEN
            WRITE(6,*) 'Something failed', ISTAT, HSTAT
            RETURN
        ENDIF
C
C-- Figure out how many bytes we actually got. This is either
C-- the number we asked for or the number sent, whichever is
C-- smaller.
C
        IF(NGOT .GT. 4*32) NGOT = 4*32
C
C-- If everything seemed to be OK we can add up the numbers
C-- returned by the host. Note that we might have to swap some
C-- bytes here.....
C--
        CALL KXSWAW(INDAT, INDAT, NGOT)
        TOT = 0
        DO 10 I=1,NGOT/4
            TOT = TOT + NGOT(I)

```

It is important to note that the data buffer being transmitted to the host is sufficiently small to fit into the 512 byte restriction. This allows us to use the KCALHO interface. Further we take care to calculate correctly the amount of data returned to the node program and to (potentially) swap bytes. Note that we have inserted calls to a byte swapping function KXSWAW in this example - if the host and node byte orderings are the same these calls are unnecessary and can be removed.

HOST INTERFACE

The previous sections described the interface to the system from the perspective of the node program. As well as incorporating one or more of the node system calls in your program you must also arrange for the host program to be linked with your host routines.

This procedure is made most tricky for a FORTRAN programmer by the fact that, since the cubix file server program is written in C its interface to any user written host routines must also be in C. This restriction, however, only refers to the top level routine called by the cubix server. If supported by the compilers used you may write the majority of your host routine in FORTRAN and somehow link it to the rest of the server process. This procedure is beyond the scope of this document and you should consult compiler documentation for details.

The linking of host routines into the cubix server process is controlled by the source file `userlink.c` supplied with the "Cubix user-link kit". This file contains an array of function pointers the indices of which correspond to the `func` argument passed to `callhost` or `starthost`. By default the top part of this file contains the following

```
static int user_no_op();

int (*user_funcs[])() = {
    user_no_op,
};
```

This code shows that a single host function is defined by default: `user_no_op`. This function doesn't actually do anything and is merely provided as a place holder to simplify the introduction of new user routines. If, for example, two additional user functions are required called, `search_DB` and `sort_DB` for example, we could modify the above part of the `userlink.c` file to read

```
extern int search_DB(), sort_DB();

int (*user_funcs[])() = {
    searchDB,
    sort_DB,
```

};

Notice that we elected to delete the `user_no_op` function and made the two new routines take indices 0 and 1. Also note that we changed the definition from `static` to `extern` since these routines are probably defined outside the `userlink.c` file.

Having initialized the data structures used by *Cubix* to find user host routines it remains only to discuss the calling sequence used when invoking them.

When a user host routine is called it is passed three arguments and is expected to return an integer value. The three arguments passed to the user routine are:

- A pointer to the buffer containing data sent to the host as the `odat` argument to `KSTRHO` or `KCALHO`.
- The number of bytes contained in this buffer. This value will be the same as that specified as `OCNT` when calling the host routine from the nodes.
- A pointer to an integer which should be set to the number of bytes to be returned to the node program as `IDAT`. This data should be placed in the buffer pointed to by the first argument, overwriting whatever values were sent there from the nodes. The value written to this argument will be returned to the node program through the `SENT` argument of `KCALHO` or `KRETHO`.

As an example, therefore, the skeleton of the `search_DB` function should be similar to

```
int search_DB(buffer, in_bytes, pout_bytes)
char *buffer;
int in_bytes;
int *pout_bytes;
{
    .....

    *pout_bytes = ...;
    return ...;
}
```

Notice that we finish the function by making sure that the `pout_bytes` argument is initialized. Finally we return a value which will be passed to the node program through the `STAT` argument.

DIAGNOSTICS

The node routines described here indicate error conditions by returning -1. Possible error conditions are as follows:

- ETOOBIG** An attempt was made to either send too much data to the host or return too much to the nodes. The maximum amount of data that can be transmitted through the system invocation mechanism is 512 bytes.

KCALHO

EBADPTR The FUNC argument indicated a function with an index outside those defined in the host's function table.

It is important to note that if an error occurs in a call to KSTRHO no call to KRETHO should be made.

SEE ALSO

KCBXSY.

NAME

KCBXSY - Specify synchronous or asynchronous system calls

SYNOPSIS

```
INTEGER FUNCTION KBCXSY (FLAG)
INTEGER FLAG
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

This call provides a system override which controls the overall synchronous or asynchronous behavior of *Cubix* system calls.

By default the system is in "synchronous mode" which means that all function calls must be made loosely synchronously. Furthermore each node must address its system requests to the same system console. (Note, however, that asynchronous I/O is still supported in this mode on a file by file basis).

Calling KCBXSY with a zero argument places the system in asynchronous mode. All further operating system requests are made on a node by node, first come-first served, basis. In this mode any node may address any host processor with impunity but the responsibility for maintaining "sensible" ordering lies with the user. Note that I/O requests will, however, still occur with the synchronization model implied by the unit's "mode" - i.e., a file which has been placed in *multi* mode with the KMULTI system call will continue to operate this way even after the KCBXSY system call has placed the system in its asynchronous state.

The value returned is the previous "synchronization mode" flag which can be used as the argument to subsequent calls to KCBXSY.

EXAMPLE

The asynchronous mode is rather difficult to control not the least because the inherent asynchronicity introduced into applications make them harder to debug. It can, however, be useful in system with multiple consoles, each under the control of a different group of nodes. In the following example we suppose that nodes with even processor numbers should communicate with host number IHOST+1 while those with odd processor numbers remain connected to the main system console.

```
PROGRAM ASYNC
INTEGER NDDATA(4)
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Setup Express and its common block.
C
CALL KXINIT
C
```

KCBXSY

```
C-- Get processor numbers.
C
    CALL KXPARA(NDDATA)
C
    CALL KCBXSY(0)
    IF (MOD(NDDATA(1),2) .EQ. 0) THEN
        CALL KCONND(IHOST+1)
    ENDIF
C
C-- The lowest numbered node now introduce themselves,
C-- asynchronously, on their respective consoles.
C
    IF (NDDATA(1) .LT. 2) THEN
        WRITE(6,*) 'Hello world'
    ENDIF
```

SEE ALSO

KCONND

NAME

KCLIP - Enable/Disable clipping.

SYNOPSIS

```
SUBROUTINE KSETCL(X0, Y0, X1, Y1)
REAL X0, Y0, X1, Y1

SUBROUTINE KENDCL
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These two calls are used to enable and disable the clipping primitives. The KSETCL routine defines a two-dimensional clipping window relative to either the default user coordinate range or that set by the most recent call to KSPACE. Further lines, points, markers and polygons will be "clipped" relative to this window and portions lying outside the indicated range will be removed.

The KENDCL routine disables the clipper.

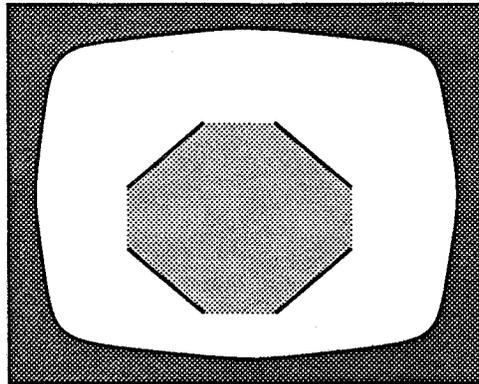
It is important to note that clipping is performed with respect to each "vport" and that the clipping window is specific to the active viewport when KSETCL is invoked. Each call to KSETVTP alters the clipping window to that associated with the particular "vport" selected.

Note that while clipping is typically expensive this process is supported on the nodes of a distributed machine rather than on the graphics device itself. As a result all clipping is performed in parallel leading to increased performance.

EXAMPLE

In the absence of the call to KSETCL the following code would draw a diamond shaped polygon on the display surface. After clipping only a portion of the figure remains.

```
CALL KSPACE(0.,0.,4., 4.)
CALL KSETCL(1., 1., 3., 3.)
C
CALL KINITP(1, 1)
CALL KPANLP(0., 2.)
CALL KPANLP(2., 4.) CALL KPANLP(4., 2.)
CALL KPANLP(2., 0.)
CALL KENDPA
C
CALL KSENDP
```



SEE ALSO

KPANEL, KMOVE, KMARKE, KCONT

NAME

KCNTOR - Contouring functions

SYNOPSIS

```
SUBROUTINE KCNTOR (FUNC, GX, GY, LEVMIN, LEVMAX, NLEV, PANELS)
REAL FUNC, LEVMIN, LEVMAX
INTEGER GX, GY, NLEV, PANELS
EXTERNAL FUNC
```

```
SUBROUTINE KINITL (FUNC, GX, GY, LEVEL, PANELS)
REAL FUNC, LEVEL
INTEGER GX, GY, PANELS
EXTERNAL FUNC
```

```
INTEGER FUNCTION KGETPT (PX, PY)
REAL PX, PY
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

The KCNTOR routine makes a two-dimensional contour map of the supplied function FUNC. Rather than provide an array of values specifying the function to be contoured a function is supplied which will be called repeatedly with pairs of integer arguments representing the position at which a value is required. The range of values is specified by the GX and GY arguments; the user function will be called as FUNC (I, J) with I and J in the ranges $0 \leq I < GX$ and $0 \leq J < GY$.

Contours are drawn at NLEV positions equally spaced between LEVMIN and LEVMAX. Optionally LEVMIN and LEVMAX can both be set to zero in which case suitable values are selected internally.

The final argument, PANELS, selects the type of contouring to be performed. If non-zero then the contours will be drawn as filled polygons while a zero value selects the more conventional style in which the contours are represented by lines. In the case of filled regions the lowest NLEV indices of the color map will be used. The KRAINB and KGREYS functions can be used to re-map the appropriate color table entries.

Since this routine calls KUSEND internally it must be called in all nodes together. Failure to do so will result in communication deadlock. Note, however, that no internode communication is done in performing the contouring. It is the responsibility of the user to distribute boundary values to processors that require them before calling the KCNTOR routine.

The contouring utility described here assumes that the data lie in a rectangular domain - i.e., that the mesh underlying the data is a Cartesian grid. In order to contour data specified in other coordinate systems, such as polar coordinates, the lower level KINITL and KGETPT

routines are available.

The former specifies a function to be contoured and a range of I and J values just as in KCNTOR. The LEVEL argument selects the contouring level and the PANELS argument indicates the style of contouring to be performed. The interpretation of this value is not so straightforward as in the KCNTOR routine. Essentially the purpose is to control exactly what type of points are returned by the KGETPT function. The allowed values and their interpretations are

- PANELS = 0 All interior points are returned. The surrounding box is treated as a true rectangle and only the vertices are returned. This option is designed for simple line contouring of rectangular regions.
- PANELS = 1 The contour map is cut into horizontal strips and coordinates are returned in such a way that the resulting polygonal regions are simply connected. The bounding box is treated as in option 0. Designed for color fill panels.
- PANELS = 2 The interior points are treated as in option 1 but the boundary is also divided into many points which are returned individually. This option is designed for cases where the actual domain to be contoured is not rectangular and hence the boundary values need to be transformed in some manner.

The KGETPT function is used, once a contour has been initialized, to return coordinates which lie on the contour. As well as returning an (X,Y) coordinate pair under the supplied pointers the returned value indicates the nature of the returned point as follows

- STATUS = 0 This contour level is finished. Ignore returned coordinates.
- STATUS = 1 The coordinates are valid for the current contour.
- STATUS = 2 A segment of the current contour line is finished. Ignore the coordinates returned and call KGETPT again in which case it will either return 0 indicating that no more points exist at this contour level or 1 indicating that another disjoint piece of the current contour exists.

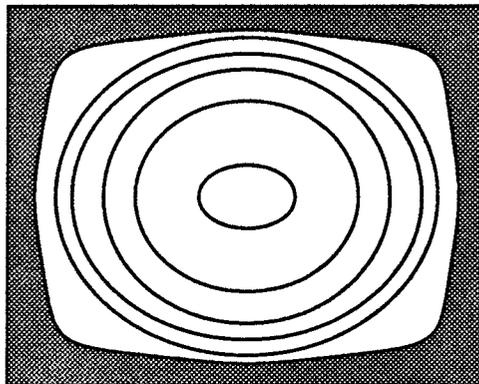
A complete example of the use of these functions to contour a function supplied in polar coordinates is shown in the *Plotix* documentation.

EXAMPLE

The following code demonstrates the elementary use of the contouring function.

```
SUBROUTINE DOCONT
REAL CIRCLS
EXTERNAL CIRCLS
C
CALL KCNTOR(CIRCLS, 10, 10, 0.0, 25.0, 6, 0)
RETURN
```

```
      END
C
C-- This is the function that will be called by the
C-- contouring utility.
C
      REAL FUNCTION CIRCLS(I,J)
      INTEGER I,J
C
      REAL X0, Y0
C
      X0 = FLOAT(I-5)
      Y0 = FLOAT(J-5)
      CIRCLS = X0*X0 + Y0*Y0
      RETURN
      END
```

**SEE ALSO**

KCOLOR, KGREYS, KRAINB.

NAME

KCOLOR - Change color attribute of graphical objects.

SYNOPSIS

```
SUBROUTINE KCOLOR (INDEX)
INTEGER INDEX
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

This routine modifies the color used in drawing all subsequent lines and markers. The **INDEX** argument is typically device-dependent but one can safely use the following on "color" devices.

0	Background color for device. ("Black")
1	Foreground color for device. ("White")
2	Red.
3	Green.
4	Blue.
5	Cyan.
6	Purple.
7	Yellow.

Monochrome devices, by default, support only two of these indices, 0 and 1.

The background color is often useful for selectively erasing previous symbols.

This function interacts with the **KGREYS** and **KRAINB** routines providing full color on devices capable of supporting such models.

EXAMPLE

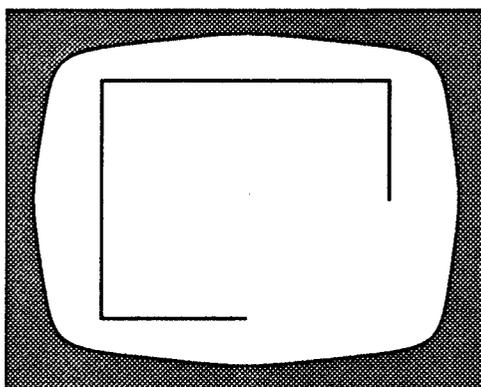
The following code defines an 8 x 8 coordinate system and draws a simple box in the foreground color. It then overwrites the lower right hand corner of the box in the background color, erasing part of the image.

```
CALL KSPACE (0., 0., 8., 8.)
C
CALL KCOLOR (1)
CALL KMOVE (1., 1.)
CALL KCONT (7., 1.)
CALL KCONT (7., 7.)
CALL KCONT (1., 7.)
CALL KCONT (1., 1.)
C
CALL KCOLOR (0)
CALL KMOVE (4., 1.)
```

```
CALL KCONT(7.,1.)  
CALL KCONT(7.,4.)
```

C

```
CALL KSENDP
```



SEE ALSO

KCONT, KMOVE, KLINEM, KRAINB, KGREYS

NAME

KCONND - Redirect system calls.

SYNOPSIS

```
SUBROUTINE KCONND (NODE)
INTEGER NODE
```

DOMAIN

Available to node programs using the *Cubix* file server only and which are compiled with either the *Cubix* or *Plotix* libraries

DESCRIPTION

This function is provided to support systems with more than one attached host. By default all *Cubix* system calls are directed to the processor which originally loaded and executed the user application. On occasion, however, it may be necessary to perform certain system tasks on other nodes in the system.

The `console_node` function has as its argument a processor number. All further (non-I/O) operating system requests will be directed to this node. To obtain suitable node numbers for use in this call we take the host identifier from the configuration utility, `cnftool`, and OR in the highest bit. If `cnftool` designated a particular host as "H1" then the appropriate node number to use is `0x8001`.

It is important to note that all file related I/O is *always* directed to the node which contained the file when it was opened, independent of the status of the KCONND function.

EXAMPLE

Let us assume that three "host" processors are attached to our system. The first is the original system console which can be addressed through the value `IHOST` from the `XPRESS` common block. The others have the identifiers `H1` and `H2` as defined in the system configuration utility, `cnftool`. The following code executes a rather simple operation; opening two files, the first on the system default "HOST" and the second on host number 1.

```
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Setup Express and its common block.
C
CALL KXINIT
C
CALL KCONND(IHOST)
OPEN(UNIT=1, FILE='host1.dat', STATUS='unknown')
C
C-- This number is 8001 in hexadecimal - i.e., Host 1
C
CALL KCONND(32769)
OPEN(UNIT=2, FILE='host2.dat', STATUS='OLD')
```

CALL KCONND (IHOST)

Note that different nodes are allowed to maintain distinct consoles with these calls although one must then use asynchronous requests to avoid deadlock.

SEE ALSO

KCBXSY.

KCONT

NAME

KCONT - Move and draw a line.

SYNOPSIS

```
SUBROUTINE KCONT (X, Y)
REAL X, Y
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

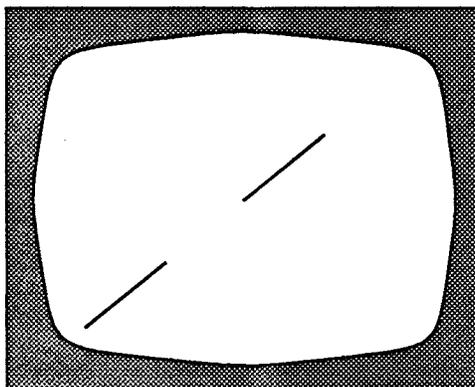
DESCRIPTION

Moves the current plotting position to (X,Y) and draws a line in the current color from the previous plotting position. X and Y are specified relative to the coordinate system defined by the most recent call to KSPACE.

EXAMPLE

The following code draws a broken diagonal line across the display surface.

```
CALL KSPACE (0.,0.,4., 4.)
C
CALL KMOVE (0.,0.)
CALL KCONT (1.,1.)
CALL KMOVE (2.,2.)
CALL KCONT (3.,3.)
C
CALL KSENDP
```



SEE ALSO

KMOVE KCOLOR, KLINEM

NAME

KCPCP, KCPELT - Dump communication profile data.

SYNOPSIS

```
SUBROUTINE KCPCP
```

```
  SUBROUTINE KCPELT (FNAME)  
  CHARACTER*80 FNAME
```

DOMAIN

KCPCP may only be called in the host processor and KCPELT may only be called in the nodes.

DESCRIPTION

These routines are used to dump the communication profiling data collected with the KCPCP functions. For each call to KCPELT on the nodes there must be a call to KCPCP in the host processor. The profiling data will be written to a file on the host with the name FNAME supplied in the node program.

Each call to KCPELT turns off the communication profiler and resets its internal counters so that further profiling starts from the zero state. This allows distinct communication profiles to be obtained for different regions of an application.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the communication profile.

1. Host Program

```
      PROGRAM HSTPRE  
      C  
      C-- Set up Express.  
      C  
      CALL KXINIT  
      C  
      C-- Allocate nodes, load programs.  
      C  
      ...  
      C  
      C-- Execute algorithm and then dump data to "phase1.cprof".  
      C  
      ...  
      C  
      CALL KCPCP  
      STOP
```

END

2. Node Program

```
PROGRAM NODPRF
C
C-- Set up Express.
C
    CALL KXINIT
C
C-- Start off profiler.
C
    CALL KCPON
C
C-- Application code .....
C
    ...
C
C-- Application complete, dump data with KCPCP/ELT.
C
    CALL KCPELT('phase1.cprof')
C
    STOP
END
```

Notice that we can selectively profile pieces of code. In this mode it makes sense to dump out the profile data independently to separate files for simplicity in later analysis.

SEE ALSO

ctool (command), KCPROF, KCPEND

NAME

KCPINQ, KCPEND - Manipulate communication profiler. under *Cubix*

SYNOPSIS

INTEGER FUNCTION KCPINQ()

SUBROUTINE KCPEND

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

These routines provide a simple control interface to the communication profiler for applications running in the *Cubix* environment.

KCPINQ returns an integer value representing the state of the “-m” runtime switch on the *cubix* command line. This can be used to conveniently enable/disable the profiling system at runtime. Consider a typical command

```
cubix -n 4 toyland 1024 1024 <noddy.dat
```

Since no “-m” switch is present a call to KCPINQ will return zero. If we modify the above command to

```
cubix -mce -n 4 toyland 1024 1024 <noddy.dat
```

then the return value would be 1 since the character ‘c’ appears in the monitoring switch, “-m”.

KCPEND is used to dump profiling data to the host computer file system. A file called “cprof.out” is created for later analysis with the *ctool* utility. In addition the profiler is disabled and its internal state reset to zero so that further profiling leads to distinct, non-overlapping data.

The operating system itself performs a check for the communication monitoring switch in the *cubix* command and, if present, turns on the profiler with a call to KCPON. It also arranges to call KCPEND at program termination with STOP. As a result a typical *Cubix* application need contain no explicit calls to the communication profiling routines - they are all made by the kernel. The only case in which such calls are needed is when more careful control is required over the profiler and the data it dumps.

EXAMPLE

The following code is a skeleton of that which might be used to control the communication profiler in a *Cubix* application.

```
PROGRAM CBXPRF
C
C-- Initialize Express.
```

```
C
    CALL KXINIT
C
C-- Start off profiler. This code is not strictly
C-- necessary since it is equivalent to the check
C-- made automatically by Express.
C
    ISTAT = KCPINQ()
    IF(ISTAT .NE. 0) THEN
        CALL KCPON
    ENDIF
C
    ...
C
C-- Program over, dump data again and exit. Again
C-- this code is superfluous since it duplicates the
C-- action of Express.
C
    IF(ISTAT .NE. 0) THEN
        CALL KCPEND
    ENDIF
    STOP
    END
```

Notice that we can selectively profile pieces of code. In this mode it makes sense to dump out the profile data independently to separate files for simplicity in later analysis. The calls to rename in the above are necessary to prevent the second call to KCPINQ from overwriting the file created by the first call.

SEE ALSO

ctool, KCPCP, KCPROF

NAME

KCPON, KCPOFF - Control communication profiler.

SYNOPSIS

SUBROUTINE KCPON

SUBROUTINE KCPOFF

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

KCPON is used to enable and start the communication profiler. After this call all subsequent calls to the communication system result in entries being made in an internal log-file. KCPOFF reverses this process - until a subsequent call to KCPON no communication profiling will be performed.

For applications which have user programs running in the host computer the profiler is initially off and must be explicitly enabled with a call to KCPON. For applications running in the *Cubix* environment the initial state of the profiler is controlled by a runtime switch in the *cubix* command. (See KCPEND).

The log of profiling information is written to the host file system with KCPCP or KCPEND.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the communication profiler.

```
PROGRAM PRFTST
C
C-- Start off profiler.
      CALL KCPON
C
C-- Application Phase 1., profiler running.
C
      ...
C
C-- Phase 1 complete, dump data with KCPCP/ELT or KCPEND.
C
      ...
C
C-- Application Phase 2., profiler turned off by
C-- previous call to KCPCP/ELT or KCPEND.
C
      ...
C
```

KCPROF

```
C-- Application phase 3., turn on profiler again.  
C  
    CALL KCPON  
    ...  
C  
C-- Program over, dump data again and exit. The STOP statement  
C-- will take care of dumping data to the host automatically.  
C  
    ...  
    STOP  
    END
```

Notice that we can selectively profile pieces of code. In this mode it makes sense to dump out the profile data independently to separate files for simplicity in later analysis.

SEE ALSO

ctool (command), KCPCP, KCPEND

NAME

KDISND - Specify alternate display surface and server.

SYNOPSIS

```
SUBROUTINE KDISND (NODE)
INTEGER NODE
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

This routine enables applications to select between output devices. By default all graphical output is sent to the console file server to be distributed to display devices. This routine allows alternate destinations to be specified, possibly connected to multiple display devices.

KDISND redirects all future graphical commands to the specified node which is assumed to be running a suitable server process.

Any of the *Plotix* functions may be assumed to be present on any graphics node attached to the system. If, however, multiple displays are in use then only asynchronous mode routines should be used since sending data to multiple host servers cannot satisfy any of the synchronization models. This can be achieved with any of the routines KAOPEN, KASEND, KAGIN, etc. or a suitable call to the *Cubix* function KCBXSY.

EXAMPLE

The following code divides the processors into two equal sized groups which direct their graphical I/O to different servers. The lowest numbered node in each group then initializes the device with a call to KAOPEN.

```

      INTEGER NDDATA(4)
      INTEGER GBUFFER(2048)
C
C-- Get parameters of system at runtime. Redirect output
C-- to servers attached to hosts 0 and 1. Note that these hosts
C-- have processor numbers 8000 and 8001 in hexadecimal which
C-- correspond to these decimal values.
C
      CALL KXPARA(NDDATA)
      IF(NDDATA(1) .LT. NDDATA(2)/2) THEN
          CALL KDISND(32768)
      ELSE
          CALL KDISND(32769)
      ENDIF
C
C-- Now setup the hardware by calling the asynchronous
```

KDISND

C-- "open".

C

CALL KAOPEN(GBUFFR, 8192)

SEE ALSO

KCONND.

NAME

KDOTEX - Draw text with complex alignment.

SYNOPSIS

```
SUBROUTINE KDOTEX(TEXT, X, Y, ANGLE, HJUST, VJUST)
CHARACTER*80 TEXT
REAL X, Y
INTEGER ANGLE, HJUST, VJUST
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

This routine draws the characters contained in the TEXT string at the position (X,Y). The text is rotated through ANGLE degrees. The two "justification" parameters are used to position the string with respect to the indicated coordinates as follows

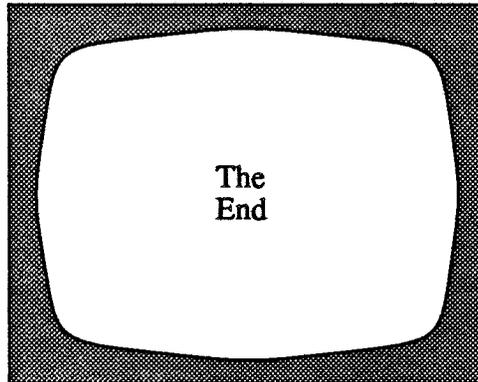
- HJUST = -1 Text is positioned entirely to the right of (X,Y).
- HJUST = 0 The text is centered about (X,Y).
- HJUST = 1 The text is placed entirely to the left of the indicated point.
- VJUST = -1 The text lies totally above (X,Y).
- VJUST = 0 Text is centered vertically on (X,Y).
- VJUST = 1 Text lies below (X,Y).

Using various combinations of these parameters is it possible to align text in fairly arbitrary manners. Using the particular combination HJUST = VJUST = 0 allows one to draw "markers" from the ASCII character set.

EXAMPLE

The following calls are used to position the phrase "The End" around a particular point.

```
CALL KDOTEX('The', X, Y, 0, 0, -1)
CALL KDOTEX('End', X, Y, 0, 0, 1)
```



WARNING

The current plotting position is undefined after this call. In order to perform reliable graphical operations KMOVE should be used before any further drawing is performed.

SEE ALSO

KMARKE, KLABEL

NAME

KEPCP, KEPELT - Dump event log data.

SYNOPSIS

SUBROUTINE KEPCP

SUBROUTINE KEPELT (FNAME)

CHARACTER*80 FNAME

DOMAIN

KEPCP may only be called in the host processor while KEPELT may only be called in the nodes.

DESCRIPTION

These routines are used to dump the event profiling data collected with the KEPROF functions. For each call to KEPELT on the nodes there must be a call to KEPCP in the host processor. The profiling data will be written to a file on the host with the name FNAME supplied in the node program.

Each call to KEPELT turns off the profiler and resets its state so that future profiling commands begin with the system in its initial state.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the event profiler.

1. Host Program

```
PROGRAM HSTEPR
C
C-- Start up Express.
C
CALL KXINIT
C
C-- Allocate node, load programs, etc..
C
...
C
C-- Dump out profile data to "eprof.out".
C
CALL KEPCP
STOP
END
```

2. Node Program

```
PROGRAM NDEPR
C
  INTEGER LOGBUF(2048), LABBUF(256)
C
  REAL ENERGY, RESID, GRIND, CRUNCH
  INTEGER ITER, I
C
C-- Initialize profiler, make labels for indices 1-3.
C-- Start running.
C
  CALL KEPINI(LABBUF, 1024, LOGBUF, 8192)
  CALL KEPLAB(1, 'Outer loop', 'Iteration %d')
  CALL KEPLAB(2, 'After crunch', 'Energy = %d')
  CALL KEPLAB(3, 'Inner loop', 'resid = %d')
  CALL KEPON
C
C-- Compute, compute, compute .....
C
  ....
C
C-- Program over, dump profile data and exit.
C
  CALL KEPELT('eprof.out')
  STOP
  END
```

The strange looking strings passed as the third argument to the KEPLAB functions are actually going to be passed to the C string formatting routine `sprintf`. All that is really important for this application is that the characters “%d” will be replaced by the decimal value supplied as the last argument to `KEPADD`.

Note that these functions may be called repeatedly - the only constraint is that each call to `KEPELT` in the nodes must have a corresponding call to `KEPCP` in the host, and each call to `KEPELT` in the nodes must be made “loosely synchronously”.

SEE ALSO

`etool` (command), `KEPCP`, `KEPEND`

NAME

KEPINQ, KEPEND - Manipulate Event profile under *Cubix*

SYNOPSIS

INTEGER FUNCTION KEPINQ ()

SUBROUTINE KEPEND

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

These routines provide a simple control interface to the event profiler for applications running in the *Cubix* environment.

KEPINQ returns an integer value representing the state of the “-me” runtime switch on the *cubix* command line. This can be used to conveniently enable/disable the profiling system at runtime. Consider a typical command

```
cubix -n 4 toyland 1024 1024 <noddy.dat
```

Since no “-m” switch is present a call to KEPINQ will return zero. If we modify the above command to

```
cubix -mce -n 4 toyland 1024 1024 <noddy.dat
```

then the return value would be 1 since the character ‘e’ appears in the monitoring switch, “-m”.

KEPEND is used to finally dump profiling data to the host computer file system. A file called “*eprof.out*” is created for later analysis with the *etool* utility. In addition the profiler is turned off and the internal state reset to its initial, zeroed, condition.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the event profiler.

```
PROGRAM EPRTST
C
  INTEGER LOGBUF(2048), LABBUF(256)
C
  REAL ENERGY, RESID, GRIND, CRUNCH
  INTEGER ITER, I
C
C-- Start up Express.
C
  CALL KXINIT
C
```

KEPINQ

```
C-- Setup profiler and make labels for indices. If
C-- asked to do so at runtime start the thing up.
C
    CALL KEPINI(LABBUF, 1024, LOGBUF, 8192)
    CALL KEPLAB(1, 'Outer loop', 'Iteration %d')
    CALL KEPLAB(2, 'After crunch', 'Energy = %d')
    CALL KEPLAB(3, 'Inner loop', 'resid = %d')
    ISTAT = KEPINQ()
    IF(ISTAT .NE. 0) THEN
        CALL KEPON
    ENDIF
C
C--Start application code, then go into main loop.
C
    ...
    DO 10 ITER=1,100
        CALL KEPADD(1, ITER)
C
        ENERGY = CRUNCH(ITER)
        CALL KEPADD(2, INT(ENERGY))
C
        DO 20 I=1,4
            RESID = GRIND(ENERGY)
            CALL KEPADD(3, RESID)
        20 CONTINUE
    10 CONTINUE
C
C-- Program over; dump data to host for later analysis.
C
    CALL KEPEND
    STOP
    END
```

The strange looking strings passed as the third argument to the KEPLAB functions are actually going to be passed to the C string formatting routine `sprintf`. All that is really important for this application is that the characters “%d” will be replaced by the decimal value supplied as the last argument to KEPADD.

