
CONCEPTUAL

The
e
t
w
o
r
k

o
r
e
c
t
n
e
s
s

and

e
r
f
o
r
m
a
n
c
e

e
s
t
i
n
g

a
n
g
u
a
g
e

Scott Pakin, pakin@lanl.gov

Copyright © 2009, Los Alamos National Security, LLC

This document describes CONCEPTUAL version 1.0.

Table of Contents

1	Introduction	1
1.1	Motivation	1
1.2	Limitations	2
1.3	Typesetting conventions	2
2	Installation	3
2.1	<i>configure</i>	5
2.2	<i>make</i>	7
2.3	<i>make install</i>	11
3	Usage	14
3.1	The CONCEPTUAL GUI	14
3.1.1	Components	15
3.1.2	Menu bar	17
3.1.3	Command bar	17
3.2	Compiling CONCEPTUAL programs	18
3.3	Supplied backends	20
3.3.1	The <i>c_seq</i> backend	21
3.3.2	The <i>c_mpi</i> backend	22
3.3.3	The <i>c_udgram</i> backend	23
3.3.4	The <i>c_trace</i> backend	23
3.3.5	The <i>c_profile</i> backend	26
3.3.6	The <i>interpret</i> backend	27
3.3.7	The <i>stats</i> backend	31
3.3.8	The <i>picl</i> backend	33
3.3.9	The <i>latex_vis</i> backend	34
3.3.10	The <i>dot_ast</i> backend	38
3.4	Running CONCEPTUAL programs	41
3.5	Interpreting CONCEPTUAL log files	44
3.5.1	Log-file format	44
3.5.2	'ncptl-logextract'	48
3.5.3	'ncptl-logmerge'	70
3.5.4	'ncptl-logunmerge'	74
4	Grammar	77
4.1	Primitives	77
4.2	Expressions	79
4.2.1	Arithmetic expressions	79
4.2.2	Built-in functions	81
4.2.3	Aggregate expressions	90
4.2.4	Aggregate functions	91
4.2.5	Relational expressions	91

4.3	Task descriptions	92
4.3.1	Restricted identifiers	93
4.3.2	Source tasks	93
4.3.3	Target tasks	94
4.4	Communication statements	94
4.4.1	Message specifications	94
4.4.2	Sending	98
4.4.3	Receiving	100
4.4.4	Awaiting completion	101
4.4.5	Multicasting	101
4.4.6	Reducing	102
4.4.7	Synchronizing	103
4.5	I/O statements	103
4.5.1	Utilizing log-file comments	103
4.5.2	Writing to standard output	104
4.5.3	Writing to a log file	104
4.6	Other statements	106
4.6.1	Resetting counters	106
4.6.2	Asserting conditions	107
4.6.3	Delaying execution	107
4.6.4	Touching memory	108
4.6.5	Reordering task IDs	110
4.6.6	Injecting arbitrary code	111
4.7	Complex statements	112
4.7.1	Combining statements	112
4.7.2	Iterating	113
4.7.3	Binding variables	117
4.7.4	Conditional execution	119
4.7.5	Grouping	119
4.8	Header declarations	120
4.8.1	Language versioning	120
4.8.2	Command-line arguments	120
4.9	Complete programs	121
4.10	Summary of the grammar	123
5	Examples	129
5.1	Latency	129
5.2	Hot potato	129
5.3	Hot spot	130
5.4	Multicast trees	131
5.5	Calling MPI functions	133

6	Implementation	135
6.1	Overview	135
6.2	Backend creation	136
6.2.1	Hook methods	137
6.2.2	A minimal C-based backend	139
6.2.3	Generated code	141
6.2.4	Internals	142
6.3	Run-time library functions	145
6.3.1	Constants, variables, and data types	145
6.3.2	Initialization functions	146
6.3.3	Memory-allocation functions	147
6.3.4	Message-buffer manipulation functions	148
6.3.5	Time-related functions	150
6.3.6	Log-file functions	152
6.3.7	Random-task functions	154
6.3.8	Task-mapping functions	154
6.3.9	Queue functions	155
6.3.10	Unordered-set functions	156
6.3.11	Language-visible functions	157
6.3.12	Finalization functions	160
7	Tips and Tricks	161
7.1	Using out-of-bound task IDs to simplify code	161
7.2	Proper use of conditionals	162
7.3	Memory efficiency	163
7.4	Cross-compilation	164
7.5	Implicit dynamic-library search paths	165
7.6	Running without installing	166
7.7	Reporting configuration information	166
8	Troubleshooting	167
8.1	Problems with configure	167
8.1.1	Interpreting configure warnings	167
8.1.2	'PRId64 is not a valid printf conversion specifier'	168
8.1.3	Header is 'present but cannot be compiled'	169
8.1.4	Slow 'checking the maximum length of command line arguments...'	169
8.1.5	'configure' is slow	169
8.1.6	Problems with 'C compiler used for Python extension modules'	170
8.1.7	Manual configuration	170
8.2	Problems with make	171
8.2.1	'Too many columns in multitable item'	171
8.2.2	Can't find 'compiler_version.h'	171
8.2.3	'could not read symbols'	171
8.2.4	Incorrect tools/flags are utilized	172

8.2.5	Compaq compilers on Alpha CPUs	172
8.2.6	‘undefined type, found ‘DEFINE_____’	173
8.2.7	<i>makehelper.py config</i> fails	173
8.2.8	Building on problematic platforms	173
8.3	Problems running	174
8.3.1	‘cannot open shared object file’	174
8.3.2	Miscellaneous mysterious hangs or crashes	175
8.3.3	Extremely noisy measurements	176
8.3.4	Keeping programs from dying on a signal	176
8.3.5	‘Unaligned access’ warnings	176
8.3.6	‘Unable to determine the OS page size’	177
8.3.7	Invalid timing measurements	177
8.3.8	‘TeX capacity exceeded’	177
8.3.9	Bad bounding boxes from <i>latex_vis</i>	178
8.4	When all else fails	178
Appendix A Reserved Words		179
A.1	Keywords	179
A.2	Predeclared variables	184
Appendix B Backend Developer’s Reference		
.....		185
B.1	Method calls	185
B.2	C hooks	187
B.3	Event types	190
B.4	Representing aggregate functions	191
Appendix C Environment Variables		192
Referenced Applications		194
License		195
Index		196

1 Introduction

This document presents a simple, special-purpose language called `CONCEPTUAL`. `CONCEPTUAL` is intended for rapidly generating programs that measure the performance and/or test the correctness of networks and network protocol layers. A few lines of `CONCEPTUAL` code can produce programs that would take significantly more effort to write in a conventional programming language.

`CONCEPTUAL` is not merely a language specification. The `CONCEPTUAL` toolset includes a compiler, run-time library, and associated utility programs that enable users to analyze network behavior quickly, conveniently, and accurately.

1.1 Motivation

A frequently reinvented wheel among network researchers is a suite of programs that test a network's performance. A problem with having umpteen versions of performance tests is that it leads to a variety in the way results are reported; colloquially, apples are often compared to oranges. Consider a bandwidth test. Does a bandwidth test run for a fixed number of iterations or a fixed length of time? Is bandwidth measured as ping-pong bandwidth (i.e., $2 \times \text{message length} \div \text{round-trip time}$) or unidirectional throughput (N messages in one direction followed by a single acknowledgement message)? Is the acknowledgement message of minimal length or as long as the entire message? Does its length contribute to the total bandwidth? Is data sent unidirectionally or in both directions at once? How many warmup messages (if any) are sent before the timing loop? Is there a delay after the warmup messages (to give the network a chance to reclaim any scarce resources)? Are receives nonblocking (possibly allowing overlap in the NIC) or blocking?

The motivation behind creating `CONCEPTUAL`, a simple specification language designed for describing network benchmarks, is that it enables a benchmark to be described sufficiently tersely as to fit easily in a report or research paper, facilitating peer review of the experimental setup and timing measurements. Because `CONCEPTUAL` code is simple to write, network tests can be developed and deployed with low turnaround times—useful when the results of one test suggest a following test that should be written. Because `CONCEPTUAL` is special-purpose its run-time system can perform the following functions, which benchmark writers often neglect to implement:

- logging information about the environment under which the benchmark ran: operating system, CPU architecture and clock speed, timer type and resolution, etc.
- aborting a program if it takes longer than a predetermined length of time to complete
- writing measurement data and descriptive statistics to a variety of output formats, including the input formats of various graph-plotting programs

`CONCEPTUAL` is not limited to network performance tests, however. It can also be used for network verification. That is, `CONCEPTUAL` programs can be used to locate failed links or to determine the frequency of bit errors—even those that may sneak past the network's CRC hardware.

In addition, because `CONCEPTUAL` is a very high-level language, the `CONCEPTUAL` compiler's backend has a great deal of potential. It would be possible for the backend to produce a variety of target formats such as Fortran + MPI, Perl + sockets, C + a network

vendor’s low-level messaging layer, and so forth. It could directly manipulate a network simulator. It could feed into a graphics program to produce a space-time diagram of a CONCEPTUAL program. The possibilities are endless.

1.2 Limitations

Although CONCEPTUAL can express a wide variety of race-free communication patterns it cannot currently express data-dependent communication. For example, CONCEPTUAL cannot express a master-worker pattern in which a master task sends a message to a worker task as a reaction to that particular worker’s sending of a message to the master. Such a communication pattern is not independent of the order in which messages happen to arrive from the workers. Similarly, CONCEPTUAL cannot use run-time performance data to guide its operations. It is therefore not currently possible to express a communication benchmark which repeats until the standard error of some performance metric drops below a given threshold. These limitations may be lifted in a future release of the system.

1.3 Typesetting conventions

The following table showcases the typesetting conventions used in this manual to attribute various meanings to text. Note that not all of the conventions are typographically distinct.

<code>-a</code>	
<code>--abcdef</code>	command-line options (e.g., <code>-C</code> or <code>--help</code>)
<code>ABCDEF</code>	environment variables (e.g., <code>PATH</code>)
<code><abcdef></code>	nonterminals in the CONCEPTUAL grammar (e.g., <code><ident></code>)
<code>abcdef</code>	commands to enter on the keyboard (e.g., <code>make install</code>)
<code>'abcdef'</code>	file and directory names (e.g., <code>'conceptual.pdf'</code>)
<code>ABCDEF</code>	CONCEPTUAL keywords (e.g., <code>RECEIVE</code>)
<code>abcdef</code>	variables, constants, functions, and types in any language (e.g., <code>bit_errors</code> or <code>gettimeofday()</code>)
<code>abcdef</code>	metasyntactic variables and formal function parameters (e.g., <code>fan-out</code>)
<code>'abcdef'</code>	snippets of code, command lines, files, etc. (e.g., <code>'10 MOD 3'</code>)

2 Installation

CONCEPTUAL uses the GNU Autotools (Autoconf, Automake, and Libtool) to increase portability, to automate compilation, and to facilitate installation. As of this writing, CONCEPTUAL has passed *make check* (see [Section 2.2 \[make\], page 7](#)) on the following platforms:

Architecture	OS	Compiler	
IA-32	Linux	'gcc'	(GNU)
		'icc'	(Intel)
		'openc++'	(Open64)
	FreeBSD	'gcc'	(GNU)
	OpenBSD	'gcc'	(GNU)
	NetBSD	'gcc'	(GNU)
	Solaris	'gcc'	(GNU)
		'cc'	(Sun)
Syllable	'gcc'	(GNU)	
Windows (via Cygwin)	'gcc'	(GNU)	
x86-64	Linux	'gcc'	(GNU)
		'pgcc'	(PGI)
		'pathcc'	(PathScale)
	Catamount	'gcc'	(GNU)
		'pgcc'	(PGI)
IA-64	Linux	'gcc'	(GNU)
		'ecc'	(Intel)
PowerPC	Linux	'gcc'	(GNU)
		'xlc'	(IBM)
	AIX	'gcc'	(GNU)
		'xlc'	(IBM)
	MacOS X	'gcc'	(GNU)
BLRTS	'xlc'	(IBM)	
Cell (Power) Cray X1	Linux	'gcc'	(GNU)
	UNICOS/mp	'cc'	(Cray)
UltraSPARC	Solaris	'gcc'	(GNU)
		'cc'	(Sun)
MIPS	IRIX	'gcc'	(GNU)
		'cc'	(MIPSpro)
Alpha	Linux	'gcc'	(GNU)
		'ccc'	(HP)
	Tru64	'gcc'	(GNU)
		'cc'	(HP)
ARM	Linux	'gcc'	(GNU)

In its simplest form, CONCEPTUAL installation works by executing the following commands at the operating-system prompt:

```
./configure
make
make install
```

(‘configure’ is normally run as `./configure` to force it to run from the current directory on the assumption that ‘.’ is not in the executable search path.) We now describe those three installation steps in detail, listing a variety of customization options for each step.

2.1 *configure*

‘configure’ is a Bourne-shell script that analyzes your system’s capabilities (compiler features, library and header-file availability, function and datatype availability, linker flags for various options, etc.) and custom-generates a ‘Makefile’ and miscellaneous other files. ‘configure’ accepts a variety of command-line options. `./configure --help` lists all of the options. The following are some of the more useful ones:

`--disable-shared`

CONCEPTUAL normally installs both static and dynamic libraries. While dynamic libraries have a number of advantages they do need to be installed on all nodes that run the compiled CONCEPTUAL programs. If global installation is not convenient/feasible, `--disable-shared` can be used to force static linking of executables. Note, however, that ‘libncptlmodule.so’, the Python interface to the CONCEPTUAL run-time library, needs to be built as a shared object so that it can be loaded dynamically into a running Python interpreter. `--disable-shared` inhibits the compilation and installation of ‘libncptlmodule.so’.

`--prefix=directory`

`make install` normally installs CONCEPTUAL into the ‘/usr/local’ directory. The `--prefix` option instructs ‘configure’ to write a ‘Makefile’ with a different installation directory. For example, `--prefix=/local/encap/conceptual-1.0` will cause CONCEPTUAL’s files to be installed in ‘/local/encap/conceptual-1.0/bin’, ‘/local/encap/conceptual-1.0/include’, etc.

`--with-ignored-libs=lib1,lib2,...`

In some circumstances it may be necessary to prevent CONCEPTUAL from using certain libraries even when `./configure` detects them and believes them to be usable. The `--with-ignored-libs` configuration option forces `./configure` to ignore one or more specified libraries. Only the base name of each library should be used; omit directory names, the ‘lib’ prefix (on Unix-like systems), and the file suffix. For example, to disable the use of ‘/usr/local/lib/libpapi.a’ you should specify `--with-ignored-libs=papi`.

`--without-fork`

`./configure` detects automatically if your system provides a working `fork()` function. However, it cannot detect if `fork()` correctly spawns a child process but corrupts the parent’s memory map while doing so, as is the case when using some InfiniBand software stacks. The `--without-fork` option inhibits the use

of `fork()` and well as functions that implicitly invoke `fork()` such as `system()` and `popen()`.

`--with-gettimeofday`

The `CONCEPTUAL` run-time library is able to use any of a variety of platform-specific microsecond timers to take timing measurements. (See [Section 6.3.5 \[Time-related functions\]](#), page 150, for a complete list.) The `--with-gettimeofday` option forces the run-time library to utilize instead the generic C `gettimeofday()` function. This can be useful in the rare, but not impossible, case that a quirk in some particular platform misleads one of `CONCEPTUAL`'s other timers. The ‘`validatetimer`’ utility (see [\[Validating the `coNcEPTuaL` timer\]](#), page 10) can help determine whether `--with-gettimeofday` is necessary.

`--with-mpi-wtime`

On some systems the most accurate timer available is provided by the `MPI_Wtime()` function in the MPI library. The `--with-mpi-wtime` option forces the run-time library to measure elapsed time using `MPI_Wtime()` instead of any of the other available timers. (See [Section 6.3.5 \[Time-related functions\]](#), page 150, for a complete list). The ramifications of `--with-mpi-wtime` are threefold:

1. The option requires that you link all `CONCEPTUAL` programs against an MPI library and run them like any other MPI program. (You may need to set `CPPFLAGS`, `LIBS`, `LDFLAGS`, or some of the other command-line variables described below.)
2. `MPI_Wtime()` may be *less* accurate than some of the other timers available to `CONCEPTUAL`. In many MPI implementations, `MPI_Wtime()` simply invokes `gettimeofday()`, for instance.
3. Although this is a rare problem, it may not be safe to invoke `MPI_Wtime()` without first invoking `MPI_Init()`. Fortunately, proper juxtaposition of the two functions is not a concern for the `CONCEPTUAL` C+MPI backend (see [Section 3.3.2 \[The `c_mpi` backend\]](#), page 22), which ensures that `MPI_Init()` is invoked before `MPI_Wtime()`.

In short, you should specify `--with-mpi-wtime` only if you have good reason to believe that `MPI_Wtime()` is likely to produce the most accurate timing measurements on your system.

`CC=C compiler`

‘`configure`’ automatically searches for a C compiler to use. To override its selection, assign a value to `CC` on the command line. For example, `./configure CC=ecc` will cause `CONCEPTUAL` to be built with ‘`ecc`’.

`CFLAGS=C compiler flags`

`LDFLAGS=linker flags`

`CPPFLAGS=C preprocessor flags`

`LIBS=extra libraries`

Like `CC`, these variables override the values determined automatically by ‘`configure`’. As an illustration, `./configure CPPFLAGS="-DSPECIAL`

```
-I/home/pakin/include/special -I." CFLAGS="-O3 -g -Wall -W"
LDFLAGS=--static LIBS="-lz /usr/lib/libstuff.a" assigns values to all
four variables.
```

```
MPICC=C compiler
MPICPPFLAGS=C preprocessor flags
MPILDFLAGS=extra linker flags
MPILIBS=extra libraries
```

These variables are analogous to *CC*, *CPPFLAGS*, *LDFLAGS*, and *LIBS*, respectively. The difference is that they are not used to build the `CONCEPTUAL` run-time library but rather to build user programs targeted to the C+MPI compiler backend. For example, if your MPI installation lacks an ‘`mpicc`’ script, you may need to specify extra header files and libraries explicitly: `./configure MPICPPFLAGS="-I/usr/lib/mpi/include" MPILIBS="-lmpich"`.

As a rather complex illustration of how some of the preceding options (as well as a few mentioned by `./configure --help`) might be combined, the following is how `CONCEPTUAL` was once configured to cross-compile from a Linux/PowerPC build machine to a prototype of the BlueGene/L supercomputer (containing, at the time, 2048 embedded PowerPC processors, each executing a minimal run-time system, BLRTS). IBM’s ‘`xlc`’ compiler was accessed via a wrapper script called ‘`mpcc`’.

```
./configure CFLAGS="-g -O -qmaxmem=64000" CC=/bgl/local/bin/mpcc
CPP="gcc -E" --host=powerpc-ibm-linux-gnu --build=powerpc-unknown-linux-gnu
--with-alignment=8 --with-gettimeofday --prefix=/bgl/bgguest/LANL/ncptl
MPICC=/bgl/local/bin/mpcc CPPFLAGS=-I/BlueLight/floor/bglsys/include
```

It’s always best to specify environment variables as arguments to `./configure` because the ‘`configure`’ script writes its entire command line as a comment to ‘`config.log`’ and as a shell command to ‘`config.status`’ to enable re-running `./configure` with exactly the same parameters.

When `./configure` finishes running it outputs a list of the warning messages that were issued during the run. If no warnings were issued, `./configure` will output ‘`Configuration completed without any errors or warnings.`’. Warnings are also written to ‘`config.log`’ and can therefore be redisplayed at any time by executing a shell command such as `grep WARNING config.log`.

2.2 make

Running `make` by itself will compile the `CONCEPTUAL` run-time library. However, the ‘`Makefile`’ generated by ‘`configure`’ can perform a variety of other actions, as well:

`make check`

Perform a series of regression tests on the `CONCEPTUAL` run-time library. This is a good thing to do after a `make` to ensure that the run-time library built properly on your system. When `make check` finishes it summarizes the test results. The following output signifies a successful completion of `make check`:

```
=====
All 21 tests passed
=====
```

The total number of tests performed depends upon the way that `CONCEPTUAL` was configured. `CONCEPTUAL` components that could not be built are not tested.

If any tests behave unexpectedly it may be possible to gain more information about the source of the problem by re-running `make check` with the `DEBUG` environment variable set to a non-empty value:

```
env DEBUG=1 make check
```

Tests can also be re-run individually:

```
cd tests
env DEBUG=1 runtime_random
```

make clean

make distclean

make maintainer-clean

make clean deletes all files generated by a preceding *make* command. *make distclean* deletes all files generated by a preceding `./configure` command. *make maintainer-clean* delete all generated files. Run *make maintainer-clean* only if you have fairly recent versions of the GNU Autotools (Autoconf 2.53, Automake 1.6, and Libtool 1.4) because those are needed to regenerate some of the generated files. The sequence of operations to regenerate all of the configuration files needed by `CONCEPTUAL` is shown below.

```
libtoolize --force --copy
aclocal
autoheader
automake --add-missing --copy
autoconf
```

make install

Install `CONCEPTUAL`, including the compiler, run-time library, header files, and tools. *make install* is described in detail in [Section 2.3 \[make install\]](#), [page 11](#).

make uninstall

Remove all of the files that *make install* installed. Most of the top-level directories are retained, however, as ‘`make`’ cannot guarantee that these are not needed by other applications.

make info

make pdf

make docbook

Produce the CONCEPTUAL user's guide (this document) in, respectively, Emacs info format, PDF format, or DocBook format. The resulting documentation ('conceptual.info*', 'conceptual.pdf', or 'conceptual.xml') is created in the 'doc' subdirectory.

make ncptl-logextract.html

CONCEPTUAL comes with a postprocessor called 'ncptl-logextract' which facilitates extracting information from CONCEPTUAL-produced log files. The complete 'ncptl-logextract' documentation is presented in [Section 3.5.2 \[ncptl-logextract\], page 48](#). As is readily apparent from that documentation, 'ncptl-logextract' supports an overwhelming number of command-line options. To make the 'ncptl-logextract' documentation more approachable, the *make ncptl-logextract.html* command creates a dynamic HTML version of it (and stores in the 'doc' subdirectory). The result, 'ncptl-logextract.html', initially presents only the top level of the 'ncptl-logextract' option hierarchy. Users can then click on the name of a command-line option to expand or contract the list of suboptions. This interactive behavior makes it easy for a user to get more information on some options without being distracted by the documentation for the others.

make empty.log

Create an empty log file called 'empty.log' which contains a complete prologue and epilogue but no data. This is convenient for validating that the CONCEPTUAL run-time library was built using your preferred build options.

make stylesheets

CONCEPTUAL can automatically produce stylesheets for a variety of programs. These stylesheets make keywords, comments, strings, and other terms in the language visually distinct from each other for a more aesthetically appealing appearance. Currently, *make stylesheets* produces a L^AT_EX 2_ε package ('ncptl.sty'), an a2ps style sheet ('ncptl.ssh'), an Emacs major mode ('ncptl-mode.el'/'ncptl-mode.elc'), a Vim syntax file ('ncptl.vim'), and a Source-highlight language definition ('ncptl.lang'). Note that the 'Makefile' currently lacks provisions for installing these files so whichever stylesheets are desired will need to be installed manually. Stylesheet installation is detailed in [\[Installing stylesheets\], page 11](#).

make modulefile

The Environment Modules package facilitates configuring the operating-system shell for a given application. The *make modulefile* command creates a 'conceptual_1.0' modulefile that checks for conflicts with previously loaded CONCEPTUAL modulefiles then sets the *PATH*, *MANPATH*, and *LD_LIBRARY_PATH* environment variables to values appropriate values as determined by 'configure' (see [Section 2.1 \[configure\], page 5](#)).