Notice that the KEPADD and KEPLAB calls are completely safe even if KEPINQ returns 0 and the profiler is not enabled.

SEE ALSO

etool (command), KEPCP, KEPROF

NAME

KEPON, KEPOFF, KEPINI, KEPLAB, KEPADD - Event driven profiler.

SYNOPSIS

SUBROUTINE KEPON

SUBROUTINE KEPOFF

SUBROUTINE KEPINI(LABBUF, LABSIZ, LOGBUF, LOGSIZ)
INTEGER LABBUF(*), LABSIZ, LOGBUF(*), LOGSIZ

KEPLAB(INDEX, TITLE, FORMAT)
INTEGER INDEX
CHARACTER*80 TITLE, FORMAT

KEPADD(INDEX, DATUM)
INTEGER INDEX, DATUM

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

These routines make up the interface to the user specified event driven profiling facility. KEPON and KEPOFF enable and disable the system respectively. While disabled no events are logged even if calls are made to KEPADD.

The routine KEPINI must be called before any of the other profiling calls. The arguments indicate two buffers to be used for "title" and "event" entries which must be supplied by the calling program. Each "entry" corresponds to a single call to the KEPLAB and KEPADD subroutines. As a guide to the amount of space which should be provided the current overheads for log entries and labels are 12 and 68 bytes respectively. Note that the LOGSIZ and LABSIZ arguments should be given in bytes.

KEPADD is the heart of the event system. It makes a new entry in the log file. Three items are logged; the event "index" and "datum" as given in the function call and the time at which the call is made. The INDEX argument is used to differentiate between events at the highest level. This index corresponds to an optional TITLE string in a call to KEPLAB. The DATUM argument is used to identify events at the lowest level. This will be used in conjunction with the FORMAT argument supplied to a call to KEPLAB.

The function KEPLAB is used to facilitate event recognition when the log-file is subsequently analyzed. Its use is optional. If no calls to KEPLAB are made then events will be identified by their "INDEX" argument in the subsequent analysis and the DATUM value will be assumed to be an integer. Making a call such as

```
CALL KEPLAB(3, 'After return from crunch', 'Energy = %d')
```

builds in extra information. Together with the event "INDEX" a legend will be presented

which connects type 3 with the string "After return from crunch". Further, when the value of the DATUM argument is shown it will be formatted according to the format string - a typical result would be

```
Energy = 23
```

Note that the FORTRAN string is interpreted according to the conventions associated with the C function `printf`. At its simplest this merely means that text is printed as entered and the special string '%d' is replaced with an integer value.

The profiler is initially off and must be explicitly enabled with calls to `KEPINI` and `KEPON`. In the *Cubix* environment this is facilitated by the use of the `KEPINQ` function to query the state of the run-time switches given to the `cubix` command.

The log of profiling information is written to the host file system with `KEPCP` or `KEPEND`.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the event profiler.

```
PROGRAM EPRTST
C
REAL ENERGY, RESID, GRIND, CRUNCH
INTEGER ITER, I
INTEGER LOGBUF(2048), LABBUF(256)
C
C-- Start up Express.
C
CALL KXINIT
C
C-- Start profiler, make labels for indices 1-3.
C
CALL KEPINI(LABBUF, 1024, LOGBUF, 8192)
CALL KEPLAB(1, 'Outer loop', 'Iteration %d')
CALL KEPLAB(2, 'After crunch', 'Energy = %d')
CALL KEPLAB(3, 'Inner loop', 'resid = %d')
CALL KEPON
C
C-- Start application code, then go into main loop.
C
...
DO 10 ITER=1,100
CALL KEPADD(1, ITER)
ENERGY = CRUNCH(ITER)
CALL KEPADD(2, INT(ENERGY))
DO 20 I=1,4
RESID = GRIND(ENERGY)
CALL KEPADD(3, INT(RESID))
```

```
20     CONTINUE
10     CONTINUE
C
C-- Program over, dump profile data and exit.
C
      ...
      STOP
      END
```

The insertion of events like these above can provide significant information about an application. The time between events 1 and 2, for example, indicates the duration of a call to the CRUNCH function. Similar information is available about GRIND from events 2 and 3, averaged over the four calls per iteration. The auxiliary DATUM fields will show the interaction between the variables and the program execution rate. It may also show up bugs and/or unexpected behavior which could be the key to understanding the failings of a particular parallelization scheme.

SEE ALSO

etool (command), KEPCP, KEPEND

KEPTOG

NAME

KEPTGI, KEPTOG - Calculate program statistics.

SYNOPSIS

```
SUBROUTINE KEPTGI (TOGGLE, LABEL)
INTEGER TOGGLE (16)
CHARACTER*80 LABEL
```

```
SUBROUTINE KEPTOG (TOGGLE)
INTEGER TOGGLE (16)
```

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

These routines allow selective analysis of particular sections of code. By surrounding code segments with calls to KEPTOG one can obtain statistics relating to the number of times the particular code section was called and the average and total times spent in these sections. The data is collected in exactly the same manner as the "event profiling" information obtained through calls to KEPADD. The same commands are available to dump the profiling data and/or rename the file containing it as are used by the other "KEPROF" utilities.

Each "toggle" data structure must be initialized with a call to KEPTGI before it can be used for data collection. This function expects to be passed an array of integers and a string that will later be used to identify the collected statistics when analyzed with `etool`.

The log of profiling information is written to the host file system with KEPCP or KEPEND.

EXAMPLE

The following example demonstrates the use of the "toggle" ideas.

```
PROGRAM TOGTST
INTEGER LPTOG (16), GRNTOG (16)
REAL*4 ENERGY, GRIND
INTEGER ITER, I
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Initialize Express and its common block.
C
CALL KXINIT
C
C-- Initialize toggle data structures.
C
CALL KEPTGI(LPTOG, 'Main iteration loop')
CALL KEPTGI(GRNTOG, 'Calls to GRIND')
```

```
C
C-- Start application code, then go into main loop.
C
      ...
C
      DO 10 ITER = 1, 100
        CALL KEPTOG(LPTOG)
C
C-- Other processing going on here....
C
      .....
C
      DO 20 I = 1, 4
        CALL KEPTOG(GRNTOG)
        ENERGY = GRIND(I)
        CALL KEPTOG(GRNTOG)
20    CONTINUE
      CALL KEPTOG(LPTOG)
10  CONTRINUE
C
C-- Dump data to host for later analysis...
C
      .....
C
      STOP
      END
```

The “toggle” data will be stored in a file with the name “eprof.out” (unless overridden by some other function call) together with the normal “event” data which may have also been collected with calls to KEPADD.

To analyze this data we execute the “etool” command with

```
etool -p -t
```

This combination of switches both suppresses the normal graphical output and also restricts

attention to the "toggle" data. the output for the above example might appear as follows

Node 0

Description	Total	#Calls	Avge.	Var.
Main iteration loop	478.32	100	4.78	.28
Calls to GRIND	363.96	400	0.91	.03

Node 1

Description	Total	#Calls	Avge.	Var.
Main iteration loop	478.32	100	4.78	.28
Calls to GRIND	363.96	400	0.91	.03

etc...

For each node is displayed the list of initialized toggles together with the number of times each code section was used, the total time elapsed in this section, the average time per call and the variance of these times. Using this information it is possible to build up a very accurate picture of the performance of a parallel program.

SEE ALSO

etool (command), KEPCP, KEPROF, KEPEND.

NAME

KERASE, KAERAS - Clear the display surface.

SYNOPSIS

SUBROUTINE KERASE

SUBROUTINE KAERAS

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These routines are used to clear the display surface. KERASE must be called at the same time in all nodes while KAERAS may be called asynchronously in any node at any time. Note that this latter option can cause rather unpredictable results unless used with some care since 1000 nodes all calling KAERAS makes the screen "blink" rather a lot!

It is important to note that neither of these routines flushes the graphics buffer to the output device. Rather they just reset the internal data structures to reflect empty buffers. All data that is required to appear, however briefly, on the display surface must be flushed explicitly with one of the KSEND routines before calling KERASE.

Hardcopy devices handle these functions in device specific ways. Some, for example, can only print on a single sheet at once and so the KERASE commands are handled by switching to a new output file. Eventually several files may be printed one at a time. Others, such as laser printers, merely switch to new pages.

EXAMPLE

The following is typical of the general use of the KERASE functions.

```
C
C-- Start up graphics system.
C
      ISTAT = KOPENP(GBUFFER, 8192)
      IF(ISTAT .LT. 0) THEN
          WRITE(6,*) 'Failed to start up graphics'
          STOP
      ENDIF
C
C-- Grind away .... graphics, graphics, graphics....
C
      .....
C
C-- Finished with first image, erase and go again.
C
```

KERASE

CALL KERASE

SEE ALSO

KSENDP

NAME

KEXEC - Overlay a node application.

SYNOPSIS

```
SUBROUTINE KEXEC (PROG)
CHARACTER*80 PROG
```

```
SUBROUTINE KAEXEC (PROG)
CHARACTER *80 PROG
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

This routine overlays (and hence terminates) the calling program by loading a new application from the file PROG. The new routine immediately begins execution from its main entry point. Unless an error occurs the call to KEXEC will not return.

By default KEXEC causes the overlay to occur in all nodes. It must, therefore, be a loosely synchronous operation. If, however, the default *Cubix* mode is asynchronous then each node performs the overlaying operation independently. This is also the case for the KAEXEC system call.

Since memory is not re-initialized across calls to these routines it is possible to share large blocks of data in each node. To do this it is merely necessary for the data to be placed in a region of memory where none of the intended programs will overwrite it.

EXAMPLE

The following code section causes the program PASS2 to be loaded on top of the currently executing routine.

```
PROGRAM DOEXEC
CHARACTER*80 PRGNAM
PARAMETER (PRGNAM='pass2')
C
C-- Start up Express.
C
CALL KXINIT
C
C-- Execute the first phase of the program and then EXEC the
C-- second phase to overlay this one.
C
CALL PHASE1
CALL KEXEC (PRGNAM)
C
C-- If we get here then something really bad happened...
```

KEXEC

```
C
  WRITE(6,*) 'Returned from EXEC ..... Hmummmmm'
  STOP
  END
```

WARNINGS

This function is only available to programs running with the *Cubix* file server. If you are running with a host program the same effect can be achieved by simply using the KXLOAD routines to download another node application.

NAME

KFLUSH - Flush I/O buffers.

SYNOPSIS

```

SUBROUTINE KFLUSH (UNIT)
INTEGER UNIT

```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

Under *Cubix* I/O is buffered which means that characters are stored up in internal data structures and then emitted in large packets to improve efficiency. In this case it is not always clear what will appear in a file when the buffers are automatically flushed by the operating system and so KFLUSH is provided to force this process to happen under the users control.

It is important to note that NO automatic flushing is ever performed on "multi" mode files - KFLUSH is required to show any output from these files.

EXAMPLES

The following code segment demonstrates the effect of the KFLUSH call on a file in "multi" mode.

```

PROGRAM MULTI
C
C-- Setup Express and its common block.
C
CALL KXINIT
C
WRITE(6,*) 'Hello world'
CALL KMULTI(6)
C
WRITE(6,*) 'This is one of the processors...'
C*****
WRITE(6,*) '...and still is'
C
CALL KSINGL(6)
WRITE(6,*) '... that's all for now folks!'
C
STOP
END

```

When executed on four processors this would produce the output

KFLUSH

```
Hello world
This is one of the processors...
    ...and still is
This is one of the processors...
    ...and still is
This is one of the processors...
    ...and still is
This is one of the processors...
    ...and still is
    ... that's all for now folks!
```

We can understand this output by noting that the `KMULTI` and `KSINGL` system calls both force an implicit `KFLUSH` operation. As can be seen the two `WRITE` statements separated by the comment line with all the asterisks were buffered - when the `KSINGL` call finally flushed their buffers the two lines came out together in each node. If we replace the line of asterisks with a call to `KFLUSH` then the output would have been:

```
Hello world
This is one of the processors...
    ...and still is
    ...and still is
    ...and still is
    ...and still is
    ... that's all for now folks!
```

in which we can see that every line is sent out as it is written.

SEE ALSO

`KMULTI`

NAME

KGETHO - Determine host specific characteristics

SYNOPSIS

```
SUBROUTINE KGETHO (NODE, BUFFER)
INTEGER NODE
CHARACTER*80 BUFFER
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

This routine is used to determine host-specific characteristics. The current implementation is restricted to returning, in BUFFER, the name of the operating system running on the host whose node identified is NODE. Up to 80 characters of this information will be transferred to the indicated buffer, any extra will be discarded.

No attempt is made to differentiate between minor versions of operating systems, or between the various "unix-like" machines.

EXAMPLE

In the following code we determine the type of operating system running on our standard host in order to find the character used to separate components of filenames. Since FORTRAN is notoriously non-standard in the area of character and string variables we assume, for the purposes of this code segment, that a routine called ISTREQ exists which compares two strings and returns zero if they are the same.

```
CHARACTER*1 FUNCTION GETSEP (NODE)
INTEGER NODE
CHARACTER*80 OSBUF

CALL KGETHO (NODE, OSBUF)
IF (ISTREQ (OSBUF, 'unix') .EQ. 0) GETSEP = '/'
ELSE IF (ISTREQ (OSBUF, 'dos') .EQ. 0) GETSEP = '\'
ELSE IF (ISTREQ (OSBUF, 'macintosh')) GETSEP = ':'
ELSE IF (ISTREQ (OSBUF, 'vms')) GETSEP = '.'
ELSE
    WRITE (6,*) 'Unrecognized OS: ', OSBUF
    GETSEP = 0
ENDIF
RETURN
END
```

KGETHO

SEE ALSO

KCBXSY

NAME

KGIN, KAGIN - Graphical input operations

SYNOPSIS

INTEGER FUNCTION KGIN (BUTTON, PX, PY)

INTEGER BUTTON

REAL PX, PY

INTEGER FUNCTION KAGIN (BUTTON, PX, PY)

INTEGER BUTTON

REAL PX, PY

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These routines are used to perform graphical input operations usually termed "locator input". Upon execution a cursor appears on the screen and is positioned and triggered in a device specific manner. After triggering the KGIN calls return and the specific trigger and position are returned.

The KGIN routine must be called in each processor simultaneously while the KAGIN function may be executed by any processor at any time. In this latter case it is the responsibility of the user to ensure that sufficient information is present to allow the operator to know which processor is requesting input. Further, no flushing is performed by these functions. It is up to the user to ensure that the display surface actually contains up-to-date data before requesting graphical input.

The coordinates returned to the user are expressed relative to those set up by the last call to KSPACE in each processor. Further a status value is returned to indicate the result of the KGIN operation. A negative value is returned by devices which are not capable of performing input. A zero return value implies that the KGIN operation completed successfully but that the cursor position was outside the window selected by the most recent call to KVPOR or KSETVP in this processor. A positive return means that the coordinates selected lay within the processor window. This last mechanism can be used to select processors with a mouse, for example.

EXAMPLE

In the following we assign different halves of the display to two processors; node 0 gets the left half and node 1 the right. We then use the input routines to select one or the other node for some processing task.

```
INTEGER NDDATA (4)
REAL X, Y
INTEGER STAT, KEY
```

C

KGIN

```
C-- Divide up the screen on the basis of processor
C-- number.
C
  CALL KXPARA(NDDATA)
C
  IF(NDDATA(1) .EQ. 0) THEN
    CALL KVPORT(0., 0., .5, 1.)
  ELSE
    CALL KVPORT(.5, 0., 1., 1.)
  ENDIF
C
C-- Now assign coordinates. Each processor's window
C-- will be mapped individually to the unit square.
C
  CALL KSPACE(0.,0.,1.,1.)
C
10  KEY = 1
    ISTAT = KGIN(KEY, X, Y)
    IF(ISTAT .LT. 0) GOTO 20
    IF(ISTAT .GT. 0) THEN
      CALL GRIND(X,Y)
    ENDIF
    IF(KEY .NE. 0) GOTO 10
20  CONTINUE
```

Having set up the windows and coordinate systems we loop until the KEY parameter is returned as zero and the processor whose region we indicated with the mouse calls the GRIND subroutine with the selected points as arguments. Note that we can perform similar operations on more processors by using the KXGRID routines to set up and coordinate the distribution of processors to screen areas.

SEE ALSO

KVPORT.

NAME

KGREYS, KAGREY - Change color attributes.

SYNOPSIS

```
SUBROUTINE KGREYS (FROM, TO)
INTEGER FROM, TO
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

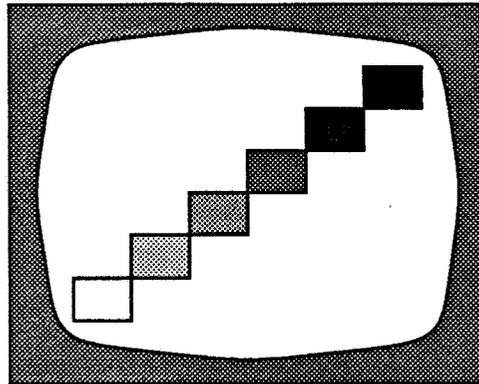
These routines change the association of color indices to device colors used by *Plotix*. By default a limited color map is used which can be extended with the KGREYS and KRAINB function calls.

KGREYS extends the *Plotix* color map by adding a smoothly varying grey-scale between the two selected values. The lower value will be white and the upper black. The number of distinct grey levels available is hardware dependent but in any case *Plotix* will map the indicated range in as smooth a manner as possible.

EXAMPLE

The following code draws a set of 6 boxes of varying grey shades along the diagonal of the screen image.

```
      INTEGER I
      REAL V
C
      CALL KSPACE (0., 0., 6., 6.)
C
      CALL KGREYS (10, 15)
C
      DO 10 I = 1, 6
         V = I
         CALL KBOX (V, V, V+1., V+1., 10+i, 0)
10    CONTINUE
      CALL KSENDP
```



SEE ALSO

KCOLOR, KRAINB

NAME

KLABEL - Add text.

SYNOPSIS

```
SUBROUTINE KLABEL(TEXT, X, Y)
CHARACTER*80 TEXT
REAL X, Y
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

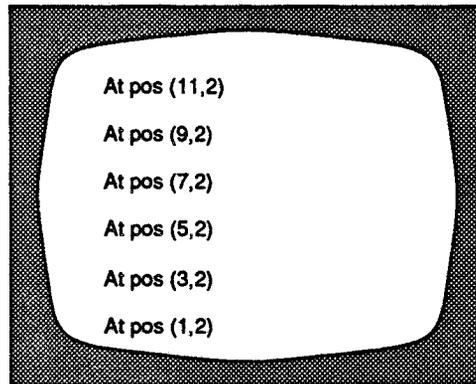
DESCRIPTION

This routine draws the characters contained in the TEXT string at the position (X,Y). The first character of the string is placed above and to the right of the indicated point. Other methods of justification can be obtained with the KDOTEX function.

EXAMPLE

The following code defines a 12 x 12 coordinate system and writes a string at several positions on the screen.

```
      INTEGER I
      CHARACTER*80 S
C
C-- Define square coordinate system.
C
      CALL KSPACE(0.,0.,12., 12.)
C
      DO 10 I=1,12,2
          WRITE(S,20) I
20      FORMAT(1X, 'At pos(', I3, ',', I3, ')')
          CALL KLABEL(S, 2.0, FLOAT(I))
10     CONTINUE
C
      CALL KSENDP
```



WARNING

The current plotting position is undefined after this call. In order to perform reliable graphical operations **KMOVE** should be used before any further drawing is performed.

SEE ALSO

KDOTEX, **KMARKE**

NAME

KLINEM - Modify drawing style for lines

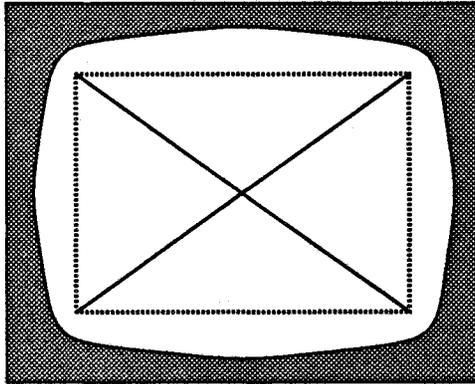
SYNOPSISSUBROUTINE KLINEM(INDEX)
INTEGER INDEX**DOMAIN**This routine may only be called in programs compiled with the *Plotix* libraries.**DESCRIPTION**

Modifies the style in which all further lines are drawn. The INDEX argument is an integer which specifies, in a device dependent manner, the actual linestyle to use. The value 0 will always create solid lines.

EXAMPLE

The following code defines a 10 x 10 coordinate system and draws a box with a dashed edge and a solid diagonal.

```
      CALL KSPACE(0.,0.,10., 10.)  
C  
      CALL KLINEM(1)  
      CALL KMOVE(1., 1.)  
      CALL KCONT(9.,1.)  
      CALL KCONT(9.,9.)  
      CALL KCONT(1.,9.)  
      CALL KCONT(1.,1.)  
C  
      CALL KLINEM(0)  
      CALL KCONT(9.,9.)  
      CALL KMOVE(1.,9.)  
      CALL KCONT(9.,1.)  
C  
      CALL KSENDP
```



SEE ALSO

KCONT, KCOLOR, KMOVE

NAME

KMARKE - Draw marker symbol.

SYNOPSIS

```
SUBROUTINE KMARKE (SYMBOL, X, Y, SIZE)
INTEGER SYMBOL
REAL X, Y, SIZE
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

This routine draws a marker symbol at the position (X,Y) expressed relative to the coordinate system most recently defined with the KSPACE function. The marker is drawn with the given SIZE, expressed in the same units as the coordinates. The SYMBOL argument is used to distinguish the various markers as follows

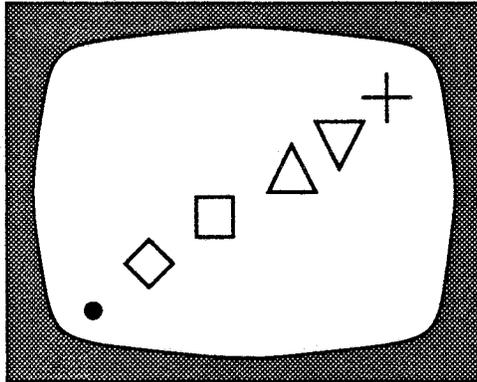
0	point
1	diamond
2	square
3	triangle
4	inverted triangle
5	cross
6	plus
7	star

Some attempt is made to compensate for the fact that "squares" look bigger than "triangles" - the SIZE argument is not strictly interpreted as the height of the triangle, for example.

EXAMPLE

The following code defines an 9 x 9 coordinate system and draws different marker symbols along the diagonal.

```
      INTEGER I
C
      CALL KSPACE(0.,0.,8.,8.)
C
      DO 10 I = 1,6
          CALL KMARKE(I, FLOAT(i+1), FLOAT(i+1), .5)
10    CONTINUE
      CALL KSENDP
```



SEE ALSO

KLABEL

NAME

KMOVE - Move without drawing.

SYNOPSIS

```
SUBROUTINE KMOVE (X, Y)
REAL X, Y
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

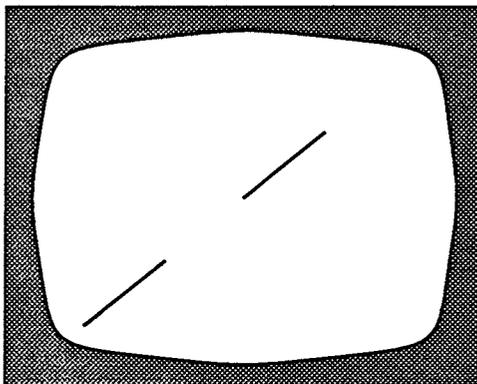
DESCRIPTION

Moves the current plotting position to (X,Y). Nothing is drawn on the display surface. X and Y are specified relative to the coordinate system defined by the most recent call to KSPACE.

EXAMPLE

The following code draws a broken diagonal line across the display surface.

```
CALL KSPACE (0., 0., 4., 4.)
C
CALL KMOVE (0., 0.)
CALL KCONT (1., 1.)
CALL KMOVE (2., 2.)
CALL KCONT (3., 3.)
C
CALL KSENDP
```

**SEE ALSO**

KCONT, KCOLOR, KLINEM

NAME

KMRD2D, KMWT2D - Read/write two dimensional data sets.

SYNOPSIS

```
INTEGER FUNCTION KMRD2D (UNIT, BUF, TOTCOL, TOTROW, ITEMSZ,  
                        COL0, COL1, ROW0, ROW1, SKIP)  
INTEGER UNIT, TOTCOL, TOTROW, ITEMSZ  
INTEGER COL0, COL1, ROW0, ROW1  
INTEGER SKIP, BUF (*)
```

```
INTEGER FUNCTION KMWT2D (UNIT, BUF, TOTCOL, TOTROW, ITEMSZ,  
                        COL0, COL1, ROW0, ROW1, SKIP)  
INTEGER UNIT, TOTCOL, TOTROW, ITEMSZ  
INTEGER COL0, COL1, ROW0, ROW1  
INTEGER SKIP, BUF (*)
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

These functions provide a primitive interface to a two-dimensional file access mode for *Cubix* programs. The basic idea is that data sets decomposed over a two dimensional array of processors can be read and written with a single function call.

Data is either read to or written from the array pointed to by BUF and consists of some number of "items" each of size ITEMSZ. This concept is used instead of the more obvious "byte" notation so that the other arguments to these functions may be assigned as row or column indices.

The disk data set is treated as an array of TOTROW by TOTCOL items of which a subset is to be read or written by each node. The particular piece of the global data set required by a given node is specified by the ROW0, ROW1, COL0 and COL1 arguments which are *inclusive* parameters indexed from zero - the specification

```
COL0 = 0  
COL1 = 9  
ROW0 = 0  
ROW1 = 9
```

would access the 10 x 10 block in the upper left hand corner of the array.

The SKIP parameter specifies the offset in the BUF array between successive "row" entries again in "items". This can be used to leave a boundary strip around the edge of the data as is common in two dimensional decompositions and is illustrated in the example below.

EXAMPLE

Suppose we have a two dimensional array of integers of size NX by NY which we wish to

decompose over the processors. The following code can be used to setup the decomposition with the KXGRID functions.

```

PROGRAM TEST
COMMON/XPRESS/
C
  INTEGER NDIM(2)
  INTEGER GBLSZ(2), LCLSZ(2), LCLSTR(2)
  INTEGER NDDATA(4)
C
C-- Set up Express and its common block.
C
  CALL KXINIT
C
C-- Get runtime parameters, processor number etc..
C
  CALL KXPARA(NDDATA)
C
C-- Divide the nodes up among the two dimensions of the
C-- data and initialize the KXGRID system.
C
  CALL KXGDSP(NDDATA(2), 2, NDIM)
  ISTAT = KXGDIN(2, NDIM)
C
C-- Figure out how much of the data fits in each node.
C
  GBLSZ(1) = NX
  GBLSZ(2) = NY
  CALL KXGDSI(NDDATA(1), GBLSZ, LCLSZ, LCLSTR)
C

```

Notice how we use KXGDSP to evenly divide up the number of processors between the data dimensions and KXGDSI to divide up the array between the processors. The parameters returned by KXGDSI can be used to read in a two-dimensional data set as follows. (We assume that UNIT is a file descriptor corresponding to some previously opened file and that the global variables defined in the previous program fragment are still available.