Normally, ‘`conceptual_1.0`’ should be installed in the system’s module path (as described by the `MODULEPATH` environment variable). However, users without administrator access can still use the `CONCEPTUAL` modulefile as a convenient mechanism for properly setting all of the environment variables needed by `CONCEPTUAL`:

```
make modulefile
module load ./conceptual_1.0
```

See the ‘`module`’ man page for more information about modules.

make dist

Package together all of the files needed to rebuild `CONCEPTUAL`. The resulting file is called ‘`conceptual-1.0.tar.gz`’ (for this version of `CONCEPTUAL`).

make all

Although `all` is the default target it can also be specified explicitly. Doing so is convenient when performing multiple actions at once, e.g., *make clean all*.

make tags

Produce/update a ‘`TAGS`’ file that the Emacs text editor can use to find function declarations, macro definitions, variable definitions, `typedefs`, etc. in the `CONCEPTUAL` run-time library source code. This is useful primarily for developers wishing to interface with the `CONCEPTUAL` run-time library. Read the Emacs documentation for *M-x find-tag* for more information.

make gui

Compile the `CONCEPTUAL` GUI, producing ‘`ncptlGUI-1.0.jar`’. Note that compilation requires both a Java compiler (e.g., `javac`) and the Jython Python-to-Java compiler (`jythonc`). Unfortunately, at the time of this writing (January 2009), `jythonc`’s future is uncertain (cf. <http://www.jython.org/Project/jythonc.html>). Hence, *make gui* has been tested only with `jythonc` version 2.2.x, not any later versions.

Validating the coNCePTuaL timer

make automatically builds a program called ‘`validatetimer`’. ‘`validatetimer`’ helps validate that the real-time clock used by the `CONCEPTUAL` run-time library accurately measures wall-clock time. The idea is to compare `CONCEPTUAL`’s timer to an external clock (i.e., one not associated with the computer). Simply follow the program’s prompts:

```
% validatetimer
Press <Enter> to start the clock ...
Press <Enter> again in exactly 60 seconds ...

coNCePTuaL measured 60.005103 seconds.
coNCePTuaL timer error = 0.008505%
```


If the difference between CONCEPTUAL's timer and an external clock is significant, then performance results from CONCEPTUAL—and possibly from other programs, as well—should not be trusted. Note that only extreme differences in timings are significant; there will always be *some* error caused by human response time and by system I/O speed. In the case that there *is* an extreme performance difference,¹ the `--with-gettimeofday` option to `'configure'` (see Section 2.1 [configure], page 5) may be a viable workaround.

`'validatetimer'` takes an optional command-line argument, which is the number of seconds of wall-clock time to expect. The default is `'60'`. Larger numbers help amortize error; smaller numbers enable the program to finish sooner.

2.3 `make install`

The CONCEPTUAL compiler and run-time library are installed with `make install`. Although `'configure'` can specify the default installation directory (see Section 2.1 [configure], page 5), this can be overridden at `make install` time in one of two ways. `make DESTDIR=prefix install` prepends `prefix` to every directory when installing. However, the files are installed believing that `DESTDIR` was not specified. For example, `make DESTDIR=/mnt install` would cause executables to be installed into `'/mnt/usr/local/bin'`, but if any of these are symbolic links, the link will omit the `'/mnt'` prefix.

The second technique for overriding installation directories is to specify a new value for `'prefix'` on the command line. That is, `make prefix=/opt/ncptl install` will install into `'/opt/ncptl/bin'`, `'/opt/ncptl/include'`, `'/opt/ncptl/man'`, etc., regardless of the `--prefix` value given to `'configure'`. CONCEPTUAL's `'Makefile'` provides even finer-grained control than that. Instead of—or in addition to—specifying a `prefix` option on the command line, individual installation directories can be named explicitly. These include `bindir`, `datadir`, `libdir`, `includedir`, `infodir`, `mandir`, `pkgdatadir`, `pythondir`, and many others. Scrutinize the `'Makefile'` to find a particular directory that should be overridden.

The remainder of this section presents a number of optional installation steps that add CONCEPTUAL support to a variety of third-party software packages.

Installing stylesheets

The `make stylesheets` command (see Section 2.2 [make], page 7) produces a variety of stylesheets for presenting CONCEPTUAL code in a more pleasant format than ordinary, monochromatic text. Stylesheets must currently be installed manually as per the following instructions:

`'ncptl.sty'`

`'ncptl.sty'` is typically installed in `'texmf/tex/latex/misc'`, where `texmf` is likely to be `'/usr/local/share/texmf'`. On a Web2c version of \TeX the command `kpsewhich -expand-var='$TEXMFLOCAL'` should

¹ To date, extreme performance differences have been observed primarily on PowerPC-based systems. The PowerPC cycle counter is clocked at a different rate from the CPU speed, which may confuse CONCEPTUAL. The run-time library compensates for this behavior on all tested platforms (see Chapter 2 [Installation], page 3), but the user should nevertheless make sure to run `'validatetimer'` to verify that CONCEPTUAL's timer is sufficiently accurate.

output the correct value of *texmf*. In most T_EX distributions the filename database needs to be refreshed after a new package is installed. See <http://www.tex.ac.uk/cgi-bin/texfaq2html?label=instpackages> for more information. ‘ncpt1.sty’ is merely a customization of the ‘listings’ package that defines a new language called ‘ncpt1’. See the ‘listings’ documentation for instructions on typesetting source code.

‘ncpt1.ssh’

Running `a2ps --list=defaults` outputs (among other things) the a2ps library path. ‘ncpt1.ssh’ should be installed in one of the ‘sheets’ directories listed there, typically ‘/usr/share/a2ps/sheets’.

‘ncpt1-mode.el’

‘ncpt1-mode.elc’

‘ncpt1-mode.el’ and ‘ncpt1-mode.elc’ belong in a local Elisp directory that is part of the Emacs load-path, e.g., ‘/usr/share/emacs/site-lisp’. The following Elisp code, which belongs in ‘~/emacs’ for GNU Emacs or ‘~/Xemacs/init.el’ for XEmacs, makes Emacs set `ncpt1-mode` whenever opening a file with extension ‘.ncpt1’:

```
(autoload 'ncpt1-mode "ncpt1-mode"
  "Major mode for editing coNCePTuaL programs." t)
(add-to-list 'auto-mode-alist ('("\\.ncpt1$" . ncpt1-mode))
```

Syntax highlighting should be enabled by default. If it isn't, the Emacs command `M-x font-lock-mode` should enable it for the current buffer.

‘ncpt1.vim’

Vim's syntax-file directory may be named after the Vim version, e.g., ‘/usr/share/vim/vim61/syntax’ for Vim 6.1. Put ‘ncpt1.vim’ there. To associate ‘.ncpt1’ files with CONCEPTUAL code, the following lines need to be added to Vim's ‘filetype.vim’ file somewhere between the ‘augroup filetypepedetect’ line and the ‘augroup END’ line:

```
" coNCePTuaL
au BufNewFile,BufRead *.ncpt1          setf ncpt1
```

‘ncpt1.lang’

Source-highlight stores all of its helper files in a single directory, typically ‘/usr/share/source-highlight’. Put ‘ncpt1.lang’ there. To associate ‘.ncpt1’ files with CONCEPTUAL code you will also need to add the following line to the ‘lang.map’ file in the same directory:

```
ncpt1 = ncpt1.lang
```

SLOCCount

SLOCCount is a utility that counts the number of lines of code in a file, excluding blank lines and comments. SLOCCount supports a variety of programming languages and it is straightforward to get it to support CONCEPTUAL, as well. The procedure follows the “Adding support for new languages” section of the SLOCCount manual:

1. Create an ‘ncpt1_count’ script with the following contents:

```
#!/bin/sh

generic_count "#" $@
```

2. Mark the script executable and install it somewhere in your executable search path.
3. Edit SLOCCount’s ‘break_filelist’ Perl script to include the following association in the %file_extensions hash:

```
"ncpt1" => "ncpt1",    # coNcEPTuaL
```

pkg-config

The pkg-config utility helps ensure that programs are given appropriate compiler and linker flags to use a particular package’s C header files and libraries. CONCEPTUAL’s ‘configure’ script (see [Section 2.1 \[configure\], page 5](#)) automatically produces a pkg-config configuration file for the CONCEPTUAL header file (‘ncpt1.h’) and run-time library (‘libncpt1’). This configuration file, ‘ncpt1.pc’, should be installed in one of the directories searched by pkg-config (‘/usr/lib/pkgconfig’ on some systems). Once ‘ncpt1.pc’ is installed, pkg-config can be used to compile C programs that require the CONCEPTUAL header file and link programs that require the CONCEPTUAL run-time library, as is shown in the following example:

```
cc ‘pkg-config --cflags ncpt1’ -c myprog.c
cc -o myprog myprog.o ‘pkg-config --libs ncpt1’
```

3 Usage

CONCEPTUAL is more than just a language; it is a complete toolset which consists of the following components:

- the CONCEPTUAL language (see [Chapter 4 \[Grammar\]](#), page 77)
- a compiler and run-time library for CONCEPTUAL programs (see [Section 3.2 \[Compiling coNcEPTuaL programs\]](#), page 18)
- a set of compiler backends that can generate code for a variety of languages and communication layers (see [Section 3.3 \[Supplied backends\]](#), page 20)
- utilities to help analyze the results (see [Section 3.5.2 \[ncptl-logextract\]](#), page 48)

This chapter explains how to compile and run CONCEPTUAL programs and how to interpret the log files they output.

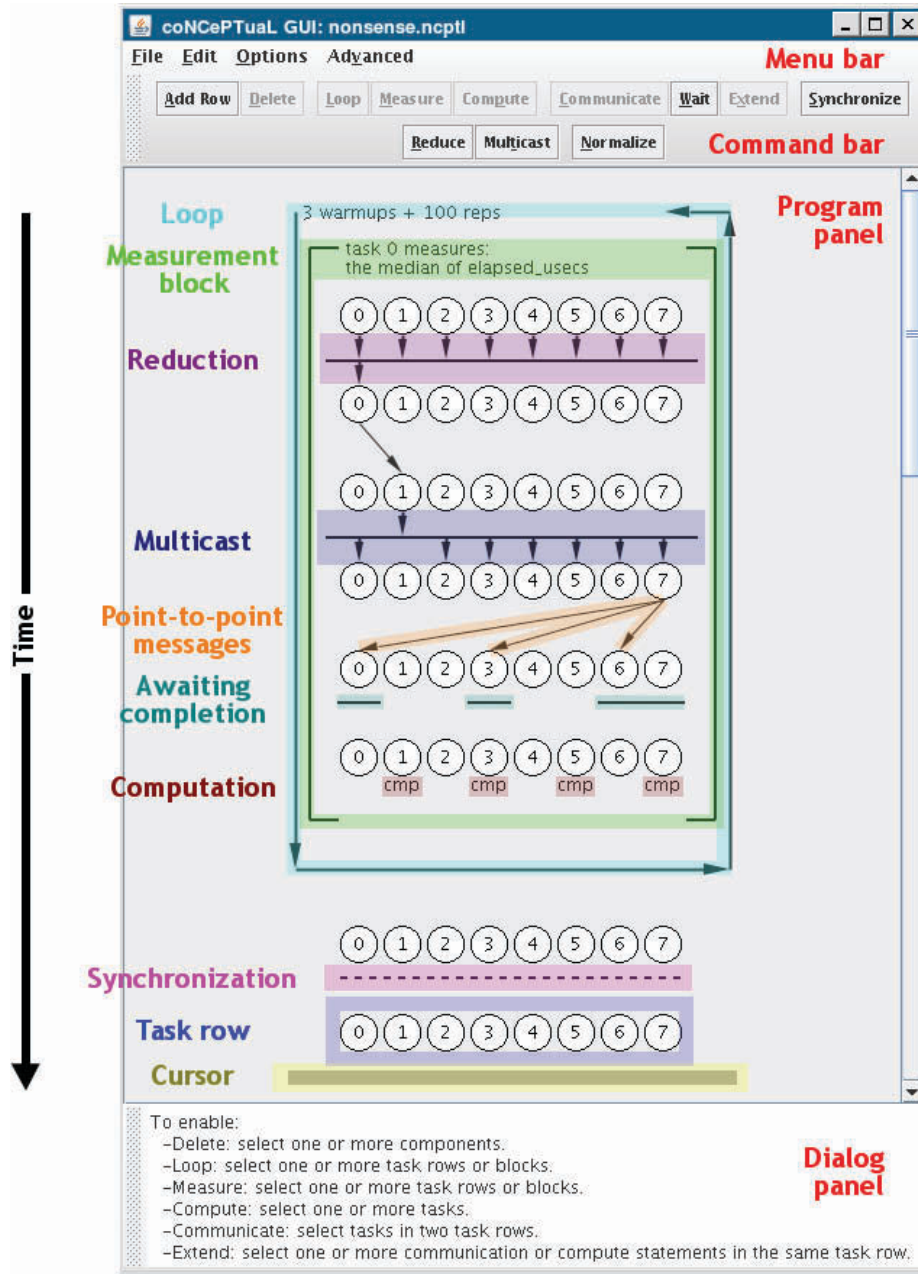
3.1 The coNcEPTuaL GUI

The CONCEPTUAL graphical user interface (GUI) is the easiest way to get started with CONCEPTUAL. Instead of writing code in the CONCEPTUAL language (documented in detail in [Chapter 4 \[Grammar\]](#), page 77), a user merely *draws* a communication pattern using the mouse, and the CONCEPTUAL GUI automatically produces a CONCEPTUAL program from that illustration. The generated program can then be compiled just like a hand-coded CONCEPTUAL program as per the instructions in [Section 3.2 \[Compiling coNcEPTuaL programs\]](#), page 18.

The CONCEPTUAL GUI is written in Java and therefore requires a Java virtual machine (JVM) to run. However, the CONCEPTUAL GUI is quite portable and should run identically on every platform for which a JVM exists. The following is a typical command for launching the CONCEPTUAL GUI from the command line:

```
java -jar ncptlGUI-1.0.jar
```

The CONCEPTUAL GUI is split into two main panels. The *program panel* displays the components that make up the program and how they interact with each other. The *dialog panel* displays fields for setting the options of selected components. Furthermore, a *menu bar* provides access to various GUI-wide operations, and a *command bar* includes buttons for creating new components in the program panel.



3.1.1 Components

Components are graphical representations of coNCEPTUAL objects and operations used by the coNCEPTUAL GUI to specify programs. Components are added to the program panel by clicking on their corresponding buttons in the command bar. To select a component in the program panel, simply left-click on it. If the component has parameters that can be edited, a dialog will appear in the dialog panel. Multiple components can be selected by dragging the mouse over target components or holding down Ctrl as you left-click on several components.

The coNCEPTUAL GUI lets a user manipulate the following components:

task/task row

Tasks represent operational units in CONCEPTUAL programs and are analogous to a process or thread in a parallel program. For example, a CONCEPTUAL operation like point-to-point communication has a source task that specifies how the message is sent and a target task that specifies how the message is received. Tasks are displayed graphically as numbered circles in the CONCEPTUAL GUI. A task row represents the total number of tasks available for operations at each step of a CONCEPTUAL program. See [Section 4.3 \[Task descriptions\]](#), page 92, for information on how to specify subsets of a program's tasks.

communication message

Messages between tasks are displayed as directed edges (arrows) from source task to target task in the CONCEPTUAL GUI. Messages are added to the program panel via the command bar or by dragging the mouse from a source task to a target task. Communication messages in the CONCEPTUAL GUI correspond to the SEND and RECEIVE statements in the CONCEPTUAL language (see [Section 4.4.2 \[Sending\]](#), page 98, and [Section 4.4.3 \[Receiving\]](#), page 100).

awaiting completion

Messages that are sent/received asynchronously must eventually be waited on. Awaits message completion is displayed as a solid line under the associated tasks in the CONCEPTUAL GUI. Awaiting completion in the CONCEPTUAL GUI correspond to the AWAIT COMPLETION statement in the CONCEPTUAL language (see [Section 4.4.4 \[Awaiting completion\]](#), page 101).

loop

One can add a loop around selected components to repeat the corresponding CONCEPTUAL operations. Loops in the CONCEPTUAL GUI correspond to the FOR statement in the CONCEPTUAL language (see [Section 4.7.2 \[Iterating\]](#), page 113).

measurement block

One can log timing or other measurements of CONCEPTUAL operations by placing them in a measurement block. Measurement blocks in the CONCEPTUAL GUI correspond to the LOGS statement in the CONCEPTUAL language (see [Section 4.5.3 \[Writing to a log file\]](#), page 104).

computation/sleeping

Artificial computation (really a spin loop) and sleeping, both of which delay the program for a given length of time, can be performed on a set of tasks. Computation is shown with 'cmp' under a task in the CONCEPTUAL GUI, and sleeping is shown with 'slp'. Computation/sleeping in the CONCEPTUAL GUI corresponds to the COMPUTE and SLEEP statements in the CONCEPTUAL language (see [Section 4.6.3 \[Delaying execution\]](#), page 107).

multicasting

A multicast operation sends a message from a source task to multiple target tasks. Multicasting in the CONCEPTUAL GUI corresponds to the MULTICAST statement in the CONCEPTUAL language (see [Section 4.4.5 \[Multicasting\]](#), page 101).

reduction A reduction operation combines messages from multiple source tasks to a single target task. Reduction in the CONCEPTUAL GUI corresponds to the REDUCE statement in the CONCEPTUAL language (see [Section 4.4.6 \[Reducing\]](#), [page 102](#)).

synchronization

Barrier synchronization forces a set of tasks to wait until each task in the set reaches the synchronization point before any task in the set proceeds past the synchronization point. Synchronization is displayed as a dotted line under the associated tasks in the CONCEPTUAL GUI. Synchronization in the CONCEPTUAL GUI corresponds to the SYNCHRONIZE statement in the CONCEPTUAL language (see [Section 4.4.7 \[Synchronizing\]](#), [page 103](#)).

3.1.2 Menu bar

The *File* menu, which appears only when the CONCEPTUAL GUI is granted access to the filesystem, contains *New*, *Open*, *Save*, *Save As*, *Print*, and *Quit* commands that exhibit the expected behavior. Programs are saved as CONCEPTUAL source code that can then be compiled with the CONCEPTUAL compiler as per [Section 3.2 \[Compiling coNCEPTuaL programs\]](#), [page 18](#). *Print* prints a graphical view of the program as it appears on screen. (See [Section 3.3.9 \[The latex_vis backend\]](#), [page 34](#), for a more sophisticated way to produce graphical views of CONCEPTUAL programs.)

The *Edit* menu provides the usual *Cut*, *Copy*, *Paste*, and *Undo* commands.

Options→*Settings* opens a dialog in the dialog panel for setting the number of tasks in a task row. The default number of tasks in a task row is 16.

Advanced→*Add conditional* adds a conditional statement to a program at the current cursor position in the program panel. A dialog in the dialog panel will open for entering the conditional expression. A placeholder expression ‘1 = 1’ is set by default.

Advanced→*Command line options* opens a dialog in the dialog panel for adding command-line options to a program. A placeholder ‘reps’ (‘number of repetitions’) option is set by default when this command is selected.

3.1.3 Command bar

The following buttons appear on the command bar:

Add Row Insert a new, empty task row at the cursor position.

Delete Delete the selected components.

Loop Add a loop around the selected components.

Measure Add a measurement block around the selected components.

Compute Make the selected tasks “compute” or sleep for a length of time.

Communicate

Add point-to-point communication between selected tasks (different task rows). This can also be achieved by dragging an arrow from a source task to a target task.

Wait Make the selected tasks (same task row) or all tasks in the row above the cursor wait for all outstanding messages sent or received asynchronously to complete.

- Extend* Extend a communication or computation pattern across an entire task row.
- Synchronize*
Synchronize the selected tasks (same task row) or all tasks in the task row above the cursor.
- Reduce* Reduce data from the selected tasks in one task row to the selected tasks in the next task row. With no selection, reduce data from all tasks above the cursor to task 0 below the cursor.
- Multicast* Multicast data from the selected tasks in one task row to the selected tasks in the next task row. With no selection, multicast data from task 0 above the cursor to all tasks below the cursor.
- Normalize* Put the program into standard form as it will appear when translated into CONCEPTUAL code. After normalization, components are drawn as early in time as possible without changing the program’s semantics.

3.2 Compiling coNCePTuaL programs

The CONCEPTUAL compiler is called ‘ncpt1’ and is, by default, installed into ‘/usr/local/bin’. Executing `ncpt1 --help` produces a brief usage string:

```
Usage: ncpt1 [--backend=<string>] [--quiet] [--no-link | --no-compile]
        [--keep-ints] [--lenient] [--filter=<sed expr>] [--output=<file>]
        <file.ncpt1> | --program=<program>
        [<backend-specific options>]
```

The usage string is followed by a list of installed backends.

The following list describes each compiler option in turn:

--backend (abbreviation: *-b*)

Specify the module that CONCEPTUAL should use as the compiler backend. `ncpt1` must be told which backend to use with either `--backend=backend` or by setting the environment variable `NCPTL_BACKEND` to the desired backend. Running `ncpt1 --help` lists the available backends. Most CONCEPTUAL backends are code generators. For example, `c_mpi` causes ‘ncpt1’ to compile CONCEPTUAL programs into C using MPI as the communication library. However, a backend need not generate code directly—or at all. The `c_trace` backend (see [Section 3.3.4 \[The c_trace backend\], page 23](#)), for instance, supplements the code generated by another backend by adding tracing output to it.

`ncpt1` searches for backends first using `NCPTL_PATH`, an environment variable containing a colon-separated list of directories (default: empty); then, in the directory in which CONCEPTUAL installed all of its Python files; and finally, in the default Python search path. Non-directories (e.g., the ‘.zip’ archives used in newer versions of Python) are not searched.

If no backend is specified, `ncpt1` runs the given CONCEPTUAL program through the lexer, parser, and semantic analyzer but does not generate an output file.

--quiet (abbreviation: *-q*)

The **--quiet** option tells ‘ncpt1’ and the chosen backend to output minimal status information.