```

C
C-- Read data into nodes, no overlap allowed.
C
  SUBROUTINE RDDATA(UNIT)
  INTEGER UNIT
C
  ISTAT = KMRD2D(UNIT, DATA, NX, NY, 4,

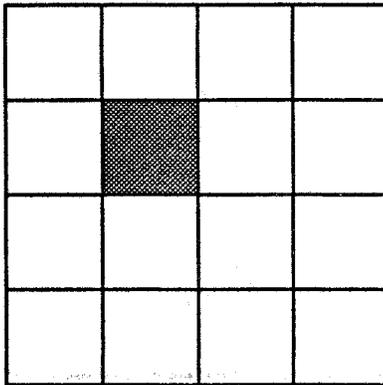
```

```

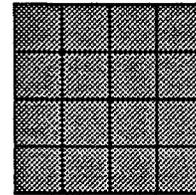
*      LCLSTR(1), LCLSTR(1)+LCLSZ(1)-1,
*      LCLSTR(2), LCLSTR(2)+LCLSZ(2)-1,
*      LCLSZ(1))
      RETURN
      END

```

This strategy uses the values returned by KXGDSI to figure out exactly which data to request from the input data set. In this case each node gets a distinct piece of data, divided as evenly as possible between the processors but with no overlap and no space for any. The mapping is as shown in the following figure.



Processor decomposition



Data array

A common situation is that in which the input data set is required to be read into the center of a block which contains, around its edges, space for one or more entries from a neighbor node. This is a common situation in image processing, for example, where some local convolution is to be applied. To achieve this effect with the above parameters we change the call to KXRD2D as follows:

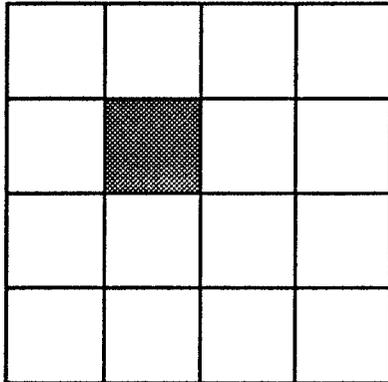
```

C
C-- Read data into nodes, overlap allowed but not
C-- performed.
C
      SUBROUTINE RDDATA(UNIT)
      INTEGER UNIT
C
      ISTAT = KMRD2D(UNIT, DATA(LCLSZ(1)+4), NX, NY, 4
* LCLSTR(1), LCLSTR(1)+LCLSZ(1)-1,
* LCLSTR(2), LCLSTR(2)+LCLSZ(2)-1,
* LCLSZ(1)+2)
      RETURN

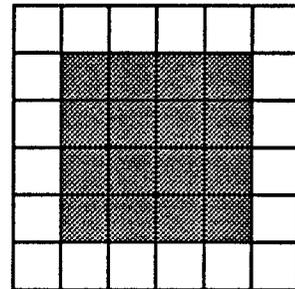
```

 END

This call performs the mapping shown in the next figure. Note that the `skip_dist` parameter has been modified to place a gap around each "row" of the data with one space at the beginning and one at the end. This would be suitable for a nearest neighbor interaction in which a single strip of data is required from each neighbor node.



Processor decomposition



Data array

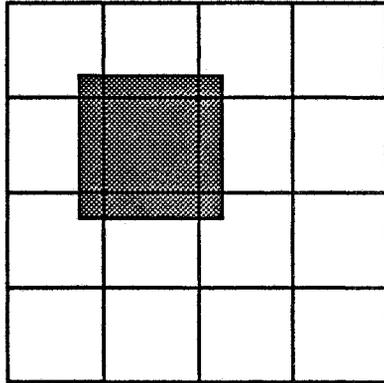
A last option which is interesting is one in which the data being read is overlapped at the time it is originally taken from the data set. This is merely a variation on the last call which provided space for the overlapped data but did not initialize it. The call required to read in overlapping data is as follows

```

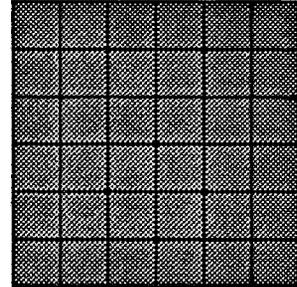
C
C-- Read data into nodes with overlapping strip one
C-- "item" wide.
C
  SUBROUTINE RDDATA(UNIT)
  INTEGER UNIT
C
  ISTAT = KXR2D(UNIT, DATA, NX, NY, 4,
  * LCLSTR(1)-1, LCLSTR(1)+LCLSZ(1),
  * LCLSTR(2)-1, LCLSTR(2)+LCLSZ(2),
  * LCLSZ(1)+2)
  RETURN
  END

```

This mapping is shown in the next figure.



Processor decomposition



Data array

WARNING

Reading and writing unformatted files is complicated by the fact that the host and nodes of the parallel processing system may not have the same type of processor (CPU) and may not share the same byte ordering properties. An example might be a Sun workstation hosting a transputer or NCUBE machine. In this case the host processor is a Motorola based system which has the most significant byte at the lowest memory address while the nodes have the opposite ordering. To cover these cases *Express* provides a set of byte swapping primitives: *KXSWAP*.

RETURN VALUE

KMRD2D returns the number of bytes read, or -1 upon unrecoverable errors. Similarly *KMWT2D* returns the number of bytes written by the calling node or -1 upon disastrous errors.

SEE ALSO

KMREAD, *KMWRT*, *KXSWAP*

NAME

KMREAD - Read independent data into each node.

SYNOPSIS

```
INTEGER FUNCTION KMREAD (UNIT, BUF, LENGTH, ORDER)
INTEGER UNIT LENGTH, ORDER, BUF (*)
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

KMREAD reads unformatted data into the nodes from the file indicated by the UNIT argument. Independent data is read into each node; the LENGTH arguments need not all be the same.

The call to KMREAD must be made loosely synchronously in all nodes.

The ORDER argument determines in what order the data from the input file are to be placed in the nodes. The simplest case, obtained by setting ORDER = NORDER, (Defined in the XPRESS common block), is for the input to appear in order of increasing processor number so that node 0 receives the first block followed by node 1 and so on. Other cases are obtained by setting the value to be an integer between 0 and the number of processors. The node which specified ORDER = 0 receives the first block and then the node which gave ORDER = 1 and so on. Note that it is an error if a value between 0 and the number of processors is not specified in some node. This condition is indicated by KMREAD returning -1.

EXAMPLE

Suppose that we have decomposed our domain into a two dimensional mesh with NX and NY processors in the two dimensions. If we now want to read data blocks in the conventional manner for such a grid - i.e., along the rows, then the KXGRID routines of *Express* can be used as follows

```
INTEGER FUNCTION BLKRD (UNIT, BLOCK, BLKSIZ, NX, NY)
REAL BLOCK (*)
INTEGER NX, NY, UNIT, BLKSIZ
INTEGER NDDATA (4)
C
C-- Assume that KXINIT has been called elsewhere.
C
NPROCS (1) = NX
NPROCS (2) = NY
ISTAT = KXGDIN (2, NPROCS)
ISTAT = KXGDCO (NDDATA (1), COORD)
C
BLKRD = KMREAD (UNIT, BLOCK, BLKSIZ,
```

KMREAD

```
*                                COORD(2)*NX + COORD(1))
  RETURN
  END
```

This will order the input according to the row and column coordinates of the processor in the two dimensional mesh.

WARNING

Reading and writing unformatted files is complicated by the fact that the host and nodes of the parallel processing system may not have the same type of processor (CPU) and may not share the same byte ordering properties. An example might be a Sun workstation hosting a transputer or NCUBE machine. In this case the host processor is a Motorola based system which has the most significant byte at the lowest memory address while the nodes have the opposite ordering. To cover these cases *Express* provides a set of byte swapping primitives: *KXSWAP*.

RETURN VALUE

KMREAD returns the number of bytes read, or -1 upon unrecoverable errors. A return value of zero indicates an "end of file" condition.

SEE ALSO

KMWRT, KXSWAP

NAME

KMULTI, KSINGL ISMULT, ISASY, KORDER - Parallel I/O characteristics.

SYNOPSIS

SUBROUTINE KMULTI (UNIT)
INTEGER UNIT

SUBROUTINE KSINGL (UNIT)
INTEGER UNIT

SUBROUTINE KASYNC (UNIT)
INTEGER UNIT

INTEGER FUNCTION ISMULT (UNIT)
INTEGER UNIT

INTEGER FUNCTION ISASYN (UNIT)
INTEGER UNIT

SUBROUTINE KORDER (UNIT, ORDER)
INTEGER UNIT, ORDER

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

These routines provide an interface to the parallel features of buffered file I/O. As well as their usual characteristics *Cubix* "UNIT"s are either in singular, multiple or asynchronous mode. This mode determines the exact behavior of READ and WRITE calls on that file with regard to the distribution of data.

If a file is in singular mode then any read operation on it must be made loosely synchronously and exactly the same data is transferred to each node. Similarly, write operations must be made loosely synchronously and only node zero actually transmits any data to the file. This has the effect of allowing all nodes to apparently write but only one copy appears in the output file. In this mode it is an error if any node attempts to read or write different data from the others. This error normally causes internode communication to "hang" or may occasionally cause the *cubix* program to abort with status -1.

In multiple mode read requests are satisfied from the file independently. Thus each node can read its own data. Output requests can also be made independently with each node writing its own data to the file. Note that in multiple mode no implicit flushing of buffers is performed and it is the responsibility of the user to call KFLUSH in order to cause data to appear in the indicated file.

In asynchronous mode, I/O requests are handled independently on the processors on which they occur. No interprocessor synchronization is performed. Each processor maintains its

own state variables recording the last byte it read or wrote in the file, and each request to read or write implicitly returns the file to that location before proceeding.

The routines KMULTI, KSINGL and KASYNC switch files between multiple, singular and asynchronous modes. KMULTI puts a file into the multiple mode and KSINGL restores it to singular mode. KASYNC places a file in asynchronous mode. All three flush any data in the file's buffers prior to the call, and all must be made loosely synchronously in all nodes.

By default both input and output operations on "multi" mode files occur in order of increasing processor number - i.e., node 0 gets the first crack followed by node 1, node 2 and so on. The KORDER system call is available to alter this default. The first argument indicates the unit for which a new ordering is desired and the second is an integer in the range 0... nprocs-1. Further "multi" mode operations on this file will result in the processor which specified order=0 being first, followed by that which gave order=1 etc. From this it should be obvious that the ORDER parameters given in the call to KORDER must form a permutation of the set {0, ..., nprocs-1} - i.e., each value must be specified exactly once in one of the nodes. Failure to observe this rule results in deadlock whenever I/O is attempted on the affected stream. (Examples of the use of this parameter in the lower-level KMREAD and KMWRIT system calls can be found on the respective manual pages).

The routine ISMULT returns 1 if its argument is in multiple mode and zero otherwise.

The routine ISASY returns 1 if its argument is in asynchronous mode, and zero otherwise.

EXAMPLES

The following code segment demonstrates the effect of the KMULTI call

```
PROGRAM MULTI
C
C-- Setup Express and its common block.
C
  CALL KXINIT
C
  WRITE(6,*) 'Hello world' CALL KMULTI(6)
  WRITE(6,*) 'This is one of the processors...'
  CALL KSINGL(6)
  WRITE(6,*) '... that's all for now folks!'
C
  STOP
END
```

When executed on four processors this would produce the output

```
Hello world
This is one of the processors...
This is one of the processors...
This is one of the processors...
```

This is one of the processors...
... that's all for now folks!

showing that only one line of output results from each call to WRITE in single mode while each processor generates its own output while the file is in multi-mode.

Asynchronous mode typically arises in one of two situations. Either a code is truly asynchronous - its behavior is too unpredictable in advance to allow use of the synchronous I/O modes or one might want to use this mode for reporting runtime errors that may only occur within a single node.

SEE ALSO

KMREAD, KMWRIT, KFLUSH

NAME

KMWRIT - Write independent data from each node.

SYNOPSIS

```
INTEGER FUNCTION KMWRIT(UNIT, BUF, LENGTH, ORDER)
INTEGER UNIT, LENGTH, ORDER, BUF(*)
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

KMWRIT writes data from the nodes to the file indicated by the UNIT argument. Independent data is written from each node; the LENGTH arguments need not all be the same.

The call to KMWRIT must be made loosely synchronously in all nodes.

The ORDER argument determines in what order the data from the various nodes are to be placed in the output file. The simplest case, obtained by setting ORDER = NORDER, is for the output to appear in order of increasing processor number. (NORDER is to be found in the XPRESS common block set up by the call to KXINIT.) Other cases are obtained by setting the value to be an integer between 0 and the number of processors. First in the output appears the data from the node which specified ORDER = 0 then that from the node with ORDER = 1 and so on. Note that it is an error if a value between 0 and the number of processors is not specified in some node. This condition is indicated by KMWRIT returning -1.

EXAMPLE

Suppose that we have decomposed our domain into a two dimensional mesh with NX and NY processors in the two dimensions. If we now want to write out data blocks in the conventional manner for such a grid - i.e., along the rows then the KXGRID routines of *Express* can be used as follows

```
INTEGER FUNCTION BLKWT(UNIT, BLOCK, BLKSIZ, NX, NY)
REAL BLOCK(*)
INTEGER NX, NY, UNIT, BLKSIZ
INTEGER NDDATA(4)

C
C-- Assume that KXINIT has been called elsewhere.
C
  NPROCS(1) = NX
  NPROCS(2) = NY
  ISTAT = KXGDIN(2, NPROCS)
  ISTAT = KXGDCO(NDDATA(1), COORD)

C
  BLKRD = KMWRIT(UNIT, BLOCK, BLKSIZ,
```

```
*                                COORD (2) *NX + COORD (1) )
  RETURN
  END
```

This will order the output according to the blocks in the two dimensional grid.

WARNING

Reading and writing unformatted files is complicated by the fact that the host and nodes of the parallel processing system may not have the same type of processor (CPU) and may not share the same byte ordering properties. An example might be a Sun workstation hosting a transputer or NCUBE machine. In this case the host processor is a Motorola based system which has the most significant byte at the lowest memory address while the nodes have the opposite ordering. To cover these cases *Express* provides a set of byte swapping primitives: KXSWAP .

RETURN VALUE

KMWRIT returns the number of bytes written, or -1 upon unrecoverable errors.

SEE ALSO

KMREAD, KXSWAP

NAME

KOPENP, KAOPEN, KCLOSP - Begin and terminate graphics system.

SYNOPSIS

INTEGER FUNCTION KOPENP (BUFFER, SIZE)
INTEGER BUFFER(*), SIZE

INTEGER FUNCTION KAOPEN (BUFFER, SIZE)
INTEGER BUFFER(*), SIZE

SUBROUTINE KCLOSP

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These routines initialize and terminate the graphics system.

One of the KOPEN functions must be the first *Plotix* function called in any graphics application. The arguments denote the internal buffer to be used for storing graphical information between calls to the KSENDP functions. The array must be provided by the user and its size (in bytes) indicated through the second argument. As a guide to appropriate sizes a call to KMOVE or KCONT requires 5 bytes.

KAOPEN performs the same function as KOPENP but asynchronously - that is any node may call this routine independent of the others with no synchronization constraints.

KOPENP returns a status code indicating the success or failure of the setup procedures. Negative values indicate errors and it is unwise to proceed if an error condition exists since terminals, for example, may be sent into strange states.

The last graphical routine to be called by an application should be KCLOSP. This serves to close any open files and reset interactive devices to their normal states.

Both KOPENP and KCLOSP must be called loosely synchronously in all nodes, while KAOPEN may be called independently at any time by any node.

EXAMPLE

The following skeleton code should provide the basis for all graphics applications.

```
PROGRAM GRAFIX
  INTEGER GBUFF (2048)
C
C-- Start up Express.
C
  CALL KXINIT
C
```

```
C-- Set up graphics.
C
      ISTAT = KOPENP(GBUFF, 8192)
      IF(ISTAT .LT. 0) THEN
          WRITE(6,*) 'Failed to init graphics system'
          STOP
      ENDIF
C
C-- Application code .....
C
      .....
C
C-- Application finished, clear up graphics system.
C
      CALL KCLOSP
      STOP
      END
```

SEE ALSO

KSENDP

NAME

KPANEL - Draw and fill polygons.

SYNOPSIS

```
SUBROUTINE KINITP (COLOR, EDGE)
INTEGER COLOR, EDGE
```

```
SUBROUTINE KPANLP (X, Y)
REAL X, Y
```

```
SUBROUTINE KENDPA
```

```
SUBROUTINE KPOLGN (NPTS, XPTS, YPTS, COLOR, EDGE)
INTEGER NPTS, COLOR, EDGE
REAL XPTS (*), YPTS (*)
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These routines are used to draw and fill polygonal regions of the display surface. KPOLGN, the most straightforward of the routines takes two arrays each containing NPTS values as the x and y coordinates of the vertices of the polygon to be drawn. The coordinates need not close - the first and last points are joined by the system. The resulting polygon will be filled according to the COLOR argument and will have its outline drawn in the current color if EDGE is non-zero.

Positive values of COLOR translate into solid colors in the same manner as the arguments to the line color primitive, KCOLOR. Negative values yield device dependent shading patterns.

All coordinates are expressed relative to the most recent call to KSPACE.

An alternative interface to the polygon routines is provided by KINITP, KPANLP and KENDPA. The first routine initializes the system so that the following polygon will be drawn and filled according to the COLOR and EDGE arguments, interpreted as above. This routine must be called to initialize each polygon. Successive calls to KPANLP then add vertices to the current polygon and the figure is closed and filled by the KENDPA call. This interface is often superior to KPOLGN since it does not have the memory overhead of storing points in arrays.

Note that filling with COLOR = 0 and EDGE = 0 results in a "selective erase" - specific areas of the screen can be erased.

EXAMPLE

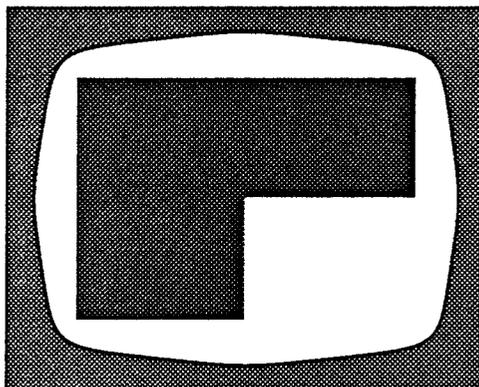
The following code draws a simple box in the foreground color using the KPOLGN primitive and then takes a "bite" out of it with the alternate routines by drawing and filling

in the background color.

```

REAL XPTS(4), YPTS(4)
DATA XPTS/1., 9., 9., 1./
DATA YPTS/1., 1., 9., 9./
C
CALL KSPACE(0.,0.,10., 10.)
C
CALL KPOLGN(4, XPTS, YPTS, 1, 1)
C
C-- Now draw a polygon filled in the background color
C
CALL KINITP(0, 0)
CALL KPANLP(5., 1.)
CALL KPANLP(9., 1.)
CALL KPANLP(9., 5.)
CALL KPANLP(5., 5.)
CALL KENDPA
C
CALL KSENDP

```



SEE ALSO

KBOX, KCOLOR

KPLOTH

NAME

KPLOTH - Analyze usage of system buffers.

SYNOPSIS

INTEGER FUNCTION KPLOTH

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

Graphics commands are buffered internally on each node until flushed by one of the KSENDP commands. This necessitates assigning a fixed size buffer for graphics. In order to "tune" the size of this buffer and ensure that neither graphics gets lost nor too much memory is devoted to this system the function KPLOTH returns the "high water mark" from the graphics system - i.e., the maximum number of bytes that were present between any two calls to the KSENDP primitives. Using this function allows the user to exactly determine system memory requirements.

EXAMPLE

Assuming that the buffer size is currently set to 8192 bytes the following code might be used to warn of impending overflows.

```
C
C-- Make display "current".
C
    CALL KUSEND
C
C-- Set up for asynchronous I/O since not all node might
C-- have overflowed.
C
    CALL KASYNC(6)
    IHWM = KPLOTH()
    IF(IHWM .GT. 8000) THEN
        WRITE(6,*) 'Warning: graphics buffer tight'
    ENDIF
```

Notice that we use the asynchronous *Cubix* mode for the warning message since it is not guaranteed that all processors will have filled their buffers to the same extent.

SEE ALSO

KOPENP, KSENDP

NAME

KPXGOP, KPXSOP - Manipulate hardware dependencies in *Plotix* programs.

SYNOPSIS

```
INTEGER FUNCTION KPXGOP (OPTION, VALUE)
CHARACTER*80 OPTION
INTEGER*4 VALUE
```

```
INTEGER FUNCTION KPXSOP (OPTION, VALUE)
CHARACTER*80 OPTION
INTEGER*4 VALUE
```

DOMAIN

These routines may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

Plotix attempts to provide device-independent graphical capabilities to *Express* programs. Due to the simple nature of the underlying graphics model this can usually be achieved to a large degree. To deal with those cases where either more capabilities are available or where more information is required about a particular *Plotix* implementation these two functions are provided.

KPXGOP accesses the value of some internal property described by the character string `OPTION` and stores it under the supplied pointer variable. The particular values of `OPTION` supported on any particular device vary according to device capabilities and are listed in the device specific section of the *Plotix* chapter of the User's Guide. If the indicated property is not supported on the device in use -1 is returned.

The opposite function is provided to set internal state of some *Plotix* system with KPXSOP. This routine takes a character representation of the required property and a single 32-bit value to which the indicated property will be set. If the named property is not supported on the device in use -1 will be returned.

When successful both routines return 0.

EXAMPLE

The following code segment initializes a *Plotix* system and also attempts to perform the following three tasks:

- Inquire how many distinct colors are available for drawing lines.
- Request output in "landscape" rather than the default "portrait" mode.
- Install a named "redraw" function which will be used in windowing versions of *Plotix* to repaint the screen under certain well-defined circumstances.

Note that any one of these requests may fail because the device currently in use may not be able to support them. In the code segment below we imagine that the calling program is able to deal with such failures without having to tell the user. In other situations we could look

for a 0 return value from the calls to indicate failure and issue diagnostics.

```
SUBROUTINE GPHINI(PBUFFR, PSIZE, NCOLS)
INTEGER*4 PBUFFR(*), PSIZE, NCOLS
EXTERNAL IPAIN
CHARACTER*80 REDRAW, LANSCP, NLCOLS
PARAMETER (REDRAW='redraw', LANSCP='landscape',
$          NLCOLS='nlcolors')

  ISTAT = KPXSOP(REDRAW, IPAIN)
  ISTAT = KPXSOP(LANSCP, 0)

  ISTAT = KOPENP(PBUFFR, PSIZE)
  IF(ISTAT .LT. 0) THEN
    WRITE(6,*) 'Failed to initialize graphics'
    STOP
  ENDIF

  ISTAT = KPXGOP(NLCOLS, NCOLS)
  IF(ISTAT .LT. 0) THEN
    WRITE(6,*) 'No data - assuming monochrome'
    NCOLS = 2
  ENDIF
  RETURN
END
```

We make the calls to KXPSOP before the call to KOPENP while the call to KPXGOP follows it. This is common practice - in many *Plotix* implementations the call to KOPENP is responsible for setting up a lot of the default behavior of the system and so it makes sense to make our preferences known before starting the system. This is one of the few cases in which KOPENP should *not* be the first call made to *Plotix*. Similarly we wait until after the device has been initialized before asking how many colors are available. This allows for systems which must be initialized before they can know how many colors are available.

SEE ALSO

KOPENP

NAME

KPROFI - Low level execution profiler

SYNOPSIS

```
SUBROUTINE KPROFI (BUFFER, BUFLen, START, SCALE)
INTEGER BUFFER(*), BUFLen, START, SCALE
```

DOMAIN

KPROFI may only be called in the nodes.

DESCRIPTION

This routine serves to initialize the execution profiler. Every few milliseconds the program counter of the user application is examined and a histogram entry in the memory area denoted by BUFFER is incremented. The size of the histogram area is BUFLen bytes.

In order to decide which histogram entry to increment a "mapping function" is applied to the program counter discovered by the system. First START is subtracted and then the result is multiplied by SCALE and divided by 0x10000 (Hexadecimal) - i.e., the complete mapping is

$$\text{BIN} = (\text{PC} - \text{START}) * \text{SCALE} / 0\text{x}10000$$

The overall effect of the SCALE parameter is to map groups of adjacent program locations into the same histogram bin. The value SCALE = 0x10000 maps every program location into a separate histogram bin, SCALE = 0x8000 maps each pair of locations into a single bin, SCALE = 0x4000 every group of four, and so on.

Using combinations of the BUFLen, START and SCALE parameters it is possible to allocate various memory ranges to be profiled. Note that no errors are incurred if the range is not large enough resulting in a calculated BIN which is out of the histogram range. In this case a special "misses" counter is incremented. This latter feature also provides some diagnostic information concerning the success of the profiling attempt - if an incorrect profiling range is selected most of the histogram entries will be in the "miss" bin allowing easy diagnosis.

KPROFI does not enable the profiler. An explicit call to KXPON must be made to begin gathering profile data.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the execution profile.

```
PROGRAM XPRTST
C
INTEGER PRFBUF(2048), PRFSCL
PARAMETER (PRFSCL = 8192)
C
```

KPROFI

```
C-- This is the name given to a particular routine in
C-- the program which is known to occur low in memory.
C-- This information can usually be obtained from a
C-- "linker map".
C
    EXTERNAL F_MAIN
C
C-- Start Express.
C
    CALL KXINIT
C
    CALL KPROFI(PRFBUF, 8192, F_MAIN, PRFSCL)
    CALL KXPON
C
C-- Application code, profiler running....
C
    ...
```

The choice of the `START` argument is most conveniently made in conjunction with the "linker map" provided by the compiler. This usually contains a list of the addresses of all the functions in an application and can be used to find the smallest.

SEE ALSO

`xtool (command)`, `KXPROF`

NAME

KRAINB - Change color attributes.

SYNOPSIS

SUBROUTINE KRAINB (FROM TO)
INTEGER FROM, TO

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

This routine changes the association of color indices to device colors used by *Plotix*. By default a limited color map is used which can be extended with the KRAINB and KGREYS function calls.

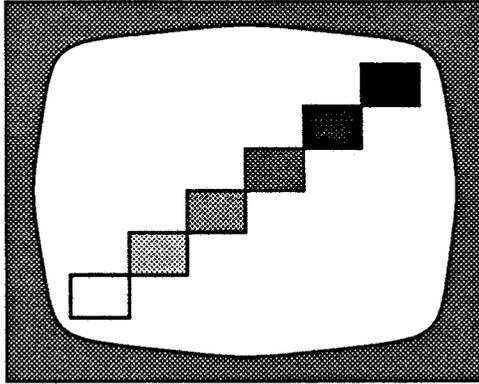
KRAINB extends the *Plotix* color map by adding a smoothly varying color spectrum between the indicated values. The "rainbow" starts with red and varies, with full saturation and value, through the different hues; red, magenta, blue, yellow, cyan and back to red. The number of distinct colors available is hardware dependent but in any case *Plotix* will map the indicated range in as smooth a manner as possible.

On devices incapable of providing color output this function is treated exactly as a call to KGREYS.

EXAMPLE

The following code draws a set of 6 boxes of varying colors along the diagonal of the screen image. Since the manual is printed on a monochrome device the output is exactly as if the call to KRAINB were replaced with one to KGREYS!

```
      INTEGER I
      REAL V
C
      CALL KSPACE(0., 0., 6., 6.)
C
      CALL KRAINB(10, 15)
C
      DO 10 I = 1, 6
          V = I
          CALL KBOX(V, V, V+1.0, V+1.0, 10+I, 0)
10  CONTINUE
      CALL KSENDP
```



SEE ALSO

KCOLOR, KGREYS

NAME

KREAD, KWRITE - Read or write unformatted data.

SYNOPSIS

```
INTEGER FUNCTION KREAD (UNIT, BUFFER, LENGTH)
INTEGER UNIT LENGTH, BUF (*)
```

```
INTEGER FUNCTION KWRITE (UNIT, BUFFER, LENGTH)
INTEGER UNIT LENGTH, BUF (*)
```

DOMAIN

This routine may only be called in programs compiled with the *Cubix* or *Plotix* libraries.

DESCRIPTION

KREAD reads unformatted data into the nodes from the file indicated by the UNIT argument. LENGTH bytes are read and placed in the indicated BUFFER. KWRITE performs the opposite function - LENGTH bytes are transmitted to the indicated UNIT from the BUFFER.

If the indicated file UNIT is in "singl" mode then KREAD and KWRITE must be called loosely synchronously in all nodes. If the file is in "async" mode or *Cubix* has been switched to asynchronous mode with a call to KCBXSY then independent calls to these functions may be made in each node.

These functions provide the fastest but least portable interface to unformatted Fortran file I/O. They make direct calls to the C routines read and write circumventing any intermediate buffering or file transfer protocols.

RETURN VALUE

KREAD returns the number of bytes read, or -1 upon unrecoverable errors. A return value of zero indicates an "end of file" condition. KWRITE returns the number of bytes written or -1 upon unrecoverable errors.

WARNING

Reading and writing unformatted files is complicated by the fact that the host and nodes of the parallel processing system may not have the same type of processor (CPU) and may not share the same byte ordering properties. An example might be a Sun workstation hosting a transputer or NCUBE machine. In this case the host processor is a Motorola based system which has the most significant byte at the lowest memory address while the nodes have the opposite ordering. To cover these cases *Express* provides a set of byte swapping primitives: KXSWAP.

SEE ALSO

KMREAD, KMWRIT, KXSWAP

NAME

KSENDP - Flush graphical data to display surface.

SYNOPSIS

SUBROUTINE KSENDP

SUBROUTINE KUSEND

SUBROUTINE KASEND

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

In the implementation of *Plotix* for parallel computers output is “buffered”. This means that each KMOVE, KCONT, KPANLP, etc. command merely stores its parameters in an area of memory rather than immediately attempting to draw the associated object. This strategy is dictated by the fact that typical parallel computers have large computing power but little I/O bandwidth. As a result it makes no sense to send lots of small messages about graphical objects to the device since this would result in spending all ones time communicating rather than computing. Instead we store up a large number of objects and then send them all at once.

This method introduces the “flushing” concept to the graphical system. No data actually appears on the display surface until one of the three KSENDP commands is executed. The differences between the three commands are typified by the following observations of common situations

- | | |
|--------|---|
| KSENDP | All the nodes have been simultaneously drawing the same part of an image. This situation is quite common - it costs nothing to duplicate the same sequential effort in all processors. All nodes make the call to KSENDP together but the data is only flushed to the display once. |
| KUSEND | The nodes have been working separately on their own pieces of the image and are now ready to ship it out to the display. All processors call KUSEND together and the set of objects from each node appear in order of increasing processor number. |
| KASEND | The nodes are working totally independently and asynchronously. A particular node wishes to send some data to the display and has no way of knowing the status of the other processors. Any node may call KASEND at any time. |

The effect of these calls is to empty the buffer on the calling node ready for more graphical objects.

The buffer size required for graphical objects varies quite significantly from application to application. In some codes it may be possible and efficient to call the KSENDP functions

quite regularly and so only a small buffer is required. Others may operate for long periods without flushing data and, as a result, need large buffers. The size of the graphics buffer is set in the call to KOPENP.

EXAMPLE

The following code segment illustrates one of the less obvious bugs possible under *Plotix*. We use the system calls to draw a "menu" and then accept selections from it with KGIN.

```
C
C-- Demo of PLOTIX code --- incorrect!!
C
    CALL KSPACE(0., 0., 4., 4.)
C
C-- Draw simple menu on left hand edge of display.
C
    CALL KLABEL('QUIT', .1, .5)
    CALL KLABEL('ITERATE', .1, 1.5)
    CALL KLABEL('RESET', .1, 2.5)
    CALL KLABEL('OUTPUT', .1, 3.5)
C
C-- Use KGIN to get user option from menu.
C
    ISTAT = GIN(KEY, X, Y)
    OPTION = INT(Y)
```

The error here is that data is not flushed before the call to KGIN. As a result the user is asked to make a selection from an invisible menu. Not very friendly. The solution is, however, very straightforward; insert a call to KSENDP before the call to KGIN. Note that this illustrates another aspect of the flushing commands - since all processors have been drawing the same thing and we only want to see one copy of it on the display the appropriate flushing function is KSENDP.

SEE ALSO

KOPENP

KSPACE

NAME

KSPACE - Define user coordinate system.

SYNOPSIS

```
SUBROUTINE KSPACE (LOWX, LOWY, HIGHX, HIGHY)
REAL LOWX, LOWY, HIGHX, HIGHY
```

```
SUBROUTINE KORTHO (LOWX, LOWY, HIGHX, HIGHY, JUSTFY)
REAL LOWX, LOWY, HIGHX, HIGHY
INTEGER JUSTFY
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These routines define a coordinate system to be mapped onto the current window. By default all plotting commands take place in a coordinate system which has (0., 0.) at its lower left corner and (1., 1.) at the upper right. After this call all future plotting commands, including the input request commands, will operate in the new coordinate system.

While the KSPACE routine covers the entire viewport with the selected coordinate range the KORTHO function can be used to preserve the aspect ratio of the indicated coordinate system. A mapping is created so that objects will actually appear with the correct shape independent of the specific characteristics of a particular output device - circles will be circular not elliptical.

Since a correctly normalized region may not completely fill the current viewport the JUSTFY parameter is used to indicate exactly where the region should be placed. The value -1 implies that the new region should be placed either to the left or at the bottom of the viewport while +1 indicates the right or top. A zero value centers the region within the viewport.

EXAMPLE

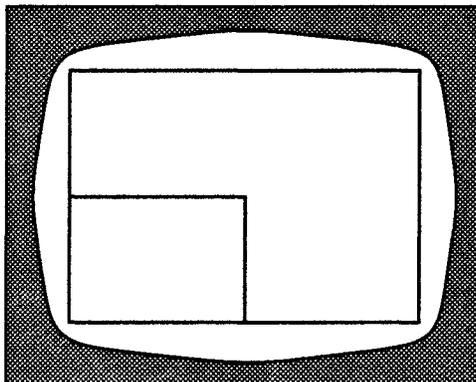
This example shows the effect of KSPACE transformations on simple objects. The routine MYBOX draws a unit square on the screen.

```
      SUBROUTINE MYBOX
C
      CALL KMOVE (0., 0.)
      CALL KCONT (1., 0.)
      CALL KCONT (1., 1.)
      CALL KCONT (0., 1.)
      CALL KCONT (0., 0.)
      RETURN
```

END

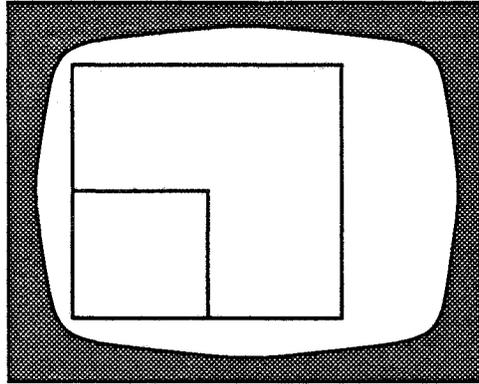
To see the effect of the KSPACE call consider the following sequence

```
C
C-- Default coords ==> full screen "square".
C
    CALL MYBOX
C
C-- Redefine coordinate system to make square fill only
C-- one quadrant of the display.
C
    CALL KSPACE(0., 0., 2., 2.)
    CALL MYBOX
    CALL KSENDP
```



As can be seen the resulting "square" is not! To correct this we could instead use the KORTHO function as shown below. Note that we chose the justification that the used area should be to the left of the viewport.

```
CALL MYBOX
CALL KORTHO(0., 0., 2., 2., -1)
CALL MYBOX
CALL KSENDP
```



In Appendix C is presented a complete example program in which the `KXGRID` routines are used to map processors to their own individual windows on the display surface and `KSPACE` is used to map each individual processor region to its own coordinate range. Note that it is possible to have different coordinate ranges in separate processors.

SEE ALSO

`KVPORT`

NAME

KVPORT, KSETVP - Specify area of display to hold image.

SYNOPSIS

```
INTEGER FUNCTION KVPORT (LOWX, LOWY, HIGHX, HIGHY)
REAL LOWX, LOWY, HIGHX, HIGHY
```

```
SUBROUTINE KSETVP (WINDOW)
INTEGER WINDOW
```

DOMAIN

This routine may only be called in programs compiled with the *Plotix* libraries.

DESCRIPTION

These routines allocate and activate/deactivate certain area of the display surface. The supplied parameters are expressed as fractions of the total view surface so that the default $0.0 < x < 1.0$ and $0.0 < y < 1.0$ is the entire display. By selecting smaller regions in x and y it is possible to confine an image to a smaller region of the display. This is useful if the final image is required to have a certain aspect ratio or in parallel processing applications where each processor is to be assigned a piece of the view surface.

Plotix allows several viewports to be present on the same display surface. Each is indicated by a number returned by the corresponding call to KVPORT and is selected by a call to KSETVP. Note that each viewport or window has its own coordinate range specified by a call to KSPACE and that clipping is performed independently in each window. Further, since the call to KVPORT selects the new viewport a call to KSPACE to set up a coordinate system must come after the corresponding call to KVPORT.

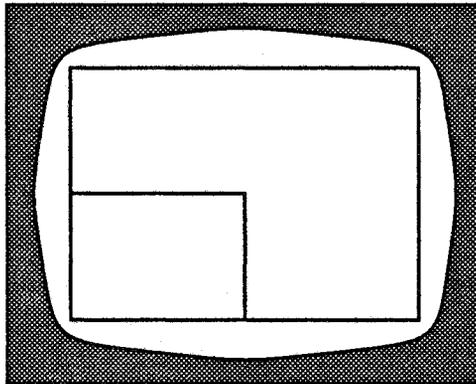
EXAMPLE

This example shows the effect of KVPORT transformations on simple objects. The routine MYBOX draws a unit square on the screen.

```
      SUBROUTINE MYBOX
C
      CALL KMOVE (0., 0.)
      CALL KCONT (1., 0.)
      CALL KCONT (1., 1.)
      CALL KCONT (0., 1.)
      CALL KCONT (0., 0.)
      RETURN
      END
```

To see the effect of the KVPORT call consider the following sequence

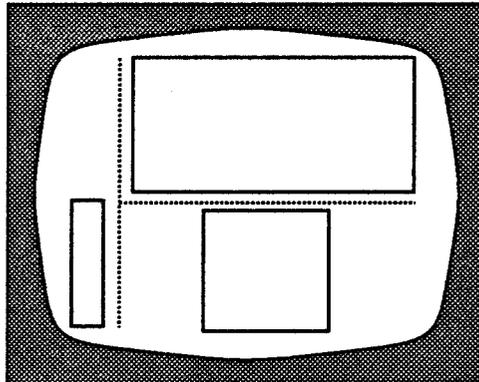
```
C
C-- Default coordinates ==> full screen "square".
C
    CALL MYBOX
C
C-- Define (and implicitly activate) new viewport.
C
    IVP = VPORT(0., 0., .5, .5)
    CALL MYBOX
    CALL KSENDP
```



To see the effect of multiple viewports consider the following code segment. We create three windows. The left window has a call to KSPACE which means that the "box" fills only the bottom part of the viewport. The second window has no call to KSPACE so its coordinate range will have the usual default. The last window uses KORTHO to make a viewport with the correct aspect ratio - the square actually comes out square!

```
    INTEGER LFTWIN, TOPWIN, LOWWIN
C
C-- Left window, scaling range (0,0) -> (1,2)
C
    LFTWIN = KVPORT(0.0, 0.0, 0.2, 1.0)
    CALL KSPACE(0.0, 0.0, 1.0, 2.0)
C
C-- Top window, default scaling range (0,0) --> (1,1)
C
    TOPWIN = KVPORT(0.2, 0.5, 1.0, 1.0) C
C-- Lower window, scaled (0,0) --> (1,1), correct
C-- aspect ratio.
C
    LOWWIN = KVPORT(0.2, 0.0, 1.0, 0.5)
```

```
        CALL KORTHO(0.0, 0.0, 1.0, 1.0, 0)
C
C-- Set up windows, draw the squares .....
C
        CALL KSETVP(LFTWIN)
        CALL MYBOX
C
        CALL KSETVP(TOPWIN)
        CALL MYBOX
C
        CALL KSETVP(LOWWIN)
        CALL MYBOX
C
        CALL KSENDP
```

**SEE ALSO**

KSPACE

KXACCS

NAME

KXACCS - Share a processor group with another process

SYNOPSIS

```
INTEGER FUNCTION KXACCS (DEVICE, NNODES)
CHARACTER*80 DEVICE
INTEGER NNODES
```

DOMAIN

Host processor only.

DESCRIPTION

This routine provides a “brute-force” mechanism by which a host program can obtain access to every node in the network irrespective of whether or not that node is currently executing a program - even if allocated to another user. This is often useful for providing overall system monitoring or when only a single application is to run on the entire network.

The first argument specifies the particular parallel computer to which access is desired and is interpreted in the same manner as the corresponding argument to KXOPEN. The last argument is returned to the caller containing the number of nodes in the system.

RETURN VALUE

The value returned by KXACCS is the *processor group index* which must be used in future references to the shared processors. If some error occurs or nodes are not accessible to the host processor -1 is returned.

WARNINGS

Communicating with shared processor groups is complicated by interactions between source and type fields specified using the NOCARE wildcard. This situation can be eased somewhat through the KXTYPE mechanisms which restrict the ranges indicated by wildcard values. It should further be noted that subsequent to this call the host must communicate with the processors using the node numbers indicated by `cnftool` rather than according to the logical mapping which results from KXOPEN or KXSHAR.

SEE ALSO

KXOPEN, KXSHAR, KXTYPE.

NAME

KXBREA - Halt program at breakpoint

SYNOPSIS

SUBROUTINE KXBREA

DOMAIN

This routine may only be called in node programs.

DESCRIPTION

The `exbreak` function causes the program to halt as though it had encountered a breakpoint of the type normally associated with the debugger, `ndb`. Examination of the process state with `ndb` will show the process to be in state `Breakpoint`.

SEE ALSO

KXPAUS.

NAME

KXBRD - Interprocessor broadcast.

SYNOPSIS

```
INTEGER FUNCTION KXBRD (BUFFER, ORIGIN, NBYTES,  
                        NNODES, NODEL, TYPE)  
INTEGER BUFFER (*)  
INTEGER ORIGIN, NBYTES, NNODES, NODEL, TYPE
```

DOMAIN

KXBRD may be called in both host and node processors.

DESCRIPTION

KXBRD is used to perform broadcasting operations among the processors.

The broadcast starts from processor ORIGIN which attempts to broadcast the NBYTES of data in the indicated BUFFER. The processors to which the broadcast will be sent are indicated by the NNODES and NODEL arguments in the following way: NODEL is an array of processor numbers which should receive the broadcast message. NNODES is the number of elements in the array. Further the special value NNODES = IALNOD (defined in the XPRESS common block set up by the KXINIT function) indicates that the broadcast should go to all processors. In this case the value of NODEL is ignored. Receiving nodes deposit the incoming data at BUFFER, up to a maximum of NBYTES.

The broadcast operation carries a "type" field in common with all other communication primitives so that overlapping broadcasts may be distinguished. This parameter is supplied as the argument TYPE and may be any positive quantity. It is illegal to use the NOCARE value from the XPRESS common block for this field

A call to KXBRD in the originating node must have corresponding calls to KXBRD in all target nodes. A corresponding call in other nodes which are not target nodes is not necessary, but will be handled without error. All calls must specify the same values of the ORIGIN, NNODES and NODEL arguments or communication deadlock will occur. A receiving node must specify NBYTES greater than or equal to that specified in the originating node. When NODEL is used the contents arrays must be exactly identical in each processor. The ORIGIN may or may not appear in the NODEL, at the convenience of the calling routine. When no errors occur, the value returned is the number of bytes written by the originating node, or the number read by a receiving node.

EXAMPLE

In the following code we use the KXGRID tools to find the processor number of the processor at the origin of a three dimensional processor decomposition. This processor then broadcasts a set of data values to all other nodes.

```
PROGRAM MYTEST
```

```
C
```

```
        INTEGER NPROCS(3), COORD(3), CORNER, TYPE
C
C-- This is the EXPRESS common block.
C
        COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
        DATA TYPE/33/
C
C-- Set up Express and initialize its common block.
C
        CALL KXINIT
C
C-- Initiate a three-dimensional decomposition of
C-- eight processors.
C
        NPROCS(1) = 2
        NPROCS(2) = 2
        NPROCS(3) = 2
        ISTAT = KXGDIN(3, NPROCS)
        IF(ISTAT .LT. 0) THEN
            STOP
        ENDIF
C
C-- Now find the processor in the (0,0,0) spot in the user
C-- topology.
C
        COORD(1) = 0
        COORD(2) = 0
        COORD(3) = 0
        CORNER = KXGDPR(COORD)
C
        ISTAT = KXBROD(DATBUF, CORNER, 32*4, IALNOD, 0, TYPE)
C
        .....
```

DIAGNOSTICS

If any error occurs in KXBROD -1 is returned. Possible sources of error are: an illegal buffer, a preposterous value of NBYTES or invalid values of ORIGIN, NNODES or NODEL. If no error occurs the number of bytes broadcast is returned in the originating processor and the number read in the receiving nodes. An error condition is also indicated in any node which reads less bytes than were originally transmitted by the originating processor.

SEE ALSO

KXCOMB, KXCONC

NAME

KXCH - Hardware dependent communication primitives.

SYNOPSIS

SUBROUTINE KXCHON (CHAN)
INTEGER CHAN

SUBROUTINE KXCHOF (CHAN)
INTEGER CHAN

INTEGER FUNCTION KXCHRD (CHAN, BUFFER, NBYTES)
INTEGER CHAN, BUFFER (*), NBYTES

INTEGER FUNCTION KXCHWT (CHAN, BUFFER, NBYTES)
INTEGER CHAN, BUFFER (*), NBYTES

DOMAIN

These routines are available only to node programs. Their availability is further subject to hardware restrictions on the system in use.

DESCRIPTION

These routines implement a message passing strategy which directly accesses the hardware present on the parallel processing system in use. As such their use is highly non-portable. Since, however, these routines have a very trivial syntax they can provide communication at the full speed of the underlying hardware. In most cases this affects the asymptotic communication rate only slightly but may reduce the start-up time (latency) by as much as an order of magnitude. They are most applicable, therefore, when the application needs to send many short messages.

Before attempting to use the message passing routines KXCHOF must have been called for every channel on which the low level functions will be used. this function serves to disable the normal *Express* processing for that channel. Note that the user is responsible for ensuring that no internode communication traffic will be disrupted by the sudden removal of one of the message paths normally used by *Express*. In practice this usually means that the application should force a synchronization through some operation before disabling any of the communication channels. Note that while a channel is disabled none of the higher level *Express* functions may be used. In particular this means that the debugger, *ndb*, will be unable to operate.

KXCHON performs the opposite function, causing *Express* to once again become active on the indicated channel. Again it is the user's responsibility to ensure that no *Express* messages are transmitted along channels that are still disabled.

The channel read function, KXCHRD, reads NBYTES bytes of data into the supplied BUFFER from the channel indicated by the CHAN argument. It will not return until exactly NBYTES have been read. The node from which data is read depends on the interpretation of the CHAN argument, which is hardware dependent.

Similarly the KXCHWT function sends NBYTES bytes of data into the channel indicated by

the CHAN argument. The data to be transmitted is taken from the user supplied BUFFER. This function will not return until all data has been read by a corresponding call to KXCHRD.

EXAMPLES

The following schematic code shows a typical sequence involving the KXCH primitives. We assume that some routine, NEARST requires heavy internode message traffic between processors directly connected to each other in hardware. As such they are able to make use of the KXCH functions.

```

C
C-- Assume that we can work here with the full Express system.
C
          .....
C
C-- For the next function we will disable Express.
C
      .   CALL KXSYNC
          DO 10 I=1,NCHANS
              CALL KXCHOF(I-1)
      10  CONTINUE

          CALL NEARST
C
C-- Assume that this routine terminates fully synchronized
C-- so that we can enable Express.
C
          DO 20 I=1,NCHANS
              CALL KXCHON(I-1)
      20  CONTINUE
C
C-- Proceed with Express functioning.....
C

```

Notice that we have used a variable NCHANS to indicate how many channels should be modified. The value of this variable is also somewhat machine dependent - on a transputer system it might be four for all the hardware links, for example, while on a hypercube it will usually be the base 2 logarithm on the number of nodes.

WARNING

These routines perform extremely hardware dependent operations and as such should be used with caution. The "nearest neighbor" communication model that they represent has however, been shown by a number of researchers to be adequate (if not optimal) for a wide class of algorithms. (An excellent reference is the book "Solving Problems on Concurrent Processors" by G.C.Fox *et al.*, published by Prentice-Hall, 1988.)

If these routines seem appropriate for your algorithm we suggest that the full *Express* routines be used during development, since this enables use of the other system tools such as the debugger, and then these routines be substituted in the final product.

SEE ALSO

exread, exwrite, exsync.

NAME

KXCHAN, KXVCHA - Synchronous scalar/vector exchange primitive.

SYNOPSIS

INTEGER FUNCTION

```
      KXCHAN (IBUF, ILEN, ISRC, ITYPE, OBUF, OLEN, ODEST, OTYPE)  
INTEGER IBUF (*), OBUF (*)  
INTEGER ILEN, ISRC, ITYPE, OLEN, ODEST, OTYPE
```

INTEGER FUNCTION

```
      KXVCHA (IBUF, ISIZE, IOFF, IITEMS, ISRC, ITYPE,  
             OBUF, OSIZE, OOFF, OITEMS, ODEST, OTYPE)  
INTEGER IBUF (*), OBUF (*)  
INTEGER ISIZE, IOFF, IITEMS, ISRC, ITYPE,  
OSIZE, OOFF, OITEMS, ODEST, OTYPE
```

DOMAIN

These functions may be called in either host or node processors.

DESCRIPTION

These functions are used to implement "synchronous" communication between two processors; a call to KXCHAN in one processor will not return until the corresponding call has been made in the sending and receiving processors.

This function essentially performs a similar task to successive calls to KXWRIT and KXREAD - i.e., data is first sent and then read from (possibly) different nodes. The advantage of this function is that its extra constraint (synchronicity) allows optimizations to be made for both speed and reliability. The former can be achieved because data transmission in the two directions can be overlapped while the latter is enhanced because low level "handshaking" can be performed to ensure that no intermediate buffers overflow. A second advantage is that the exchange of information can be considered to be simultaneous - the user is freed from any worry about which node should read first and which write. As a result these functions should be preferred to the analogous pair of KXREAD/KXWRIT operations whenever the synchronous constraint can be met.

KXCHAN causes OLEN bytes of data to be sent to the node denoted by ODEST in a message of type OTYPE. The data is taken from memory at OBUF. It is not guaranteed that OLEN will be read by the reading processor; the actual number of bytes read depends on the number specified in the corresponding call to KXCHAN. If no error occurs, the actual number of bytes written is returned to the calling program. The interpretation of the destination and type fields is exactly as in KXWRIT. Note that this includes the fact that neither ODEST nor OTYPE may take the special NOCARE value.

KXCHAN also causes at most ILEN bytes of data to be read from the source denoted by ISRC from a message of a type matching ITYPE. The data is placed in memory at IBUF. It is not guaranteed that ILEN bytes will be read; the actual number of bytes read depends on the number written by the transmitting processor. If no error occurs, the actual number

of bytes read is returned to the calling program. The interpretation of the ISRC and ITYPE arguments is exactly as in KXREAD.

A call to KXCHAN must be complemented by calls to KXCHAN in the processors denoted by ODEST and ISRC in order to prevent communication deadlock. Similarly the message types in these processors must be compatible.

Note that the exchange of data is conceptually simultaneous - data is written to the output processor at the same time as it is received from the sender. This allows, for example, the buffer arguments to be identical. The kernel maintains the integrity of the data and handles any read/write synchronization problems.

The above discussion holds equally well for the KXVCHA function. The difference between the two is analogous to the difference between KXREAD and KXVREA. While the former is used to transmit contiguous blocks of memory the latter is able to send messages made up of several disjoint memory areas.

The arguments to KXVCHA are interpreted in the same way as their counterparts in KXVREA. The message is specified by defining a number of "objects" to be sent. Each is of length SIZE bytes and is separated from the next by OFFSET bytes. In total ITEMS objects will be transmitted. This description applies to both the input and output arguments of KXVCHA.

EXAMPLE

Consider a simple model of a two-dimensional terminal screen. We assume that the data currently displayed is represented by an 80 x 24 array of characters. Using the KXGRID and KXCHAN primitives it is easy to construct routines which, for example, scroll the data in different directions when decomposed in parallel.

```
PROGRAM MYTEST
PARAMETER (IHORIZ = 0, IVERT = 1)
INTEGER NPROCS(2)
C
C-- The amount of the display in each node is found by
C-- decomposing the 80 x 24 total over the processors.
C
LOGICAL*1 SCREEN(20,12)
C
C-- Set up Express
C
CALL KXINIT
C
NPROCS(IHORIZ) = 4
NPROCS(IVERT) = 2
ISTAT = KXGDIN(2, NPROCS)
IF(ISTAT .LT. 0) THEN
STOP
ENDIF
```

The macros IHORIZ and IVERT are defined for our convenience and just serve to label the two axes on the screen. We assign four processors to the horizontal dimension and two to the vertical. (A more flexible assignment scheme is easily devised using the KXPARA and KXGDSP system calls to determine at runtime the number of processors available.)

Now consider a simple scrolling operation in which data is to be passed to the right. We need to figure out the processor numbers necessary to communicate in this direction using KXGDNO.

```

      INTEGER NDDATA(4)
      INTEGER RECNUM(2), PERBC(2), TYPE
      INTEGER UNODE, DNODE, LNODE, RNODE
      DATA TYPE/12/
C
C-- Get runtime parameters, number of nodes, etc..
C
      CALL KXPARA(NDDATA)
C
C-- Turn off periodic boundary conditions.
C
      PERBC(1) = 0
      PERBC(2) = 0
      CALL KXGDBC(PERBC)
C
      UNODE = KXGDNO(NDDATA(1), IVERT, 1)
      DNODE = KXGDNO(NDDATA(1), IVERT, -1)
      LNODE = KXGDNO(NDDATA(1), IHORIZ, -1)
      RNODE = KXGDNO(NDDATA(1), IVERT, 1)

```

Note that we have made the additional step of dealing with the boundaries of the screen correctly. If a processor is on the extreme left edge of the display and it tries to communicate with a processor to its left then the value of LNODE will be correctly assigned the value NONODE which will, in turn, direct KXCHAN to omit communication with this non-existent processor.

Now in order to “scroll” the data over to the left we merely use the following call to KXCHAN.

```

      ISTAT = KXCHAN(SCREEN, 12, LNODE, TYPE,
                   SCREEN(20, 1), 12, RNODE, TYPE)

```

Notice that at no point in these calculations did the topology of the hardware enter. Everything is specified in the user domain - i.e., screen coordinates, and KXGRID and KXCHAN do the rest. Notice the appearance of the “magic” number 12 in the above call. To arrive at this value we divided the height of the screen (24) by the number of processors in that direction (2). We could do much better by using the KXGDSI function which would also allow the possibility of changing the number of processors allocated to each dimension

at runtime.

In order to scroll data vertically instead of horizontally we would just use the call

```
ISTAT = KXVCHA (SCREEN, 1, 12, 20, DNODE, TYPE,  
              SCREEN(1,12), 1, 12, 20, UNODE, TYPE)
```

The arguments here are arrived at in a similarly simple manner. If each processors piece of the display surface is 20 x 12 then we need to take every twelfth byte when we scroll upward. Also there are twenty bytes to transmit. (Again this can be made more flexible using the KXGDSI function.)

DIAGNOSTICS

If any error occurs in KXCHAN or KXVCHA -1 is returned. Possible sources of error are: an illegal source or destination, an illegal buffer or a preposterous value of length, size, offset or item arguments. If no error occurs KXCHAN returns the number of bytes read and KXVCHA the number of items read.

SEE ALSO

KXREAD, KXWRIT, KXVREA, KXGRID, KXPARA

NAME

KXCLOS - Deallocate processors.

SYNOPSIS

```
SUBROUTINE KXCLOS (PGIND)
INTEGER PGIND
```

DOMAIN

Available to host programs only.

DESCRIPTION

This routine is used to terminate a connection between the host and a processor group.

This routine should be called at the end of an application's use of a processor group to ensure that system resources are correctly reset. The sole argument, PGIND, is the *Processor group index* originally returned by the KXOPEN call.

EXAMPLES

The following schematic code should be the general template of any host process which allocates and uses processor groups.

```
PROGRAM MYTEST
INTEGER PG1
CHARACTER*80 DEVICE
PARAMETER (DEVICE='/dev/transputer')

C
C-- Set up Express.
C
    CALL KXINIT
C
C-- Allocate a processor group.
C
    PG1 = KXOPEN(DEVICE, 4, NOCARE)
    IF(PG1 .LT. 0) THEN
        WRITE(6,*) 'Failed to allocate nodes'
        STOP
    ENDIF
C
C-- Load progs, send/rec messages to groups of processors. C
    ...
C
C-- Program finished. Clean up by deallocating processors.
C
    CALL KXCLOS(PG1)
    STOP
```

KXCLOS

END

SEE ALSO

KXOPEN, KXSHAR, KXINIT

NAME

KXCOMB - Node data compaction.

SYNOPSIS

```
INTEGER FUNCTION KXCOMB(BUFFER, FUNC, SIZE, ITEMS,  
                        NNODES, NODEL, TYPE)  
INTEGER BUFFER(*)  
INTEGER FUNC, SIZE, ITEMS, NNODES, NODEL, TYPE
```

DOMAIN

KXCOMB may be called in the node processors only.

DESCRIPTION

This routine is used to performing “combining” operations on data within the node processors. An example of such an operation is the sum of a set of values distributed over the nodes of the parallel machines. Other example combining functions are products, maximum and minimum functions.

NITEMS of data, each of size SIZE and taken from BUFFER are individually combined across the specified node processors. The final vector of results will overwrite BUFFER in each node.

In the nodes the user-supplied combining function will be called for each of the NITEMS to be combined. In each case three arguments will be supplied to the routine. The first two are the objects to be combined and the third is the SIZE argument supplied in the call to KXCOMB. The combining function’s responsibility is to perform whatever operation is required and write the result over the first operand. The value returned by the combining function is used to detect errors in the KXCOMB routine. If a negative value is returned the current call to KXCOMB is aborted and an error returned to its caller.

The nodes involved in the combining operation are specified by the NNODES and NODEL arguments. The latter is an array of processor numbers listing those nodes on which the combining operation is to take place. NNODES is the number of elements in this list. The special value NNODES = IALNOD is allowed and performs the combining operation on all processors. In this case the value of NODEL is ignored. IALNOD is found in the XPRESS common block setup by the call to KXINIT.

The TYPE argument is used to specify a “type” for the combine operation. This is used to distinguish between potentially overlapping communication operations. Any positive value is legal - the special NOCARE value may not be used in this function.

All processors involved in the combining operation must call KXCOMB together with identical values for both NNODES and NODEL - otherwise communication deadlock will occur.

EXAMPLE

The first example merely calculates the global sum of the components of a vector distributed over all processors. We assume that each processor contains NPTS values.

```
PROGRAM MYTEST
C
  INTEGER NPTS, I, TYPE
  EXTERNAL SUMUP
  REAL RESULT
  DATA TYPE/37/
  COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Set up Express and initialize the common block.
C
  CALL KXINIT
C
C-- First compute the subtotal in each node.
C
  RESULT = 0.0
  DO 10 I=1,NPTS
    RESULT = RESULT + VALUE(I)
10  CONTINUE
C
C-- Now combine these values with the sum_up function.
C
  ISTAT = KXCOMB(RESULT, SUMUP, 4, 1, IALNOD, 0, TYPE)
C
  .....
C
  STOP
  END

  INTEGER FUNCTION SUMUP(V1, V2, SIZE)
  REAL V1, V2
  INTEGER SIZE
C
  V1 = V1 + V2
  SUMUP = 0
  RETURN
  END
```

Notice how the combining function replaces its first argument with the result of the combining operation and returns zero to indicate successful combination.

In the second example processors 0 thru 4 have obtained a vector of four floating point values, MYVEC. The purpose of the call is to calculate, for each array slot, the maximum value distributed over the nodes.

```

      INTEGER I, NNODES, NODEL(5), TYPE
      EXTERNAL MAXFLT
      REAL RESVEC(4)
      DATA TYPE/48/

C
C-- Setup nodelist array to specify combining nodes.
C
      NNODES = 5
      DO 10 I = 1,NNODES
          NODEL(I) = I-1
10  CONTINUE
C
C-- Combine values with the maxflt function.
C
      ISTAT = KXCOMB(MYVEC, MAXFLT, 4, 4, NNODES, NODEL, TYPE)
C
      .....
C
      STOP
      END

      INTEGER FUNCTION MAXFLT(V1, V2, SIZE)
      REAL V1, V2
      INTEGER SIZE
C
      IF (V2 .GT. V1) THEN
          V1 = V2
      ENDIF
      MAXFLT = 0
      RETURN
      END

```

This example points out another important point. The purpose of the NITEMS field is to allow multiple data items to be combined AND LEFT SEPARATE - it is not to perform "on-node" combination operations before the global operation. This is the reason why we explicitly coded the subtotal operation in the first example. In this last example four values will be left in each node corresponding to the maximum value among all processors of MYVEC(1), the maximum of MYVEC(2) and so on.

DIAGNOSTICS

If any error occurs in KXCOMB -1 is returned. Possible sources of error are: an illegal BUFFER, preposterous values of NITEMS or SIZE and the return of a negative value from the combining function. If no error occurs the number of items combined is returned.

WARNING

The combining function must be both commutative and associative in order to give results which are independent of the underlying hardware topology. If we denote the operation of the combining function on two elements A and B by A op B then the requirements can be written as

Commutativity: $A \text{ op } B = B \text{ op } A$

Associativity: $(A \text{ op } B) \text{ op } C = A \text{ op } (B \text{ op } C)$

Useful functions which satisfy these constraints are: addition, multiplication, maximum, minimum, logical AND, logical OR, logical XOR. Operations which do not satisfy the constraint are: subtraction ($3 - 1 \neq 1 - 3$) and division ($4 / 2 \neq 2 / 4$)

A particularly unpleasant problem concerns the use of this function with floating point numbers. Because of the intricacies of rounding and truncation while performing floating point operations it cannot be guaranteed that all nodes will have *exactly* the same result after a KXCOMB operation such as addition. While the difference will be (at worst) microscopic it can occasionally be sufficient to cause *Cubix* programs to abort while printing values in "singl" mode. At present no solution for this problem is known.

BUGS

There is an implementation specific upper limit on the SIZE of each individual item that can be combined. In most applications this should be an unimportant restriction.

SEE ALSO

KXBROD, KXCONC

NAME

KXCONC - Concatenate data from nodes.

SYNOPSIS

```
INTEGER FUNCTION KXCONC (MYBUF, MYBYTE, RESBUF, RESSZ,  
                        SIZES, NNODES, NODEL, TYPE)  
INTEGER MYBUF (*), RESBUF (*)  
INTEGER MYBYTE, RESSZ, SIZES (*), NNODES, NODEL (*), TYPE
```

DOMAIN

KXCONC may be called in only the node processors.

DESCRIPTION

This routine is used to collect and concatenate data in a set of node processors.

Each node contributes MYBYTE bytes of data from the array MYBUF to be placed in each node's RESBUF. The individual blocks of data are sorted into order of increasing processor number and placed in the RESBUF buffer, separated by RESSZ bytes. If any node contributes more than RESSZ bytes to the global vector then the excess are discarded. The amount of data contributed by each processor is stored in the appropriate slot of the SIZES array.

The group of nodes participating in the concatenation operation is specified by the NNODES and NODEL arguments. The latter is an array of processor numbers; NNODES specifies the number of processors in the list. If NNODES has the special value IALNOD then the concatenation is performed by all nodes irrespective of the value of the NODEL argument. The IALNOD parameter is to be found in the XPRESS common block setup by the call to KXINIT.

The TYPE argument is used to specify a "type" for the concatenation operation which will distinguish it from other active communication. Any positive value may be supplied - the special NOCARE value may not be used.

All nodes involved in the concatenation operation must call KXCONC together and with identical values for the NNODES and NODEL arguments or communication deadlock will occur.

EXAMPLE

Consider a simple case with four processors which have buffers as follows

```
Processor 0: INTEGER MYBUF (4) = {12, 13, 14}  
Processor 1: INTEGER MYBUF (4) = {32, 33, 34}  
Processor 2: INTEGER MYBUF (4) = {52, 53, 54}  
Processor 3: INTEGER MYBUF (4) = {72, 73, 74}
```

In the simplest case we can concatenate all four buffers with the code

```
ISTAT = KXCONC (MYBUF, 3*4, IBUF, 3*4, SIZES, IALNOD, 0, TYPE)
```

which would result in each processor obtaining the following result in IBUF

```
IBUF = {12, 13, 14, 32, 33, 34, 52, 53, 54, 72, 73, 74}
```

and the value returned by the call would be 12 times the size of an INTEGER in each node.

Another simple case is obtained by sending different amounts of data from each processor. Consider the following code:

```

      INTEGER SIZES(4), TYPE, NDDATA(4)
      DATA TYPE/56/
C
      CALL KXPARA(NDDATA)
      ISTAT = KXCONC(MYBUF, 4*NDDATA(1), IBUF, 3*4, SIZES,
$                IALNOD, 0, TYPE)

```

In this case the "result" buffers on each node would be

```
IBUF = {0, 0, 0, 32, 0, 0, 52, 53, 0, 72, 73, 74}
SIZES = {0, 4, 8, 12}
```

where we have assumed 4-byte integers. If we now change the arguments again by reducing the RESSZ parameter to 8 then the resulting buffers would be

```
IBUF = {0, 0, 32, 0, 52, 53, 72, 73}
SIZES = {0, 4, 8, 8}
```

In each node the call to KXCONC would now return an error, to reflect the fact that node 3 attempted to send more data than was to be read.

In the final example we perform the concatenation only in processors 0,1 and 3.