--no-link (abbreviation: *-c*)

By default, ‘ncpt1’ instructs the backend to compile and link the user’s CONCEPTUAL program into an executable file. **--no-link** tells the backend to skip the linking step and produce only an object file.

--no-compile (abbreviation: *-E*)

By default, ‘ncpt1’ instructs the backend to compile and link the user’s CONCEPTUAL program into an executable file. **--no-compile** tells the backend to skip both the compilation and the linking step and to produce only a source file in the target language.

--keep-ints (abbreviation: *-K*)

CONCEPTUAL backends normally delete any files created as part of the compiling or linking process. **--keep-ints** tells ‘ncpt1’ and the chosen backend to preserve their intermediate files.

--lenient (abbreviation: *-L*)

The **--lenient** option tells the compiler to permit certain constructs that would otherwise result in a compilation error. First, using the same command-line option (either the long or short variant) for two different variables normally generates an ‘Option *opt* is multiply defined’ error. (See [Section 4.8.2 \[Command-line arguments\]](#), page 120, for a description of how to declare command-line options in CONCEPTUAL.) **--lenient** tells the CONCEPTUAL compiler to automatically rename duplicate options to avoid conflicts. Only the option strings can be renamed; the programmer must still ensure that the option variables are unique. Second, using a variable without declaring it normally produces an error message at compile time. Passing **--lenient** to ‘ncpt1’ tells the compiler to automatically generate a command-line option for each missing variable. This is convenient when entering brief programs on the command line with **--program** (described below) as it can save a significant amount of typing.

--filter (abbreviation: *-f*)

The **--filter** option applies a ‘sed’-style substitution expression to the backend-translated code (e.g., a ‘.c’ file output by the `c_udgram` backend or a ‘.tex’ file output by the `latex_vis` backend) before the backend compiles it. The **--filter** option can be used multiple times on the command line; filters are applied in the order specified. Substitution expressions must be of the form `s/pattern/replacement/flags`, although the ‘/’ characters can be replaced by any other character. *pattern* is a regular expression; *replacement* is an optional replacement string; and, *flags* is a sequence of zero or more modifiers from the set {‘i’, ‘l’, ‘m’, ‘s’, ‘u’, ‘x’}, as described in the Python Library Reference. For example, ‘i’ means to perform a case-insensitive substitution. In addition, the ‘g’ flag performs a global search-and-replace instead of replacing only the first occurrence of *pattern*. An important difference between **--filter** and ‘sed’ is that omitting the ‘g’ flag instructs

`--filter` to make at most one substitution *total* while it instructs ‘sed’ to make at most one substitution *per line*.

`--output` (abbreviation: `-o`)

‘ncpt1’ normally writes its output to a file with the same base name as the input file (or ‘a.out’ if the program was specified on the command line using `--program`). `--output` lets the user specify a file to which to write the generated code.

`--program` (abbreviation: `-p`)

Because CONCEPTUAL programs can be quite short `--program` enables a program to be specified in its entirety on the command line. The alternative to using `--program` is to specify the name of a file containing a CONCEPTUAL program. By convention, CONCEPTUAL programs have a ‘.ncpt1’ file extension.

The following—to be entered without line breaks—is a sample command line:

```
ncpt1 --backend=c_mpi --output=sample.c --program='Task 0 sends a 0
byte message to task 1 then task 1 sends a 0 byte message to task 0
then task 0 logs elapsed_usecs/2 as "Startup latency (usecs)".'
```

‘ncpt1’ stops processing the command line at the first unrecognized option it encounters. That option and all subsequent options—including those which ‘ncpt1’ would otherwise process—are passed to the backend without interpretation by ‘ncpt1’. Furthermore, the `--` (i.e., empty) option tells ‘ncpt1’ explicitly to stop processing the command line at that point. For example, in the command `ncpt1 --lenient myprogram.ncpt1 -- --help`, ‘ncpt1’ will process the `--lenient` option but will pass `--help` to the backend even though ‘ncpt1’ has its own `--help` option.

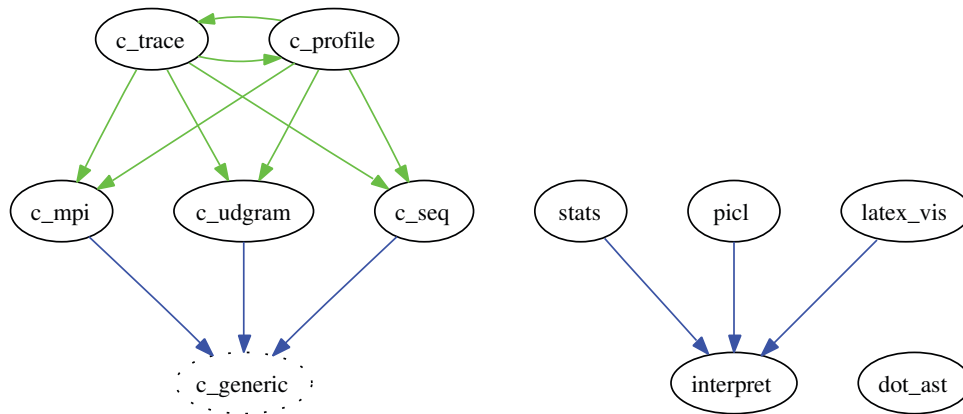
3.3 Supplied backends

The CONCEPTUAL 1.0 distribution includes the following compiler backends:

- `c_seq` Generate ANSI C code with no communication support.
- `c_mpi` Generate ANSI C code with calls to the MPI library for communication.
- `c_udgram` Generate ANSI C code that communicates using Unix-domain (i.e., local to a single machine) datagram sockets.
- `c_trace` Instrument a C-based backend either to include a call to `fprintf()` before every program event or to utilize the ‘curses’ library to display graphically the execution of a selected task.
- `c_profile` Instrument a C-based backend either to write event timings and tallies to each log file or to the standard error device if the CONCEPTUAL program doesn’t use a log file.
- `interpret` Interpret a CONCEPTUAL program, simulating any number of processors and checking for common problems such as deadlocks and mismatched sends and receives.

- stats** Output statistics of a program’s execution—message tallies, byte counts, communication peers, network bisection crossings, event tallies, etc.
- picl** Output in the PICL trace format a logical-time trace of a CONCEPTUAL program’s communication pattern.
- latex_vis** Use L^AT_EX to produce an Encapsulated PostScript visualization of a program’s communication pattern.
- dot_ast** Output a program’s parse tree in the Graphviz dot format.

CONCEPTUAL employs a highly modular software structure for its backends. Many of the backends listed above are built atop other backends. The following figure illustrates the current set of dependencies:



Some dependencies are defined as a static characteristic of a backend. For example, the `c_mpi` backend is hardwired to derive some of its functionality from `c_generic`. Other dependencies are determined dynamically. For example, the `c_trace` backend must be instructed to derive its functionality from one of `c_profile`, `c_mpi`, `c_udgram`, or `c_seq`. (See Section 3.3.4 [The `c_trace` backend], page 23, for more information on the `c_trace` backend.)

Except for `c_generic`, all of the compiler backends are described in turn in the following sections. The `c_generic` backend is unique because it is used exclusively to construct C-based backends; it does nothing by itself. See Section 6.2 [Backend creation], page 136, for a detailed description of `c_generic`.

Each backend accepts a `--help` option that explains the backend’s command-line options. The easiest way to request help from a specific backend is with `'ncpt1 --backend=backend -- --help'`. The empty `--` option, mentioned in Section 3.2 [Compiling coNCePTuaL programs], page 18, prevents the compiler from intercepting `--help` and providing the standard, backend-independent information.

3.3.1 The `c_seq` backend

The `c_seq` backend is intended primarily to provide backend developers with a minimal C-based backend that can be used as a starting point for creating new backends. See Section 6.2 [Backend creation], page 136, explains how to write backends.

3.3.2 The `c_mpi` backend

The `c_mpi` backend is CONCEPTUAL's workhorse. It generates parallel programs written in ANSI C that communicate using the industry-standard MPI messaging library.

By default, `c_mpi` produces an executable program that can be run with `'mpirun'`, `'prun'`, `'pdsh'`, or whatever other job-launching program is normally used to run MPI programs. When `'ncptl'` is run with the `--no-link` option, `c_mpi` produces an object file that needs to be linked with the appropriate MPI library. When `'ncptl'` is run with the `--no-compile` option, `c_mpi` outputs ANSI C code that must be both compiled and linked.

`c_mpi` honors the following environment variables when compiling and linking C+MPI programs: `MPICC`, `MPICPPFLAGS`, `MPICFLAGS`, `MPILDFLAGS`, `MPILIBS`. If any of these variables is not found in the environment, `c_mpi` will use the value specified/discovered at configuration time (see [Section 2.1 \[configure\], page 5](#)). `MPICC` defaults to the value of `CC`; `MPICFLAGS` defaults to the value of `CFLAGS`; the remaining variables are appended respectively to `CPPFLAGS`, `LDFLAGS`, and `LIBS`.

The following is a complete list of MPI functions employed by the `c_mpi` backend: `MPI_Allreduce()`, `MPI_Barrier()`, `MPI_Bcast()`, `MPI_Comm_rank()`, `MPI_Comm_size()`, `MPI_Comm_split()`, `MPI_Errhandler_create()`, `MPI_Errhandler_set()`, `MPI_Finalize()`, `MPI_Init()`, `MPI_Irecv()`, `MPI_Isend()`, `MPI_Recv()`, `MPI_Reduce()`, `MPI_Send()`, `MPI_Ssend()`, `MPI_Waitall()`. In addition, if `./configure` is passed the `--with-mpi-wtime` option as described in [Section 2.1 \[configure\], page 5](#), then *all* backends that utilize the CONCEPTUAL run-time library, including `c_mpi`, will use `MPI_Wtime()` for taking performance measurements.

Command-line options for `c_mpi`

When `'ncptl'` is passed `--backend=c_mpi` as a command-line option, `c_mpi` processes the following backend-specific command-line options:

`--ssend` In the generated code, `MPI_Isend()` and `MPI_Irecv()` are used for asynchronous communication and `MPI_Send()` and `MPI_Recv()` are normally used for synchronous communication. However, the `c_mpi`-specific compiler option `--ssend` instructs `c_mpi` to replace all calls to `MPI_Send()` in the generated code with calls to `MPI_Ssend()`, MPI's synchronizing send function. A program's log files indicate whether the program was built to use `MPI_Send()` or `MPI_Ssend()`.

`--reduce=MPI_Op`

The default reduction operation is `MPI_SUM` but a different operation can be specified using the `c_mpi`-specific `--reduce` compiler option. A program's log files indicate the reduction operation that was used.

Implementation of reductions

The `c_mpi` backend implements the CONCEPTUAL `REDUCE` statement (see [Section 4.4.6 \[Reducing\], page 102](#)) as follows. Many-to-one reductions are implemented with a single call to `MPI_Reduce()`. Many-to-many reductions in which the sources exactly match the targets are implemented with a single call to `MPI_Allreduce()`. All other reductions are implemented by calling `MPI_Reduce()` to reduce the data to the first target task then `MPI_Bcast()` to distribute the reduced data to the remaining targets.

Because MPI requires that the source and target buffers in an `MPI_Reduce()` or `MPI_Allreduce()` be different, the `c_mpi` backend utilizes *two* buffers when `UNIQUE` (see [Unique messages], page 96) is specified. It utilizes two *adjacent* buffers when `FROM BUFFER` or `INTO BUFFER` (see [Buffer control], page 98) is specified.

3.3.3 The `c_udgram` backend

`CONCEPTUAL` program development on a workstation is facilitated by the `c_udgram` backend. `c_udgram` runs on only a single machine but, unlike `c_seq`, supports all of `CONCEPTUAL`'s communication statements. Communication is performed over Unix-domain datagram sockets. Unix-domain datagrams are reliable and guarantee order (unlike UDP/IP datagrams) but have a maximum packet size. `c_udgram` backend write this maximum to every log file and automatically packetizes larger messages.

By default, `c_udgram` produces an executable program that can be run directly from the command line. When `'ncpt1'` is run with the `--no-link` option, `c_udgram` produces an object file that needs to be linked with the appropriate sockets library (on systems that require a separate library for socket calls). When `'ncpt1'` is run with the `--no-compile` option, `c_udgram` outputs ANSI C code that must be both compiled and linked. Like all C-based backends, `c_udgram` honors the `CC`, `CPPFLAGS`, `LDFLAGS`, and `LIBS` environment variables when compiling and linking. Values not found in the environment are taken from those specified/discovered at configuration time (see Section 2.1 [configure], page 5).

In addition to supporting the default set of command-line options, programs generated using the `c_udgram` backend further support a `--tasks` option that designates the number of tasks to use:

<code>-T, --tasks=<number></code>	Number of tasks to use [default: 1]
---	-------------------------------------

`c_udgram` programs spawn one OS-level process for each task in the program. They also create a number of sockets in the current directory named `'c_udgram_{tag}'`. These are automatically deleted if the program exits cleanly but will need to be removed manually in the case that the program is killed by a non-trappable signal.

3.3.4 The `c_trace` backend

While most `CONCEPTUAL` backends are code generators, the `c_trace` backend adds tracing code to the code produced by a code-generating backend. `c_trace` is useful as a debugging aid and as a means to help understand the control flow of a `CONCEPTUAL` program.

Command-line options for `c_trace`

When `'ncpt1'` is passed `--backend=c_trace` as a command-line option, `c_trace` processes the following backend-specific command-line options:

`--trace=backend`

Specify a backend that will produce C code for `c_trace` to trace. The `--trace` option is required to use `c_trace`; `c_trace` will issue an error message if `--trace` is not specified. The restrictions on `backend` are that it must produce C code and must be derived from the `c_generic` backend (see Section 6.2 [Backend creation], page 136). Improper backends cause `c_trace` to abort abnormally.

`--curses` Instead of injecting `fprintf()` statements into the generated C code, inject calls to the ‘`curses`’ (or ‘`ncurses`’) library to show graphically the line of code currently executing on a given processor.

Default `c_trace` tracing

Without `--curses`, `c_trace` alters the generated C code to write data like the following to the standard error device:

```
[TRACE] phys: 1 | virt: 1 | action: RECV | event: 1 / 44001 | lines: 18 - 18
[TRACE] phys: 0 | virt: 0 | action: RESET | event: 1 / 88023 | lines: 17 - 17
[TRACE] phys: 0 | virt: 0 | action: SEND | event: 2 / 88023 | lines: 18 - 18
[TRACE] phys: 0 | virt: 0 | action: RECV | event: 3 / 88023 | lines: 19 - 19
[TRACE] phys: 1 | virt: 1 | action: SEND | event: 2 / 44001 | lines: 19 - 19
[TRACE] phys: 1 | virt: 1 | action: RECV | event: 3 / 44001 | lines: 18 - 18
[TRACE] phys: 0 | virt: 0 | action: CODE | event: 4 / 88023 | lines: 20 - 21
[TRACE] phys: 0 | virt: 0 | action: RESET | event: 5 / 88023 | lines: 17 - 17
[TRACE] phys: 0 | virt: 0 | action: SEND | event: 6 / 88023 | lines: 18 - 18
[TRACE] phys: 0 | virt: 0 | action: RECV | event: 7 / 88023 | lines: 19 - 19

:
```

The format is designed to be easy to read and easy for a program to parse. Each line of trace data begins with the string ‘`[TRACE]`’ and lists the (physical) processor number, the (virtual) task ID, the action (a.k.a., event type) that is about to be performed, the current event number and total number of events that will execute on the given processor, and the range of lines of source code to which the current event corresponds. An “event” corresponds more-or-less to a statement in the CONCEPTUAL language.¹ Loops are unrolled at initialization time and therefore produce no events. [Section B.3 \[Event types\], page 190](#), lists and briefly describes the various event types.

`c_trace` tracing with ‘`curses`’

The `--curses` option enables a more interactive tracing environment. Generated programs must be linked with the ‘`curses`’ (or compatible, such as ‘`ncurses`’) library. The resulting executable supports the following additional command-line options:

```
-D, --delay=number
    delay in milliseconds after each screen update ('0'=no delay)

-M, --monitor=number
    processor number to monitor

-B, --breakpoint=number
    source line at which to enter single-stepping mode ('-1'=none; '0'=first event)
```

When the program is run it brings up a screen like the following:

¹ A more precise correspondence is to a `<simple.stmt>` in the formal grammar presented in [Chapter 4 \[Grammar\], page 77](#).

```

1. # Determine computational "noise"
2.
3. Require language version "1.0".
4.
5. accesses is "Number of data accesses to perform" and comes from
6.   "--accesses" or "-a" with default 500000.
7.
8. trials is "Number of timings to take" and comes from "--timings" or
9.   "-t" with default 1000.
10.
11. For trials repetitions {
12.   all tasks reset their counters then
13.   all tasks touch a 1 word memory region accesses times with stride 0 w
14.   all tasks log a histogram of elapsed_usecs as "Actual time (usecs)"
15. }

Phys: 0  Virt: 0  Action: RESET    Event:    1/3001

```

The program displays its source code (truncated vertically if too tall and truncated horizontally if too wide) at the top of the screen and a status bar at the bottom of the screen. As the program executes, a cursor indicates the line of source code that is currently executing. Likewise, the status bar updates dynamically to indicate the processor's current task ID, action, and event number. In 'curses' mode, the program's standard output (see [Section 4.5.2 \[Writing to standard output\], page 104](#)) is suppressed so as not to disrupt the trace display.

Programs traced with `c_trace` and the `--curses` option are made interactive and support the following (case-insensitive) keyboard commands:

- 'S' Enable single-stepping mode. While single-stepping mode is enabled the traced processor will execute only one event per keystroke from the user.
- 'space' Disable single-stepping mode. The program executes without further user intervention.
- 'D' Delete the breakpoint.
- 'Q' Quit the program. The log file will indicate that the program did not run to completion.

All other keystrokes cause the program to advance to the next event immediately.

A single breakpoint can be set using the program's `-B` or `--breakpoint` command-line option. Whenever the monitored processor reaches the source-code line at which a breakpoint has been set, it enters single-stepping mode exactly as if `S` were pressed. Setting a breakpoint at line 0 tells the program to begin single-stepping as soon as the program begins. Note that only lines corresponding to `CONCEPTUAL` events can support breakpoints.

Offline tracing with ‘curses’

The `c_trace` backend can be told to trace by writing messages to the standard error device (see [Default `c_trace` tracing], page 24) or by employing an interactive display (see [c_trace tracing with curses], page 24). These two alternatives can be combined to support offline tracing of a `CONCEPTUAL` program. The idea is to compile the program without the `--curses` option. When running the program, the standard-error output should be redirected to a file. The `ncptl-replaytrace` utility, which comes with `CONCEPTUAL`, can then be used to play back the program’s execution by reading and displaying the file of redirected trace data.

‘`ncptl-replaytrace`’ accepts the following command-line options, which correspond closely to those accepted by a program compiled with the `--curses` option to `c_trace` (see [c_trace tracing with curses], page 24):

- `--trace=file`
Specify a file containing redirected trace data. *file* defaults to the standard input device.
- `--delay=number`
Specify the delay in milliseconds after each screen update (‘0’=no delay).
- `--monitor=processor`
Specify the processor number to monitor. *processor* defaults to ‘0’.
- `--breakpoint=line`
Specify a line of source code at which to enter single-stepping mode (‘-1’=none; ‘0’=first event).

In addition, ‘`ncptl-replaytrace`’ requires that the `CONCEPTUAL` source-code file be specified on the command line, as the source code is not included in the trace data.

The interactive display presented by the offline ‘`ncptl-replaytrace`’ tool is nearly identical to that presented by a program compiled with the `--curses` option to `c_trace`. See [c_trace tracing with curses], page 24, for a usage description.

3.3.5 The `c_profile` backend

Profiling communication events can help explain surprising performance measurements. A profiling tool which automatically instruments `CONCEPTUAL` programs is simpler than manually adding `LOGS` or `OUTPUTS` statements to an existing program.

Like the `c_trace` backend (see Section 3.3.4 [The `c_trace` backend], page 23), the `c_profile` backend adds code to that produced by other backends. The `c_profile` backend accepts a single—mandatory—command-line argument, `--profile=backend`, which designates a target backend to use. The restrictions on *backend* are that it must produce C code and must be derived from the `c_generic` backend (see Section 6.2 [Backend creation], page 136). Improper backends cause `c_profile` to abort abnormally.

If the profiled `CONCEPTUAL` program produces log files, the log files will include profiling information in the epilogue, one line for each event type which was used at least once by the corresponding process. (Section B.3 [Event types], page 190, lists and briefly describes the various event types.) Each line names an event and presents the total number of microseconds spent processing that event, a tally of the number of times that event type

was executed, and the quotient of those two numbers. In addition, the amount of memory used to store the event list is also logged.

For example, the following extract from a log file indicates that the process which wrote it spent a total of 6.7 s processing 22,000 ‘RECV’ events (which includes waiting time) for an average of 303.5 μ s per ‘RECV’ event:

```
# Profile of SEND (microseconds, count, average): 5267469 22001 239.4
# Profile of RECV (microseconds, count, average): 6676521 22000 303.5
# Profile of REPEAT (microseconds, count, average): 11985167 1 11985167.0
# Profile of NEWSTMT (microseconds, count, average): 43 1 43.0
# Profile of CODE (microseconds, count, average): 5516 1 5516.0
# Profile of event memory: 528 bytes (6 events * 88 bytes/event)
```

Each ‘REPEAT’ event includes the time for all of the events that it repeats.

Although the preceding log-file excerpt indicates that a total of $22001+22000+1+1+1 = 44004$ events were executed, the ‘event memory’ line clarifies that the event list contained only 6 unique events and therefore required only 528 bytes of memory.

Profiled programs that do not produce log files write profiling information to the standard error device. Because all processes may share a single standard error device, each line of output is preceded by a processor ID as in the following example:

```
1 SEND 5527125 22000 251.2
1 RECV 6322699 22001 287.4
1 REPEAT 11894523 1 11894523.0
1 event-memory 352 4 88
0 SEND 5267469 22001 239.4
0 RECV 6676521 22000 303.5
0 REPEAT 11985167 1 11985167.0
0 event-memory 352 4 88
```

The columns are *processor ID*, *event*, *total microseconds*, *tally*, and *average microseconds* except when *event* is ‘event-memory’ in which case the columns are *processor ID*, ‘event-memory’, *total bytes*, *number of events*, and *bytes per event*. The intention is for the output to be easily parseable using tools such as ‘awk’.

3.3.6 The interpret backend

Like the `c_udgram` backend (see [Section 3.3.3 \[The c_udgram backend\]](#), page 23), the `interpret` backend is designed to help programmers ensure the correctness of CONCEPTUAL code. The `interpret` backend does not output code. As its name implies, `interpret` is an *interpreter* of CONCEPTUAL programs rather than a compiler. `interpret` exhibits the following salient features:

1. Some programs run faster than with a compiler because the interpreter does not actually send messages. `interpret` merely simulates communication. It also skips over statements such as `COMPUTES/SLEEPS` (see [Section 4.6.3 \[Delaying execution\]](#), page 107) and `TOUCHES` (see [Section 4.6.4 \[Touching memory\]](#), page 108).