```

      INTEGER NNODES, NODEL(4), TYPE, SIZES(3)
      DATA TYPE/12/
C
      NNODES = 3
      NODEL(1) = 0
      NODEL(2) = 1
      NODEL(3) = 3
C
      ISTAT = KXCONC(MYBUF, 3*4, IBUF, 3*4, SIZES,
$                NNODES, NODEL, TYPE)

```

DIAGNOSTICS

If any error occurs in KXCONC -1 is returned. Possible sources of error are: illegal values of MYBUF or RESBUF and preposterous values of MYSIZE or RESSZ. If no error occurs the total number of bytes stored in memory is returned. An error condition can also be generated if the value of RESSZ on a node is smaller than the amount of data which is being

sent by a processor (including the node itself).

SEE ALSO

KXBROD, KXCOMB

KXCUST

NAME

KXCUST - Indicate an alternative system configuration file.

SYNOPSIS

```
INTEGER FUNCTION KXCUST(FNAME)
CHARACTER*80 FNAME
```

DOMAIN

Only available to host programs.

DESCRIPTION

KXCUST indicates that *Express* should use system configuration information from the named file rather than the system default. This allows applications to maintain their own customization programs independent of any other user or system requirements.

To complete the customization process the `exinit` command has an optional argument which names the customization file which should be used while loading *Express* into the transputer system. Similarly `cubix` has an additional '-E' switch allowing an alternative file to be named at runtime. In both cases the KXCUST function is invoked with the named file as argument.

The KXCUST call must be made before any other *Express* system calls.

RETURN VALUE

The returned value indicates whether or not the indicated customization file was found. Non-zero values indicate a failure to locate the named file.

EXAMPLES

The following code fragment could be used to allow a program to use an alternative customization file.

```
PROGRAM MYCUST
C
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
CALL KXINIT
C
ISTAT = KXCUST('myfile.cst')
IF(ISTAT .NE. 0) THEN
    WRITE(6,*) 'Failed to find customization file'
    STOP
ENDIF
```

SEE ALSO

excustom (command).

KXGRID

NAME

KXGRID - Automatic decomposition tools

SYNOPSIS

```
INTEGER FUNCTION KXGDIN (GRDDIM, NPROCS)
INTEGER GRDDIM, NPROCS (*)
```

```
INTEGER FUNCTION KXGDSP (NODES, GRDDIM, NSPLIT)
INTEGER NODES, GRDDIM, NSPLIT (*)
```

```
INTEGER FUNCTION KXGDCO (PROCNO, COORD)
INTEGER PROCNO, COORD (*)
```

```
INTEGER FUNCTION KXGDPR (COORD)
INTEGER COORD (*)
```

```
INTEGER FUNCTION KXGDSI (PROCNO, GLOBAL, SIZE, START)
INTEGER PROCNO, GLOBAL (*), SIZE (*), START (*)
```

```
INTEGER FUNCTION KXGDBC (PERBC)
INTEGER PERBC (*)
```

```
INTEGER FUNCTION KXGDNO (PROCNO, DIR, DIST)
INTEGER PROCNO, DIR, DIST
```

DOMAIN

The KXGRID routines may be called from any *Express* program.

DESCRIPTION

KXGRID collectively refers to a set of utilities that perform automatic decompositions of user domains onto the underlying machine topology. A user specification for a problem domain which has the topology of a Cartesian grid in N dimensions is mapped onto the hardware topology and routines are available to enable processors (defined relative to the user topology) to communicate through the primitive system calls.

KXGDIN is the routine which performs the elementary mapping and must be called before any of the other KXGRID routines (except KXGDSP). The arguments are the number of dimensions in the user topology and the number of processors to be assigned to each dimension. If the requested topology is successfully mapped to the hardware zero is returned; otherwise the value returned is -1.

The function KXGDSP is used to divide up the NODES processors between the GRDDIM dimensions in as even a way as possible consistent with the requirement that all processors be used. The number in each dimension will be returned in the array NSPLIT. A simple example would be that of two dimensional decompositions: for eight nodes we would obtain an 4 x 2 decomposition while nine processors yields 3 x 3.

Having set up the KXGRID system in this way the other function calls are available to inquire about specific details of the decomposition. Particularly useful is information concerning where, in the user defined topology, a certain processor is to be found. The KXGDCO function call takes a processor number as argument and returns the coordinates in the Cartesian grid of this processor. The inverse transformation is provided by the function KXGDPR which takes as arguments an array of coordinates and returns the processor number of the node at that position in the user grid.

The interface to the underlying communication structure is provided by the KXGDNO and KXGDBC functions. The arguments to the former are a processor number, a direction in the user grid and a distance. The returned value is a "NODE" suitable for use in calls such as KXCHAN and KXVCHA which contains the necessary information for communication in that direction. The distance parameter specifies the offset from the current node in the direction indicated so that a value of +1 implies the next node along the positive axis while -1 indicates the next node in the negative direction. Magnitudes greater than 1 are also possible and correspond to multiple hops in the given direction.

The KXGDBC function is provided to alter the boundary conditions at the edges of the user domain. By default KXGDNO assumes that boundaries are connected periodically so that the processor to the "left" of the leftmost is the one on the extreme right hand edge. To suppress this feature one uses KXGDBC. The sole argument is an array of integers, one for each dimension in the user domain. A non-zero value indicates that this dimension is to be considered periodic while a zero value causes KXGDNO to return a NONODE at the boundary.

The last function in the KXGRID collection, KXGDSI, is used to distribute an array over the user grid. The first argument is again a processor number and the second is an array containing the global sizes of the array to be decomposed. After the call the third argument will be an array containing the number of entries in each dimension of the array which lie in the processor specified. The final argument will be an array containing the global index that corresponds to an index of zero in the local array.

A final point to note is that these routines are very useful in conjunction with the low level I/O primitives KMREAD, KMWRIT, KMRD2D and KMWT2D which require arguments easily calculated by the KXGRID functions.

EXAMPLE

As a simple example consider a problem involving two dimensional images to be executed on eight processors. A suitable call to initialize the system might be

```

      INTEGER NPROCS(2), TYPE PARAMETER (IHORIZ=0, IVERT=1)
      DATA TYPE/33/
C
      NPROCS(IHORIZ) = 4
      NPROCS(IVERT) = 2
C
      ISTAT = KXGDIN(2, NPROCS)
      IF(ISTAT .LT. 0) THEN

```

```

      STOP
    ENDIF

```

The macros IHORIZ and IVERT are defined for our convenience and just serve to label the two axes in the grid. We assign four processors to the horizontal dimension and two to the vertical. (A more flexible assignment scheme is easily devised using the KXPARA system call to determine at runtime the number of processors available.)

Now consider a simple scrolling operation in which data is to be passed to the right. We need to figure out the processor numbers of the appropriate nodes in order to communicate in this direction. The simple thing to do in this case is to use KXGDNO to calculate the appropriate values. However, one must first consider the boundary values; What should happen when data is scrolled off the right hand edge of the display? The two options are to have it appear on the left hand edge, or to disappear completely. We adopt the latter approach which entails altering the default assumption of KXGDNO that boundaries are periodic. The following code uses KXGDBC to override this default and KXGDNO to assign suitable processor values for the four directions we will be interested in.

```

      INTEGER PERBC (2)
      INTEGER NDDATA (4)
C
      CALL KXPARA (NDDATA)
C
      PERBC (1) = 0
      PERBC (2) = 0
      CALL KXGDBC (PERBC)
C
      UNODE = KXGDNO (NDDATA (1), IVERT, 1)
      DNODE = KXGDNO (NDDATA (1), IVERT, -1)
      LNODE = KXGDNO (NDDATA (1), IHORIZ, -1)
      RNODE = KXGDNO (NDDATA (1), IHORIZ, 1)

```

Now all the “nodes” are valid. If a processor is on the extreme left edge of the domain and it tries to communicate with a processor to its left then the value of LNODE has been correctly assigned the value NONODE which will, in turn, direct the communication system to omit communication with this non-existent processor. Note how simple it would be to adopt the alternative strategy and have data scroll off the right edge and re-appear on the left. We simply omit the call to KXGDBC (or else change the zero values to ones) and the correct values would be returned.

To show the actual use of these processor numbers assume that we wish to “scroll” 512 bytes along to the right. In each processor the data is to be taken from an array OBUF and the data coming in from the left is to be read into an array IBUF. The following call to KXCHAN is all that is required

```

      ISTAT = KXCHAN (IBUF, 512, LNODE, TYPE,

```

```
      OBUF, 512, RNODE, TYPE)
```

Notice that at no point in these calculations did the topology of the hardware enter. Everything is specified in the user domain - i.e., that of the image, and KXGRID does the rest.

To demonstrate the use of the KXGDSI function assume that the image to be "scrolled" is not 1024 bytes tall as was implicitly assumed in the previous code (We scrolled 512 bytes left in each processor and there are two processors in the vertical direction for a total of 1024 bytes.) Instead we will make the strange choice of an image which is 767 bytes high, and 1024 bytes wide. The KXGDSI routine can then be used to tell us how many elements are in each processor through the following code

```
      INTEGER GLOBAL(2), SIZES(2), START(2)
C
C-- Decompose the array over the processor ring.
C
      GLOBAL(IHORIZ) = 1024
      GLOBAL(IVERT) = 767
      ISTAT = KXGDSI(NDDATA(1), GLOBAL, SIZES, START)
```

At the completion of this call the values SIZES(1) and SIZES(2) contain the sizes of the subregions assigned to each processor. Further, the values START(1) and START(2) contain the horizontal and vertical index of the first byte that is stored in this processor. In the case described here every processor would have the value 256 for SIZES(1) since the horizontal size is divided exactly by the number of processors in that direction. In the vertical direction, however, the division does not work out correctly and so the processors whose responsibility is the lower half of the display would have SIZES(2) = 384 while those in the upper half would have 383. Similarly, the processors in the upper half have START(2) = 0 while those in the lower half have START(2) = 384. The modified call to KXCHAN which scrolls the data to the right is

```
      ISTAT = KXCHAN(IBUF, SIZES(2), LNODE, TYPE,
                   OBUF, SIZES(2), RNODE, TYPE)
```

RETURN VALUE

If any error occurs in the KXGRID routines they return -1. Particular errors include failing to call KXGDIN before using the other functions and a failure of KXGDIN to match the user requested topology onto that of the hardware.

SEE ALSO

KXPARA, KMREAD, KMRD2D, KORDER

NAME

KXHAND - Asynchronous message handler.

SYNOPSIS

```
INTEGER FUNCTION KXHAND (FUNC, SRC, TYPE)
INTEGER FUNC, SRC, TYPE
EXTERNAL FUNC
```

DOMAIN

KXHAND may be called in the node processors only.

DESCRIPTION

This routine is used to initialize a “handler” for messages of certain types and sources. The idea is that whenever a message arrives that matches the SRC and TYPE parameters the user-supplied procedure FUNC is invoked to process the data. This process occurs immediately upon receipt of the message with as little overhead as possible and can be used to implement a totally asynchronous processing style in which messages can be handled transparently without the intervention of the main application code.

The FUNC is invoked immediately a message has arrived in the internal node buffers with the following arguments

```
FUNC (PTR, LENGTH, SRC, TYPE)
INTEGER PTR (*)
INTEGER LENGTH, SRC, TYPE
```

i.e., it looks just like a call to KXREAD. Note however, that the supplied PTR argument actually points to a buffer within the *Express* kernel. If the application needs to keep the message for later processing memory must be allocated and the buffer copied. Otherwise the data becomes unavailable when the user function completes.

The SRC and TYPE fields reflect the actual source and type of the message being handled in cases where “NOCARE” values were originally supplied to the KXHAND function.

The user supplied function must return an integer value to its caller. This value will determine the future behavior of the system; a negative value will terminate the association between the message source/type and the function while positive (and zero) values maintain the *status quo*. In this way it is possible to have a message handler that is invoked only once, several times until a particular message arrives, or permanently.

EXAMPLE

The following example shows how this function can be used to implement a global, “read only” memory. A handler is set up which intercepts all messages of type MEMRD and responds by sending back a message containing the memory requested. Obviously one could implement a writable shared memory in a similar manner although problems concerning mutual exclusion would probably have to be addressed.

```

PROGRAM MYTEST
PARAMETER (MEMRD=10, MEMDAT=11)
EXTERNAL MEMHND
INTEGER MEMTYP, SRC
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Setup Express and initialize its common block.
C
CALL KXINIT
C
C-- Allow anyone to send memory requests.
C
MEMTYP = MEMRD
SRC = NOCARE
ISTAT = KXHAND(MEMHND, SRC, MEMTYP)
CALL KXSYNC
.....
C-- This is the function that fields requests for memory.
C-- The first argument will point to an array containing
C-- the address and number of bytes to read.
C-- NOTE: we return 0 so that the handler continues to
C-- operate.
C
INTEGER FUNCTION MEMHND(REQ, LENGTH, SRC, TYPE)
INTEGER REQ(2)
INTEGER LENGTH, SRC, TYPE
C
INTEGER RTYPE
PARAMETER (MEMRD=10, MEMDAT=11)
DATA RTYPE /MEMDAT/
C
ISTAT = KXWRIT(REQ(1), REQ(2), SRC, RTYPE)
MEMHND = 0
RETURN
END

```

Having set up this message handler we can access memory on another node by simply sending a message of type MEMRD. Notice that the message handler sends back the data in a message of a different type that it read. This is an important point - if the routine adopted the simpler strategy of returning the same type message as it received then that message would be trapped by the message handler on the original node and treated as a memory request. In this way an infinite chain of requests would be generated!

The following routine reads LENGTH bytes of memory from processor NODE and stores it

in the specified BUFFER. The routine returns the number of bytes read.

```
      INTEGER FUNCTION RDMEM(NODE, ADDR, LENGTH, BUFFER)
      INTEGER NODE, ADDRESS, LENGTH, BUFFER(*)
C
      PARAMETER (MEMRD=10, MEMDAT=11)
      INTEGER REQ(2), STYPE, RTYPE
C
      REQ(1) = ADDR
      REQ(2) = NODE
C
      STYPE = MEMRD
      RTYPE = MEMDAT
C
      ISTAT = KXWRIT(REQ, 8, NODE, STYPE)
      RDMEM = KXREAD(BUFFER, LENGTH, NODE, RTYPE)
      RETURN
      END
```

This function forms the basis of an extremely elegant multitasking system under *Express* which is discussed in more detail in the accompanying manual, "Multitasking under *Express*".

DIAGNOSTICS

If the kernel is unable to install the message handler -1 is returned. Otherwise the return value will be 0.

WARNING

The current implementation restricts the length of a message that can be sent to a handler to the "packet size" as specified in the customization procedure, *excustom*.

SEE ALSO

KXREAD, KXRECV

NAME

KXINIT - Start Express system.

SYNOPSIS

SUBROUTINE KXINIT

DOMAIN

Available to host and node programs.

DESCRIPTION

This routine **MUST** be the first *Express* routine called in both host and node programs. It serves to initialize the internal state of *Express* and also to set up a common block containing useful parameters for use by application codes.

The system common block has the name XPRESS and is defined as follows:

```
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
INTEGER NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
```

The various variables in this block are as follows

NO CARE	Used to indicate that a message should be read from any node or that its type is of no concern. Note that you should not attempt to send a message to destination NO CARE or with type NO CARE.
NORDER	Used by <i>Cubix</i> programs to indicate the default ordering of output in "multi"-mode I/O.
NONODE	Used by the KXGRID utilities to indicate that no processor lies in the indicated position of the user topology.
IHOST	Special "node" value used to send messages to or receive them from the host processor which loaded the node program.
IALNOD	Used in KXBROD, KXCOMB, KXCONC etc. operations to indicate that all nodes should be involved in a particular communication operation.
IALPRC	Used in KXBROD to indicate that the host processor should be included as a recipient of a broadcast message.

EXAMPLE

The following schematic code should be the general template of any host or node program which uses *Express*.

```
PROGRAM MYTEST
INTEGER PG1
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Set up Express
```

KXINIT

```
C
    CALL KXINIT
C
C-- Start application code.....
C
    ....

    CALL KXCLOS(PG1)
    STOP
    END
```

SEE ALSO

KXOPEN, KXSHAR.

NAME

KXLOAD - Load a program.

SYNOPSIS

```
INTEGER FUNCTION KXLOAD (PGIND, NDPROG)
INTEGER PGIND
CHARACTER*80 NDPROG
```

DOMAIN

KXLOAD may only be called in the host computer.

DESCRIPTION

KXLOAD loads the program NDPROG into a set of processors previously allocated with KXOPEN. The PGIND argument is the *processor group index* returned by the KXOPEN call.

The KXLOAD function provides the simplest interface to allocating processors and loading application programs. A single application code is loaded into all processors. The alternative call KXPLOA is provided if different programs are to be loaded into different processors.

EXAMPLES

The following code loads a program (called MYPROG) into four processors.

```
PROGRAM EXPTST
INTEGER PGIND
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
CHARACTER*80 DEVICE, PROG
PARAMETER (DEVICE='/dev/ncube', PROG='myprog')
C
C-- Set up Express and its common block.
C
CALL KXINIT
C
ISTAT = KXOPEN(DEVICE, 4, NOCARE)
IF(ISTAT .LT. 0) THEN
WRITE(6,*) 'Failed to allocate processors'
STOP
ENDIF
C
ISTAT = KXLOAD(PGIND, PROG)
C
...
```

KXLOAD

DIAGNOSTICS

If any error occurs during loading -1 is returned. Possible sources of error are: an illegal value of PGIND or the failure of the system to allocate a the correct number of processors. Errors are also returned if a routine fails to find an appropriate executable to load or if a communication error occurs during loading.

SEE ALSO

KXPLOA

NAME

KXOPEN - Allocate a processor group.

SYNOPSIS

```
INTEGER FUNCTION KXOPEN (DEVICE, NODES, START)
CHARACTER*80 DEVICE
INTEGER NODES, START
```

DOMAIN

Only available to host programs.

DESCRIPTION

KXOPEN allocates a processor group containing NODES processors in the device pointed to by the character string DEVICE.

The NODES argument indicates the number of nodes to be allocated and the last argument optionally requests a specific set of nodes within the parallel machine. The default value NOCARE allows any group of nodes to be selected.

The KXOPEN call must be used before attempting to access any processor group.

RETURN VALUE

The returned value is a *processor group index* which must be used in all further references to the allocated processors. In cases where no processor group of the appropriate size is available or some other hardware error occurs the value returned is -1.

EXAMPLES

The following code allocates a group of 4 processors anywhere in the parallel machine.

```
PROGRAM MYTEST
INTEGER PGIND
COMMON/XPRESS/NO CARE,NORDER, NONODE, IHOST, IALNOD, IALPRC
CHARACTER*80 DEVICE
PARAMETER (DEVICE='/dev/transputer')
C
C-- Setup Express and its common block.
C
CALL KXINIT
C
PGIND = KXOPEN (DEVICE, 4, NOCARE)
IF (PGIND .LT. 0) THEN
WRITE (6,*) 'Failed to allocate processors'
STOP
ENDIF
```

KXOPEN

If we wanted to allocate a particular set of processors in the machine then we could replace, for example, the NOCARE argument in the KXOPEN call:

```
ISTAT = KXOPEN('/dev/transputer', 4, 8)
```

which attempts to allocate nodes 8 thru 11. In this case it is even more important that the value returned by KXOPEN be checked since there is a larger chance of failure.

WARNINGS

In current implementations the DEVICE parameter will be one of

/dev/transputer	Transputer based devices
/dev/ncube	NCUBE systems.
/dev/symult	Symult S2010, 68000 based nodes
/dev/symfpa	Symult S2010, Weitek VFPA nodes

Note that this list is not necessarily exhaustive. It is complete at the time of writing but may be extended at future dates.

Some systems are unable to support the `start_node` argument to this function.

SEE ALSO

KXSHAR, KXLOAD, KXREAD, KXTEST, KXTYPE, KXWRIT

NAME

KXPARA - Runtime parameters.

SYNOPSIS

```
SUBROUTINE KXPARA (NDDATA)
INTEGER NDDATA (*)
```

DOMAIN

KXPARA may be called in either the host or node processors.

DESCRIPTION

This routine is used when an application program requires to know the details of its runtime environment. The information available and its correspondence to the array elements returned is

- NDDATA (1) Processor number of the calling node. Nodes are numbered consecutively from (and including) 0.
- NDDATA (2) Number of processors allocated in this processor group.
- NDDATA (3) Specifies the *processor group index* containing this node.
- NDDATA (4) Specifies the *process identifier* of the process making the call.

The last two pieces of information are currently unused.

The use of this information and the KXGRID utilities is the key to writing reconfigurable applications since they allow the program to adapt to different processor configurations at runtime.

EXAMPLE

Assume that we wish to use the KXGRID tools to map the parallel machine to a two dimensional mesh of processors. The following code supplies the necessary parameters to the KXGRID routines.

```
PROGRAM MYTEST
  INTEGER NDDATA (4)
  INTEGER NPROCS (2)
C
C-- Set up Express.
C
  CALL KXINIT
C
C-- Get runtime parameters.
C
  CALL KXPARA (NDDATA)
C
C-- Divide up processors in two dimensional mesh. Set up
```

KXPARA

C-- the KXGRID routine with this decomposition.

C

```
      ISTAT = KXGDSP (NDDATA(2), 2, NPROCS)
      ISTAT = KXGDIN(2, NPROCS)
      IF (ISTAT .LT 0) THEN
          STOP
      ENDIF
```

Note that we use the KXGDSP function to divide up the processors between the physical dimensions.

SEE ALSO

KXGRID

NAME

KXPAUS - Arrange for programs to be loaded "stopped".

SYNOPSIS

SUBROUTINE KXPAUS

DOMAIN

Only available to host programs.

DESCRIPTION

This routine is used to control the initial state of a program or programs being loaded into groups of processors. By default node programs start immediately. If KXPAUS is used before the appropriate KXLOAD call then the programs will halt at their first instruction after loading. This is useful when using the debugger, ndb, since it allows the user to control the entire course of execution by setting breakpoints etc.

EXAMPLE

Consider the case where debugging is occasionally required. The following code segment illustrates the use of KXPAUS to load programs in a stopped state if the number of processors entered is negative. Otherwise programs will be loaded in the (default) running state.

```
PROGRAM MYTEST
INTEGER NNODES, PGIND
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
CHARACTER*80 DEVICE, PRGNAM
PARAMETER (DEVICE='/dev/transputer', PRGNAM='myprog')
C
C--Set up Express and its common block.
C
CALL KXINIT
C
WRITE(6,*) 'How many nodes ? (Negative ==> stopped)'
READ(5,*) NNODES

IF(NNODES .LT. 0) THEN
    NNODES = -NNODES
    CALL KXPAUS
ENDIF
ISTAT = KXOPEN(DEVICE, NNODES, NOCARE)
IF(PGIND .LT. 0) THEN
    WRITE(6,*) 'Failed to allocate nodes'
    STOP
ENDIF
```

KXPAUS

```
C  
C-- Finally load application program into nodes.  
C  
    ISTAT = KXLOAD(PGIND, PRGNAM)
```

SEE ALSO

KXOPEN, KXSHAR.

NAME

KXPCP, KXPELT - Dump execution profile data.

SYNOPSIS

```
SUBROUTINE KXPCP
```

```
  SUBROUTINE KXPELT(FNAME)
```

```
  CHARACTER*80 FNAME
```

DOMAIN

KXPCP may only be called in the host processor while KXPELT may only be called in the nodes.

DESCRIPTION

These routines are used to dump the execution profile data collected with the KXPROF functions. For each call to KXPELT on the nodes there must be a call to KXPCP in the host processor. The profiling data will be written to a file on the host with the name FNAME supplied in the node program.

In addition to dumping out the profile data KXPELT also turns off the profiler and resets its internal state so that further invocations of the execution profiler will begin from the zero state and hence be totally independent.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the execution profiler.

1. Host Program

```
      PROGRAM HSTXPR
      C
      C-- Start Express.
      C
            CALL KXINIT
      C
      C-- Allocate nodes, load programs.
      C
            ...
      C
      C-- Execute algorithm .....
      C
            ...
      C
      C-- Dump profiling data.
      C
            CALL KXPCP
```

```
C
  STOP
  END
```

2. Node Program

```
PROGRAM NODXPR
C
  INTEGER PRFBUF(2048), PRFSCL
  PARAMETER (PRFSCL = 8192)
C
C-- This is the name of a function found to live at the
C-- low end of memory. This information can usually be
C-- found in the "linker map".
C
  EXTERNAL F_MAIN
C
C-- Start off profiler.
C
  CALL KPROFI(PRFBUF, 8192, F_MAIN, PRFSCL)
  CALL KXPON
C
C-- Application code, profiler running.
C
  ...
C
C-- Program over, dump data and exit.
C
  CALL KXPELT('xprof.out')
  STOP
  END
```

SEE ALSO

xtool (command), KPROFI, KXPROF, KXPEND

NAME

KXPINQ, KXPEND - Manipulate execution profiler under *Cubix*.

SYNOPSIS

INTEGER FUNCTION KXPINQ()

SUBROUTINE KXPEND

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

These routines provide a simple control interface to the execution profiler for applications running in the *Cubix* environment.

KXPINQ returns an integer value representing the state of the “-mx” runtime switch on the cubix command line. This can be used to conveniently enable/disable the profiling system at runtime. Consider a typical command

```
cubix -n 4 toyland 1024 1024 <noddy.dat
```

Since no “-m” switch is present a call to KXPINQ will return zero. If we modify the above command to

```
cubix -mcxe -n 4 toyland 1024 1024 <noddy.dat
```

then the return value would be 1 since the character ‘x’ appears in the monitoring switch, “-m”.

KXPEND is used to finally dump profiling data to the host computer file system. A file called “xprof.out” is created for later analysis with the xtool utility. In addition the profiler is disabled and its initial state reset to zero. This allows distinct phases of an application to be profiled totally independently.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the execution profiler.

```

PROGRAM XPRST
C
  INTEGER PRFBUF(2048), PRFSCL
C
C-- This value is 0x2000 (hexadecimal)
C
  PARAMETER (PRFSCL = 8192)
C
C-- This is the name of a function in the program, low
C-- in memory. A suitable candidate can usually be found

```

KXPEND

```
C-- by looking through the "linker map".
C
    EXTERNAL F_MAIN
C
C-- Start up Express.
C
    CALL KXINIT
C
C-- Start up profiler if user selected -mx option.
C
    ISTAT = KXPINQ()
    IF(ISTAT .NE. 0) THEN
        CALL KPROFI(PRFBUF, 8192, F_MAIN, PRFSCL)
        CALL KXPON
    ENDIF
C
C-- Execute application code with profiler running.
C
    ...
C
C-- Program over, dump data and exit.
C
    IF(ISTAT .NE. 0) THEN
        CALL KXPEND
    ENDIF
    STOP
    END
```

SEE ALSO

xtool(command), KPROFI, KXPCP, KXPROF

NAME

KXPLOA - Load a program into individual nodes.

SYNOPSIS

```
INTEGER FUNCTION KXPLOA(PGIND, PROG, NODE)
INTEGER PGIND, NODE
CHARACTER*80 PROG
```

DOMAIN

This routine may only be called in the host processor.

DESCRIPTION

KXPLOA provides a complementary interface to the KXLOAD routine for loading programs into groups of processors. Instead of loading the entire array with a single node program this routine allows different applications to be loaded into individual nodes of the machine.

In each case a previous call to KXOPEN must have allocated a set of processors into which we are attempting to load programs. The *processor group index* returned by this call must be supplied to the KXPLOA functions as the argument PGIND.

Having allocated a group of nodes user applications are loaded with the KXPLOA primitive which loads the named code into the processor specified by the NODE argument. The special value IALNOD defined in the XPRESS common block specifies that all processors are to be loaded with the same item.

Before execution of the node program can begin calls must be made to the KXMAIN function.

EXAMPLES

The following calls allocate, load and start a program in four processors

```
PROGRAM MYTEST
INTEGER PGIND
COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
CHARACTER*80 DEVICE, PRGNAM
PARAMETER (DEVICE='/dev/transputer', PRGNAM='myprog')
C
C-- Initialize Express and its common block.
C
CALL KXINIT
C
PGIND = KXOPEN(DEVICE, 4, NOCARE)
IF(PGIND .LT. 0) THEN
WRITE(6,*) 'Failed to allocate nodes'
STOP
ENDIF
```

```
C
  ISTAT = KXPLOA(PGIND, PRGNAM, IALNOD)
  CALL KXMAIN(PGIND, IALNOD)
```

...

Note that the particular arguments chosen here make this code functionally equivalent to a call to KXLOAD.

In the following example we load the programs "prog1" into nodes 0 through 3 and "prog2" into nodes 4 through 15 of a sixteen processor group.

```
PROGRAM MYTEST
  INTEGER PGIND
  COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
  CHARACTER*80 DEVICE, PROG1, PROG2
  PARAMETER (DEVICE='/dev/transputer')
  PARAMETER (PROG1='prog1', PROG2='prog2')
C
C-- Initialize Express and its common block.
C
  CALL KXINIT
C
  PGIND = KXOPEN(DEVICE, 4, NOCARE)
  IF(PGIND .LT. 0) THEN
    WRITE(6,*) 'Failed to allocate nodes'
    STOP
  ENDIF
C
  DO 10 I=0, 3
    ISTAT = KXPLOA(PGIND, PROG1, I)
10  CONTINUE
  DO 20 I=4, 15
    ISTAT = KXPLOA(PGIND, PROG2, I)
20  CONTINUE
C
  CALL KXMAIN(PGIND, IALNOD)
```

...

DIAGNOSTICS

KXPLOA returns zero upon successful loading of the executable program. If the executable file is not found, or is invalid in some way the value -1 is returned.

SEE ALSO

KXLOAD, KXSTAR, KXMAIN

KXPROF

NAME

KXPON, KXPOFF - Control execution profiler.

SYNOPSIS

```
SUBROUTINE KXPON
```

```
SUBROUTINE KXPOFF
```

DOMAIN

These routines may only be called from the nodes.

DESCRIPTION

KXPON is used to enable and start the execution profiler which must have been previously initialized with a call to KPROFI. Subsequently a periodically scheduled event occurs which causes the program counter of the user application to be "logged" in an internal structure. KXPOFF reverses this process - until a subsequent call to KXPON no execution profiling will be performed.

The profiler is initially off and must be explicitly enabled with calls to KPROFI and KXPON.

The log of profiling information is written to the host file system with KXPCP or KXPEND.

EXAMPLE

The following code is a skeleton of that which might typically be used to control the execution profiler.

```
PROGRAM XPTEST
C
  INTEGER PRFBUF(2048), PRFSCL
C
C-- This value is 0x2000 (hexadecimal)
C
  PARAMETER(PRFSCL = 8192)
C
C-- This is the name of a function found to be in
C-- low memory by perusing the "linker map".
C
  EXTERNAL F_MAIN
C
C-- Start Express.
C
  CALL KXINIT
C
C-- Start off profiler.
C
```

```
        CALL KPROFI (PRFBUF, 8192, F_MAIN, PRFSCL)
        CALL KXPON
C
C-- Application code, profiler running.
C
        ...
C
C-- Program finishes. Dump data and exit.....
C
        ...
        STOP
        END
```

SEE ALSO

xtool (command), KPROFI, KXPCP, KXPEND

NAME

KXREAD - Read a message

SYNOPSIS

```
INTEGER FUNCTION KXREAD (BUF, LENGTH, SRC, TYPE)
INTEGER LENGTH, SRC, TYPE
INTEGER BUF (*)
```

DOMAIN

KXREAD is available to both host and node programs with identical calling sequences.

DESCRIPTION

This call is used to read messages in the *Express* system. This routine provides the simplest interface to the message system - a blocking read; the function only returns when a suitable message has been read.

The accepted message is read into the buffer pointed to by the BUF argument and is truncated to size LENGTH bytes if necessary. The source and type of the message to be read can be specified by the arguments SRC and TYPE as described below.

This routine blocks until a message with suitable parameters has been received.

OPTIONS

Under *Express* messages have both destinations and types which are used by reading processes to distinguish between various available messages. A message will only be read if it matches, in both source and type, the parameters supplied in the read call. However, several options are available to allow the user extra flexibility. Both source and type fields are treated equivalently at this level so the following discussion applies equally to both.

SRC = NOCARE A message will be read from any node. The particular node will be indicated by modifying the value SRC. The value NOCARE is to be found in the XPRESS common block setup by the call to KXINIT.

SRC = number Any positive numeric value will restrict attention to messages with that particular source.

These same considerations apply to the type, TYPE, except that the interpretation of the wildcard value, NOCARE, is subject to modification through the KXTYPE system calls.

The special value IHOST is used by nodes wishing to send messages to the host processor. (This value is also to be found in the XPRESS common block.)

RETURN VALUE

The value returned is the length of the received message, after any necessary truncation has been performed. If some sort of hard error occurs then -1 is returned.

EXAMPLES

In the following examples we consider a case in which the following four messages have arrived on our node in the order given.

1. Source 1	Type 12	Length 32
2. Source HOST	Type 2	Length 512
3. Source 1	Type 15	Length 1024
4. Source 2	Type 0	Length 0

The simplest case is where both source and type are explicitly stated as in the call

```
ISTAT = KXREAD (BUFFER, 512, 1, 15)
```

In this case message three will be accepted for reading. Note, however, that the actual message is longer than the request length so only the first 512 bytes will be read and the rest discarded. The returned value, ISTAT will be 512.

The next example uses the wildcard value, NOCARE, to read a message but retain information about its source.

```
INTEGER SOURCE
SOURCE = NOCARE
ISTAT = KXREAD (BUFFER, 512, SOURCE, 0)
```

In this case the type is explicitly given and so message 4 will be read. The returned value will be 0, the length of the message read and the SOURCE variable will contain 2, the source of the message.

In the last example a wildcard value is given for the type field.

```
SOURCE = 1
TYPE = NOCARE
ISTAT = KXREAD (BUFFER, 512, SOURCE, TYPE)
```

In this case the source is given explicitly and the type allowed to take any value. With the parameters shown message 1 will be read and the value 12 stored in the TYPE variable. 32 bytes will be copied into the user buffer and the same value returned as STAT. Note that types are subject to extra processing through the KXTYPE commands. If we had specifically excluded type 12 from consideration then message 3 would have been read instead since it has the correct source and has not been excluded. If we had excluded both types 12 and 15 then the call to KXREAD would block until a more suitable message arrived.

WARNINGS

Types are restricted to be positive integers less than 16384. Other message types are reserved for use within the *Express* kernel.

One very common “bug” concerns the use of the NOCARE parameter when reading messages. Consider a situation where one needs to loop over all processors reading a single message from each, in any order. The following code is incorrect:

```
C-- Attempt to read from each node, in any order
```

KXREAD

```
C-- INCORRECT CODE
C
  TYPE = 124
  NODE = NOCARE
  DO 10 I = 1, NPROCS
    ISTAT = KXREAD (BUF, 128, NODE, TYPE)
  10 CONTINUE
```

The error in this code lies in the fact that the receipt of the first message in the loop overwrites the value of the `NODE` variable. As a result the second call to `KXREAD` attempts to read from the same node that responded in the first cycle rather than any node as was desired. The simple solution to the problem is to move the assignment `NODE = NOCARE` variable inside the loop.

SEE ALSO

`KXOPEN`, `KXSHAR`, `KXTEST`, `KXWRIT`, `KXTYPE`, `KXGRID`.

NAME

KXRCV - Non-blocking read function.

SYNOPSIS

```

INTEGER FUNCTION
      KXRCV(BUFFER, LENGTH, SRC, TYPE, STATUS)
INTEGER LENGTH, SRC, TYPE, STATUS
INTEGER BUFFER(*)

```

DOMAIN

KXRCV may be called in only the node processors.

DESCRIPTION

This function provides a non-blocking read function for *Express* messages. It is intended for use in applications such as "double-buffering" in which one wishes to process some data while waiting for another message to arrive.

When called it looks for a message in the buffers that match the supplied SRC and TYPE parameters. If such a message exists it is read as though by a normal call to KXREAD and the STATUS value will contain the message length.

If no message exists which matches the requested parameters the value -1 is written under the STATUS flag and the function immediately returns to its caller. When a message of the correct type and source subsequently arrives it will be read into memory at the address BUFFER and the length will be written under the STATUS variable replacing the -1. The SRC and TYPE variables will also be updated at that time to reflect the newly read message.

The interpretation of the first four arguments is exactly as in the corresponding call to KXREAD. The last argument, STATUS, is a mechanism by which one can poll for the arrival of the requested message; while negative, no message has been received.

EXAMPLE

The following example is a sketch of a typical "double-buffered" application. We assume that processor SOURCE is sending messages of type PROCES which must be passed to the function GRIND for processing. When all messages for such treatment have been received a message of type DONE will be sent. We assume that each of the PROCES messages will be of no more than 1024 bytes.

```

      SUBROUTINE DOGRIN(NODE)
      INTEGER NODE
C
      INTEGER PROCES, DONE
      PARAMETER (PROCES=10, DONE=11)
C
      INTEGER BUFFER(1024,2)
      INTEGER STOP, TYPE, THIS, NEXT

```

```
      INTEGER STAT(2)
C
      COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
      STOP = 0
      THIS = 1
      NEXT = 2
C
C-- Get first buffer, blocking read this time. We assume
C-- that someone else called KXINIT to set up the XPRESS
C-- common block.
C
      TYPE = NO CARE
      STAT(THIS) = KXREAD(BUFFER(1,THIS), 1024, NODE, TYPE)

      STOP = 0
20  CONTINUE

      IF(TYPE .NE. DONE) THEN
          TYPE = NO CARE
          ISTAT = KXRCV(BUFFER(1,NEXT), 1024, NODE, TYPE,
$                   STAT(NEXT))
      ELSE
          STOP = 1
      ENDIF

      CALL GRIND(BUFFER(1,THIS), STAT(THIS))
C
C-- If we've not finished then now is the time to poll for
C-- the next buffer to arrive.
C
      IF(STOP .EQ. 0) THEN
30  CONTINUE
          IF(STAT(NEXT) .LT. 0) GOTO 30
          NEXT = MOD(NEXT, 2) + 1
          THIS = MOD(THIS, 2) + 1
      ENDIF

      IF(STOP .EQ. 0) GOTO 20
      RETURN
      END
```

There are several points to note in this code. We assume that we must process the buffer with the DONE type - this saves us a message since we can send valid data and still use the type field to convey the important information. We also save the length of the message we are going to process in the STAT variable - this could be important in the GRIND function.

Note that it would be dangerous to use a single variable here since it would get overwritten whenever the second buffer arrived - possibly before the call to GRIND had been passed the value. Finally note that we have to keep setting TYPE = NOCARE since its value is overwritten whenever a message comes. Failing to do this is quite a common error and would result in the failure to read the DONE message.

RETURN VALUE

This function returns zero unless insufficient memory is available to register the read function. In this case -1 is returned.

SEE ALSO

KXREAD, KXHAND

NAME

KXSEM - Various semaphore operations.

SYNOPSIS

INTEGER FUNCTION KXSEMI (SEMPTR)

INTEGER SEMPTR (2)

SUBROUTINE KXSEMW (SEMPTR)

INTEGER SEMPTR (2)

SUBROUTINE KXSEMS (SEMPTR)

INTEGER SEMPTR (2)

DOMAIN

This routine may only be called in node programs.

DESCRIPTION

These routines implement a semaphore mechanism essential to protect critical sections of code in a multitasking environment.

KXSEMI initializes a new semaphore in the array SEMPTR and sets it so that the first call to KXSEMW will not block. If the initialization attempt fails -1 will be returned, otherwise the result will be 0.

Each call to KXSEMW checks the status of the associated semaphore. If locked the calling task sleeps until another process unlocks the semaphore. While sleeping no CPU time is expended allowing other tasks to proceed.

The KXSEMS call unlocks the indicated semaphore allowing other processes to enter a critical section of code.

EXAMPLE

The following code could be used to implement a global shared memory system for a distributed memory machine. We will assume that the data being accessed is such that only one process can be allowed access at any one time. This would be the case where, say, extended records are being written to memory in which case the integrity of any particular record is crucial. We would not, for example, allow two processes to both write records simultaneously since they may each write half leaving inconsistent data.

To implement these ideas we need to register a message handler which will field the read/write requests. For simplicity we will use only one handler for both purposes and let the data sent indicate the requested operation. We will encode the various requests in a 3 element integer array with the elements identified as follows:

ARRAY (1) A code value which is either MEMRD for read requests or MEMWT
 for writes.

ARRAY (2) A memory address for the read/write operation.

ARRAY (3) A number of bytes to be read or written.

The necessary message handler is as follows

```

      INTEGER FUNCTION MEMHND (PTR, LENGTH, SRC, TYPE)
      INTEGER PTR(3), LENGTH, SRC, TYPE
      INTEGER RTYPE
C
      COMMON/MEMORY/MEMRD, MEMWT, MEMACC, MEMRSP, MEMDAT, MEMSEM
      INTEGER MEMSEM(2)
C
C-- Block other users from entering this section of code
C-- while we're doing things.
C
      CALL KXSEMW (MEMSEM)
C
C-- Check: is this a read request ?
C
      IF (PTR(1) .EQ. MEMRD) THEN
          RTYPE = MEMRSP
          ISTAT = KXWRIT (PTR(2), PTR(3), SRC, RTYPE)
      ENDIF
C
C-- Is it a write request ?
C
      IF (PTR(1) .EQ. MEMWT) THEN
          RTYPE = MEMDAT
          ISTAT = KXREAD (PTR(2), PTR(3), SRC, RTYPE)
      ENDIF
C
C-- Release the semaphore.
C
      CALL KXSEMS (MEMSEM)
C
      MEMHND = 0
      RETURN
      END

```

We have assumed in the above code that the call to KXHAND which sets up this handler is made elsewhere. Similarly the MEMSEM semaphore should be allocated before any use will be made of this routine.

To use these routines it is merely necessary to add the following calls.

```
C
C-- Function to read global memory using the message handler
C-- installed above.
C
  INTEGER FUNCTION RDMEM(BUFFER, LENGTH, NODE, ADDR)
  INTEGER BUFFER(*), LENGTH, NODE, ADDR
C
  INTEGER MSG(3)
  INTEGER TYPE, RTYPE
C
  COMMON/MEMORY/MEMRD, MEMWT, MEMACC, MEMRSP, MEMDAT, MEMSEM
  INTEGER MEMSEM(2)
C
  TYPE = MEMACC
  RTYPE = MEMRSP
C
C-- Build array to make memory request.
C
  MSG(1) = MEMRD
  MSG(2) = ADDR
  MSG(3) = LENGTH
C
  ISTAT = KXWRIT(MSG, 3*4, NODE, TYPE)
  RDMEM = KXREAD(BUFFER, LENGTH, NODE, RTYPE)
  RETURN
  END
C
C-- Function to write global memory using the message
C-- handler installed above.
C
  INTEGER FUNCTION WTMEM(BUFFER, LENGTH, NODE, ADDR)
  INTEGER BUFFER(*), LENGTH, NODE, ADDR
C
  INTEGER MSG(3)
  INTEGER TYPE, RTYPE
C
  COMMON /MEMORY/MEMRD, MEMWT, MEMACC, MEMRSP, MEMDAT, MEMSEM
  INTEGER MEMSEM(2)
C
  TYPE = MEMACC
  RTYPE = MEMRSP
C
C-- Build array to make memory request.
C
  MSG(1) = MEMWT
```

```
MSG(2) = ADDR
MSG(3) = LENGTH
C
  ISTAT = KXWRIT(MSG, 3*4, NODE, TYPE)
  RDMEM = KXWRIT(BUFFER, LENGTH, NODE, RTYPE)
  RETURN
  END
```

Notice that several potential improvements could be made to this code. In particular we could speed up the writing process by sending short amounts of data in the same message as invokes the MEMHND handler. (KXHAND can only deal with messages up to the system packet size so any extra could be sent in a second message.) A further bottleneck is due to the fact that we have a single semaphore protecting a large memory space on each node. It might be more practical to have separate semaphores protecting disjoint areas of memory so that fewer processes would have to "sleep".

SEE ALSO

KXSLEE, KXHAND

KXSEND

NAME

KXSEND - Non-blocking write function.

SYNOPSIS

```
INTEGER FUNCTION KXSEND (BUFFER, LENGTH, SRC, TYPE, STATUS)
INTEGER BUFFER (*), LENGTH, SRC, TYPE, STATUS
```

DOMAIN

KXSEND may be called in only the node processors.

DESCRIPTION

This function provides a non-blocking write function for *Express* messages. It is intended for use in applications such as "double-buffering" in which one wishes to process some data while waiting for another message to arrive or be sent.

This routine provides a mechanism by which a node can transmit a message and then carry on processing regardless of whether or not the message has actually been sent. Upon return from the kernel the STATUS variable is set to -1. When the message is finally processed this value will be changed to the number of bytes sent. Until this has happened the user should (probably) not alter the data in the message since it is unknown which bytes have been transmitted to the receiving node and which have yet to be sent.

The interpretation of the first four arguments is exactly as in the corresponding call to KXWRIT. The last argument, STATUS, is a mechanism by which one can poll for the final dispatch of the requested message; while negative, the message has still to be sent.

EXAMPLE

The following example is a sketch of a typical "double-buffered" application. We assume that processor SOURCE is sending messages of type PROCES which must be passed to the function GRIND for processing. When all messages for such treatment have been received a message of type DONE will be sent. We assume that each of the PROCES messages will be of no more than 1024 bytes.

```
      SUBROUTINE DOGRIN (SRC, DEST)
      INTEGER DEST, SRC
C
      INTEGER PROCES, DONE
      PARAMETER (PROCES=10, DONE=11)
C
      INTEGER BUFFER(1024, 3)
      INTEGER STOP, THIS, LAST, NEXT
      INTEGER STAT(3), TYPE(3)
C
      COMMON /XPRESS/ NOCARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
      STOP = 0
      LAST = -1
```

```
        THIS = 1
        NEXT = 2
C
C-- Get first buffer, blocking read this time. We assume
C-- that someone else called KXINIT to set up the XPRESS
C-- common block.
C
        TYPE(THIS) = NOCARE
        STAT(THIS) = KXREAD(BUFFER(1,THIS), 1024,
$                               SRC, TYPE(THIS))

        STOP = 0
20    CONTINUE

        IF(TYPE(THIS) .NE. DONE) THEN
            TYPE(NEXT) = NOCARE
            IF(LAST .GE. 0 THEN
80                CONTINUE
                    IF(STAT(LAST) .LT. 0) GOTO 80
            ENDIF
            ISTAT = KXRCV(BUFFER(1,NEXT),1024, SRC, TYPE(NEXT),
$                               STAT(NEXT))
        ELSE
            STOP = 1
        ENDIF

        CALL GRIND(BUFFER(1,THIS), STAT(THIS))
        ISTAT = KXSEND(BUFFER(1,THIS),STAT(THIS),
$                               DEST,TYPE(THIS),STAT(THIS))
C
C-- If we've not finished then now is the time to poll for
C-- the next buffer to arrive.
C
        IF(STOP .EQ. 0) THEN
30            CONTINUE
                IF(STAT(NEXT) .LT. 0) GOTO 30
                LAST = MOD(LAST, 3) + 1
                NEXT = MOD(NEXT, 3) + 1
                THIS = MOD(THIS, 3) + 1
        ENDIF

        IF(STOP .EQ. 0) GOTO 20
        RETURN
    END
```

There are several points to note in this code. We assume that we must process the buffer with the `DONE` type - this saves us a message since we can send valid data and still use the type field to convey the important information. We also save the length of the message we are going to process in the `STAT` variable - this could be important in the `GRIND` function. Note that it would be dangerous to use a single variable here since it would get overwritten whenever the second buffer arrived - possibly before the call to `GRIND` had been passed the value. Finally note that we have to keep setting `TYPE = NOCARE` since its value is overwritten whenever a message comes. Failing to do this is quite a common error and would result in the failure to read the `DONE` message.

RETURN VALUE

This function returns zero unless insufficient memory is available to register the write function. In this case -1 is returned.

SEE ALSO

`KXWRIT`, `KXHAND`, `KXRECV`

NAME

KXSHAR - Share a processor group with another process

SYNOPSIS

```
INTEGER FUNCTION KXSHAR(DEVICE, PID, NODES)
CHARACTER*80 DEVICE
INTEGER PID, NODES
```

```
INTEGER FUNCTION KXPID(UNIXID)
INTEGER UNIXID
```

DOMAIN

Host processor only.

DESCRIPTION

The KXSHAR routine allows two or more host processes to share access to the same processor group. The first argument, *DEVICE*, specifies which array contains the processor group to be shared and is interpreted exactly as in the KXOPEN call. The process ID of the process with which the processor group is to be shared must be specified by *PID*. Upon return the number of nodes in the shared processor group is written under the value *NODES*.

The most reliable source of information about process ID's is provided by the KXOPEN system call which reports the appropriate value. Similar information is often available from the *exstat* command. On UNIX machines the function KXPID is available whose argument is the UNIX process ID. The returned value is the *Express* process I.D, suitable for giving to the KXSHAR function.

RETURN VALUE

The value returned by KXSHAR is the *processor group index* which must be used in future references to the shared processors.

If the indicated process has terminated or is not using any processors itself the value -1 is returned.

EXAMPLE

The following code would be used if a second process wished to share the processor group currently assigned to the process with *process-ID* 349.

```
PROGRAM MYTEST
  INTEGER NNODES, MSGTYP, MSGSRC, PGIND
  COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
  CHARACTER*80 DEVICE
  PARAMETER (DEVICE='/dev/transputer')
C
C-- Set up Express and its common block.
```

KXSHAR

```
C
CALL KXINIT
C
PGIND = KXSHAR(DEVICE, KXPID(349), NNODES)
IF(PGIND .LT. 0) THEN
    WRITE(6,*) 'Failed to share, job may have ended'
    STOP
ELSE
    WRITE(6,*) 'Sharing ', NNODES, ' processors'
ENDIF
C
C-- Successfully shared nodes, restrict wildcard message
C-- types and start reading.
C
CALL KXINCT(123, 125)
C
MSGTYP = NOCARE
MSGSRC = NOCARE
ISTAT = KXREAD(BUFFER, 512, MSGSRC, MSGTYP)
```

Note that having successfully shared the nodes with process 349 we use the KXTYPE functions to restrict attention to the message types from 123 to 125. This allows us the freedom to use the wildcard NOCARE values in reading without clashing with the process whose nodes we are sharing.

WARNINGS

Communicating with shared groups of nodes is complicated by interactions between source and type fields specified using the NOCARE wildcard. This situation can be eased somewhat through the KXTYPE mechanisms which restrict the ranges indicated by wildcard values.

SEE ALSO

KXOPEN, KXTYPE

NAME

KXSLEE - Pause process.

SYNOPSIS

```
SUBROUTINE KXSLEE (MSECS)
INTEGER MSECS
```

DOMAIN

This routine may only be called in node programs

DESCRIPTION

This routine is used when a process needs to wait for an event without using CPU resources. The supplied argument is the minimum time to wait in microseconds. This routine should be used in multitasking applications where one task needs to wait for an event which will potentially be generated by another task on this node.

EXAMPLE

The following code makes use of the KXSLEE function to implement a global semaphore - i.e., a semaphore that can be used from any node. For definiteness we assume that the physical semaphore is located on node 0. In this node we should register the following function with a call to KXHAND.

```

      INTEGER FUNCTION SEMHND (MSG, LENGTH, NODE, TYPE)
      INTER MSG, LENGTH, NODE, TYPE
C
      COMMON /SEMDAT/WAIT, SIGNAL, OPEN, CLOSED, SEMREQ, SEMRSP
      INTEGER WAIT, SIGNAL, OPEN, CLOSED, SEMREQ, SEMRSP
      COMMON /SEMSYS/ GBLSEM, GBLFLG
      INTEGER GBLSEM(2), GBLFLG
C
      INTEGER RESP, RTYPE
C
      RTYPE = SEMRSP
C
C-- There are two types of requests which basically
C-- correspond to the KXSEMW and KXSEMS calls on local
C-- semaphores.
C
      IF (MSG .EQ. WAIT) THEN
          CALL KXSEMW (GBLSEM)
          IF (GBLFLG .EQ. OPEN) THEN
              GBLFLG = CLOSED
              RESP = OPEN
          ELSE

```

```
        RESP = CLOSED
    ENDIF
    CALL KXSEMS (GBLSEM)
    ISTAT = KXWRIT (RESP, 4, NODE, RTYPE)
ENDIF
C
IF (MSG .EQ. SIGNAL) THEN
    CALL KXSEMW (GBLSEM)
    GBLFLG = OPEN
    CALL KXSEMS (GBLSEM)
ENDIF
C
SEMHND = 1
RETURN
END
```

Note that we implement the global semaphore with a simple variable, GBLFLG to which access is restricted with the local semaphore, GBLSEM. If the semaphore is "locked" a message is sent back to the requesting node indicating that it should sleep. The code which implements the "signal" and "wait" requests for this global semaphore is shown below. For simplicity we do not show the code which initializes the local semaphores or sets up the message handler.

```
    SUBROUTINE GBLSIG
C
    INTEGER MSG, DEST, TYPE
    COMMON /SEMDAT/WAIT, SIGNAL, OPEN, CLOSED, SEMREQ, SEMRSP
    INTEGER WAIT, SIGNAL, OPEN, CLOSED, SEMREQ, SEMRSP
C
    MSG = SIGNAL
    DEST = 0
    TYPE = SEMREQ
    ISTAT = KXWRIT (MSG, 4, DEST, TYPE)
    RETURN
    END

    SUBROUTINE GBLWAT
C
    INTEGER MSG, DEST, TYPE, RTYPE, STATUS
    COMMON /SEMDAT/WAIT, SIGNAL, OPEN, CLOSED, SEMREQ, SEMRSP
    INTEGER WAIT, SIGNAL, OPEN, CLOSED, SEMREQ, SEMRSP
C
    MSG = WAIT
    DEST = 0
    TYPE = SEMREQ
```

```
        RTYPE = SEMRSP
C
        STATUS = CLOSED
C
10  ISTAT = KXWRIT(MSG, 4, DEST, TYPE)
    ISTAT = KXREAD(STATUS, 4, DEST, RTYPE)
    IF(STATUS .EQ. CLOSED) THEN
        CALL KXSLEE(10)
        GOTO 10
    ENDIF
    RETURN
END
```

The important point to note in this code is the call to KXSLEE in the last routine. This allows other processes on a node to proceed even though the calling process is blocked waiting for the global semaphore.

SEE ALSO

KXSEM, KXHAND

NAM

KXSTAR - Start execution of program

SYNOPSIS

```
SUBROUTINE KXSTAR(PGIND, NODE)
  INTEGER PGIND, NODE
```

```
SUBROUTINE KXMAIN(PGIND, NODE)
  INTEGER PGIND, NODE
```

DOMAIN

Available to host processes only.

DESCRIPTION

These routines begin execution of a program previously loaded into a node with the KXPLOA system call. Programs loaded with KXLOAD do not need to use these calls.

The special value `NODE = IALNOD` may be specified to perform the action on all allocated nodes. This value is defined in the XPRESS common block set up by the call to KXINIT.

EXAMPLE

The following example shows the correct use of KXSTAR and KXMAIN to begin execution of a job successfully loaded into the nodes.

```
PROGRAM MYTEST
  INTEGER PGIND
  COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
  CHARACTER*80 DEVICE, PRGNAM
  PARAMETER (DEVICE='/dev/ncube', PRGNAM='nobby')
C
C-- Set up Express and its common block.
C
  CALL KXINIT
C
  PGIND = KXOPEN(DEVICE, 4, NOCARE)
  IF(PGIND .LT. 0) THEN
    WRITE(6,*) 'Failed to allocate processors'
    STOP
  ENDIF
C
C-- Load program into processor group using index returned.
C
  ISTAT = KXPLOA(PGIND, PRGNAM, IALNOD)
  IF(ISTAT .LT. 0) THEN
    WRITE(6,*) 'Failed to load application'
```

```
        STOP
    ENDIF
C
C-- Start application running.
C
        CALL KXSTAR(PGIND, IALNOD)
        CALL KXMAIN(PGIND, IALNOD)
```

Note that these calls can be used to explicitly control when a process begins executing. It may be important, for example, that certain actions be performed on the host before execution begins. In this case the "start" calls can be deferred until an appropriate time.

SEE ALSO

KXLOAD, KXPLOA, KXOPEN, KXPAUS

NAME

KXSWAB, KXSWAW, KXSWAD - Byte swapping routines

SYNOPSIS

```
SUBROUTINE KXSWAB (FROM, TO, NBYTES)
INTEGER*2 FROM (*), TO (*)
INTEGER NBYTES
```

```
SUBROUTINE KXSWAW (FROM, TO, NBYTES)
INTEGER*4 FROM (*), TO (*)
INTEGER NBYTES
```

```
SUBROUTINE KXSWAD (FROM, TO, NBYTES)
REAL*8 FROM (*), TO (*)
INTEGER NBYTES
```

DOMAIN

These routines may be called in any *Express* program.

DESCRIPTION

An unfortunate problem with many parallel processing systems is that the host machines and node processors have different CPU types. It is often then the case that the binary representation of various data types is different. Typical examples are Sun workstations hosting transputer or NCUBE systems. The former has a Motorola CPU with the most significant byte of a word having the lowest memory address while the node processors store the least significant byte first.

To aid with these problems *Express* provides a set of byte swapping primitives for transforming data between "big" and "little" endian machines. Each routine has a pair of pointer arguments which denote the buffers from which data should be taken and into which it should be placed after swapping. These two pointers may be the same. The last argument, in each case, is the number of bytes in the buffer to be swapped. This should be a multiple of the size of each item being swapped.

The three routines each serve a different swapping style as follows:

KXSWAB	Swaps adjacent bytes in 2-byte quantities
KXSWAW	Reverses the bytes in 4-byte quantities - i.e., the original order {0,1,2,3} becomes {3,2,1,0}.
KXSWAD	Reverses the bytes in 8-byte quantities - the original order {0,1,2,3,4,5,6,7} becomes {7,6,5,4,3,2,1,0}.

Note that these routines are sufficient to transform data items between Motorola byte ordered machines (Sun workstations, etc.) and INTEL byte ordered machines (NCUBE, transputers, etc.)

EXAMPLE

When necessary, byte swapping typically occurs in one of two places depending on the programming model in use.

In “Host-node” programs it is typical to have to swap all data items that are transmitted to or received from the nodes. The issue of which processor should perform the byte swapping is one of pure convenience - either the host or the nodes can swap the bytes. Often this decision is made according to who has to further use the data being swapped - the following code fragment represents a typical bug

```

C
C-- Byte swapping in a "host-node" program - INCORRECT
C
      SUBROUTINE ITERAT(NTIMES)
      INTEGER*4 NTIMES
      INTEGER I, TYPE
      COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC

      TYPE = 123
      CALL KXSWAW(NTIMES, NTIMES, 4)
      ISTAT = KXBROD(NTIMES, IHOST, 4, IALNOD, 0, TYPE)
C
C-- This is a BUG ..... ntimes just had its bytes swapped!
C
      DO 10 I=1,NTIMES
          ....
10  CONTINUE

```

This code shows some typical features in a byte-swapping environment. The “bug” in the above code is that the host program swaps the bytes in the NTIMES value and sends it to the nodes (correct) but then attempts to use the value in the following loop - *without swapping the bytes back*. As a result the loop will probably run for an extremely long time!

Among several possible “fixes” are:

- Adding another call to KXSWAW after the call to KXBROD to restore the NTIMES variable to its proper state.
- Making a temporary variable, swapping NTIMES into it and sending the temporary value to the nodes.
- Having the nodes do the byte swapping in this case.

Cubix programs should only have byte swapping problems when performing binary I/O. Regular text files should pose no problems since the internal protocols take care of all appropriate byte swapping. Arguments to system calls that will be sent to another host are also byte swapped automatically. For binary files, however, the problem remains and the only viable solution seems to be the insertion of many calls to the appropriate swapping

KXSWAP

routine.

SEE ALSO

KMREAD, KMRD2D, KMWRT.

NAME

KXSYNC - Synchronization primitive

SYNOPSIS

SUBROUTINE KXSYNC

DOMAIN

KXSYNC may only be called from the nodes.

DESCRIPTION

This routine is used to implement synchronization points in applications. It is guaranteed that no processor will proceed past the call to KXSYNC until all are ready to do so. Furthermore the processors emerge from the KXSYNC calls on their respective nodes as synchronized as can be arranged.

A call to KXSYNC in one processor must be complemented by a call to KXSYNC in all other processors.

EXAMPLE

In the following code we assume that it is important that all processors be synchronized between two phases of an algorithm.