2. `interpret` can simulate massively parallel computer systems from a single process.
3. As `interpret` runs it checks for common communication errors such as deadlocks, asynchronous sends and receives that are never completed, and blocking operations left over at the end of the program (which would likely cause hung tasks under a real messaging layer).

The drawbacks are that `interpret` is slow when interpreting control-intensive programs and that timing measurements are not indicative of any real network. (The `interpret` backend utilizes logical time rather than physical time.) `interpret` is intended primarily as a development tool for helping ensure the correctness of `CONCEPTUAL` programs.

The `interpret` backend accepts all of the command-line options described in [Section 3.4 \[Running `coNCePTuaL` programs\]](#), [page 41](#), plus the following three options:

<code>-H, --hierarchy=<string></code>	Latency hierarchy as a comma-separated list of <code>task_factor:latency_delta</code> pairs [default: "tasks:1"]
<code>-M, --mcastsync=<number></code>	Perform an implicit synchronization after a multicast (0=no; 1=yes) [default: 0]
<code>-T, --tasks=<number></code>	Number of tasks to use [default: 1]

Normally, the `interpret` backend assigns unit latency to every communication operation. The `--hierarchy` option can make communication with distant tasks observe more latency than communication with nearby tasks. An explanation of the argument to `--hierarchy` is presented in [\[Task latency hierarchies\]](#), [page 30](#).

A multicast operation (see [Section 4.4.5 \[Multicasting\]](#), [page 101](#)) is normally treated as multiple point-to-point operations with the same send time. The `--mcastsync` option instructs the `interpret` backend to perform an implicit barrier synchronization at the end of the multicast.

The `--tasks` option specifies the number of tasks to simulate. Because this number can be quite large the `NCPTL_LOG_ONLY` environment variable (see [Appendix C \[Environment Variables\]](#), [page 192](#)) may be used to limit the set of processors that are allowed to create log files. That way, if task 0 is the only task out of thousands that logs any data, `NCPTL_LOG_ONLY` can specify that only one log file will be produced, not thousands. By default, all processors create a log file.

All other command-line arguments are passed to the program being interpreted.

The `--output` option described in [Section 3.2 \[Compiling `coNCePTuaL` programs\]](#), [page 18](#), has special meaning to the `interpret` backend. When `--output` is used, `interpret` dumps a list of events to the specified file after a successful run. For example, the `CONCEPTUAL` program 'ALL TASKS t ASYNCHRONOUSLY SEND A 384 BYTE MESSAGE TO TASK t XOR 2 THEN ALL TASKS AWAIT COMPLETION' results in the following event dump:

```
Task 0 posted a NEWSTMT at time 0 and completed it at time 0
Task 0 posted a RECEIVE at time 0 and completed it at time 0
Task 0 posted a SEND at time 1 and completed it at time 1
Task 0 posted a WAIT_ALL at time 2 and completed it at time 2
Task 1 posted a NEWSTMT at time 0 and completed it at time 0
Task 1 posted a RECEIVE at time 0 and completed it at time 0
```

```

Task 1 posted a SEND at time 1 and completed it at time 1
Task 1 posted a WAIT_ALL at time 2 and completed it at time 2
Task 2 posted a NEWSTMT at time 0 and completed it at time 0
Task 2 posted a RECEIVE at time 0 and completed it at time 0
Task 2 posted a SEND at time 1 and completed it at time 1
Task 2 posted a WAIT_ALL at time 2 and completed it at time 2
Task 3 posted a NEWSTMT at time 0 and completed it at time 0
Task 3 posted a RECEIVE at time 0 and completed it at time 0
Task 3 posted a SEND at time 1 and completed it at time 1
Task 3 posted a WAIT_ALL at time 2 and completed it at time 2

```

As an example of the `interpret` backend’s usage, here’s how to simulate 100,000 processors communicating in a simple ring pattern:

```

% ncptl --backend=interpret --lenient --program='All tasks t send
  nummsgs 1024 gigabyte messages to task t+1 then task num_tasks-1
  sends nummsgs 1024 gigabyte messages to task 0.' --tasks=100000
  --nummsgs=5

```

The preceding command ran to completion in under 5 minutes on a 1.5 GHz Xeon uniprocessor workstation—not too bad considering that 488 petabytes of data are transmitted on the program’s critical path.

The `interpret` backend is especially useful for finding communication-related program errors:

```

% ncptl --backend=interpret --quiet --program='All tasks t send
  a 10 doubleword message to task (t+1) mod num_tasks.' --tasks=3
<command line>: The following tasks have deadlocked: 0 --> 2 --> 1
  --> 0

```

Deadlocked tasks are shown with ‘-->’ signifying “is blocked waiting for”. In the preceding example, all receives are posted before all sends. Hence, task 0 is blocked waiting for task 2 to send it a message. Task 2, in turn, is blocked waiting for task 1 to send it a message. Finally, task 1 is blocked waiting for task 0 to send it a message, which creates a cycle of dependencies.

The `interpret` backend can find other errors, as well:

```

% ncptl --backend=interpret --quiet --program='All tasks t
  asynchronously send a 10 doubleword message to task (t+1) mod
  num_tasks.' --tasks=4
<command line>: The program ended with the following leftover-event
  errors:
* Task 0 posted an asynchronous RECEIVE which was never waited for
* Task 0 posted an asynchronous SEND which was never waited for
* Task 0 sent a message to task 1 which was never received
* Task 1 posted an asynchronous RECEIVE which was never waited for
* Task 1 posted an asynchronous SEND which was never waited for
* Task 1 sent a message to task 2 which was never received
* Task 2 posted an asynchronous RECEIVE which was never waited for
* Task 2 posted an asynchronous SEND which was never waited for
* Task 2 sent a message to task 3 which was never received

```

- * Task 3 posted an asynchronous RECEIVE which was never waited for
- * Task 3 posted an asynchronous SEND which was never waited for
- * Task 3 sent a message to task 0 which was never received

(A message received ASYNCHRONOUSLY is not considered received until after the corresponding AWAITs COMPLETION; hence, all of the ‘was never received’ messages listed above.)

```
% ncptl --backend=interpret --quiet --program='Task 0 sends a 40
kilobyte message to unsuspecting task 1 then task 0 receives a 40
kilobyte message from task 1.' --tasks=2
```

<command line>: The program ended with the following leftover-event errors:

- * Task 0 sent a message to task 1 which was never received
- * Task 1 terminated before satisfying task 0’s RECEIVE operation

In short, it is well worth testing the correctness of new CONCEPTUAL programs with `interpret` before performing timing runs with one of the message-passing backends.

Task latency hierarchies

The `interpret` backend normally simulates a flat network in which communication between any two tasks takes unit latency. However, the `--hierarchy` option lets one specify a hierarchy of latencies: latency l_1 within a set of t_1 tasks, latency $l_1 + l_2$ within a set of $t_1 \cdot t_2$ tasks, latency $l_1 + l_2 + l_3$ within a set of $t_1 \cdot t_2 \cdot t_3$ tasks, and so forth.

The argument to `--hierarchy` is a comma-separated list of $\langle \text{task factor} : \text{latency delta} \rangle$ pairs. The *latency delta* component is optional and defaults to ‘1’. If the list ends with ‘...’, the final $\langle \text{task factor} : \text{latency delta} \rangle$ pair is repeatedly indefinitely. All *task factor* values must be positive integers, and all *latency delta* values must be nonnegative integers.

As an example, `--hierarchy=4` (or `--hierarchy=4:1`) partitions the program’s tasks into sets of four with one unit of additional latency to communicate with another set. Hence, tasks 0, 1, 2, and 3 can communicate with each other in unit time, as can tasks 4, 5, 6, and 7. However, communication between any task in the first set and any task in the second set (e.g., between tasks 3 and 4) takes an additional unit of time for a total latency of two units.

As a more complex example, consider a cluster of symmetric multiprocessors (SMPs) in which each SMP comprises two processor sockets with each socket containing a quad-core processor (four CPUs). Further assume that the SMPs are networked together via a fat-tree network consisting of 12-port switches. In this example, let’s say that it takes unit latency to communicate within a socket, two units of latency to communicate with another socket in the same SMP, and an additional three units of latency to traverse each switch. This configuration can be specified to the `interpret` backend as `--hierarchy="4:1, 2:1, 12:3, 12:6, ..."` (or simply `--hierarchy=4,2,12:3,12:6, ...`). With this setting, task 0, for instance, communicates with tasks 1–3 in 1 time unit, with tasks 4–7 in 2 time units, with tasks 8–95 in 5 time units (one switch crossing), with tasks 96–1151 in 11 time units (three switch crossings—a level 0 switch, a level 1 switch, and another level 0 switch), with tasks 1152–13,823 in 17 time units (five switch crossings—of switch levels 0, 1, 2, 1, and 0), and so forth.

3.3.7 The stats backend

The `stats` backend outputs various statistics about a `CONCEPTUAL` program. It can be used to help verify a program's correctness or merely to search for interesting patterns within a complex communication pattern. The following is some sample output from a program compiled with `stats`:

```

Execution parameters
-----
Number of processors:      16
Random-number seed:      -1727114895
Command line:             ncptl --backend=stats hycom.ncptl --xdim=4 --ydim=4 --tasks=16 --iter=10
Timestamp:               Wed Jan  4 17:32:04 2006

Message traffic
-----
Total messages sent:      710
Total bytes sent:        1454080
Unique message sizes sent: 2048

Per-processor message traffic
-----
Processors sending a total of 61440 bytes:  0
Processors sending a total of 81920 bytes:  1-3, 12-15
Processors sending a total of 102400 bytes: 4-11
Processors receiving a total of 61440 bytes: 1-3, 13-15
Processors receiving a total of 81920 bytes: 5-7, 9-11
Processors receiving a total of 122880 bytes: 12
Processors receiving a total of 143360 bytes: 4, 8
Processors receiving a total of 184320 bytes: 0
Processors sending a total of 30 messages:   0
Processors sending a total of 40 messages:  1-3, 12-15
Processors sending a total of 50 messages:  4-11
Processors receiving a total of 30 messages: 1-3, 13-15
Processors receiving a total of 40 messages: 5-7, 9-11
Processors receiving a total of 60 messages: 12
Processors receiving a total of 70 messages: 4, 8
Processors receiving a total of 90 messages: 0

Processor SEND-event peers
-----
Processors posting SEND events to offsets {-12, -4, +1, +3}: 12
Processors posting SEND events to offsets {-8, -4, +1, +3, +4}: 8
Processors posting SEND events to offsets {-4, -3, -1}: 15
Processors posting SEND events to offsets {-4, -3, -1, +4}: 7, 11
Processors posting SEND events to offsets {-4, -2, -1, +1}: 14
Processors posting SEND events to offsets {-4, -2, -1, +1, +4}: 6, 10
Processors posting SEND events to offsets {-4, -1, +1}: 13
Processors posting SEND events to offsets {-4, -1, +1, +4}: 5, 9
Processors posting SEND events to offsets {-4, +1, +3, +4}: 4
Processors posting SEND events to offsets {-3, -1, +4}: 3
Processors posting SEND events to offsets {-2, -1, +1, +4}: 2
Processors posting SEND events to offsets {-1, +1, +4}: 1
Processors posting SEND events to offsets {+1, +3, +4}: 0

Network bisection crossings
-----
Bisection messages:      100
Bisection bytes:        204800

```

```

Event tallies
-----
Total number of LOG events:      10
Total number of NEWSTMT events:  48
Total number of RECEIVE events: 710
Total number of RESET events:    10
Total number of SEND events:     710
Total number of WAIT_ALL events: 480

Per-processor event sets
-----
Processors executing only {LOG, NEWSTMT, RECEIVE, RESET, SEND, WAIT_ALL}: 0
Processors executing only {NEWSTMT, RECEIVE, SEND, WAIT_ALL}:             1-15

Per-processor event tallies
-----
Processors executing 10 LOG events:      0
Processors executing 3 NEWSTMT events:  0-15
Processors executing 30 RECEIVE events:  1-3, 13-15
Processors executing 40 RECEIVE events:  5-7, 9-11
Processors executing 60 RECEIVE events:  12
Processors executing 70 RECEIVE events:  4, 8
Processors executing 90 RECEIVE events:  0
Processors executing 10 RESET events:    0
Processors executing 30 SEND events:     0
Processors executing 40 SEND events:     1-3, 12-15
Processors executing 50 SEND events:     4-11
Processors executing 30 WAIT_ALL events: 0-15

```

`stats` is derived from `interpret` (see [Section 3.3.6 \[The `interpret` backend\]](#), page 27) and therefore supports the `interpret` backend's `--tasks` and `--mcastsync` options as well as the standard options described in [Section 3.4 \[Running `coNCePTuaL` programs\]](#), page 41. However, because `stats` does not produce log files, the `--logfile` option is absent. `stats` additionally supports the following three command-line options:

<code>-E, --expand-lists=<number></code>	0=collapse lists of numbers into ranges; 1=show all numbers [default: 0]
<code>-F, --format=<string></code>	Output format, either "text", "excelcsv", or "sep:<string>" [default: "text"]
<code>-X, --exclude=<string></code>	Name of a category or individual field to exclude from output [default: ""]

The `--expand-lists` option tells the `stats` backend whether it should output sets of numbers as comma-separated ranges (e.g., '1-3, 12-15') or as individual numbers (e.g., '1, 2, 3, 12, 13, 14, 15'). The `--format` option specifies the output format. By default, it outputs human-readable text, as shown in the above example. However, `--format=excelcsv` outputs the data in comma-separated value format suitable for loading directly into Microsoft Excel. A more general form of computer-parseable output is specified with the 'sep' sub-option, which takes a separator string as a sub-option. For example, `--format=sep:="#"` produces output like the following:

```

"Event tallies"#"Total number of LOG events"#10
#"Total number of NEWSTMT events"#48

```

```

#"Total number of RECEIVE events"#710
#"Total number of RESET events"#10
#"Total number of SEND events"#710
#"Total number of WAIT_ALL events"#480

```

The difference between `--format=excelcsv` and `--format=sep:","` is that the former employs a number of bizarre special cases when quoting strings so as to coerce Excel into properly processing the file. The latter simply inserts a comma between columns with no special string processing other than preceding the characters `'` and `\` with `\`.

The `stats` backend issues a `Column separator sep appears within a field` warning if the separator string appears in any of the output fields. (Note that it will always appear in the `Command line` line because the `--format` line is specified on the command line.) The idea is to warn the user that automatic parsing of the output (e.g., using `awk`) may need to be careful when processing lines containing such fields.

As can be seen from the sample output, `stats` outputs a lot of information. The `--exclude` option—which can be specified repeatedly on the command line—provides the user with fine-grained control over which output to suppress. For example, `--exclude="Total number of NEWSTMT events"` tells `stats` not to output the line with that key. Similarly, `--exclude="Processor SEND-event peers"` eliminates an entire category of information.

3.3.8 The `picl` backend

PICL, the Portable Instrumented Communication Library, defines a trace file format that records an execution trace of a message-passing application. In particular, MPICL is an MPI-specific implementation of PICL that facilitates instrumenting MPI programs. Normally, one passes a PICL file to a performance-visualization tool such as ParaGraph, which renders the trace data using any of a number of graphical views.

Although MPICL is compatible with the `c_mpi` backend (see [Section 3.3.2 \[The `c_mpi` backend\]](#), page 22), `CONCEPTUAL` provides a `picl` backend which produces PICL files directly. The intention is to use the PICL format to represent an idealized view of a communication pattern rather than to store a time-sensitive execution trace. For example, timing events recorded by `picl` occur at logical times instead of physical times and statements such as `COMPUTES/SLEEPS` (see [Section 4.6.3 \[Delaying execution\]](#), page 107) and `TOUCHES` (see [Section 4.6.4 \[Touching memory\]](#), page 108) take either zero time or unit time, depending upon a command-line option. By providing an idealized view of a communication pattern, `picl` abstracts away timing artifacts so events that one would expect to occur simultaneously are written to the trace file as being exactly simultaneous.

As an example, the following is the PICL output produced by passing `picl` the `CONCEPTUAL` program `'TASK 0 SENDS A 32 KILOBYTE MESSAGE TO TASK 1 THEN TASK 1 SENDS A 256 BYTE MESSAGE TO TASK 0.'`:

```

-3 -901 0.00000 0 0 0
-3 -901 0.00000 1 1 0
-3 -21 0.00001 0 0 4 2 32768 1 1 1
-4 -21 0.00001 0 0 0
-3 -52 0.00001 1 1 3 2 1 0 0
-3 -52 0.00002 0 0 3 2 1 1 1

```

```

-4 -52 0.00002 1 1 4 2 32768 1 0 0
-3 -21 0.00002 1 1 4 2 256 1 0 0
-4 -21 0.00002 1 1 0
-4 -52 0.00003 0 0 4 2 256 1 1 1
-4 -901 0.00003 1 1 0
-4 -901 0.00004 0 0 0

```

See the PICL manual, *A new PICL trace file format* (ORNL/TM-12125), for a detailed description of the various fields used in the preceding trace file.

ParaGraph is a visualization tool that reads PICL trace files and graphically displays/animates the corresponding execution in a variety of formats. Running ParaGraph on output from `picl` makes it easy for a CONCEPTUAL programmer to explain complex communication patterns to other people.

`picl` is derived from `interpret` (see Section 3.3.6 [The `interpret` backend], page 27) and therefore supports the `interpret` backend’s `--tasks` and `--mcastsync` options as well as the standard options described in Section 3.4 [Running `coNcEPTuaL` programs], page 41. However, because `picl` does not produce log files, the `--logfile` option is absent. `picl` additionally supports the following two command-line options:

<code>-A, --all-events=<number></code>	0=include only communication events; 1=include all events [default: 0]
<code>-F, --frequency=<number></code>	PICL event frequency (Hz) [default: 100000]

By default, only communication events are written to the trace file. If `--all-events` is set to ‘1’ then all events are written to the trace file. Events not related to communication are defined to take unit time. Regardless of the setting of `--all-events`, the `OUTPUTS` statement (see Section 4.5.2 [Writing to standard output], page 104) writes a PICL ‘`tracemsg`’ event to the trace file. A user can tell ParaGraph to pause visualization at ‘`tracemsg`’ events, making it possible to isolate key components of a CONCEPTUAL program’s execution.

Because PICL events are marked with physical time (a floating-point number of seconds) but `picl` uses exclusively logical time, `picl` needs to associate a (fabricated) physical time with each logical time. The `--frequency` option specifies that mapping. By default, `picl` pretends that each unit of logical time corresponds to 1/100000 of a second (i.e, 10 microseconds) of physical time. As an example, the default time step in ParaGraph is 1 microsecond; this can be specified to `picl` with `--frequency=1E6`.

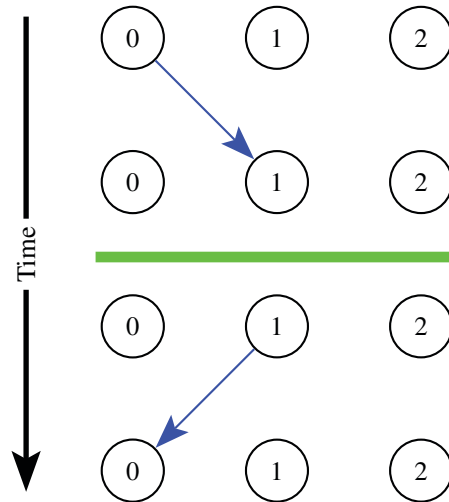
One of the goals of CONCEPTUAL is to facilitate the explanation of network performance tests. The `picl` backend aids in the explanation by making it easy to show graphically how tasks communicate with each other in an arbitrary CONCEPTUAL program.

3.3.9 The `latex_vis` backend

The `latex_vis` backend visualizes a program’s communication pattern as an Encapsulated PostScript time-space diagram. The backend thereby provides a static counterpart to the animated visualizations made possible by the `picl` backend and the ParaGraph tool (see Section 3.3.8 [The `picl` backend], page 33). `latex_vis` outputs `LATEX` code with calls to the `PSTricks` package (available from <http://www.tug.org/applications/PSTricks/> but shipped with most `LATEX` distributions) to draw the figure, then runs `latex` (or whatever

the *LATEX* environment variable is set to) to convert the `.tex` file to DVI format and *dvips* (or whatever the *DVIPS* environment variable is set to) to convert the `.dvi` file into EPS format. Finally, `latex_vis` runs the graphic through Ghostscript (*gs* or whatever the *GS* environment variable is set to) to tighten the graphic's bounding box.² `latex_vis` requires that *L^AT_EX* and *PSTricks* be installed in order to produce its visualizations. However, it can still run with `--no-compile` (see Section 3.2 [Compiling coNcEPTuaL programs], page 18) to produce a `.tex` file which can later be run manually through *latex*.

As an example, the following is the default `latex_vis` output from the `coNcEPTuaL` program `'TASK 0 SENDS A 1 MEGABYTE MESSAGE TO TASK 1 THEN ALL TASKS SYNCHRONIZE THEN TASK 1 SENDS A 3 KILOBYTE MESSAGE TO TASK 0'` when run with three tasks:



Currently, the output diagram does not indicate message size but this may change in a future release of `latex_vis`.

The `latex_vis` backend has a number of uses. First, it can be used to illustrate a nontraditional communication pattern for a presentation, research paper, or technical report. Second, it can be used as a teaching aid to demonstrate common communication patterns (e.g., a butterfly pattern) to students. Third, it can be used as a debugging aid to ensure that a `coNcEPTuaL` program is, in fact, performing the expected communication operations.

`latex_vis` is derived from `interpret` (see Section 3.3.6 [The `interpret` backend], page 27) and therefore supports the `interpret` backend's `--tasks` and `--mcastsync` options as well as the standard options described in Section 3.4 [Running coNcEPTuaL programs], page 41. However, because `latex_vis` does not produce log files, the `--logfile` option is absent. `latex_vis` additionally supports the following five command-line options:

² If Ghostscript is not installed or fails to run, `latex_vis` issues a warning message, not an error message. The figure is still usable without a tight bounding box and a loose bounding box can be corrected manually by editing the `'%BoundingBox:'` line in the generated EPS file.

-A, --annotate=<string>	Annotation level (0=no annotations; 1=annotate communication events; 2=annotate all events; "<event>..."=annotate only the specified events) [default: "0"]
-B, --binary-tasks=<number>	Display task numbers in binary rather than decimal (0=decimal; 1=binary) [default: 0]
-E, --every-event=<number>	Events requiring nonzero time to complete (0=only communication events; 1=every event) [default: 0]
-G, --stagger=<number>	Number of points by which to stagger overlapping arrows [default: 2]
-L, --source-lines=<number>	Associate source-code line numbers with each event annotation (0=no; 1=yes) [default: 0]
-R, --arrow-width=<string>	Python expression to map m, representing a message size in bytes, to an arrow width in points [default: "1"]
-Z, --zero-latency=<number>	Depict communication as having zero latency (0=unit latency; 1=zero latency) [default: 0]

In the preceding list a “point” refers to a PostScript point. \TeX calls these “big points” and defines $1\text{ bp} \equiv 1/72''$)

The `--annotate` option places adjacent to appropriate nodes in the output diagram a list of textual annotations which indicate the communication operations that were posted or completed and the non-communication operations that were executed by the corresponding task at the corresponding time. Some sample annotations are ‘Post SEND’, ‘Complete SEND’, and ‘Execute OUTPUT’. With `--annotate=1`, only communication events are annotated; with `--annotate=2`, all events are annotated;³ otherwise, a list of specific events can be annotated. For example, `--annotate=SYNC,MCAST` causes only the ‘SYNC’ and ‘MCAST’ events to be annotated. Compiling a `CONCEPTUAL` program with the `latex_vis` backend and the `--keep-ints` options (see [Section 3.2 \[Compiling coNCePTuaL programs\]](#), page 18) produces a ‘.tex’ file which contains in the prologue comments a list of all events used by a program that were and were not annotated.

Task numbers are normally shown in base ten. The `--binary-tasks` option causes task numbers to be output in base two and using the same number of bits for each task. For example, a 3-task visualization with `--binary-tasks` would number the tasks ‘00’, ‘01’, and ‘11’.

Because `latex_vis` is intended primarily for visualizing communication patterns, by default only communication operations take nonzero time to complete. The `--every-event` option indicates that all events—including local events such as ‘SLEEP’, ‘COMPUTE’, ‘LOG’, and ‘OUTPUT’—should be deemed to complete in unit time.

³ To avoid the confusion of annotating what is essentially a “do nothing” event, `latex_vis` does not annotate the ‘NEWSTMT’ event, which is injected at the beginning of each top-level statement.

`latex_vis` output distinguishes coincident arrows (consider the phrase ‘TASK 0 ASYNCHRONOUSLY SENDS 5 MESSAGES TO TASK 1’) by staggering them slightly. The `--stagger` option specifies the number of PostScript points by which to stagger each overlapping arrow with a default of 2 points. Large stagger values are more visually distinctive while small stagger values allow arrows to drift less from their associated nodes.

The `--source-lines` option, when used with `--annotate`, augments each event annotation with the corresponding line(s) of source code which produced that event. This feature improves the utility of `latex_vis` as a debugging aid for CONCEPTUAL programs.

By default, all arrows which indicate message transmissions are drawn with equally thick line widths. The `--arrow-width` option lets you specify a Python function to map a message size in bytes, m , to a line width in PostScript points, with the default being 1 point. As a simple example, `--arrow-width=2` doubles the arrow width for all message (i.e., it represents the constant function $f : m \mapsto 2$). A more typical example would be `--arrow-width="log10(m)"` (representing $f : m \mapsto \log_{10}(m)$), which causes a tenfold increase in message size to yield a unit increase in line width. The argument to `--arrow-width` can be any Python expression using either built-in functions or functions from the Python `math` module. See the [Python library reference](#) for details.

Normally, messages are considered to require one time unit to travel from source to destination. The `--zero-latency` option shows messages as being received in the same timestep as they were sent. Some communication patterns are more aesthetically pleasing when drawn this way.

Further customizations

In addition to the options described above, the front end’s `--filter` option (see [Section 3.2 \[Compiling coNCePTuaL programs\]](#), page 18) is a useful mechanism for customizing the formatting of the communication diagram. For example, specifying `--filter="s/rowsep=30bp/rowsep=1.5in/g"` increases the separation between rows from 30 bp (where 1 bp \equiv 1/72") to 1.5". See the [PSTricks documentation](#) for more information about the PSTricks commands used in the generated ‘.tex’ files.

To facilitate the use of `--filter`, the `latex_vis` backend uses a helper macro (`\viscolor`) to define colors. `\viscolor` takes an argument of the form ‘`name=color`’ and defines a macro `\namecolor` which expands to `color`. To further facilitate the use of `--filter`, the L^AT_EX code generated by the `latex_vis` backend contains a number of strategically placed placeholder comments of the form ‘% PLACEHOLDER: `tag`’. A `--filter` command can thereby insert code into the document by replacing an appropriate ‘PLACEHOLDER’ line. In alphabetical order, the currently defined placeholder tags are ANNOTATIONS, COLORS, COMMUNICATION, DEADLOCK, DOCUMENT, END, NODESHAPE, PACKAGES, PSMATRIX, TEXTOEPS, and TIMELINE. Look through any `latex_vis`-generated ‘.tex’ file to see where these are situated. As an example, the following command-line options define a “chartreuse” color then change the color used to indicate point-to-point messages from blue to chartreuse:

```
--filter="s/% PLACEHOLDER: COLORS/\newrgbcolor{chartreuse}{0.5 1 0}/"
--filter="s/sendrecv=blue/sendrecv=chartreuse/"
```

(According to the [PSTricks documentation](#), the predefined colors are red, green, blue, cyan, magenta, and yellow, and the predefined grayscales are black, darkgray, gray, lightgray, and white.)

Tasks which are active at a given time are drawn using the `\task` macro, which takes the task number as an argument. Tasks which are idle at a given time are drawn using the `\idle` macro, which is initially defined to be the same as `\task`. The following “cookbook” examples showcase some of the ways that the power of \LaTeX and PSTricks can be exploited to display idle tasks in a variety of different styles (best used with `--annotate=2`):

omitting idle tasks

```
--filter="s/\let\idle=\task/\newcommand*{\idle}[1]{[mnode=R]}/"
```

showing each idle task as a gray dot

```
--filter="s/\let\idle=\task/\newcommand*{\idle}[1]{[mnode=dot,
linecolor=gray]}/"
```

drawing idle tasks with a dotted circle instead of a solid circle

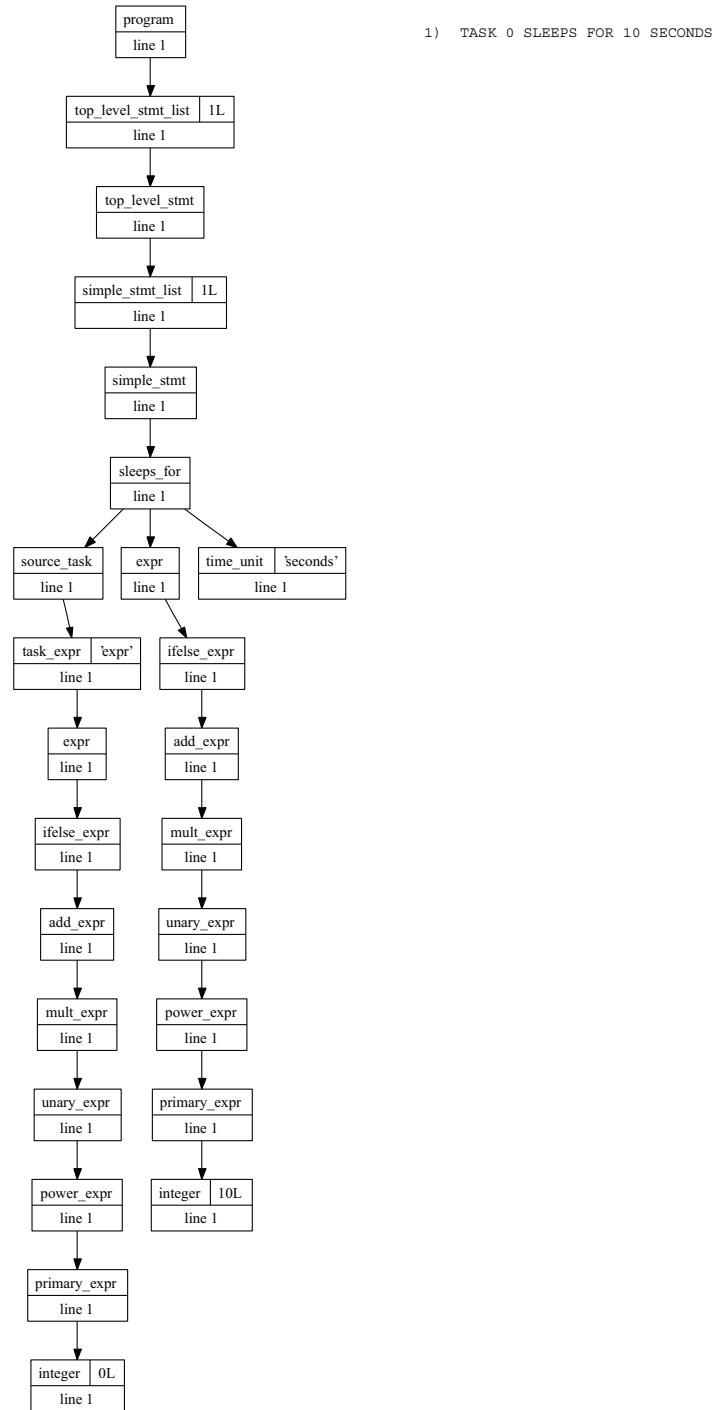
```
--filter="s/\let\idle=\task/\newcommand*{\idle}[1]
{[linestyle=dotted]\task{#1}}/"
```

As an alternative to replacing the `\let\idle=\task` binding, the preceding substitutions can be expressed as the insertion of a \LaTeX `\renewcommand`. That is, idle tasks can also be omitted by specifying `--filter="s/% PLACEHOLDER: NODESHAPE/\renewcommand*{\idle}[1]{[mnode=R]}/"`.

In short, the `latex_vis` backend produces highly customizable illustrations of communication patterns. Because `latex_vis` produces commented \LaTeX code, any customization not provided through the use of `--filter` or one of the backend-specific command-line options is easily performed by compiling with `--keep-ints` (see [Section 3.2 \[Compiling coNCePTuaL programs\]](#), page 18) and editing the generated \LaTeX code.

3.3.10 The dot_ast backend

`dot` is a format for describing graphs in terms of their edges and vertices. The tools in the Graphviz suite typeset `dot` files in a variety of output formats and using a variety of graph-layout algorithms. `coNCePTuaL`'s `dot_ast` backend outputs in `dot` format the abstract-syntax tree corresponding to a given `coNCePTuaL` program. As an example, `dot_ast` renders the one-line `coNCePTuaL` program `'TASK 0 SLEEPS FOR 10 SECONDS.'` as follows:



`dot_ast` is expected to be of particular use to backend developers, who can use it to help prioritize the methods that need to be implemented (i.e., implementing first the AST node types needed by in a trivial program, then those needed by successively more complex programs).

The `dot_ast` backend accepts the following options from the `'ncpt1'` command line:

`--format=dot_format`

The programs in the Graphviz suite can output graphs in a variety of formats such as PostScript, SVG, and PNG. By default, the `dot_ast` backend outputs PostScript. The `--format` option specifies an alternate format to use. At the time of this writing, the Graphviz programs support the following formats: `canon`, `cmap`, `dot`, `fig`, `gd`, `gd2`, `gif`, `hpgl`, `imap`, `ismap`, `jpeg`, `jpg`, `mif`, `mp`, `pcl`, `pic`, `plain`, `plain-ext`, `png`, `ps`, `ps2`, `svg`, `svgz`, `vrml`, `vtx`, `wbmp`, and `xdot`. See the Graphviz documentation for more information about these formats.

`--node-code=characters`

To facilitate associating nodes in the AST with fragments of the `CONCEPTUAL` program being graphed, the `dot_ast` backend provides a `--node-code` option that labels each node with the fragment of code to which it corresponds. The argument to `--node-code` is a number of characters at which to truncate the code fragment or `-1` to inhibit truncation. (The purpose of truncation is to prevent excessively large nodes from disturbing the graph layout. The `'program'` node, for example, includes the complete program source code if not truncated.)

`--extra-dot=dot_code`

The `dot_ast` backend's `--extra-dot` option enables the user to inject arbitrary dot code into the generated file. For example, specifying `--extra-dot="node [shape=Mrecord]"4` tells `dot` to use draw nodes as rounded rectangles and specifying `--extra-dot='edge [color="green"]'` colors all edges green. `--extra-dot` can be specified repeatedly on the command line; `dot_ast` concatenates all of the extra dot code with intervening semicolons.

`--no-lines`

By default, each AST node indicates the lines in the program's source code to which it corresponds. The `--no-lines` option suppresses the outputting of source-code line numbers.

`--no-attrs`

Every node in the AST has a type. Some nodes additionally have an attribute. `dot_ast` normally outputs attributes but `--no-attrs` prevents `dot_ast` from doing so.

`--no-source`

The complete source code corresponding to the AST is included in the generated dot graph unless `--no-source` is specified on the command line.

The `DOT` environment variable names the Graphviz program that `dot_ast` should run on the generated code. If `DOT` is not set, `dot_ast` uses whatever value was specified/discovered at configuration time (see [Section 2.1 \[configure\], page 5](#)), with the default being `'dot'`. By default, `dot_ast` produces dot code and runs this through the designated Graphviz program to produce a PostScript file (or whatever format is named by the `--format` option). If `'ncpt1'` is run with either the `--no-link` or `--no-compile` options, it produces a dot file that should be run manually through `'dot'` or another Graphviz tool.

⁴ `dot_ast` automatically places a semicolon after the extra dot code.

3.4 Running coNcEPTuaL programs

CONCEPTUAL programs can be run like any other program built with the same compiler and communication library. For example if a program ‘myprog’ was built with CONCEPTUAL’s C+MPI backend, the program might be run with a command like `mpirun -np nodes myprog` or `prun -Nnodes myprog` or `pdsh -w node_list myprog`. The important point is that job launching is external to CONCEPTUAL. A CONCEPTUAL program is oblivious to whether it is being run with a single thread on each multiprocessor node or with one thread on each CPU, for example. However, CONCEPTUAL log files do include the host name in the prologue comments (see [Section 3.5.1 \[Log-file format\], page 44](#)) so job-launching parameters can potentially be inferred from those.

CONCEPTUAL programs automatically support a “help” option. This is usually specified as `--help` or `-?`, depending on which option-parsing library ‘configure’ configured in. (See [Section 2.1 \[configure\], page 5](#).) The output of running `myprog --help` most likely looks something like this:

```
Usage: myprog [OPTION...]
  -C, --comment=<string>      Additional commentary to write to the log
                              file, @FILE to import commentary from FILE,
                              or !COMMAND to import commentary from COMMAND
                              (may be specified repeatedly)
  -L, --logfile=<string>      Log-file template [default: "a.out-%p.log"]
  -N, --no-trap=<string>      List of signals that should not be trapped
                              [default: ""]
  -S, --seed=<number>         Seed for the random-number generator
                              [default: 0]

Help options:
  -?, --help                  Show this help message
  --usage                     Display brief usage message
```

Although a CONCEPTUAL program can specify its own command-line options (see [Section 4.8.2 \[Command-line arguments\], page 120](#)), a few are provided by default. In addition to `--help` these include `--comment`, `--logfile`, `--no-trap`, and `--seed`.

`--comment`

`--comment` makes it possible to add arbitrary commentary to a log file. This is useful for incorporating information that CONCEPTUAL would be unable to (or simply does not currently) determine on its own, for example, `--comment="Last experiment before upgrading the network device driver"`. Two special cases are supported:

1. If the comment string begins with ‘@’ then the remainder of the string is treated as a filename. Each line of the corresponding file is treated as a separate comment string. Hence, if the file ‘sysdesc.txt’ contains the lines ‘Using FooBarNet’ and ‘Quux is enabled’, then specifying `--comment=sysdesc.txt` is similar to specifying both `--comment="Using FooBarNet"` and `--comment="Quux is enabled"`.

2. If the comment string begins with ‘!’ then the remainder of the string is treated as a shell command. The command is executed and each line of its output is treated as a separate comment string. For example, `--comment='!lspci | grep -i net'` executes ‘lspci’, extracts only those lines containing the string ‘net’, and makes log-file comments out of the result. Note that `--comment='!command'` differs from `--comment="command"` in that the former causes *command* to be executed individually by each process in the program while the latter executes *command* only once and only before launching the program. Also note that ‘!’ must be escaped in ‘csh’ and derivative shells (i.e., `--comment='\!command'`).

As an example, the command line arguments `--comment="This is a simple comment."` `--comment=!lspci` `--comment=@/proc/version` `--comment="This is another simple comment."` produce log-file lines like the following:

```
# User comment 1: This is a simple comment.
# Output of 'lspci', line 1: 00:00.0 Host bridge: Intel Corp. 82860 860 (Wombat) Chipset Host
# Output of 'lspci', line 2: 00:01.0 PCI bridge: Intel Corp. 82850 850 (Tehama) Chipset AGP B
# Output of 'lspci', line 3: 00:02.0 PCI bridge: Intel Corp. 82860 860 (Wombat) Chipset AGP B
# Output of 'lspci', line 4: 00:1e.0 PCI bridge: Intel Corp. 82801BA/CA/DB PCI Bridge (rev 04)
# Output of 'lspci', line 5: 00:1f.0 ISA bridge: Intel Corp. 82801BA ISA Bridge (LPC) (rev 04)
# Output of 'lspci', line 6: 00:1f.1 IDE interface: Intel Corp. 82801BA IDE U100 (rev 04)
# Output of 'lspci', line 7: 00:1f.2 USB Controller: Intel Corp. 82801BA/BAM USB (Hub (rev 04)
# Output of 'lspci', line 8: 00:1f.3 SMBus: Intel Corp. 82801BA/BAM SMBus (rev 04)
# Output of 'lspci', line 9: 00:1f.4 USB Controller: Intel Corp. 82801BA/BAM USB (Hub (rev 04)
# Output of 'lspci', line 10: 00:1f.5 Multimedia audio controller: Intel Corp. 82801BA/BAM AC
# Output of 'lspci', line 11: 01:00.0 VGA compatible controller: nVidia Corporation NV11 [GeF
# Output of 'lspci', line 12: 02:1f.0 PCI bridge: Intel Corp. 82806AA PCI64 Hub PCI Bridge (r
# Output of 'lspci', line 13: 03:00.0 PIC: Intel Corp. 82806AA PCI64 Hub Advanced Programmabl
# Output of 'lspci', line 14: 04:0b.0 Ethernet controller: 3Com Corporation 3c905C-TX/TX-M [T
# Output of 'lspci', line 15: 04:0d.0 Multimedia audio controller: Creative Labs SB Live! EMU
# Output of 'lspci', line 16: 04:0d.1 Input device controller: Creative Labs SB Live! MIDI/Ga
# Contents of /proc/version, line 1: Linux version 2.4.20-28.7 (bhcompile@porky.devel.redhat.
# User comment 2: This is another simple comment.
```

To facilitate log-file parsing, all colons in the name of an ‘@’-file or ‘!’-command are written as periods. This affects only the display, not the file to read or command to execute.

Be careful when using shell backquotes with `--comment` (e.g., as in `--comment="who"`). Different shells have different ways of handling newlines output by a backquoted command. In some shells (e.g., ‘bash’), the CONCEPTUAL program sees the entire output with embedded newlines as a single argument; in others (e.g., ‘tcsh’), each line of output constitutes a separate command-line argument. Because CONCEPTUAL programs currently ignore arguments which do not begin with ‘--’, the comment written to the log file in the latter case terminates at the first newline. In virtually all shells, the double quotes around the backquoted command are needed to prevent the shell from splitting arguments at word boundaries. As a consequence, `--comment='who'` logs only the first word output by the *who* command.

--logfile

--logfile specifies a template for naming log files. Each task maintains a log file based on the template name but with ‘%p’ replaced with the processor number, ‘%r’ replaced with the run number (the smallest nonnegative integer that produces a filename which does not already exist), and ‘%%’ replaced with a literal “%” character. The program outputs an error message and aborts if the log-file template does not contain at least one ‘%p’. The only exception is that an empty template (i.e., **--logfile=""**) inhibits the production of log files entirely.

Like C’s `printf()` function, a numeric field width can occur between the ‘%’ and the conversion specifier (the ‘p’ or ‘r’ in the case of **--logfile**). The field is padded on the left with spaces to the given width. More practically, if the field width begins with the number ‘0’ the field is padded on the left with zeroes to the given width. For example, specifying **--logfile="mydata-%03p.log"** on the command line produces log files named ‘mydata-000.log’, ‘mydata-001.log’, ‘mydata-002.log’, ‘mydata-003.log’, and so forth.

--no-trap

--no-trap specifies a list of signals or ranges of signals that should not be trapped. For example, **--no-trap=10-12,17** prevents signals 10, 11, 12, and 17 from being trapped.⁵ Signals can also be referred to by name, with or without a ‘SIG’ prefix. Also, names and numbers can be freely mixed. Hence, **--no-trap=10-12,INT,17,SIGSTOP,SIGCONT** is a valid argument to a CONCEPTUAL program. Because signal reception can adversely affect performance, CONCEPTUAL’s default behavior is to terminate the program on receipt of a signal. However, some signals may be necessary for the underlying communication layer’s proper operation. **--no-trap** enables such signals to pass through CONCEPTUAL untouched. (Some signals, however, are needed by CONCEPTUAL or by a particular backend and are always trapped.)

--seed

--seed (which selects a different default value on each run) is used in any program that utilizes the RANDOM TASK construct (see [Section 4.7.3 \[Binding variables\]](#), page 117) or that sends message WITH VERIFICATION (see [Section 4.4.1 \[Message specifications\]](#), page 94).

If the LOGS statement (see [Section 4.5.3 \[Writing to a log file\]](#), page 104) is used anywhere in a CONCEPTUAL program, then *all* processes write a log file. This is done because CONCEPTUAL log files—even those which contain no measurement data—include a wealth of important information in prologue comments, as described [Section 3.5 \[Interpreting coN-CePTuaL log files\]](#), page 44. As a consequence, a program run with thousands of processes produces thousands of log files. If process 0 is the only process which logs actual data, the ‘ncptl-logmerge’ script (see [Section 3.5.3 \[ncptl-logmerge\]](#), page 70) can merge these log files into a single, more manageable, file. For situations in which it is unreasonable for every process to write a log file (e.g., if the filesystem is unable to handle large numbers of simultaneous file creations, the `NCPTL_LOG_ONLY` environment variable lets the user limit

⁵ On some platforms, these signals correspond to SIGUSR1, SIGSEGV, SIGUSR2, and SIGCHLD, respectively.

the set of processes which produce log files. `NCPTL_LOG_ONLY` accepts a comma-separated list of dash-separated process ranges such as '0-3,12-16,24,25,32-48'. Only processes included in the list produce log files.