```
C
C-- PHASE 1. of application .....
C
          .....

C
C-- Before beginning second phase make sure all processors
C-- in sync.
C
          CALL KXSYNC
C
C-- PHASE 2. of application - all processors synchronized.
C
          .....
```

Another good place for this function is after installing message handlers with the KXHAND system call. Synchronizing all processors is a good idea since it prevents any one processor sending a message to another which has yet to install its signal handler.

SEE ALSOKXCHAN

NAME

KXTEST - Test for an incoming message, non-blocking

SYNOPSIS

```
INTEGER FUNCTION KXTEST(SRC, TYPE)
  INTEGER SRC, TYPE
```

DOMAIN

KXTEST is available to both host and node programs. The calling sequence is identical in both cases.

DESCRIPTION

This function looks for an incoming message in a non-blocking fashion. It is intended for use in implementing strategies which require non-blocking read capabilities. The arguments SRC and TYPE are interpreted just as in the KXREAD call with the same wildcard interpretations.

The useful feature of the "test" function is that it returns immediately indicating by the return value whether or not a message currently exists which matches the supplied parameters. If no such message is found -1 is returned. Otherwise the return value is the length of the matching message.

EXAMPLES

In the following examples we consider a case in which the following four messages have arrived on our node in the order given.

1. Source 1	Type 12	Length 32
2. Source HOST	Type 2	Length 512
3. Source 1	Type 15	Length 1024
4. Source 2	Type 0	Length 0

The simplest case is where both source and type are explicitly stated as in the call

```
SOURCE = 1
TYPE = 15
ISTAT = KXTEST(SOURCE, TYPE)
```

In this case message three will be accepted. The returned value, ISTAT will be 1024, the length of the acceptable message.

The next example uses the wildcard value, NOCARE, to look for any message but retain information about its source. (This value is defined in the XPRESS common block set up by the call to KXINIT.

```
SOURCE = NOCARE
TYPE = 0
ISTAT = KXTEST(SOURCE, TYPE)
```

In this case the type is explicitly given and so message 4 will be matched. The returned

variable will be 0, the length of the message and the SOURCE variable will contain 2, the source of the message.

In the last example a wildcard value is given for the type field.

```
SOURCE = 1
TYPE = NOCARE
ISTAT = KXTEST(SOURCE, TYPE)
```

In this case the source is given explicitly and the type allowed to take any value. With the parameters shown message 1 will be accepted and the value 12 stored in the TYPE variable. The value 32 will be returned. Note that types are subject to extra processing through the KXTYPE commands. If we had specifically excluded type 12 from consideration then message 3 would have been used instead since it has the correct source and has not been excluded. If we had excluded both types 12 and 15 then the call to KXTEST would return -1 to indicate that no suitable message had yet arrived.

RETURN VALUE

The return value is the length of the matching message or -1 if no message can be found which fits the indicated parameters.

SEE ALSO

KXOPEN, KXREAD, KXTYPE

KXTIME

NAME

KXTIME, KXTICK - Time measurement

SYNOPSIS

```
INTEGER*4 FUNCTION KXTIME ()
```

```
INTEGER FUNCTION KXTICK ()
```

DOMAIN

These functions are available to all node programs.

DESCRIPTION

KXTIME returns the number of microseconds since a fixed reference point.

KXTICK returns the number of hardware clock ticks since a fixed reference point.

Both routines are intended to be used for timing measurements. KXTIME provides measurements in convenient units but suffers from the fact that its accuracy may depend on some "unknown" constant such as the hardware's clock speed. It may further require significantly longer than KXTICK to return a result since one or more arithmetic operations will normally be required to convert the machine clock ticks to microseconds.

Note that the availability of a routine which returns time in microseconds should not be taken to imply the existence of hardware with this resolution. In most cases the hardware timers will have intervals of many microseconds.

NAME

KXINCT, KXEXCT - Include or exclude certain message types in interpreting wildcards.

SYNOPSIS

```
SUBROUTINE KXINCT (LOTYP, HITYP)
INTEGER LOTYP, HITYP
```

```
SUBROUTINE KXEXCT (LOTYP, HITYP)
INTEGER LOTYP, HITYP
```

DOMAIN

KXINCT and KXEXCT are available in both host and node processors.

DESCRIPTION

These routines are used to modify the behavior of the "NOCARE" wildcard value used in the TYPE field of the calls KXREAD, KXTEST, etc. In particular the user can specify that certain types be excluded or included among those that match the "any type" condition.

KXEXCT specifies a low and high type value defining an (inclusive) range of types which should not be considered when processing the wildcard value. All the other types will remain acceptable.

KXINCT specifies the low and high end of an (inclusive) range of types which can be accepted by the program. All other types of messages will be excluded.

These routines are of most use when two or more processes share the same processor group with the KXSHAR call or when message handlers are being used (*cf.* KXHAND). In this case the use of wildcards is dangerous, without previously calling these routines, since otherwise the recipient of any given message is unpredictable. Using these routines it is possible to allow one process access to only a restricted range of types while the other process can safely use all the other types and BOTH may still be permitted the use of wildcards.

EXAMPLES

In the following code we limit attention to types in the range 123 thru 125.

```
PROGRAM MYTEST
INTEGER MSGSRC, MSGTYP
COMMON/XPRESS/NOCARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Initialize Express and its common block.
C
CALL KXINIT
C
C-- Code to allocate nodes and load programs.
```

KXTYPE

```
C
      ...
C
C-- Restrict attention to only a small range of message
C-- types.
C
      CALL KXINCT(123,125)
C
C-- Read with wildcard values, restricted to range [123,125]
C
      MSGSRC = NOCARE
      MSGTYP = NOCARE
      ISTAT = KXREAD(BUFFER, 128, MSGSRC, MSGTYP)
```

After including only the specific types the wildcard values may be used freely but with their meanings restricted to a smaller range. In the above example the call to KXREAD will only ever read messages whose types lie in the range 123-125.

As mentioned above this technique is most useful when two or more processes wish to share access to a particular set of nodes. If the above call had been made in one process then the other one might wish to make a call such as

```
      CALL KXEXCT(99, 125)
```

in which we explicitly delete the message type range 99-125 from consideration. (This would be useful if yet another process were sharing the same nodes and using types 99-100.) All other message types will remain valid.

The include/exclude mechanism can be turned off by supplying two NOCARE arguments to the appropriate function.

SEE ALSO

KXREAD, KXTEST, KXSHAR.

NAME

KXVREA, KXVWRI - Vector read/write functions

SYNOPSIS

```
INTEGER FUNCTION
      KXVREA (BUFFER, SIZE, OFFSET, ITEMS, SRC, TYPE)
INTEGER BUFFER (*)
INTEGER SIZE, OFFSET, ITEMS, SRC, TYPE
```

```
INTEGER FUNCTION
      KXVWRI (BUFFER, SIZE, OFFSET, ITEMS, DEST, TYPE)
INTEGER BUFFER (*)
INTEGER SIZE, OFFSET, ITEMS, DEST, TYPE
```

DOMAIN

KXVREA and KXVWRI may be called in both host and node processors.

DESCRIPTION

These routines implement direct read and write functions. Additionally they allow non-contiguous blocks of data to be transmitted as a single message.

These functions correspond directly to KXREAD and KXWRIT except in the interpretation of the actual bytes to be transmitted. In the KXREAD function a single block of contiguous data is transmitted while the KXVREA functions allow messages to be built up from non-contiguous memory blocks.

The manner in which the blocks are specified to KXVWRI is as follows: ITEMS objects, each of size SIZE bytes are taken starting from BUFFER. In addition each block is separated from the next by OFFSET bytes.

The specification is similar for KXVREA except that objects are read into distinct memory blocks separated by OFFSET bytes.

In all other regards the arguments to KXVREA and KXVWRI perform just as they would in KXREAD and KXWRIT - including the restriction that neither DEST nor TYPE arguments may be NOCARE in calls to KXVWRI.

EXAMPLE

The most useful application of these functions is to deal with multi-dimensional arrays in which we are required to pass data across a dimension in which the array data is not contiguous. (In FORTRAN the first array dimension is the one that indexes contiguous memory locations). Consider an example in which we have a 10 x 10 array of values in each node corresponding to a two-dimensional image. The first dimension refers to the vertical axis while the second refers to the horizontal. (array (2, 2) is thus near the bottom left-hand corner, for example). If we now consider a simple scrolling operation in which data is to be moved from left to right then we see that the data lies correctly and a suitable call to KXREAD, for example, would be

KXVREA

```
ISTAT = KXREAD (ARRAY, 10*4, LNODE, TYPE)
```

assuming that LNODE had been correctly assigned and a suitable call to KXWRIT had been made in some processor. If the scroll were to be in the vertical direction, however, then KXREAD is not appropriate; the operation can be coded as

```
ISTAT = KXVREA (ARRAY, 4, 10*4, 10, DNODE, TYPE)
```

which specifies that each array element has the size of an INTEGER and that the total distance between elements ARRAY (I, J) and ARRAY (I, J+1) is 10 times the size of an individual element. Finally ten items should be transmitted. Notice that we can also use a call to KXVREA for the horizontal shift by merely changing the OFFSET field in the above call from 10*4 to 4. This allows the code to have a uniform structure for both axes.

DIAGNOSTICS

If any error occurs in KXVREA or KXVWRI -1 is returned. Possible sources of error are: an illegal source or destination, an illegal buffer or a preposterous value of SIZE, OFFSET or ITEMS. If no error occurs KXVREA returns the number of items read and KXVWRI the number written.

SEE ALSO

KXREAD, KXWRIT, KXCHAN.

NAME

KXWRIT - Write a message

SYNOPSIS

```

INTEGER FUNCTION KXWRIT (BUF, LENGTH, DEST, TYPE)
INTEGER LENGTH, DEST, TYPE
INTEGER BUF (*)

```

DOMAIN

KXWRIT is available to both host and node programs with identical calling sequences.

DESCRIPTION

This routine sends a message to the processor indicated by the DEST argument. The message will consist of LENGTH bytes taken from the supplied BUF pointer. The message has the type specified by the TYPE parameter which may not take the special NOCARE value from the XPRESS common block.

The special value IHOST may be used to give the host processor as destination. This value is to be found in the XPRESS common block set up by the call to KXINIT.

RETURN VALUE

KXWRIT returns the number of bytes written, or -1 upon unrecoverable errors.

EXAMPLES

The following code is used to send 15 bytes taken from the address MYBUF to processor 12. The message will have type 99.

```

      PROGRAM MYTEST
      INTEGER DEST, TYPE
C
C-- Set up Express.
C
      CALL KXINIT
C
      DEST = 12
      TYPE = 99
C
      ISTAT = KXWRIT(MYBUF, 15, DEST, TYPE)
C
      ...

```

The next code sends a 128 byte message to the host processor. The message type will be 10.

```

      PROGRAM MYTEST

```

KXWRIT

```
        INTEGER DEST, TYPE
        COMMON/XPRESS/NO CARE, NORDER, NONODE, IHOST, IALNOD, IALPRC
C
C-- Set up Express and its common block.
C
        CALL KXINIT
C
        DEST = IHOST
        TYPE = 10
C
        ISTAT = KXWRIT(DATA, 128, DEST, TYPE)