3.5 Interpreting coNCePTuaL log files

Any coNCePTuaL program that uses the LOGS keyword (see [Section 4.5.3 \[Writing to a log file\]](#), page 104) will produce a log file as it runs. The coNCePTuaL run-time library writes log files in a simple, plain-text format. In addition to measurement data, a wealth of information is stored within log-file comments. coNCePTuaL comes with a tool, 'ncptl-logextract', which can extract data and other information from a log file and convert it into any of a variety of other formats.

3.5.1 Log-file format

The coNCePTuaL run-time library writes log files in the following (textual) format:

- Lines beginning with '#' are comments.
- Columns are separated by commas.
- Strings are output between double quotes. Literal double-quotes are output as '\"' and literal backslashes are output as '\\'.

A sample log file is listed below. The log file is presented in its entirety.

```
#####
# =====
# coNCePTuaL log file
# =====
# coNCePTuaL version: 0.6.4a
# coNCePTuaL backend: c_mpi (C + MPI)
# Executable name: /home/pakin/src/coNCePTuaL/example
# Working directory: /home/pakin/src/coNCePTuaL
# Command line: ./example
# Number of tasks: 2
# Processor (0<=P<tasks): 0
# Host name: a1
# Operating system: Linux version 2.4.21-3.5qsnet (root@a31) (gcc version 2.96 20000731 (Red Hat Linux 7.
# CPU vendor: GenuineIntel
# CPU architecture: ia64
# CPU count: 2
# CPU frequency: 1300000000 Hz (1.3 GHz)
# Cycle-counter frequency: 1300000000 Hz (1.3 GHz)
# OS page size: 16384 bytes
# Physical memory: 2047901696 bytes (1.9 GB)
# Elan capability: [1965771c.6213bf5d.3a26411c.5b4c5d58] Version 10002 Type a001 Context 640.640.640 Node
# coNCePTuaL configuration: ./configure '--prefix=/tmp/ncptl' 'MPICPPFLAGS=-I/usr/local/include' 'CFLAGS=
# Library compiler+linker: /opt/intel-7.1.033/compiler70/ia64/bin/ecc
# Library compiler version: Intel(R) C++ gcc 3.0 mode [7.1]
# Library compiler options: -g -O3 -ansi_alias -ansi
# Library linker options: -lrmscall -lelan -lpopt
# Library compiler mode: LP64
# Dynamic libraries used: /usr/lib/qsnet/elan3/lib/librmscall.so.1 /usr/lib/qsnet/elan3/lib/libelan.so.1
# Microsecond timer type: gettimeofday()
# Average microsecond timer overhead: <1 microsecond
# Microsecond timer increment: 1.00466 +/- 0.123576 microseconds (ideal: 1 +/- 0)
# Minimum sleep time: 1946.6 +/- 31.5872 microseconds (ideal: 1 +/- 0)
# WARNING: Sleeping exhibits poor granularity (not a serious problem).
```

```

# WARNING: Sleeping has a large error component (not a serious problem).
# Process CPU timer: getrusage()
# Process CPU-time increment: 976.59 +/- 0.494311 microseconds (ideal: 1 +/- 0)
# WARNING: Process timer exhibits poor granularity (not a serious problem).
# Log-file template: example-%p.log
# Number of minutes after which to kill the job (-1=never): -1
# List of signals which should not be trapped: 14
# Log-file checkpointing interval: 60 seconds (i.e., 1 minute)
# MPI send routine: MPI_Send()
# MPI error checking: off
# Front-end compilation command line: ncptl --backend=c_mpi example.ncptl
# Back-end compilation command line: /usr/lib/mpi/mpi_intel/bin/mpicc -I/tmp/ncptl/include -I/usr/local/
# Log creator: Scott Pakin
# Log creation time: Mon Dec 19 12:02:18 2005
#
# Environment variables
# -----
# CVS_RSH: /usr/bin/ssh
# DISPLAY: localhost:19.0
# DYNINSTAPI_RT_LIB: /home/pakin/dyninstAPI-3.0/lib/i386-unknown-linux2.2/libdyninstAPI_RT.so.1
# DYNINST_ROOT: /home/pakin/dyninstAPI-3.0
# EDITOR: /usr/bin/emacs
# GROUP: CCS3
# HOME: /home/pakin
# HOST: a0
# HOSTNAME: a0
# HOSTTYPE: unknown
# KDEDIR: /usr
# LANG: en_US
# LD_LIBRARY_PATH: /opt/intel-7.1.033/compiler70/ia64/lib:/users/pakin/lib:/usr/lib:/usr/ccs/lib:/opt/SUNW
# LESSOPEN: |/usr/bin/lesspipe.sh %s
# LOGNAME: pakin
# LPDEST: lwy
# LS_COLORS: no=00:fi=00:di=01;34:ln=01;36:pi=40;33:so=01;35:bd=40;33;01:cd=40;33;01:or=01;05;37;41:mi=01
# MACHTYPE: unknown
# MAIL: /var/mail/pakin
# MANPATH: /opt/intel-7.1.033/compiler70/man:/home/pakin/man:/usr/man:/opt/SUNWspro/man:/usr/dt/man:/usr/
# MOZILLA_HOME: /usr/local/netscape/java
# NAME: Scott Pakin
# ORGANIZATION: Los Alamos National Lab
# OSTYPE: linux
# PATH: /opt/intel-7.1.033/compiler70/ia64/bin:./home/pakin/bin:/usr/local/bin:/usr/dt/bin:/usr/openwin/
# PRINTER: lwy
# PVM_ROOT: /usr/share/pvm3
# PVM_RSH: /usr/bin/rsh
# PWD: /home/pakin/src/coNcEPTuaL
# QTDIR: /usr/lib/qt-2.3.1
# REMOTEHOST: antero.c3.lanl.gov
# RMS_JOBID: 308
# RMS_MACHINE: a
# RMS_NNODES: 2
# RMS_NODEID: 0
# RMS_NPROCS: 2
# RMS_PROCID: 0
# RMS_RANK: 0
# RMS_RESOURCEID: parallel.318
# RMS_STOPONELANINIT: 0
# SHELL: /bin/tcsh
# SHLVL: 2

```

```

# SSH_AGENT_PID: 8586
# SSH_ASKPASS: /usr/libexec/openssh/gnome-ssh-askpass
# SSH_AUTH_SOCK: /tmp/ssh-XXfG457q/agent.8562
# SSH_CLIENT: 128.165.20.177 47456 22
# SSH_TTY: /dev/pts/10
# SUPPORTED: en_US:en
# TERM: xterm
# TZ: MST7MDT
# USER: pakin
# VENDOR: unknown
#
# coNcEPTuaL source code
# -----
#     FOR 10 REPETITIONS {
#         TASK 0 RESETS ITS COUNTERS THEN
#         TASK 0 SENDS A 0 BYTE MESSAGE TO TASK 1 THEN
#         TASK 1 SENDS A 0 BYTE MESSAGE TO TASK 0 THEN
#         TASK 0 LOGS EACH elapsed_usecs/2 AS "Latency (usecs)"
#     }
#
#####
"Latency (usecs)"
"(all data)"
192.5
7
5.5
5.5
5
5
5
5.5
5
5.5
#####
# Program exited normally.
# Log completion time: Mon Dec 19 12:02:18 2005
# Elapsed time: 0 seconds
# Process CPU usage (user+system): 0 seconds
# Number of interrupts received (all CPUs): 22
# Peak memory allocation: 1072987 bytes (1.0 MB)
#####

```

As the preceding example indicates, a log file's comment block can be divided into multiple stanzas:

- a list of $\langle key:value \rangle$ pairs that describe various characteristics of the run-time environment, including hardware and software identification, timer quality, values of command-line arguments, and a timestamp
- a dump of the environment variables active when the program ran
- the complete `coNcEPTuaL` source program

Two rows of column headers and the measurement data follow the prologue comment block. A set of epilogue comments completes the log file.

The motivation for writing so much information to the log file is to facilitate reproduction of the experiment. The ideal situation is for a third party to be able to look at a

CONCEPTUAL log file and from that, recreate the original experiment and get identical results.

Some of the comments that may benefit from additional explanation include the following:

‘Library compiler mode’

‘LP64’ means “**L**ong integers and **P**ointers contain exactly **64** bits while ordinary integers contain exactly **32** bits”. ‘ILP32’ means “ordinary **I**ntegers, **L**ong integers, and **P**ointers all contain exactly **32** bits”. The library compiler mode will be ‘nonstandard’ for any other combination of datatype sizes.

‘Average timer overhead’

During initialization, the CONCEPTUAL run-time library performs some calibration routines. Among these routines is a measurement of the quality of whatever mechanisms the library is using to measure elapsed time. In the sample log file presented above, the mechanism used was ‘inline assembly code’, meaning the run-time library reads the hardware cycle counter without going through a standard library or system call. The complete set of timer mechanisms and selection criteria is presented in [Section 6.3.5 \[Time-related functions\], page 150](#), in the documentation for the `ncptl_time()` function.

The log file then reports “average timer overhead” as the mean time between back-to-back invocations of whichever timer routine is being used. Ideally, the mean should be ‘<1 microsecond’ but this is not the case on all systems. Large values indicate that a performance penalty is charged every time a CONCEPTUAL program reads the timer.

‘Timer increment’

In addition to measuring the overhead of reading the timer, the CONCEPTUAL run-time library also measures timer accuracy. The library expects to be able to read the timer with microsecond accuracy. That is, the time reported should not increase by more than a microsecond between successive readings of the timer. To gauge timer accuracy, the library’s initialization routine performs a number of back-to-back invocations of the timer routine and reports the mean and standard deviation of the number of microseconds that elapsed between readings, discarding any deltas of zero microseconds. Ideally, the microsecond timer, when read multiple times in rapid succession, should report nonzero increments of exactly one microsecond with no variation. The log file will contain warning messages if the increment or standard deviation are excessively large, as this may indicate a large margin of error in the measurement data.

‘Process CPU-time increment’

‘Process CPU usage (user+system)’

Log files end with an epilogue section which includes ‘Process CPU usage (user+system)’, which indicates the subset of total wall-clock time for which the program was running (‘user’) or for which the operating system was running on the program’s behalf (‘system’). The log-file prologue reports as ‘Process CPU-time increment’ the resolution of the timer user to report process CPU time. Note that process CPU time is not exported to

CONCEPTUAL programs; it is therefore much less critical than the wall-clock timer and is reported primarily for informational purposes.

‘Number of interrupts received (all CPUs)’

On certain platforms, CONCEPTUAL can tally the number of CPU interrupts that were processed during the run of the program. Because a multiprocessor may migrate tasks among CPUs during their execution, a per-CPU interrupt count may have little merit. Consequently, the number reported represents the sum across all CPUs in the same node (but not across nodes). CONCEPTUAL attempts to read interrupt information from the `/proc/interrupts` file if no alternative mechanism is available. In case having a large number of processes accessing `/proc/interrupts` poses a problem the `--disable-proc-interrupts` configuration option prevents programs from accessing that file.

‘Peak memory allocation’

CONCEPTUAL programs heap-allocate memory for a variety of purposes: message buffers, event lists (see [Section 6.2.3 \[Generated code\], page 141](#)), unaggregated performance data (see [\[Computing aggregates\], page 106](#)), etc. Allocated memory which is no longer needed is returned to the heap. The total amount of allocated memory the program is holding therefore increases and decreases over time. The ‘Peak memory allocation’ comment reports the maximum amount of memory the program held at any given time. If this value nears or exceeds the value reported for ‘Physical memory’, it is possible that paging overheads may be negatively impacting some of the program’s timing measurements. On systems without demand paging, exceeding ‘Physical memory’ is likely to crash the program; a large ‘Peak memory allocation’ on a smaller run can therefore help explain why a larger run is crashing. ‘Peak memory allocation’ includes only memory allocated using the memory-allocation functions described in [Section 6.3.3 \[Memory-allocation functions\], page 147](#).

3.5.2 ‘ncptl-logextract’

To facilitate converting CONCEPTUAL log files into input data for other applications, CONCEPTUAL provides a Perl script called ‘ncptl-logextract’. ‘ncptl-logextract’ can extract the data from a log file—as well as various information that appears in the log file’s comments—into a variety of formats suitable for graphing or typesetting.

Running `ncptl-logextract --usage` causes ‘ncptl-logextract’ to list a synopsis of its core command-line options to the standard output device; running `ncptl-logextract --help` produces basic usage information; and, running `ncptl-logextract --man` outputs a complete manual page. See [\[ncptl-logextract manual page\], page 49](#), shows the ‘ncptl-logextract’ documentation as produced by `ncptl-logextract --man`.

NAME

ncptl-logextract - Extract various bits of information from a CONCEPTUAL log file

SYNOPSIS

ncptl-logextract *--usage* | *--help* | *--man*

```
ncptl-logextract  [--extract=[data|params|env|source|warnings]]  [--format=format]
[format-specific options...] [--before=string] [--after=string] [--force-merge[=number]]
[--procs=string] [--quiet] [--verbose] [--output=filename] [filename...]
```

DESCRIPTION

Background CONCEPTUAL is a domain-specific programming language designed to facilitate writing networking benchmarks and validation suites. CONCEPTUAL programs can log data to a file but in only a single file format. ‘ncptl-logextract’ extracts this log data and outputs it in a variety of formats for use with other applications.

The CONCEPTUAL-generated log files that serve as input to ‘ncptl-logextract’ are plain ASCII files. Syntactically, they contain a number of newline-separated tables. Each table contains a number of newline-separated rows of comma-separated columns. This is known generically as *comma-separated value* or *CSV* format. Each table begins with two rows of header text followed by one or more rows of numbers. Text is written within double quotes. Double-quote characters and backslashes within text are escaped with a backslash. No other escaped characters are recognized. Lines that begin with **#** are considered comments.

Semantically, there are four types of data present in every CONCEPTUAL-generated log file:

1. The complete source code of the CONCEPTUAL program that produced the log file
2. Characteristics of the run-time environment and the values of all command-line parameters
3. A list of warning messages that CONCEPTUAL issued while analyzing the run-time environment
4. One or more tables of measurement data produced by the CONCEPTUAL program

The first three items appear within comment lines. The measurement data is written in CSV format.

Extracting information from coNCePTuaL log files It is common to want to extract information (especially measurement data) from log files. For simple formatting operations, a one-line awk or Perl script suffices. However, as the complexity of the formatting increases, the complexity of these scripts increases even more. That’s where ‘ncptl-logextract’ fits in. ‘ncptl-logextract’ makes it easy to extract any of the four types of log data described above and format it in variety of ways. Although the number of options that ‘ncptl-logextract’ supports may be somewhat daunting, it is well worth learning how to use ‘ncptl-logextract’ to avoid reinventing the wheel every time a CONCEPTUAL log file needs to be processed. ‘ncptl-logextract’ takes care of all sorts of special cases that crop up when manipulating CONCEPTUAL log files.

OPTIONS

‘ncpt1-logextract’ accepts the following command-line options regardless of what data is extracted from the log file and what formatting occurs:

- h, --help**
Output the Synopsis section and the Options section then exit the program.
- m, --man** Output a complete Unix man (“manual”) page for ‘ncpt1-logextract’ then exit the program.
- e info, --extract=info**
Specify what sort of data should be extracted from the log file. Acceptable values for *info* are listed and described in the Additional Options section and include **data**, **params**, **env**, and **source**.
- f format, --format=format**
Specify how the extracted data should be formatted. Valid arguments depend upon the value passed to **--extract** and include such formats as **csv**, **html**, **latex**, **text**, and **bash**. See the Additional Options section for details, explanations, and descriptions of applicability.
- b string, --before=string**
Output an arbitrary string of text before any other output. *string* can contain escape characters such as **\n** for newline, **\t** for tab, and **** for backslash.
- a string, --after=string**
Output an arbitrary string of text after all other output. *string* can contain escape characters such as **\n** for newline, **\t** for tab, and **** for backslash.
- F [number], --force-merge[=number]**
Try extra hard to merge multiple log files, even if they seem to have been produced by different programs or in different execution environments. This generally implies padding empty rows and columns with blanks. However, if **--force-merge** is given a numeric argument, the value of that argument is used instead of blanks to pad empty locations. Note that **--force-merge** is different from **--force-merge=0** because data-merging functions (**mean**, **max**, etc.) ignore blanks but consider zeroes.
- p string, --procs=string**
When given a “merged” log file, unmerge only the data corresponding to the comma-separated processor ranges in *string*. For example, **--procs=0,16-20,25** unmerges the data for processors 0, 16, 17, 18, 19, 20, and 25. By default, ‘ncpt1-logextract’ uses all of the data from a merged log file.
- q, --quiet**
Suppress progress output. Normally, ‘ncpt1-logextract’ outputs status information regarding its operation. The **--quiet** option instruct ‘ncpt1-logextract’ to output only warning and error messages.
- v, --verbose**
Increase progress output. Normally, ‘ncpt1-logextract’ outputs basic status information regarding its operation. The **--verbose** option instruct

‘ncptl-logextract’ to output more detailed information. Each time `--verbose` is specified, the program’s verbosity increases (up to a maximum).

`-o filename, --output=filename`

Redirect the output from ‘ncptl-logextract’ to a file. By default, ‘ncptl-logextract’ writes to the standard output device.

The above is merely a terse summary of the ‘ncptl-logextract’ command-line options. The reader is directed to the Additional Options section for descriptions of the numerous ways that ‘ncptl-logextract’ can format information. Note that `--extract` and `--format` are the two most common options as they specify what to extract and how to format it; most of the remaining options in the Additional Options section exist to provide precise control over formatting details.

ADDITIONAL OPTIONS

The ‘ncptl-logextract’ command-line options follow a hierarchy. At the top level is `--extract`, which specifies which of the four types of data ‘ncptl-logextract’ should extract. Next, `--format` specifies how the extracted data should be formatted. Valid values for `--format` differ based on the argument to `--extract`. Finally, there are various format-specific options that fine-tune the formatted output. Each output format accepts a different set of options. Many of the options appear at multiple places within the hierarchy, although usually with different default values.

The following hierarchical list describes all of the valid combinations of `--extract`, `--format`, and the various format-specific options:

`--extract=data` [default]

Extract measurement data

`--format=csv` [default]

Output each table in comma-separated-value format

`--noheaders`

Do not output column headers

`--colbegin=string`

Specify the text placed at the beginning of each data column [default: “”]

`--colsep=string`

Specify the text used to separate data columns [default: “,”]

`--colend=string`

Specify the text placed at the end of each data column [default: “”]

`--rowbegin=string`

Specify the text placed at the beginning of each data row [default: “”]

`--rowsep=string`

Specify the text used to separate data rows [default: “”]

`--rowend=string`
Specify the text placed at the end of each data row
[default: “\n”]

`--hcolbegin=string`
Specify the text placed at the beginning of each header
column [default: same as `colbegin`]

`--hcolsep=string`
Specify the text used to separate header columns [de-
fault: same as `colsep`]

`--hcolend=string`
Specify the text placed at the end of each header col-
umn [default: same as `colend`]

`--hrowbegin=string`
Specify the text placed at the beginning of each header
row [default: same as `rowbegin`]

`--hrowsep=string`
Specify the text used to separate header rows [default:
same as `rowsep`]

`--hrowend=string`
Specify the text placed at the end of each header row
[default: same as `rowend`]

`--tablebegin=string`
Specify the text placed at the beginning of each table
[default: “”]

`--tablesep=string`
Specify the text used to separate tables [default: “\n”]

`--tableend=string`
Specify the text placed at the end of each table [default:
“”]

`--quote=string`
Specify the text used to begin quoted text [default:
“”]

`--unquote=string`
Specify the text used to end quoted text [default: same
as `quote`]

`--excel` Output strings in a format readable by Microsoft Excel

`--merge=function`
Specify how to merge data from multiple files [default:
“mean”]

`--showfnames=option`
Add an extra header row showing the filename the data
came from [default: “none”]

`--format=tsv`
Output each table in tab-separated-value format

`--noheaders`
Do not output column headers

`--colbegin=string`
Specify the text placed at the beginning of each data column [default: “”]

`--colsep=string`
Specify the text used to separate data columns [default: “\t”]

`--colend=string`
Specify the text placed at the end of each data column [default: “”]

`--rowbegin=string`
Specify the text placed at the beginning of each data row [default: “”]

`--rowsep=string`
Specify the text used to separate data rows [default: “”]

`--rowend=string`
Specify the text placed at the end of each data row [default: “\n”]

`--hcolbegin=string`
Specify the text placed at the beginning of each header column [default: same as `colbegin`]

`--hcolsep=string`
Specify the text used to separate header columns [default: same as `colsep`]

`--hcolend=string`
Specify the text placed at the end of each header column [default: same as `colend`]

`--hrowbegin=string`
Specify the text placed at the beginning of each header row [default: same as `rowbegin`]

`--hrowsep=string`
Specify the text used to separate header rows [default: same as `rowsep`]

`--hrowend=string`
Specify the text placed at the end of each header row [default: same as `rowend`]

`--tablebegin=string`
Specify the text placed at the beginning of each table
[default: “”]

`--tablesep=string`
Specify the text used to separate tables [default: “\n”]

`--tableend=string`
Specify the text placed at the end of each table [default:
“”]

`--quote=string`
Specify the text used to begin quoted text [default:
“”]

`--unquote=string`
Specify the text used to end quoted text [default: same
as `quote`]

`--excel` Output strings in a format readable by Microsoft Excel

`--merge=function`
Specify how to merge data from multiple files [default:
“mean”]

`--showfnames=option`
Add an extra header row showing the filename the data
came from [default: “none”]

`--format=html`
Output each table in HTML table format

`--noheaders`
Do not output column headers

`--colbegin=string`
Specify the text placed at the beginning of each data
column [default: “<td>”]

`--colsep=string`
Specify the text used to separate data columns [default:
“ ”]

`--colend=string`
Specify the text placed at the end of each data column
[default: “</td>”]

`--rowbegin=string`
Specify the text placed at the beginning of each data
row [default: “<tr>”]

`--rowsep=string`
Specify the text used to separate data rows [default:
“”]

--rowend=string
Specify the text placed at the end of each data row
[default: “</tr>\n”]

--hcolbegin=string
Specify the text placed at the beginning of each header
column [default: “<th>”]

--hcolsep=string
Specify the text used to separate header columns [de-
fault: same as `colsep`]

--hcolend=string
Specify the text placed at the end of each header col-
umn [default: “</th>”]

--hrowbegin=string
Specify the text placed at the beginning of each header
row [default: same as `rowbegin`]

--hrowsep=string
Specify the text used to separate header rows [default:
same as `rowsep`]

--hrowend=string
Specify the text placed at the end of each header row
[default: same as `rowend`]

--tablebegin=string
Specify the text placed at the beginning of each table
[default: “<table>\n”]

--tablesep=string
Specify the text used to separate tables [default: “”]

--tableend=string
Specify the text placed at the end of each table [default:
“</table>\n”]

--quote=string
Specify the text used to begin quoted text [default: “”]