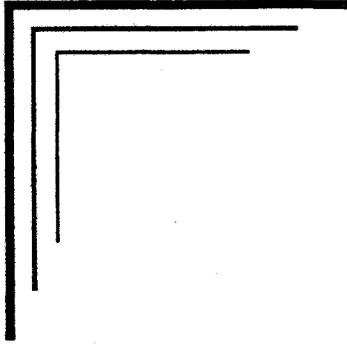
        ...
```

WARNINGS

Certain message types are restricted to the *Express* kernel. User message types must be less than 16384.

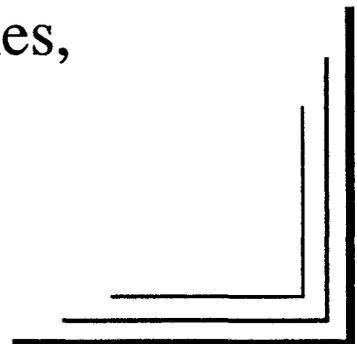
SEE ALSO

KXOPEN, KXREAD



Classification of routines

A listing of the *Express* routines,
broken down by functionality



This section lists the various functions and routines available to *Express* programs grouped according to functionality. While no exact division between routines is possible this information may serve as a useful guide to “related” functions.

User Commands (Man page)

acctool	Analyze accounting data	acctool
cnftool	Configure Transputer systems	cnftool
ctool	Analyze communication profile data	ctool
cubix	Download and execute <i>Cubix</i> programs, I/O server	cubix
etool	analyze event profile data and “toggles”	etool
excustom	Modify <i>Express</i> system parameters	excustom
exdump	Retrieve data from RAM files	exdump
exinit	Reboot and reload <i>Express</i> kernel	exinit
exreset	Reset transputer system	exinit
exstat	Display node usage information	exstat
ndb	Source level debugger	ndb
xtool	Analyze execution profile data	xtool

Compilers (Man page)

ncc	C compiler and linker for NCUBE	ncc
nf77	FORTRAN compiler and linker for NCUBE	nf77
symcc	C compiler and linker for SYMULT	symcc
symf77	FORTRAN compiler and linker for SYMULT	symf77
tcc	Logical Systems C compiler and linker for transputers	tcc
tcc3L	3L C compiler and linker for transputers	tcc3L
tfc	3L FORTRAN compiler and linker for transputers	tfc

System Initialization (Man page)

KXINIT	Start up <i>Express</i> and initialize XPRESS common block . . .	KXINIT
--------	--	--------

Processor Allocation and Control (Man page)

KXCLOS	Deallocate processor group	KXCLOS
KXLOAD	Load program into all nodes	KXLOAD
KXMAIN	Start execution of main program	KXSTAR
KXOPEN	Allocate a group of processors	KXOPEN
KXPAUS	Arrange for program to be loaded “stopped”	KXPAUS
KXPID	Translate UNIX process ID to <i>Express</i> process ID	KXSHAR

KXPLOA	Load a program into a single node	KXPLOA
KXSHAR	Share a processor group between multiple host programs . .	KXSHAR
KXSTAR	Start execution of a node program	KXSTAR

Basic Communication System (Man page)

KXEXCT	Define meaning of "read/write" wildcards	KXTYPE
KXINCT	Define meaning of "read/write" wildcards	KXTYPE
KXREAD	Read a message	KXREAD
KXTEST	Test for an incoming message - non-blocking	KXTEST
KXVREA	Read a <i>vector</i> message	KXVREA
KXVWRI	Send a <i>vector</i> message	KXVREA
KXWRIT	Send a message	KXWRIT

"Global" Communication System (Man page)

KXBROD	Interprocessor broadcast	KXBROD
KXCHAN	Synchronous multi-node data exchange	KXCHAN
KXCOMB	Apply user supplied operation to distributed data set	KXCOMB
KXCONC	Transfer distributed data to local memory	KXCONC
KXSYNC	Synchronize processors	KXSYNC
KXVCHA	Synchronous multi-node <i>vector</i> exchange	KXCHAN

Asynchronous Communication System (Man page)

KXHAND	Install asynchronous message handler	KXHAND
KXRECV	Read a message - non-blocking	KXRECV
KXSEND	Send a message - non-blocking	KXSEND

Hardware Dependent Communication System (Man page)

KXCHON	Re-enable <i>Express</i> processing on a channel	KXCH
KXCHOF	Disable <i>Express</i> processing on a channel	KXCH
KXCHRD	Read bytes from disabled channel	KXCH
KXCHWT	Write bytes to disabled channel	KXCH

Decomposition Tools (Man page)

KXGDBC	Define boundary conditions on user domain	KXGRID
KXGDCO	Determine position in user domain	KXGRID

KXGDIN	Initialize decomposition system	KXGRID
KXGDNO	Determine communication parameters from user domain . .	KXGRID
KXGDPR	Map user domain coordinates to processor number	KXGRID
KXGDSI	Distribute data among processors	KXGRID
KXGDSP	Distribute processors on user domain	KXGRID
KXPARA	Determine run-time configuration	KXPARA

Customization (Man page)

KXCUST	Modify <i>Express</i> system parameters	KXCUST
--------	---	--------

I/O (Man page)

KABORT	Immediately terminate node program	KABORT
KMULTI	Switch file I/O mode to "multi"	KMULTI
KSINGL	Switch file I/O mode to "single"	KMULTI
ISASYN	Inquire file I/O mode	KMULTI
ISMULT	Inquire file I/O mode	KMULTI
KMREAD	Read independent data into nodes	KMREAD
KMRD2D	Read two-dimensional data set into nodes	KMRD2D
KMWRT	Write independent data from node	KMWRT
KMWT2D	Write two-dimensional data set into nodes	KMRD2D
KCBXSY	Assign overall synchronous/asynchronous mode	KCBXSY

Debugging (Man page)

KABORT	Immediately terminate node program	KABORT
KXBREA	Halt program at breakpoint	KXBREA
KXPAUS	Load node program "stopped" at a breakpoint	KXPAUS

Multi-Host systems (Man page)

KCONND	Indicate an alternative host for system calls	KCONND
KDSPND	Indicate an alternative host for graphical output	KDSPND
KXACCS	Override access to all nodes in system	KXACCS
KXSHAR	Share a group of nodes with another host program	KXSHAR

Multitasking (Man page)

KEXEC	Overlay a node program with another	KEXEC
-------	---	-------

KXHAND	Install asynchronous message handler	KXHAND
KXSEMA	Allocate and initialize a semaphore	KXSEM
KXSEMF	Deallocate a semaphore structure	KXSEM
KXSEMS	Exit a critical section and "signal" any waiting processes	KXSEM
KXSEMW	Attempt to enter a critical section, sleeping if necessary	KXSEM
KXSLEE	Suspend process for indicated time	KXSLEE

Graphics

(Man page)

KAERAS	Erase display asynchronously	KERASE
KAGIN	Perform asynchronous graphical input operations	KGIN
KAOPEN	Initialize Plotix asynchronously	KOPENP
KASEND	Flush graphical data to display surface asynchronously	KSENDP
KASPEC	Inquire device aspect ratio	KASPEC
KBOX	Draw, and optionally fill, rectangle	KBOX
KCLOSP	Terminate Plotix	KCLOSP
KCOLOR	Set line drawing color	KCOLOR
KCONT	Draw visible line in current color	KCONT
KCNTOR	Draw a contour plot of a user supplied function	KCNTOR
KDISND	Indicate an alternative host for graphical output	KDISND
KDOTEX	Draw and justify text	KDOTEX
KENDCL	Disable clipping	KCLIP
KENDPA	Close and optionally fill polygon	KPANEL
KERASE	Erase display surface	KERASE
KGIN	Perform "locator" input	KGIN
KGREYS	Modify color look-up table, create greyscale	KGREYS
KINITP	Begin polygon	KPANEL
KLABEL	Draw text	KLABEL
KLINEM	Set line style	KLINEM
KMARKE	Draw marker symbol	KMARKE
KMOVE	Move current position without drawing	KMOVE
KOPENP	Initialize Plotix	KOPENP
KORTHO	Define user coordinate range	KSPACE
KPANLP	Define point in polygon	KPANEL
KPLOTH	Monitor graphics buffer usage	KPLOTH
KPOLGN	Draw polygon	KPANEL
KPXGOP	Inquire device capability	KPLXOP
KPX SOP	Set hardware-dependent graphics option	KPLXOP
KRAINB	Modify color look-up table, create HSV table	KRAINB
KSENDP	Flush graphical data to display surface synchronously	KSENDP

KSETCL	Enable clipping against rectangular region	KCLIP
KSETVP	Switch between “windows”	KSETVP
KSPACE	Define user coordinate range	KSPACE
KVPORT	Define a region of the display as a “window”	KVPORT
KUSEND	Flush independent data to display synchronously	KSENDP

Performance Analysis	(Man page)
-----------------------------	-------------------

KCPEND	Terminate communication profiler and dump data	KCPINQ
KCPINQ	Inquire setting of runtime ‘-mc’ switch	KCPINQ
KCPON	Enable communication profiler	KCPROF
KCPOFF	Disable communication profiler	KCPROF
KCPCP	Receive communication profile data in host processor	KCPCP
KCPELT	Send communication profile data to host processor	KCPCP
KEPADD	Indicate a “user” event	KEPROF
KEPEND	Terminate event profiler and dump data	KEPINQ
KEPINI	Initialize memory for event profiler	KEPROF
KEPINQ	Inquire setting of runtime ‘-me’ switch	KEPINQ
KEPLAB	Assign a label to a user specified “event”	KEPROF
KEPON	Enable event profiler	KEPROF
KEPOFF	Disable event profiler	KEPROF
KEPTOG	Enable/disable timing for a region of source code	KEPTOG
KEPTGI	Initialize memory for a “toggle”	KEPTOG
KEPCP	Receive event profile data in host processor	KEPCP
KEPELT	Send event profile data to host processor	KEPCP
KPROFI	Assign memory for execution profiler	KPROFI
KXPEND	Terminate execution profiler and dump data	KXPINQ
KXPINQ	Inquire setting of runtime ‘-mx’ switch	KXPINQ
KXPON	Enable execution profiler	KXPROF
KXPOFF	Disable execution profiler	KXPROF
KXPCP	Receive execution profile data in host processor	KXPCP
KXPELT	Send execution profile data to host processor	KXPCP

Host Interface Utilities	(Man page)
---------------------------------	-------------------

KCALHO	Call host routine from <i>Cubix</i> node program	KCALHO
KGETHO	Inquire host capabilities	KGETHO
KRETHO	Return from host routine in <i>Cubix</i> node programs	KCALHO
KSTRHO	Start host routine in <i>Cubix</i> node program	KCALHO

Utility Routines**(Man page)**

KXSWAB	Reverse bytes in 16-bit quantities	KXSWAP
KXSWAD	Reverse bytes in 64-bit quantities	KXSWAP
KXSWAW	Reverse bytes in 32-bit quantities	KXSWAP
KXTICK	Measure time in hardware "ticks"	KXTIME
KXTIME	Measure time in microseconds	KXTIME

1 Correspondence between C and FORTRAN

The first two columns of the following table list the equivalent C and FORTRAN routines. A blank entry indicates that no such routine exists.

2 Synchronization Rules

The third column of the table indicates the synchronization modes associated with each function. The various codes are:

- a These routines may be called with no regard to any synchronization constraints - any node may make such a call at any time.
- ls, all These routines must be made "loosely synchronously" in *all processors*. When a node calls one of these routines it will halt until all other nodes have called the same routine. Arguments may or may not be different in each node according to the particular function involved.
- ls, group These routines must be made "loosely synchronously" in all participating processors. Typically this means that two processors will be involved in some transaction in which case the first to arrive will halt until the others arrive at the synchronization point.
- mode The synchronization requirements of these calls depend on the global synchronization state of the system, as modified with the `syncmode` or `KCBXSY` system calls. If the global synchronization mode is "on" (the default) then these routines behave as though their synchronization constraint were "ls, all". If the global state is "off" they behave as "a".

3 Libraries and Programming Models

The last column in the table indicates the libraries and/or programming model combinations which support the named routines. These latter are coded as follows:

- h Routine is available to programs running on the host processor, linked with the *Express* library.
- n Available to programs running on the parallel computer nodes in the "Host-Node" programming style. Such programs should NOT be linked with either the *Cubix* or *Plotix* libraries.
- c These routines are part of the *Cubix* I/O library and may only be linked with programs using the *Cubix* programming model. Usually a compiler switch is available to indicate this programming model and the associated libraries.
- p These routines are to be found in the *Plotix* library which can be linked to programs running under the *Cubix* programming model. In some cases a compiler switch is available which links both the *Cubix* and *Plotix* libraries. If this is not so on your system the *Cubix* switch should be used supplemented by the pathname of the *Plotix* library.

4 NOTES

- (i) While no corresponding routine is available in FORTRAN the effect can be achieved by modifying the parameters to an OPEN statement. See the section on "open file modes" in the *Cubix* chapter for more details.
- (ii) These calls may be made asynchronously but they have no subsequent effect on the objects they access. The graphical open functions, for example, may be made asynchronously but the mode in which data is flushed to the output device is still determined by the flushing function used. Similarly a file opened with one of the asynchronous "open" functions still has as its default the "singl" access mode.
- (iii) These functions *can* be called asynchronously but will usually be used in a mode similar to "ls, group". Once invoked they leave the affected nodes in a state which will almost certainly function in an unpredictable manner until the corresponding action has been performed on other members of the group.
- (iv) These functions can be called asynchronously but should be used with extreme care when so doing. Because of their nature it is easy to introduce "race conditions" when using these routines asynchronously. In most cases it is easy (and safer) to force a synchronization after using one of these routines.
- (v) The synchronization behavior of these routines depends upon the "mode" of the associated file. For "singl" and "multi" mode files the constraint is "ls, group" while it becomes "a" for "async" mode files.

C	FORTRAN	Synchronization	Library
	KXEXIT	ls, all	n
	KXINIT	ls, all	h,n,c,p
_ex_swab	KXSWAB	a	h,n,c,p
_ex_swad	KXSWAD	a	h,n,c,p
_ex_swaw	KXSWAW	a	h,n,c,p
abort	KABORT	a	c,p
aerase	KAERAS	a	p
aexecve	KAEXEC	a	c,p
agin	KAGIN	a	p
aopen	(i)	a, (ii)	c,p
aopenpl	KAOPEN	a, (ii)	p
asendplot	KASEND	a,	p
aspect	KASPEC	a,	p
box	KBOX	a,	p
callhost	KCALHO	mode	c,p
closepl	KCLOSP	ls, all	p
color	KCOLOR	a	p
console_node	KCONND	a, (ii)	c,p
cont	KCONT	a	p
contour	KCNTOR	ls, all	p
cprof_end	KCPEND	ls, all	c,p
cprof_inq	KCPINQ	mode	c,p
cprof_off	KCPOFF	a	n,c,p
cprof_on	KCPON	a	n,c,p
cprofcp	KPCPCP		h
cprofelt	KPELT	ls, all	n
display_node	KDISND	a, (ii)	p
dotext	KDOTEX	a	p
endpanel	KENDP	a	p
eprof_add	KEPADD	a	n,c,p
eprof_end	KEPEND	ls, all	c,p
eprof_init	KEPINI	ls, all	n,c,p
eprof_inq	KEPINQ	mode	c,p
eprof_label	KEPLAB	a	n,c,p
eprof_off	KEPOFF	a	n,c,p
eprof_on	KEPON	a	n,c,p
eprof_toggle	KEPTOG	a	n,c,p
eprof_toginit	KEPTGI	a	n,c,p
eprofcp	KEPCP		h
eprofelt	KEPELT	ls, all	n
erase	KERASE	ls, all	p
exaccess	KXACCS		h
exargldl			h
exargldv			h

C	FORTRAN	Synchronization	Library
exbreak	KXBREA	a	n,c,p
exbroadcast	KXBROD	ls, group	h,n,c,p
exchange	KXCHAN	ls, group	h,n,c,p
exchanoff	KXCHOF	a, (iii)	n,c,p
exchanon	KXCHON	a, (iii)	n,c,p
exchanrd	KXCHRD	ls, group	n,c,p
exchanwt	KXCHWT	ls, group	n,c,p
exclose	KXCLOS		h
excombine	KXCOMB	ls, group	n,c,p
exconcat	KXCONC	ls, group	n,c,p
excustom	KXCUST		h
execve	KEXEC	mode	c,p
exenvld			h
exexctype	KXEXCT	a, (iv)	h,n,c,p
exgridbc	KXGDBC	a	h,n,c,p
exgridcoord	KXGDCO	a	h,n,c,p
exgridinit	KXGDIN	a	h,n,c,p
exgridnode	KXGDNO	a	h,n,c,p
exgridproc	KXGDPR	a	h,n,c,p
exgridsize	KXGDSI	a	h,n,c,p
exgridsplit	KXGDSP	a	h,n,c,p
exhandle	KXHAND	a, (iv)	n,c,p
exinctype	KXINCT	a, (iv)	h,n,c,p
exload	KXLOAD		h
exloadl			h
exloadle			h
exloadv			h
exloadve			h
exmain	KXMAIN		h
exopen	KXOPEN		h
exparam	KXPARA	a	h,n,c,p
expause	KXPAUS		h
expid	KXPID		h
expload	KXPLOA		h
exread	KXREAD	a	h,n,c,p
exreadfd			h
exreceive	KXRECV	a	n,c,p
exsemalloc	KXSEMI	a	n,c,p
exsemfree		a	n,c,p
exsemsig	KXSEMS	a	n,c,p
exsemwait	KXSEMW	a	n,c,p
exsend	KXSEND	a	n,c,p
exshare	KXSHAR		h
exsleep	KXSLEE	a	n,c,p

C	FORTRAN	Synchronization	Library
exstart	KXSTAR		h
exsync	KXSYNC	ls, all	n,c,p
extest	KXTEST	a	h,n,c,p
extick	KXTIME	a	n,c,p
extime	KXTIME	a	n,c,p
exvchange	KXVCHA	ls, group	h,n,c,p
exvread	KXVREA	a	h,n,c,p
exvwrite	KXVWRI	a	h,n,c,p
exwrite	KXWRIT	a	h,n,c,p
exwritefd			h
fasync	KASYNC	ls, all	c,p
fmulti	KMULTI	ls,all	c,p
forder	KORDER	a, (ii)	c,p
fsingl	KSINGL	ls, all	c,p
gethost	KGETHO	mode	c,p
getplxopt	KPXGOP	mode	p
getpoint	KGETPT	a	p
gin	KGIN	mode	p
greyscale	KGREYS	a	p
initlevel	KINITL	a	p
initpanel	KINITP	a	p
isasync	KISASY	a	c,p
ismulti	KISMUL	a	c,p
label	KLABEL	a	p
linemod	KLINEM	a	p
malloc_avail		a	c,p
malloc_debug		a	c,p
malloc_print		ls, group (v)	c,p
malloc_verify		a	c,p
marker	KMARKE	a	p
move	KMOVE	a	p
mread	KMREAD	ls, all	c,p
mread2d	KMRD2D	ls,all	c,p
mwrite	KMWRIT	ls,all	c,p
mwrite2d	KMWT2D	ls,all	c,p
openpl	KOPENP	mode	p
ortho_space	KORTHO	a	p
panelpoint	KPANLP	a	p
plothwm	KPLOTH	a	p
polgn	KPOLGN	a	p
profil	KPROFI	a	n,c,p
rainbow	KRAINB	a	p
ramfopen	(i)	a	c,p
rethost	KRETHO	mode	c,p

C	FORTTRAN	Synchronization	Library
sendplot	KSENDP	mode	p
setclip	KSETCL	a	p
setplxopt	KPXSOP	mode	p
setvbuf		a, (ii)	c,p
space	KSPACE	a	p
starthost	KSTRHO	mode	c,p
syncmode	KCBXSY	a, (ii)	c,p
usendplot	KUSEND	ls, all	p
vport	KVPORT	a	p
xprof_end	KXPEND	ls, all	c,p
xprof_inq	KXPINQ	mode	c,p
xprof_off	KXPOFF	a	n,c,p
xprof_on	KXPON	a	n,c,p
xprofcp	KXPCP		h
xprofelt	KXPELT	ls, all	n

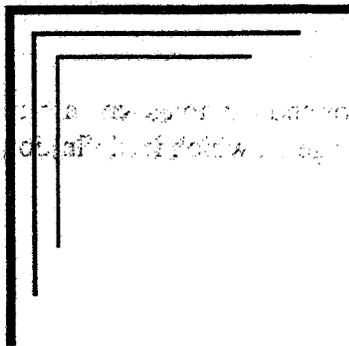
1

1. The first part of the document discusses the importance of maintaining accurate records of all activities. It emphasizes that these records are essential for ensuring accountability and transparency in the organization's operations.

2. The second part of the document outlines the specific procedures for collecting and organizing data. It details the steps involved in identifying relevant information, gathering it from various sources, and then categorizing it into a structured format for easy access and analysis.

3. The third part of the document focuses on the analysis and interpretation of the collected data. It describes how to identify trends, patterns, and anomalies within the data set, and how to use these insights to inform decision-making and strategic planning.

4. The final part of the document discusses the importance of regular reporting and communication. It stresses that clear and concise reports are necessary to keep stakeholders informed of the organization's progress and to facilitate collaborative problem-solving.



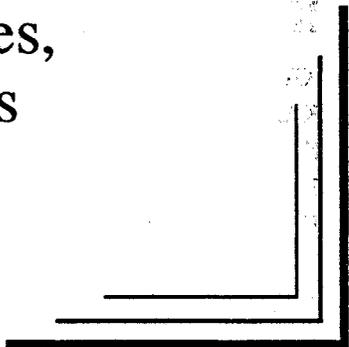
Index of Routines

and
to
of

ACCT	101
AD	102
ADJ	103
ADJ2	104
ADJ3	105
ADJ4	106
ADJ5	107
ADJ6	108
ADJ7	109
ADJ8	110
ADJ9	111
ADJ10	112
ADJ11	113
ADJ12	114
ADJ13	115
ADJ14	116
ADJ15	117
ADJ16	118
ADJ17	119
ADJ18	120
ADJ19	121
ADJ20	122
ADJ21	123
ADJ22	124
ADJ23	125
ADJ24	126
ADJ25	127
ADJ26	128
ADJ27	129
ADJ28	130
ADJ29	131
ADJ30	132
ADJ31	133
ADJ32	134
ADJ33	135
ADJ34	136
ADJ35	137
ADJ36	138
ADJ37	139
ADJ38	140
ADJ39	141
ADJ40	142
ADJ41	143
ADJ42	144
ADJ43	145
ADJ44	146
ADJ45	147
ADJ46	148
ADJ47	149
ADJ48	150
ADJ49	151
ADJ50	152
ADJ51	153
ADJ52	154
ADJ53	155
ADJ54	156
ADJ55	157
ADJ56	158
ADJ57	159
ADJ58	160
ADJ59	161
ADJ60	162
ADJ61	163
ADJ62	164
ADJ63	165
ADJ64	166
ADJ65	167
ADJ66	168
ADJ67	169
ADJ68	170
ADJ69	171
ADJ70	172
ADJ71	173
ADJ72	174
ADJ73	175
ADJ74	176
ADJ75	177
ADJ76	178
ADJ77	179
ADJ78	180
ADJ79	181
ADJ80	182
ADJ81	183
ADJ82	184
ADJ83	185
ADJ84	186
ADJ85	187
ADJ86	188
ADJ87	189
ADJ88	190
ADJ89	191
ADJ90	192
ADJ91	193
ADJ92	194
ADJ93	195
ADJ94	196
ADJ95	197
ADJ96	198
ADJ97	199
ADJ98	200
ADJ99	201
ADJ100	202

An alphabetical listing of routines,
variables, commands and macros

ADJ101
ADJ102
ADJ103
ADJ104
ADJ105
ADJ106
ADJ107
ADJ108
ADJ109
ADJ110
ADJ111
ADJ112
ADJ113
ADJ114
ADJ115
ADJ116
ADJ117
ADJ118
ADJ119
ADJ120
ADJ121
ADJ122
ADJ123
ADJ124
ADJ125
ADJ126
ADJ127
ADJ128
ADJ129
ADJ130
ADJ131
ADJ132
ADJ133
ADJ134
ADJ135
ADJ136
ADJ137
ADJ138
ADJ139
ADJ140
ADJ141
ADJ142
ADJ143
ADJ144
ADJ145
ADJ146
ADJ147
ADJ148
ADJ149
ADJ150
ADJ151
ADJ152
ADJ153
ADJ154
ADJ155
ADJ156
ADJ157
ADJ158
ADJ159
ADJ160
ADJ161
ADJ162
ADJ163
ADJ164
ADJ165
ADJ166
ADJ167
ADJ168
ADJ169
ADJ170
ADJ171
ADJ172
ADJ173
ADJ174
ADJ175
ADJ176
ADJ177
ADJ178
ADJ179
ADJ180
ADJ181
ADJ182
ADJ183
ADJ184
ADJ185
ADJ186
ADJ187
ADJ188
ADJ189
ADJ190
ADJ191
ADJ192
ADJ193
ADJ194
ADJ195
ADJ196
ADJ197
ADJ198
ADJ199
ADJ200
ADJ201
ADJ202
ADJ203
ADJ204
ADJ205
ADJ206
ADJ207
ADJ208
ADJ209
ADJ210
ADJ211
ADJ212
ADJ213
ADJ214
ADJ215
ADJ216
ADJ217
ADJ218
ADJ219
ADJ220
ADJ221
ADJ222
ADJ223
ADJ224
ADJ225
ADJ226
ADJ227
ADJ228
ADJ229
ADJ230
ADJ231
ADJ232
ADJ233
ADJ234
ADJ235
ADJ236
ADJ237
ADJ238
ADJ239
ADJ240
ADJ241
ADJ242
ADJ243
ADJ244
ADJ245
ADJ246
ADJ247
ADJ248
ADJ249
ADJ250
ADJ251
ADJ252
ADJ253
ADJ254
ADJ255
ADJ256
ADJ257
ADJ258
ADJ259
ADJ260
ADJ261
ADJ262
ADJ263
ADJ264
ADJ265
ADJ266
ADJ267
ADJ268
ADJ269
ADJ270
ADJ271
ADJ272
ADJ273
ADJ274
ADJ275
ADJ276
ADJ277
ADJ278
ADJ279
ADJ280
ADJ281
ADJ282
ADJ283
ADJ284
ADJ285
ADJ286
ADJ287
ADJ288
ADJ289
ADJ290
ADJ291
ADJ292
ADJ293
ADJ294
ADJ295
ADJ296
ADJ297
ADJ298
ADJ299
ADJ300
ADJ301
ADJ302
ADJ303
ADJ304
ADJ305
ADJ306
ADJ307
ADJ308
ADJ309
ADJ310
ADJ311
ADJ312
ADJ313
ADJ314
ADJ315
ADJ316
ADJ317
ADJ318
ADJ319
ADJ320
ADJ321
ADJ322
ADJ323
ADJ324
ADJ325
ADJ326
ADJ327
ADJ328
ADJ329
ADJ330
ADJ331
ADJ332
ADJ333
ADJ334
ADJ335
ADJ336
ADJ337
ADJ338
ADJ339
ADJ340
ADJ341
ADJ342
ADJ343
ADJ344
ADJ345
ADJ346
ADJ347
ADJ348
ADJ349
ADJ350
ADJ351
ADJ352
ADJ353
ADJ354
ADJ355
ADJ356
ADJ357
ADJ358
ADJ359
ADJ360
ADJ361
ADJ362
ADJ363
ADJ364
ADJ365
ADJ366
ADJ367
ADJ368
ADJ369
ADJ370
ADJ371
ADJ372
ADJ373
ADJ374
ADJ375
ADJ376
ADJ377
ADJ378
ADJ379
ADJ380
ADJ381
ADJ382
ADJ383
ADJ384
ADJ385
ADJ386
ADJ387
ADJ388
ADJ389
ADJ390
ADJ391
ADJ392
ADJ393
ADJ394
ADJ395
ADJ396
ADJ397
ADJ398
ADJ399
ADJ400
ADJ401
ADJ402
ADJ403
ADJ404
ADJ405
ADJ406
ADJ407
ADJ408
ADJ409
ADJ410
ADJ411
ADJ412
ADJ413
ADJ414
ADJ415
ADJ416
ADJ417
ADJ418
ADJ419
ADJ420
ADJ421
ADJ422
ADJ423
ADJ424
ADJ425
ADJ426
ADJ427
ADJ428
ADJ429
ADJ430
ADJ431
ADJ432
ADJ433
ADJ434
ADJ435
ADJ436
ADJ437
ADJ438
ADJ439
ADJ440
ADJ441
ADJ442
ADJ443
ADJ444
ADJ445
ADJ446
ADJ447
ADJ448
ADJ449
ADJ450
ADJ451
ADJ452
ADJ453
ADJ454
ADJ455
ADJ456
ADJ457
ADJ458
ADJ459
ADJ460
ADJ461
ADJ462
ADJ463
ADJ464
ADJ465
ADJ466
ADJ467
ADJ468
ADJ469
ADJ470
ADJ471
ADJ472
ADJ473
ADJ474
ADJ475
ADJ476
ADJ477
ADJ478
ADJ479
ADJ480
ADJ481
ADJ482
ADJ483
ADJ484
ADJ485
ADJ486
ADJ487
ADJ488
ADJ489
ADJ490
ADJ491
ADJ492
ADJ493
ADJ494
ADJ495
ADJ496
ADJ497
ADJ498
ADJ499
ADJ500
ADJ501
ADJ502
ADJ503
ADJ504
ADJ505
ADJ506
ADJ507
ADJ508
ADJ509
ADJ510
ADJ511
ADJ512
ADJ513
ADJ514
ADJ515
ADJ516
ADJ517
ADJ518
ADJ519
ADJ520
ADJ521
ADJ522
ADJ523
ADJ524
ADJ525
ADJ526
ADJ527
ADJ528
ADJ529
ADJ530
ADJ531
ADJ532
ADJ533
ADJ534
ADJ535
ADJ536
ADJ537
ADJ538
ADJ539
ADJ540
ADJ541
ADJ542
ADJ543
ADJ544
ADJ545
ADJ546
ADJ547
ADJ548
ADJ549
ADJ550
ADJ551
ADJ552
ADJ553
ADJ554
ADJ555
ADJ556
ADJ557
ADJ558
ADJ559
ADJ560
ADJ561
ADJ562
ADJ563
ADJ564
ADJ565
ADJ566
ADJ567
ADJ568
ADJ569
ADJ570
ADJ571
ADJ572
ADJ573
ADJ574
ADJ575
ADJ576
ADJ577
ADJ578
ADJ579
ADJ580
ADJ581
ADJ582
ADJ583
ADJ584
ADJ585
ADJ586
ADJ587
ADJ588
ADJ589
ADJ590
ADJ591
ADJ592
ADJ593
ADJ594
ADJ595
ADJ596
ADJ597
ADJ598
ADJ599
ADJ600
ADJ601
ADJ602
ADJ603
ADJ604
ADJ605
ADJ606
ADJ607
ADJ608
ADJ609
ADJ610
ADJ611
ADJ612
ADJ613
ADJ614
ADJ615
ADJ616
ADJ617
ADJ618
ADJ619
ADJ620
ADJ621
ADJ622
ADJ623
ADJ624
ADJ625
ADJ626
ADJ627
ADJ628
ADJ629
ADJ630
ADJ631
ADJ632
ADJ633
ADJ634
ADJ635
ADJ636
ADJ637
ADJ638
ADJ639
ADJ640
ADJ641
ADJ642
ADJ643
ADJ644
ADJ645
ADJ646
ADJ647
ADJ648
ADJ649
ADJ650
ADJ651
ADJ652
ADJ653
ADJ654
ADJ655
ADJ656
ADJ657
ADJ658
ADJ659
ADJ660
ADJ661
ADJ662
ADJ663
ADJ664
ADJ665
ADJ666
ADJ667
ADJ668
ADJ669
ADJ670
ADJ671
ADJ672
ADJ673
ADJ674
ADJ675
ADJ676
ADJ677
ADJ678
ADJ679
ADJ680
ADJ681
ADJ682
ADJ683
ADJ684
ADJ685
ADJ686
ADJ687
ADJ688
ADJ689
ADJ690
ADJ691
ADJ692
ADJ693
ADJ694
ADJ695
ADJ696
ADJ697
ADJ698
ADJ699
ADJ700
ADJ701
ADJ702
ADJ703
ADJ704
ADJ705
ADJ706
ADJ707
ADJ708
ADJ709
ADJ710
ADJ711
ADJ712
ADJ713
ADJ714
ADJ715
ADJ716
ADJ717
ADJ718
ADJ719
ADJ720
ADJ721
ADJ722
ADJ723
ADJ724
ADJ725
ADJ726
ADJ727
ADJ728
ADJ729
ADJ730
ADJ731
ADJ732
ADJ733
ADJ734
ADJ735
ADJ736
ADJ737
ADJ738
ADJ739
ADJ740
ADJ741
ADJ742
ADJ743
ADJ744
ADJ745
ADJ746
ADJ747
ADJ748
ADJ749
ADJ750
ADJ751
ADJ752
ADJ753
ADJ754
ADJ755
ADJ756
ADJ757
ADJ758
ADJ759
ADJ760
ADJ761
ADJ762
ADJ763
ADJ764
ADJ765
ADJ766
ADJ767
ADJ768
ADJ769
ADJ770
ADJ771
ADJ772
ADJ773
ADJ774
ADJ775
ADJ776
ADJ777
ADJ778
ADJ779
ADJ780
ADJ781
ADJ782
ADJ783
ADJ784
ADJ785
ADJ786
ADJ787
ADJ788
ADJ789
ADJ790
ADJ791
ADJ792
ADJ793
ADJ794
ADJ795
ADJ796
ADJ797
ADJ798
ADJ799
ADJ800
ADJ801
ADJ802
ADJ803
ADJ804
ADJ805
ADJ806
ADJ807
ADJ808
ADJ809
ADJ810
ADJ811
ADJ812
ADJ813
ADJ814
ADJ815
ADJ816
ADJ817
ADJ818
ADJ819
ADJ820
ADJ821
ADJ822
ADJ823
ADJ824
ADJ825
ADJ826
ADJ827
ADJ828
ADJ829
ADJ830
ADJ831
ADJ832
ADJ833
ADJ834
ADJ835
ADJ836
ADJ837
ADJ838
ADJ839
ADJ840
ADJ841
ADJ842
ADJ843
ADJ844
ADJ845
ADJ846
ADJ847
ADJ848
ADJ849
ADJ850
ADJ851
ADJ852
ADJ853
ADJ854
ADJ855
ADJ856
ADJ857
ADJ858
ADJ859
ADJ860
ADJ861
ADJ862
ADJ863
ADJ864
ADJ865
ADJ866
ADJ867
ADJ868
ADJ869
ADJ870
ADJ871
ADJ872
ADJ873
ADJ874
ADJ875
ADJ876
ADJ877
ADJ878
ADJ879
ADJ880
ADJ881
ADJ882
ADJ883
ADJ884
ADJ885
ADJ886
ADJ887
ADJ888
ADJ889
ADJ890
ADJ891
ADJ892
ADJ893
ADJ894
ADJ895
ADJ896
ADJ897
ADJ898
ADJ899
ADJ900
ADJ901
ADJ902
ADJ903
ADJ904
ADJ905
ADJ906
ADJ907
ADJ908
ADJ909
ADJ910
ADJ911
ADJ912
ADJ913
ADJ914
ADJ915
ADJ916
ADJ917
ADJ918
ADJ919
ADJ920
ADJ921
ADJ922
ADJ923
ADJ924
ADJ925
ADJ926
ADJ927
ADJ928
ADJ929
ADJ930
ADJ931
ADJ932
ADJ933
ADJ934
ADJ935
ADJ936
ADJ937
ADJ938
ADJ939
ADJ940
ADJ941
ADJ942
ADJ943
ADJ944
ADJ945
ADJ946
ADJ947
ADJ948
ADJ949
ADJ950
ADJ951
ADJ952
ADJ953
ADJ954
ADJ955
ADJ956
ADJ957
ADJ958
ADJ959
ADJ960
ADJ961
ADJ962
ADJ963
ADJ964
ADJ965
ADJ966
ADJ967
ADJ968
ADJ969
ADJ970
ADJ971
ADJ972
ADJ973
ADJ974
ADJ975
ADJ976
ADJ977
ADJ978
ADJ979
ADJ980
ADJ981
ADJ982
ADJ983
ADJ984
ADJ985
ADJ986
ADJ987
ADJ988
ADJ989
ADJ990
ADJ991
ADJ992
ADJ993
ADJ994
ADJ995
ADJ996
ADJ997
ADJ998
ADJ999
ADJ1000



Index to Routines

This index contains an alphabetical list of the various subroutines, macros and variables which may be of use to *Express* programs. Each routine has an indication of the page on which its definition and arguments can be found.

A

acctool 7

C

cnftool 9

ctool 10

cubix 12

E

etool 15

examples

 system initialization 185

excustom 16

exdump 18

exinit 21

exreset 23

exsend 214

exstat 24

I

ISASY 131

ISMULT 131

K

KABORT 63

KAERAS 105

KAEXEC 107

KAGIN 113

KAOPEN 136

KASEND 148

KASPEC 64

KASYNC 131

KBOX 66

KCALHO 68

KCBXSY 73

KCNTOR 77

KCOLOR 80

KCONND 82

KCONT 84

KCPCP 85

KCPFLT 85

KCPEND 87

KCPINQ 87

KCPOFF 89

KCPON 89

KDISND 91

KDOTEX 93

KENDCL 75

KENDPA 138

KEPADD 99

KEPCP 95

KEPELT 95

KEPEND 97

KEPINI 99

KEPINQ 97

KEPLAB 99

KEPOFF 99

KEPON 99

KEPTGI 102

KEPTOG 102

KERASE 105

KEXEC 107

KFLUSH 109

KGETHO 111

KGETPT 77	KXCUST 176
KGIN 113	KXEXCT 231
KGREYS 115	KXGDBC 178
KINITL 77	KXGDCO 178
KINITP 138	KXGDIN 178
KLABEL 117	KXGDNO 178
KLINEM 119	KXGDPR 178
KMARKE 121	KXGDSI 178
KMOVE 123	KXGDSP 178
KMRD2D 124	KXHAND 182
KMREAD 129	KXINCT 231
KMULTI 131	KXINIT 185
KMWRT 134	KXLOAD 187
KMWT2D 124	KXMAIN 222
KOPENP 136	KXOPEN 189
KORDER 131	KXPARA 191
KORTHO 150	KXPAUS 193
KPANLP 138	KXPCP 195
KPLOTH 140	KXPELT 195
KPOLGN 138	KXPEND 197
KPROFI 143	KXPID 217
KPXGOP 141	KXPINQ 197
KPXSOP 141	KXPLOA 199
KRAINB 145	KXPOFF 202
KREAD 147	KXPON 202
KRETHO 68	KXREAD 204
KSENDP 148	KXRECV 207
KSETCL 75	KXSEMI 210
KSETVP 153	KXSEMS 210
KSINGL 131	KXSEMW 210
KSPACE 150	KXSHAR 217
KSTRHO 68	KXSLEE 219
KUSEND 148	KXSTAR 222
KVPORT 153	KXSWAB 224
KWRITE 147	KXSWAD 224
KXACCS 156	KXSWAW 224
KXBREA 157	KXSYNC 227
KXBROD 158	KXTEST 228
KXCHAN 163	KXTICK 230
KXCHOF 160	KXTIME 230
KXCHON 160	KXVCHA 163
KXCHRD 160	KXVREA 233
KXCHWT 160	KXVWRI 233
KXCLOS 167	KXWRIT 235
KXCOMB 169	
KXCONC 173	

N

ndb 25

S

system initialization 185

T

tcc 39

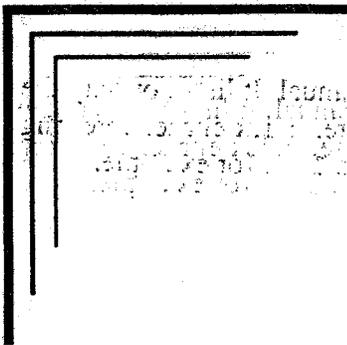
tcc31 43

tfc 47

X

xtool 50

00000000
 00000001
 00000002
 00000003
 00000004
 00000005
 00000006
 00000007
 00000008
 00000009
 0000000A
 0000000B
 0000000C
 0000000D
 0000000E
 0000000F
 00000010
 00000011
 00000012
 00000013
 00000014
 00000015
 00000016
 00000017
 00000018
 00000019
 0000001A
 0000001B
 0000001C
 0000001D
 0000001E
 0000001F
 00000020
 00000021
 00000022
 00000023
 00000024
 00000025
 00000026
 00000027
 00000028
 00000029
 0000002A
 0000002B
 0000002C
 0000002D
 0000002E
 0000002F
 00000030
 00000031
 00000032
 00000033
 00000034
 00000035
 00000036
 00000037
 00000038
 00000039
 0000003A
 0000003B
 0000003C
 0000003D
 0000003E
 0000003F
 00000040
 00000041
 00000042
 00000043
 00000044
 00000045
 00000046
 00000047
 00000048
 00000049
 0000004A
 0000004B
 0000004C
 0000004D
 0000004E
 0000004F
 00000050
 00000051
 00000052
 00000053
 00000054
 00000055
 00000056
 00000057
 00000058
 00000059
 0000005A
 0000005B
 0000005C
 0000005D
 0000005E
 0000005F
 00000060
 00000061
 00000062
 00000063
 00000064
 00000065
 00000066
 00000067
 00000068
 00000069
 0000006A
 0000006B
 0000006C
 0000006D
 0000006E
 0000006F
 00000070
 00000071
 00000072
 00000073
 00000074
 00000075
 00000076
 00000077
 00000078
 00000079
 0000007A
 0000007B
 0000007C
 0000007D
 0000007E
 0000007F
 00000080
 00000081
 00000082
 00000083
 00000084
 00000085
 00000086
 00000087
 00000088
 00000089
 0000008A
 0000008B
 0000008C
 0000008D
 0000008E
 0000008F
 00000090
 00000091
 00000092
 00000093
 00000094
 00000095
 00000096
 00000097
 00000098
 00000099
 0000009A
 0000009B
 0000009C
 0000009D
 0000009E
 0000009F
 000000A0
 000000A1
 000000A2
 000000A3
 000000A4
 000000A5
 000000A6
 000000A7
 000000A8
 000000A9
 000000AA
 000000AB
 000000AC
 000000AD
 000000AE
 000000AF
 000000B0
 000000B1
 000000B2
 000000B3
 000000B4
 000000B5
 000000B6
 000000B7
 000000B8
 000000B9
 000000BA
 000000BB
 000000BC
 000000BD
 000000BE
 000000BF
 000000C0
 000000C1
 000000C2
 000000C3
 000000C4
 000000C5
 000000C6
 000000C7
 000000C8
 000000C9
 000000CA
 000000CB
 000000CC
 000000CD
 000000CE
 000000CF
 000000D0
 000000D1
 000000D2
 000000D3
 000000D4
 000000D5
 000000D6
 000000D7
 000000D8
 000000D9
 000000DA
 000000DB
 000000DC
 000000DD
 000000DE
 000000DF
 000000E0
 000000E1
 000000E2
 000000E3
 000000E4
 000000E5
 000000E6
 000000E7
 000000E8
 000000E9
 000000EA
 000000EB
 000000EC
 000000ED
 000000EE
 000000EF
 000000F0
 000000F1
 000000F2
 000000F3
 000000F4
 000000F5
 000000F6
 000000F7
 000000F8
 000000F9
 000000FA
 000000FB
 000000FC
 000000FD
 000000FE
 000000FF



Faint, illegible text in the top-left section, possibly bleed-through from the reverse side of the page.

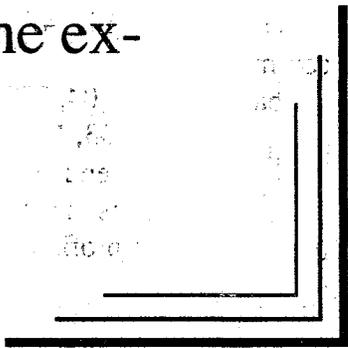
Index

Faint, illegible text in the middle-left section, possibly bleed-through from the reverse side of the page.

Faint, illegible text in the middle-left section, possibly bleed-through from the reverse side of the page.

General index to *Express* and the examples from the text

Faint, illegible text in the bottom-left section, possibly bleed-through from the reverse side of the page.



General Index

This index is the general reference for all the topics discussed in this manual. It lists not only the various functions/routines but also the examples and other points of note. Index entries referring solely to subroutines have their page numbers in typewriter font: `exwrite 178`, for example.

A

accounting 7
argc, argv 187
arguments 187
 ordering 54
 type 54
aspect ratio 64
asynchronous I/O 73, 131
automatic decomposition 178

B

binary file I/O 147
breakpoints 157, 193
buffering
 graphics 140, 148
byte swapping 224

C

clear display surface 105
clipping 75
color 80
 color maps 115, 145
communication
 basic 53–54, 204, 233, 235
 global 158, 163, 169, 173
 hardware dependent 160
 overlapped, asynchronous 214
communication profiler 10, 85, 87, 89
compilers

3L C 43

3L FORTRAN 47

Logical Systems C 39

configuration 9, 16, 82, 176
contour plots 77
coordinate systems 150, 153
customization 16, 176

D

debugging 193
 assembly code 35–37
 breakpoints 157
 interactive 25
 post mortem 18
 source code 27–35
decomposition 124, 178
disk farms 82
domain decomposition 124, 178
domains 53
DONTCARE 204
double buffering 207, 214
downloading
 data 158
 programs 187, 199

E

errors
 asynchronous runtime 63
event driven profiler 15, 95, 97, 99
event driven profiling 102

examples

argc, argv 199
 asynchronous programs 207, 214
 asynchronous system calls 73
 basic communication 205, 233, 235
 broadcast 158
 byte swapping 225
 communicating arrays 233
 communication
 hardware dependent 161
 data-base 69
 decomposition 191
 display processing 179
 DONTCARE 228
 double buffering 207, 214
 exgrid 124, 134, 158, 164
 flushing files 109
 global block 227, 228
 global communication 164, 169, 173
 global maximum 170
 global memory 182
 global semaphores 219
 global sum 169
 global to local data transformation 173
 graphics
 aspect ratio 64
 buffer control 140
 clipping 75
 color 80, 115, 145
 contouring 78
 coordinate systems 150, 153
 erase 105
 flushing 149
 hardware dependencies 141
 initialization 136
 input 113
 line drawing 84, 119, 123
 markers 121
 multiple hosts 91, 93
 polygons 66, 138
 text 93, 117, 149
 hardware dependencies 141
 host capabilities 111
 I/O modes 73
 image analysis 124
 message types 205, 228, 231

multiple hosts 82
 multitasking 107, 182, 219
 parallel I/O 109, 124, 129, 132, 134
 performance
 evaluation
 xtool 195
 performance evaluation 102
 ctool 85, 87, 89
 etool 95, 97, 100
 xtool 143, 197, 202
 preparation for debugging 193
 processor (de)allocation 167, 187, 189,
 199, 222
 processor control 222
 processor sharing 217
 program loading 187
 runtime configuration 176
 runtime errors 63, 140
 shared memory 210
 synchronization 227, 228
 user host routines 69
 wildcard processing 205
 wildcards 228

execution profiler 50, 143, 195, 197, 202

F

file I/O, parallel 131
 file server 12, 82
 flushing
 graphics 148
 flushing files 109

G

global block 227
 global communication 158, 163, 169, 173, 227
 global memory 210
 global operations 169, 173
 graphical input 113
 graphics
 buffering 140, 148
 clipping 75
 color 145
 coordinate systems 150, 153
 device dependencies 141
 initialization 136

line drawing 84, 123
 polygons 138
 symbols 121
 text 117
 graphics servers 91

H

hardware communication 160
 hardware reset 23
 hardware specific graphics 141
 help
 ndb 26
 host capabilities 111
 host programs
 interface to cubix programs 68
 Hostless programming 12

I

I/O 58–59
 non-FORTRAN behavior 109
 parallel 124, 129, 134
 unformatted, reliable 147
 I/O modes
 asynchronous 73
 IALNOD 60
 IALPRC 60
 IHOST 60
 initializing programs 59
 installation 16
 interrupt handling 182

L

libraries 53
 linestyle 119
 load individual nodes 199
 load program “stopped” 193
 loading programs 187, 199

M

message types 54, 204
 process specific 231
 restrictions 205, 236
 messages 54
 multiple host programs 156, 217
 multiple hosts 9, 23, 82, 91, 231

multitasking 56, 107, 182, 210, 219
 multiuser systems 24, 156, 217
 mutual exclusion 210

N

NOCARE 59
 node processes 24
 nodes
 allocation 189
 non-blocking communication 207, 214
 NONODE 60
 NORDER 60

O

overlapped communication 163
 overlaying programs 107

P

performance
 analysis 10, 15, 50
 ctool 85, 89
 etool 99
 evaluation
 ctool 87
 etool 95, 97
 xtool 143, 195, 197, 202
 optimization 16, 176
 performance analysis 102
 polygons 66
 process ID 217
 processor
 (de)allocation 57, 167, 189
 control 222
 synchronization 227
 program
 startup 222
 programming models 53

R

RAM files 18
 read message 204
 rebooting Express 21
 rectangle 66
 runtime configuration 191
 runtime parameters 191

S

semaphores 210
 global 219
send message 235
shared memory 210
sharing processor groups 156, 217
statistics 24, 102
suspend process 219
system constants
 IALNOD 60
 IALPRC 60
 IHOST 60
 NOCARE 59
 NONODE 60
 NORDER 60
system variables 59

T

time measurement 230

W

wildcards 54, 204, 228, 231

X

XPRESS common block 59