--unquote=string
Specify the text used to end quoted text [default: same
as `quote`]

--merge=function
Specify how to merge data from multiple files [default:
“mean”]

--showfnames=option
Add an extra header row showing the filename the data
came from [default: “none”]

`--format=gnuplot`
Output each table as a gnuplot data file

`--noheaders`
Do not output column headers

`--colbegin=string`
Specify the text placed at the beginning of each data column [default: “”]

`--colsep=string`
Specify the text used to separate data columns [default: “ ”]

`--colend=string`
Specify the text placed at the end of each data column [default: “”]

`--rowbegin=string`
Specify the text placed at the beginning of each data row [default: “”]

`--rowsep=string`
Specify the text used to separate data rows [default: “”]

`--rowend=string`
Specify the text placed at the end of each data row [default: “\n”]

`--hcolbegin=string`
Specify the text placed at the beginning of each header column [default: same as `colbegin`]

`--hcolsep=string`
Specify the text used to separate header columns [default: same as `colsep`]

`--hcolend=string`
Specify the text placed at the end of each header column [default: same as `colend`]

`--hrowbegin=string`
Specify the text placed at the beginning of each header row [default: “# ”]

`--hrowsep=string`
Specify the text used to separate header rows [default: same as `rowsep`]

`--hrowend=string`
Specify the text placed at the end of each header row [default: same as `rowend`]

`--tablebegin=string`
Specify the text placed at the beginning of each table
[default: “”]

`--tablesep=string`
Specify the text used to separate tables [default:
“\n\n”]

`--tableend=string`
Specify the text placed at the end of each table [default:
“”]

`--quote=string`
Specify the text used to begin quoted text [default:
“”]

`--unquote=string`
Specify the text used to end quoted text [default: same
as `quote`]

`--merge=function`
Specify how to merge data from multiple files [default:
“mean”]

`--shownames=option`
Add an extra header row showing the filename the data
came from [default: “none”]

`--format=octave`
Output each table as an Octave text-format data file

`--noheaders`
Do not output column headers

`--colbegin=string`
Specify the text placed at the beginning of each data
column [default: “”]

`--colsep=string`
Specify the text used to separate data columns [default:
“”]

`--colend=string`
Specify the text placed at the end of each data column
[default: “\n”]

`--rowbegin=string`
Specify the text placed at the beginning of each data
row [default: “”]

`--rowend=string`
Specify the text placed at the end of each data row
[default: “”]

--hcolbegin=string
Specify the text placed at the beginning of each header column [default: “”]

--hcolsep=string
Specify the text used to separate header columns [default: “_”]

--hcolend=string
Specify the text placed at the end of each header column [default: “”]

--hrowbegin=string
Specify the text placed at the beginning of each header row [default: “# ”]

--hrowsep=string
Specify the text used to separate header rows [default: “”]

--hrowend=string
Specify the text placed at the end of each header row [default: “\n”]

--tablebegin=string
Specify the text placed at the beginning of each table [default: “”]

--tablesep=string
Specify the text used to separate tables [default: “\n”]

--tableend=string
Specify the text placed at the end of each table [default: “”]

--quote=string
Specify the text used to begin quoted text [default: “”]

--unquote=string
Specify the text used to end quoted text [default: same as quote]

--merge=function
Specify how to merge data from multiple files [default: “mean”]

--showfnames=option
Add an extra header row showing the filename the data came from [default: “none”]

--format=custom
Output each table in a completely user-specified format

--noheaders
Do not output column headers

`--colbegin=string`
Specify the text placed at the beginning of each data column [default: `""`]

`--colsep=string`
Specify the text used to separate data columns [default: `""`]

`--colend=string`
Specify the text placed at the end of each data column [default: `""`]

`--rowbegin=string`
Specify the text placed at the beginning of each data row [default: `""`]

`--rowsep=string`
Specify the text used to separate data rows [default: `""`]

`--rowend=string`
Specify the text placed at the end of each data row [default: `""`]

`--hcolbegin=string`
Specify the text placed at the beginning of each header column [default: same as `colbegin`]

`--hcolsep=string`
Specify the text used to separate header columns [default: same as `colsep`]

`--hcolend=string`
Specify the text placed at the end of each header column [default: same as `colend`]

`--hrowbegin=string`
Specify the text placed at the beginning of each header row [default: same as `rowbegin`]

`--hrowsep=string`
Specify the text used to separate header rows [default: same as `rowsep`]

`--hrowend=string`
Specify the text placed at the end of each header row [default: same as `rowend`]

`--tablebegin=string`
Specify the text placed at the beginning of each table [default: `""`]

`--tablesep=string`
Specify the text used to separate tables [default: `""`]

```

--tableend=string
    Specify the text placed at the end of each table [default:
    “”]

--quote=string
    Specify the text used to begin quoted text [default: “”]

--unquote=string
    Specify the text used to end quoted text [default: same
    as quote]

--excel    Output strings in a format readable by Microsoft Excel

--merge=function
    Specify how to merge data from multiple files [default:
    “mean”]

--showfnames=option
    Add an extra header row showing the filename the data
    came from [default: “none”]

--format=latex
    Output each table as a LATEX tabular environment

--dcolumn
    Use the dcolumn package to align numbers on the dec-
    imal point

--booktabs
    Use the booktabs package for a more professionally
    typeset look

--longtable
    Use the longtable package to enable multi-page tables

--merge=function
    Specify how to merge data from multiple files [default:
    “mean”]

--showfnames=option
    Add an extra header row showing the filename the data
    came from [default: “none”]

--extract=params
    Extract the program’s run-time parameters and environment variables

--format=text [default]
    Output the parameters in plain-text format

--include=filename
    Read from a file the list of keys to output

--exclude=regexp
    Ignore any keys whose name matches a regular expres-
    sion

```

```

--sort      Sort the list of parameters alphabetically by key
--noenv     Exclude environment variables
--noparams  Exclude run-time parameters
--envformat=template
            Format environment variable names using the given
            template [default: "%s (environment variable)"]
--columns=number
            Output the parameters as a 1-, 2-, or 3-column table
            [default: 1]
--colsep=string
            Specify the text used to separate data columns [default:
            ": "]
--rowbegin=string
            Specify the text that's output at the start of each data
            row [default: ""]
--rowend=string
            Specify the text that's output at the end of each data
            row [default: "\\n"]
--format=dumpkeys
            Output a list of the keys only (i.e., no values)
--include=filename
            Read the list of parameters to output from a given file
--exclude=regexp
            Ignore any keys whose name matches a regular expres-
            sion
--envformat=template
            Format environment variable names using the given
            template [default: "%s (environment variable)"]
--sort      Sort the list of parameters alphabetically by key
--noenv     Exclude environment variables
--noparams  Exclude run-time parameters
--format=latex
            Output the parameters as a LATEX tabular environment
--include=filename
            Read from a file the list of keys to output
--exclude=regexp
            Ignore any keys whose name matches a regular expres-
            sion

```

```
--envformat=template
    Format environment variable names using the given
    template [default: "%s (environment variable)"]
--sort      Sort the list of parameters alphabetically by key
--booktabs  Use the booktabs package for a more professionally
            typeset look
--tabularx  Use the tabularx package to enable line wraps within
            the value column
--longtable Use the longtable package to enable multi-page tables
--noenv     Exclude environment variables
--noparams  Exclude run-time parameters

--extract=env
    Extract the environment in which the program was run
--format=sh [default]
    Use Bourne shell syntax for setting environment variables
--newlines  Separate commands with newlines instead of
            semicolons
--unset     Unset all other environment variables
--chdir     Switch to the program's original working directory

--format=bash
    Use Bourne Again shell syntax for setting environment variables
--newlines  Separate commands with newlines instead of
            semicolons
--unset     Unset all other environment variables
--chdir     Switch to the program's original working directory

--format=ksh
    Use Korn shell syntax for setting environment variables
--newlines  Separate commands with newlines instead of
            semicolons
--unset     Unset all other environment variables
--chdir     Switch to the program's original working directory
```

```
--format=csh
    Use C shell syntax for setting environment variables
    --newlines
        Separate commands with newlines instead of
        semicolons
    --unset   Unset all other environment variables
    --chdir   Switch to the program's original working directory

--format=zsh
    Use Z shell syntax for setting environment variables
    --newlines
        Separate commands with newlines instead of
        semicolons
    --unset   Unset all other environment variables
    --chdir   Switch to the program's original working directory

--format=tcsh
    Use tcsh syntax for setting environment variables
    --newlines
        Separate commands with newlines instead of
        semicolons
    --unset   Unset all other environment variables
    --chdir   Switch to the program's original working directory

--format=ash
    Use ash syntax for setting environment variables
    --newlines
        Separate commands with newlines instead of
        semicolons
    --unset   Unset all other environment variables
    --chdir   Switch to the program's original working directory

--extract=source
    Extract CONCEPTUAL source code
--format=text [default]
    Output the source code in plain-text format
--linebegin=string
    Specify the text placed at the beginning of each line
    [default: ""]
--lineend=string
    Specify the text placed at the end of each line [default:
    "\\n"]
```

```
--kwbegin=string
    Specify the text placed before each keyword [default:
    “”]

--kwend=string
    Specify the text placed after each keyword [default: “”]

--strbegin=string
    Specify the text placed before each string [default: “”]

--strend=string
    Specify the text placed after each string [default: “”]

--combegin=string
    Specify the text placed before each comment [default:
    “”]

--comend=string
    Specify the text placed after each comment [default:
    “”]

--indent=number
    Indent each line by a given number of spaces

--wrap=number
    Wrap the source code into a paragraph with a given
    character width

--extract=warnings
    Extract a list of warnings the program issued during initialization

--format=text [default]
    Output warnings in plain-text format

--listbegin=string
    Specify text to appear at the beginning of the list [de-
    fault: “”]

--listend=string
    Specify text to appear at the end of the list [default:
    “”]

--itembegin=string
    Specify text to appear before each warning [default:
    “* ”]

--itemend=string
    Specify text to appear after each warning [default:
    “\n”]

--format=html
    Output warnings as an HTML list

--listbegin=string
    Specify text to appear at the beginning of the list [de-
    fault: “<ul>\n”]
```

```

--listend=string
    Specify text to appear at the end of the list [default:
    "</ul>\n"]

--itembegin=string
    Specify text to appear before each warning [default:
    "<li>"]

--itemend=string
    Specify text to appear after each warning [default:
    "</li>\n"]

--format=latex
    Output warnings as a LATEX list

--listbegin=string
    Specify text to appear at the beginning of the list [de-
    fault: "\begin@{itemize@}\n"]

--listend=string
    Specify text to appear at the end of the list [default:
    "\end@{itemize@}\n"]

--itembegin=string
    Specify text to appear before each warning [default:
    "\item "]

--itemend=string
    Specify text to appear after each warning [default:
    "\n"]

```

The following represent additional clarification for some of the above:

- If `--indent` is specified without an argument, the argument defaults to 2.
- If `--wrap` is specified without an argument, the argument defaults to 72.
- The following are examples of the different arguments to the `--columns` option:

```
--columns=1 (default)
```

```

coNCePTuaL version: 1.0
coNCePTuaL backend: c_mpi
Average timer overhead [gettimeofday()]: <1 microsecond
Log creation time: Thu Mar 27 19:22:48 2003
Log completion time: Thu Mar 27 19:22:48 2003

```

```
--columns=2
```

```

coNCePTuaL version:                1.0
coNCePTuaL backend:                c_mpi
Average timer overhead [gettimeofday()]: <1 microsecond
Log creation time:                  Thu Mar 27 19:22:48 2003
Log completion time:                Thu Mar 27 19:22:48 2003

```

```
--columns=3
```

```

coNCePTuaL version                : 1.0

```

```

coNCePTuaL backend           : c_mpi
Average timer overhead [gettimeofday()]: <1 microsecond
Log creation time           : Thu Mar 27 19:22:48 2003
Log completion time        : Thu Mar 27 19:22:48 2003

```

- `--dumpkeys` produces suitable input for the `--include` option.
- `--exclude` can be specified repeatedly on the command line.
- `--merge` takes one of `mean` (arithmetic mean), `hmean` (harmonic mean), `min` (minimum), `max` (maximum), `median` (median), `all` (all values from each column), or `concat` (horizontal concatenation of all data), and applies the function to corresponding data values across all of the input files. `--merge` can also accept a comma-separated list of the above functions, one per data column. This enables a different merge operation to be used for each column. For example, `--merge=min,min,mean` will take the minimum value across all files of each element in the first and second columns and the arithmetic mean across all files of each element in the third column. If the number of comma-separated values differs from the number of columns and `--force-merge` is specified, ‘ncptl-logextract’ will cycle over the given values until all columns are accounted for. The `concat` merge type applies to all columns and therefore cannot be combined with any other merge type. The difference between `--merge=all` and `--merge=concat` is that the former merges three files each with columns *A* and *B* as {*A*, *A*, *A*, *B*, *B*, *B*} while the latter merges the same files as {*A*, *B*, *A*, *B*, *A*, *B*}.
- `--showfnames` prepends to each data table in the input file an extra header line indicating the log file the data was extracted from. This option makes sense only when data is being extracted and primarily when `--merge=all` is specified. `--showfnames` takes one of `none`, `all`, or `first`. The default is `none`, which doesn’t add an extra header row. `all` repeats the filename in each column of the extra header row. `first` outputs the filename in only the first column, leaving the remaining columns with an empty string. The following examples show how a sample data table is formatted with `--showfnames` set in turn to each of `none`, `all`, and `first`:

- Set to `none` (the default):

```

"Size", "Value"
1,2
2,4
3,6

```

- Set to `all` (filename repeated in each column of the first row):

```

"mydata.log", "mydata.log"
"Size", "Value"
1,2
2,4
3,6

```

- Set to `first` (filename shown only in the first column of the first row):

```

"mydata.log", ""
"Size", "Value"
1,2
2,4
3,6

```


- If `--format=params` is used with both `--longtable` and `--tabularx`, the generated table will be formatted for use with the `ltxtable` L^AT_EX package. See the `ltxtable` documentation for more information.

NOTES

If no filenames are given, `'ncptl-logextract'` will read from the standard input device. If multiple log files are specified, `CONCEPTUAL` will merge the data values and take all other information from the first file specified. Note, however, that all of the log files must have been produced by the same `CONCEPTUAL` program and that that program must have been run in the same environment. In other words, only the data values may change across log files; everything else must be invariant. See the description of `--merge` in the Additional Options section for more information about merging data values from multiple log files.

`'ncptl-logextract'` treats certain files specially:

- If `'ncptl-logextract'` is given a filename ending in `'.gz'`, `'.bz2'`, or `'.Z'` it automatically decompresses the file to a temporary location using `'gunzip'`, `'bunzip2'`, or `'uncompress'`, as appropriate, then recursively processes the decompressed file.
- If `'ncptl-logextract'` is given a filename ending in `'.tar'` or `'.zip'` it automatically extracts the file's contents to a temporary directory using `'tar'` or `'unzip'`, as appropriate, then recursively processes the temporary directory.
- If `'ncptl-logextract'` is given the name of a directory it processes all of the plain files found (recursively) beneath that directory.
- If an input file is a merged `CONCEPTUAL` log file (i.e., produced by `'ncptl-logmerge'`), `'ncptl-logextract'` automatically invokes `'ncptl-logunmerge'` to split the file into its constituent, ordinary log files then recursively processes those.

`'ncptl-logmerge'` treats filenames ending in `'.tgz'` as if they ended in `'.tar.gz'` and filenames ending in `'.taz'` as if they ended in `'.tar.Z'`.

If the argument provided to any `'ncptl-logextract'` option begins with an at sign (“@”), the value is treated as a filename and is replaced by the file's contents. To specify a non-filename argument that begins with an at sign, merely prepend an additional “@”:

`--this=that`

The option `this` is given the value “that”.

`--this=@that`

The option `this` is set to the contents of the file called `'that'`.

`--this=@@that`

The option `this` is given the value “@that”.

EXAMPLES

For the following examples, we assume that `'results.log'` is the name of a log file produced by a `CONCEPTUAL` program.

Extract the data in CSV format and write it to `'results.csv'`:

```
ncptl-logextract --extract=data results.log --output=results.csv
```

Note that `--extract=data` is the default and therefore optional:

```
ncptl-logextract results.log --output=results.csv
```

‘ncptl-logextract’ can combine data from multiple log files (using an arithmetic mean by default):

```
ncptl-logextract results-*.log --output=results.csv
```

Put the data from all of the log files side-by-side and produce a CSV file that Microsoft Excel can read directly:

```
ncptl-logextract results-*.log --output=results.csv --merge=all \
--showfnames=first --excel
```

Output the data from ‘result.log’ in tab-separated-value format:

```
ncptl-logextract --format=tsv results.log
```

Output the data in space-separated-value format:

```
ncptl-logextract --colsep=" " results.log
```

Use ‘gnuplot’ to draw a PostScript graph of the data:

```
ncptl-logextract results.log --format=gnuplot \
--before=@params.gp | gnuplot > results.eps
```

In the above, the ‘params.gp’ file might contain ‘gnuplot’ commands such as the following:

```
set terminal postscript eps enhanced color "Times-Roman" 30
set output
set logscale xy
set data style linespoints
set pointsize 3
plot "-" title "Latency"
```

(There should be an extra blank line at the end of the file because ‘ncptl-logextract’ strips off a trailing newline character whenever it reads a file using “@”.)

Produce a complete HTML file of the data (noting that `--format=html` produces only tables, not complete documents):

```
ncptl-logextract --format=html
--before='<html>\n<head>\n<title>Data</title>\n</head>\n<body>\n' \
--after='</body>\n</html>\n' results.log
```

Output the data as a \LaTeX tabular, relying on both the (standard) `dcolumn` and (non-standard) `booktabs` packages for more attractive formatting:

```
ncptl-logextract --format=latex --dcolumn --booktabs \
--output=results.tex results.log
```

Output the run-time parameters in the form “key ---> value” with all of the arrows aligned:

```
ncptl-logextract results.log --extract=params --columns=3 --colsep=" ---> "
```

Output the run-time parameters as an HTML description list:

```
ncptl-logextract results.log --extract=params --before='<dl>' \
--rowbegin='<dt>' --colsep='</dt><dd>' --rowend='</dd>\n' \
--after='</dl>\n'
```

Restore the exact execution environment that was used to produce ‘results.log’, including the current working directory (assuming that ‘bash’ is the current command shell):

```
eval 'ncptl-logextract --extract=env --format=bash \
--unset --chdir results.log'
```

Set all of the environment variables that were used to produce ‘results.log’, overwriting—but not removing—whatever environment variables are currently set (assuming that ‘tcsh’ is the current command shell):

```
eval 'ncptl-logextract --extract=env --format=tcsh results.log'
```

Extract the source code that produced ‘results.log’:

```
ncptl-logextract --extract=source results.log
```

Do the same, but indent the code by four spaces then re-wrap it into a 60-column paragraph:

```
ncptl-logextract --extract=source --indent=4 --wrap=60 results.log
```

Here are a variety of ways to express the same thing:

```
ncptl-logextract -e source --indent=4 --wrap=60 results.log
ncptl-logextract -e source --indent=4 results.log --wrap=60
cat results.log | ncptl-logextract --wrap=60 --indent=4 -e source
```

Output the source code wrapped to 72 columns, with no indentation, and formatted within an HTML preformatted-text block:

```
ncptl-logextract --extract=source --wrap --before="<PRE>\n" \
after="</PRE>\n" results.log
```

List all of the warning messages which occur in ‘results.log’:

```
ncptl-logextract --extract=warnings results.log
```

SEE ALSO

ncptl-logmerge(1), ncptl-logunmerge(1), the CONCEPTUAL User’s Guide

AUTHOR

Scott Pakin, pakin@lanl.gov

3.5.3 ‘ncptl-logmerge’

CONCEPTUAL programs produce one log file per process. An unwieldy number of files can therefore be generated on large-scale computer systems. For the case in which only a single log file contains measurement data, ‘ncptl-logmerge’ can merge a number of log files into a single file. Only lines that differ across log files are repeated, making the result fairly space-efficient. The primary advantage of ‘ncptl-logmerge’ over an archiving program such as ‘tar’ or ‘zip’ is that the output of ‘ncptl-logmerge’ is designed to be human-readable—in fact, easily readable.

‘ncptl-logmerge’ can also be used to highlight differences in log files. It is therefore an important diagnostic tool for unearthing subtle configuration discrepancies across nodes in a large-scale computer system.

Running *ncptl-logmerge --usage* causes ‘ncptl-logmerge’ to list a synopsis of its core command-line options to the standard output device; running *ncptl-logmerge --help* produces basic usage information; and, running *ncptl-logmerge --man* outputs a complete manual page. See [\[ncptl-logmerge manual page\]](#), page 71, shows the ‘ncptl-logmerge’ documentation as produced by *ncptl-logmerge --man*.

NAME

ncptl-logmerge - Merge CONCEPTUAL log files

SYNOPSIS

ncptl-logmerge *--usage* | *--help* | *--man*

ncptl-logmerge [*--output=filename*] [*--simplify*] *filename...*

DESCRIPTION

A CONCEPTUAL program produces one log file per process. For large numbers of processes the result can be unwieldy. ‘ncptl-logmerge’ combines a large set of log files into a single, merged file which can later be expanded back into its constituent log files. There are a number of restrictions on the input to ‘ncptl-logmerge’; see the Restrictions section for details.

The merged output file does not modify lines which are identical in all of the input files. Lines which do differ across input files are prefixed with the processors and processor ranges in which they appeared.

As an example, the following text was extracted from a set of 186 CONCEPTUAL log files (from a 186-processor run):

```

# Microsecond timer type: PAPI_get_real_usec()
# Average microsecond timer overhead: <1 microsecond
#[0-4,6-12,14-16,18-52,54-78,80-94,96-101,103-121,123-140,142-169,
  171-185]# Microsecond timer increment: 1 +/- 0 microseconds
  (ideal: 1 +/- 0)
#[5]# Microsecond timer increment: 1.00229 +/- 0.15854
  microseconds (ideal: 1 +/- 0)
#[13]# Microsecond timer increment: 1.00228 +/- 0.158442
  microseconds (ideal: 1 +/- 0)
#[17,79]# Microsecond timer increment: 1.00228 +/- 0.158392
  microseconds (ideal: 1 +/- 0)
#[53]# Microsecond timer increment: 1.00228 +/- 0.158409
  microseconds (ideal: 1 +/- 0)
#[102]# Microsecond timer increment: 1.00228 +/- 0.158458
  microseconds (ideal: 1 +/- 0)
#[95,122]# Microsecond timer increment: 1.00228 +/- 0.158474
  microseconds (ideal: 1 +/- 0)
#[141]# Microsecond timer increment: 1.00228 +/- 0.158491
  microseconds (ideal: 1 +/- 0)
#[170]# Microsecond timer increment: 1.00228 +/- 0.158524
  microseconds (ideal: 1 +/- 0)

```

All of the input files contained the same `Microsecond timer type` and `Average microsecond timer overhead` lines. However, the measured `Microsecond timer increment` varied across input files. While many of the processors observed an increment of `1 +/- 0`, processor 5 was alone in observing `1.00229 +/- 0.15854`; processor 13 was

alone in observing 1.00228 ± 0.158442 ; and, both processor 17 and processor 79 observed 1.00228 ± 0.158392 as the timer increment.

'ncpt1-logmerge' can also be instructed to output only the lines which differ across files. Common lines are not output. This feature is useful for discovering misconfigured nodes in a large computer system. For example, on one computer system on which CONCEPTUAL was run, five processors were running at a higher clock rate than the remainder which naturally affected performance. 'ncpt1-logmerge' can be used to help identify such outliers.

OPTIONS

'ncpt1-logmerge' accepts the following command-line options:

- u, --usage**
Output the Synopsis section then exit the program.
- h, --help**
Output the Synopsis section and the Options section then exit the program.
- m, --man** Output a complete Unix man ("manual") page for 'ncpt1-logmerge' then exit the program.
- o filename, --output=filename**
'ncpt1-logmerge' normally writes to the standard output device. The **--output** option redirects 'ncpt1-logmerge's output to a file.
- s, --simplify**
Simplify the output by including only lines which differ across input files. No data is output, only prologue and epilogue comments. **--simplify** can be specified up to four times on the command line:
 - once** Omit all comments and all lines which are identical across all input files.
 - twice** Lines which differ across *all* output files (e.g., **Processor (0<=P<tasks)**) are also omitted.
 - three times**
The amount of output is further reduced by rounding to two significant digits all numbers appearing in all input files. Doing so makes 1.10644 ± 0.593714 match 1.12511 ± 0.58829 , for example. (Both are converted to 1.1 ± 0.59 .)
 - four times** Lists of processors are replaced by the list size. For example, **#[22,67,86,430]** becomes **#[4]**.

Note that **--simplify** is intended as a diagnostic tool; files output using **--simplify** cannot be un-merged to recover the original input files.

In addition to the preceding options 'ncpt1-logmerge' requires a list of log files to merge. If a directory is specified, all of the files immediately under that directory are used. (Note that 'ncpt1-logmerge' does not descend into subdirectories, however.) Files containing lists of filenames can be specified with a leading at sign ("@"). For example, **@filelist.txt** means to read a list of filenames from **@filelist.txt**. Filenames beginning with an at sign can be specified by doubling the at sign on the command line.

DIAGNOSTICS

filename does not look like an unmerged coNcEPTuaL log file

‘ncptl-logmerge’ accepts as input only log files produced directly by a coNcEPTuaL program. It is not a general-purpose file combiner nor does it accept its own output as input. Unrecognized input files cause ‘ncptl-logmerge’ to abort with the preceding error message.

No processor number found in *filename*

‘ncptl-logmerge’ needs to map filenames to processor numbers to indicate which processors produced which lines of output. If an input file does not contain a Processor (0<=P<tasks) comment, ‘ncptl-logmerge’ aborts with the preceding error message.

EXAMPLES

Merge a set of coNcEPTuaL log files:

```
ncptl-logmerge mybenchmark-[0-9]*.log > mybenchmark-all.log
```

The following command is equivalent to the preceding one:

```
ncptl-logmerge mybenchmark-[0-9]*.log --output=mybenchmark-all.log
```

Show only “interesting” differences among the input files:

```
ncptl-logmerge --simplify --simplify mybenchmark-[0-9]*.log
```

For convenience, one can abbreviate `--simplify --simplify --simplify --simplify` to `-s -s -s -s` or even `-ssss`:

```
ncptl-logmerge -ssss mybenchmark-[0-9]*.log
```

RESTRICTIONS

The log files passed to ‘ncptl-logmerge’ are subject to the following restrictions:

- All files must be produced by the same run of the same coNcEPTuaL program.
- None of the files can have been previously merged by ‘ncptl-logmerge’ (i.e., ‘ncptl-logmerge’ can’t read its own output).
- Only the first filename passed to ‘ncptl-logmerge’ is allowed to contain data. Data from all other files is discarded with a warning message.

BUGS

‘ncptl-logmerge’ is not a particularly robust script. Specifically, it is confused when input files contain different numbers of comment lines. For example, if one input file includes more environment variables than another or issued a warning about a timer where another input file didn’t, ‘ncptl-logmerge’ will erroneously report all subsequent lines as being mismatched across input files.

SEE ALSO

ncptl-logunmerge(1), ncptl-logextract(1), the coNcEPTuaL User’s Guide

AUTHOR

Scott Pakin, pakin@lanl.gov

3.5.4 ‘ncptl-logunmerge’

The primary capability of ‘ncptl-logmerge’ (see [Section 3.5.3 \[ncptl-logmerge\], page 70](#)) is to merge multiple CONCEPTUAL log files into a more maintainable single file. ‘ncptl-logunmerge’ performs the complementary operation of splitting that merged file back into the original set of CONCEPTUAL log files.

Running `ncptl-logunmerge --usage` causes ‘ncptl-logunmerge’ to list a synopsis of its core command-line options to the standard output device; running `ncptl-logunmerge --help` produces basic usage information; and, running `ncptl-logunmerge --man` outputs a complete manual page. See [\[ncptl-logunmerge manual page\], page 75](#), shows the ‘ncptl-logunmerge’ documentation as produced by `ncptl-logunmerge --man`.

NAME

ncptl-logunmerge - Recover individual CONCEPTUAL log files from ncptl-logmerge output

SYNOPSIS

ncptl-logunmerge *--usage* | *--help* | *--man*

ncptl-logunmerge [*--logfile=template*] [*--procs=process_list*] [*--quiet*]
[*--memcache=megabytes*] *filename*

DESCRIPTION

‘ncptl-logmerge’ merges a set of CONCEPTUAL log files into a more convenient, single file. ‘ncptl-logunmerge’ performs the inverse operation, splitting a merged file into separate CONCEPTUAL log files. Specifically, unadorned comment lines such as the following are written to all log files:

```
# Executable name: /home/me/mybenchmark
```

Comment lines which specify processor ranges are written to the appropriate log files. For example, the following line—with the leading `#[35,43,89]` stripped—is written only to the log files corresponding to processes 35, 43, and 89:

```
#[35,43,89]# Minimum sleep time: 9.68 +/- 0.556776  
microseconds (ideal: 1 +/- 0)
```

Non-comment lines (i.e., measurement data) such as the following are written only to process 0’s log file:

```
"Contention factor","Msg. size (B)","1/2 RTT (us)","MB/s"  
"(all data)","(all data)","(all data)","(all data)"  
0,1048576,1368.0295,730.978389  
0,524288,691.9115,722.6357706
```

OPTIONS

‘ncptl-logunmerge’ accepts the following command-line options:

-u, --usage

Output the Synopsis section then exit the program.

-h, --help

Output the Synopsis section and the Options section then exit the program.

-m, --man Output a complete Unix man (“manual”) page for ‘ncptl-logunmerge’ then exit the program.

-L template, --logfile=template

Specify a template for the names of the generated log files. The template must contain the literal string `%p` which will be replaced by the appropriate processor number. The template may contain the literal string `%r` (run number) which will be replaced by the smallest integer which produces a filename that does not already exist. In addition, `printf()`-style field widths can be used with `%p` and `%r`. For example, `%04p` outputs the processor number as a four-digit number padded on the left with zeroes.

If `--logfile` is not specified, ‘ncptl-logunmerge’ takes the default template from the merged log file’s `Log-file template` line, discards the directory component of the filename, and uses the result as the log-file template.

`-p process_list, --procs=process_list`

Identify a subset of log files to extract from the merged log file. By default, ‘ncptl-logunmerge’ extracts all of the constituent log files. `process_list` is a comma-separated list of process number or process ranges.

`-q, --quiet`

Suppress progress output. Normally, ‘ncptl-logunmerge’ outputs status information regarding its operation. The `--quiet` option instruct ‘ncptl-logunmerge’ to output only warning and error messages.

`-M megabytes, --memcache=megabytes`

Specify the size of the in-memory file cache. By default the program keeps up to 8 MB of extracted file data resident in memory to improve performance. The `--memcache` option enables more or less data to be cached in memory. For example, `--memcache=512` specifies that 512 MB of memory should be reserved for the file cache.

In addition to the preceding options ‘ncptl-logunmerge’ requires the name of a merged CONCEPTUAL log file. If not provided, ‘ncptl-logunmerge’ reads the contents of the merged CONCEPTUAL log file from standard input.

DIAGNOSTICS

The input file does not look like a merged coNcEPTuaL log file

Only the output from ‘ncptl-logmerge’ is valid input to ‘ncptl-logunmerge’.

Unable to find a unique number of tasks in the input file

‘ncptl-logunmerge’ determines the number of files to generate from the `Number of tasks` prologue comment. If that comment does not appear or takes on different values, ‘ncptl-logunmerge’ rejects the input file.

EXAMPLES

Extract a set of CONCEPTUAL log files from a merged log file and name the extracted files ‘happy-0.log’, ‘happy-1.log’, ‘happy-2.log’, etc.:

```
ncptl-logunmerge --logfile=happy-%p.log mybenchmark-all.log
```

Extract only ‘mybenchmark-0.log’, ‘mybenchmark-50.log’, ‘mybenchmark-51.log’, ‘mybenchmark-52.log’, and ‘mybenchmark-100.log’ from ‘mybenchmark-all.log’ (assuming that ‘mybenchmark-all.log’ contains the line `# Log-file template: mybenchmark-%p.log`):

```
ncptl-logunmerge --procs=0,50-52,100 mybenchmark-all.log
```

SEE ALSO

ncptl-logmerge(1), ncptl-logextract(1), printf(3), the CONCEPTUAL User’s Guide

AUTHOR

Scott Pakin, pakin@lanl.gov

4 Grammar

The CONCEPTUAL language was designed to produce precise specifications of network correctness and performance tests yet read like an English-language document that contains a hint of mathematical notation. Unlike more traditional programming languages, CONCEPTUAL is more descriptive than imperative. There are no classes, functions, arrays, pointers, or even variable assignments (although expressions can be let-bound to identifiers).¹ The language operates primarily on integers, with support for string constants and floating-point numbers in only a few constructs. A CONCEPTUAL program merely describes a communication pattern and the CONCEPTUAL compiler generates code to implement that pattern.

As a domain-specific language, CONCEPTUAL contains primitives to send and receive messages. It is capable of measuring time, computing statistics, and logging results. It knows that it will be run in a shared-nothing SPMD² style with explicit message-passing. As a result of its special-purpose design CONCEPTUAL can express communication patterns in a clearer and terser style than is possible using a general-purpose programming language.

The CONCEPTUAL language is case-insensitive. `Hello` is the same as `HELLO` or `hello`. Furthermore, whitespace is insignificant; one space has the same meaning as multiple spaces. Comments are designated with a `#` character and extend to the end of the line.

We now describe the CONCEPTUAL grammar in a bottom-up manner, i.e., starting from primitives and working up to complete programs. Note that many of the sections in this chapter use the following syntax to formally describe language elements:

<code><nonterminal></code>	a placeholder for a list of language primitives and additional placeholders
<code>::=</code>	“is defined as”
<code>KEYWORD</code>	a primitive with special meaning to the language
<code>[. . .]</code>	optional items
<code>(. . .)</code>	grouping of multiple items into one
<code>*</code>	zero or more occurrences of the preceding item
<code>+</code>	one or more occurrences of the preceding item
<code> </code>	either the item to the left or the item to the right but not both

4.1 Primitives

At the lowest level, CONCEPTUAL programs are composed of identifiers, strings, and integers (and a modicum of punctuation). Identifiers consist of a letter followed by zero or more alphanumerics or underscores. `'potato'`, `'x'`, and `'This_is_program_123'` are all examples of valid identifiers. Identifiers are used for two purposes: variables and keywords.

¹ CONCEPTUAL is not even Turing-complete. That is, it cannot perform arbitrary computations.

² Single Program, Multiple Data

Variables—referred to in the formal grammar as *<ident>*s—can be bound but not assigned. That is, once a variable is given a value it retains that value for the entire scope although it may be given a different value within a subordinate scope for the duration of that scope. All variables are of integer type. There are a number of variables that are predeclared and maintained automatically by CONCEPTUAL. These are listed and described in [Section A.2 \[Predeclared variables\]](#), page 184. Predeclared variables can be used by CONCEPTUAL programs but cannot be redeclared; an attempt to do so will result in a compile-time error message.

Keywords introduce actions. For example, SEND and RECEIVE are keywords. (A complete list of CONCEPTUAL keywords is presented in [Section A.1 \[Keywords\]](#), page 179.) Most keywords can appear in multiple forms. For example, OUTPUT and OUTPUTS are synonymous, as are COMPUTE and COMPUTES, A and AN, and TASK and TASKS. The intention is for programs to use whichever sounds better in an English-language sentence. Keywords may not be used as variable names; an attempt to do so will cause the compiler to output a parse error.

As a special case to increase program readability, a single ‘-’ preceding a keyword is treated as a whitespace character. Hence, ‘INTEGER-SIZED PAGE-ALIGNED MESSAGE’ is equivalent to ‘INTEGER SIZED PAGE ALIGNED MESSAGE’ and ‘10 64-BYTE MESSAGES’ is equivalent to ‘10 64 BYTE MESSAGES’. However, ‘x-3’ and ‘3-x’ still represent subtraction operations.

Although identifiers are case insensitive—‘SEND’ is the same as ‘send’ is the same as ‘sEND’—to increase clarity, this manual presents keywords in uppercase and variables in lowercase.

Strings consist of double-quoted text. Within a string—and only within a string—whitespace and case are significant. In particular, a literal newline is honored but can be suppressed by preceding it with a backslash as in the following example:

```
"This string\  
contains some  
newline characters."  
⇒ This string contains some  
newline characters.
```

Use ‘\’ for a double-quote character, ‘\’ for a backslash, and ‘\n’ for a newline character. All other escape sequences produce a warning message and are discarded. As examples of valid escape-sequence usage, the string "January 2009" represents the text “January 2009” and "I store \"stuff\" in C:\\MyStuff." represents the text “I store "stuff" in C:\\MyStuff.”

Integers consist of an optional ‘+’ or ‘-’³ followed by one or more digits followed by an optional multiplier. This multiplier is unique to CONCEPTUAL and consists of one of the following four letters:

‘K’ (kilo) multiplies the integer by 1,024
‘M’ (mega) multiplies the integer by 1,048,576

³ From the lexer’s perspective, integers are always unsigned and ‘+’ or ‘-’ are merely operators (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79). The parser, however, applies unary operators to integer literals. This alteration is evident in the output of the dot_ast backend (see [Section 3.3.10 \[The dot_ast backend\]](#), page 38).

‘G’ (giga) multiplies the integer by 1,073,741,824

‘T’ (tera) multiplies the integer by 1,099,511,627,776

In addition, a multiplier can be ‘E’ (exponent) followed by a positive integer. An ‘E’ multiplier multiplies the base integer by 10 raised to the power of the alternate integer.

Some examples of valid integers include ‘2009’, ‘-42’, ‘64K’ (= 65,536), and ‘8E3’ (= 8,000).

4.2 Expressions

Expressions, as in any language, are a combination of primitives and other expressions in a semantically meaningful juxtaposition. CONCEPTUAL provides arithmetic expressions, which evaluate to a number, and relational expressions, which evaluate to either TRUE or FALSE. In addition, CONCEPTUAL provides the notion of an *aggregate expression*, which represents a function (e.g., statistical mean) applied to every value taken on by an arithmetic expression during the run of a program.

4.2.1 Arithmetic expressions

CONCEPTUAL supports a variety of arithmetic expressions. The following is the language’s order of operations from highest to lowest precedence:

unary ‘+’, ‘-’, NOT, ‘*<function>*(*<expr>*, ...)', ‘REAL(*<expr>*)’

power ‘**’

multiplicative ‘*’, ‘/’, MOD, ‘<<’, ‘>>’, ‘&’

additive ‘+’, ‘-’, ‘|’, XOR

conditional ‘*<expr>* IF *<rel_expr>* OTHERWISE *<expr>*’

In addition, as in most programming languages, parentheses can be used to group subexpressions.

The ‘&’ (“and”), ‘|’ (“or”), XOR, and NOT operators perform bitwise, not logical, operations. That is, they accept numerical arguments, not truth-value arguments. Hence, for example, ‘3 | 5’ is equal to ‘7’.

‘<<’ and ‘>>’ are bit-shift operators. That is, ‘a << b’ is the CONCEPTUAL equivalent of the mathematical expression $a \times 2^b$ and ‘a >> b’ is the CONCEPTUAL equivalent of the mathematical expression $a \div 2^b$. Consequently, negative values of b are valid and correspond to a shift in the opposite direction. (In contrast, C and Perl treat a negative shift amount as a—usually large—unsigned number and Python raises a `ValueError` exception on negative shifts.)

MOD is a modulo (i.e., remainder) operator: ‘10 MOD 3’ returns ‘1’. ‘MOD’ is guaranteed to return a nonnegative remainder. Hence, ‘16 MOD 7’ and ‘16 MOD -7’ both return ‘2’ even though ‘-5’ is also mathematically valid. Similarly, ‘-16 MOD 7’ and ‘-16 MOD -7’ both return ‘5’ even though ‘-2’ is also mathematically valid.

The function calls allowed in ‘*<function>*(*<expr>*, ...)’ are listed and described in [Section 4.2.2 \[Built-in functions\], page 81](#). All functions take one or more arithmetic expressions

as an argument. The operator ‘*’ represents multiplication; ‘/’ represents division; and ‘**’ represents exponentiation (i.e., ‘x ** y’ $\equiv [x^y]$). Note that 0^y generates a run-time error for $y \leq 0$.

A conditional expression ‘ $\langle expr1 \rangle$ IF $\langle rel_expr \rangle$ OTHERWISE $\langle expr2 \rangle$ ’ evaluates to $\langle expr1 \rangle$ if the relational expression $\langle rel_expr \rangle$ evaluates to TRUE and $\langle expr2 \rangle$ if $\langle rel_expr \rangle$ evaluates to FALSE.⁴ Relational expressions are described in Section 4.2.5 [Relational expressions], page 91. As some examples of conditional expressions, ‘666 IF 2+2=5 OTHERWISE 777’ returns ‘777’ while ‘666 IF 2+2=4 OTHERWISE 777’ returns ‘666’.

All operations proceed left-to-right except power and conditional expressions, which proceed right-to-left. That is, ‘4-3-2’ means $(4 - 3) - 2$ but ‘4**3**2’ means $4^{(3^2)}$. Similarly, ‘2 IF p=0 OTHERWISE 1 IF p=1 OTHERWISE 0’ associates like ‘2 IF p=0 OTHERWISE (1 IF p=1 OTHERWISE 0)’, not like ‘(2 IF p=0 OTHERWISE 1) IF p=1 OTHERWISE 0’.

Evaluation contexts

CONCEPTUAL normally evaluates arithmetic expressions in “integer context”, meaning that each subexpression is truncated to the nearest integer after being evaluated. Hence, ‘24/5*5’ is ‘20’, not ‘24’, because ‘24/5*5’ = $[24 \div 5] \times 5 = 4 \times 5 = 20$. There are a few situations, however, in which CONCEPTUAL evaluates expressions in “floating-point context”, meaning that no truncation occurs:

- within an OUTPUTS statement (see Section 4.5.2 [Writing to standard output], page 104)
- within a LOGS statement (see Section 4.5.3 [Writing to a log file], page 104)
- within a BACKEND EXECUTES statement (see Section 4.6.6 [Injecting arbitrary code], page 111)
- within a range in a FOR EACH statement (see [Range loops], page 113) when CONCEPTUAL is unable to find an arithmetic or geometric progression by evaluating the component $\langle expr \rangle$ s in integer context

Within any of the preceding statements, the expression ‘24/5*5’ evaluates to 24. Furthermore, the expression ‘24/5’ evaluates to 4.8, which is a number that can’t be entered directly in a CONCEPTUAL program. (The language supports only integral constants, as mentioned in Section 4.1 [Primitives], page 77.)

The CONCEPTUAL language provides a special form called REAL which resembles a single-argument function. When evaluated in floating-point context, REAL returns its argument evaluated normally, as if ‘REAL’ were absent. When evaluated in integer context, however, REAL evaluates its argument in floating-point context and then rounds the result to the nearest integer. As an example, ‘9/2 + 1/2’ is ‘4’ in integer context because ‘9/2 + 1/2’ = $[9/2] + [1/2] = 4 + 0 = 4$. However, ‘REAL(9/2 + 1/2)’ is ‘5’ in integer context because ‘REAL(9/2 + 1/2)’ = $[9/2 + 1/2 + 0.5] = [5 + 0.5] = 5$.

Formal grammar for arithmetic expressions

For completeness, the following productions formalize the process by which CONCEPTUAL parses arithmetic expressions:

$\langle expr \rangle ::= \langle cond_expr \rangle$

⁴ It is therefore analogous to ‘ $\langle rel_expr \rangle ? \langle expr1 \rangle : \langle expr2 \rangle$ ’ in the C programming language.

```

<cond_expr> ::= <add_expr> IF <rel_expr> OTHERWISE <add_expr>
<add_expr> ::= <mult_expr>
              | <add_expr> '+' <mult_expr>
              | <add_expr> '-' <mult_expr>
              | <add_expr> '|' <mult_expr>
              | <add_expr> XOR <mult_expr>
<mult_expr> ::= <unary_expr>
              | <mult_expr> '*' <unary_expr>
              | <mult_expr> '/' <unary_expr>
              | <mult_expr> MOD <unary_expr>
              | <mult_expr> '>>' <unary_expr>
              | <mult_expr> '<<' <unary_expr>
              | <mult_expr> '&' <unary_expr>
<power_expr> ::= <primary_expr> ['**' <unary_expr>]
<unary_expr> ::= <power_expr>
              | <unary_operator> <unary_expr>
<unary_operator> ::= '+' | '-' | NOT
<primary_expr> ::= '(' <expr> ')'
              | <ident>
              | <integer>
              | <func_name> '(' <enumerated_exprs> ')'
              | REAL '(' <expr> ')'
<enumerated_exprs> ::= <expr> [',' <expr>]*

```

4.2.2 Built-in functions

In addition to the operators described in [Section 4.2.1 \[Arithmetic expressions\]](#), page 79, CONCEPTUAL contains a number of built-in functions that perform a variety of arithmetic operations that are often found to be useful in network correctness and performance testing codes. These include simple functions that map one number to another as well as a set of topology-specific functions that help implement communication across various topologies, specifically n -ary trees, meshes, tori, and k -nomial trees. CONCEPTUAL currently supports the following functions:

```

<func_name> ::= ABS | BITS | CBRT | FACTOR10 | LOG10 | MAX | MIN | ROOT | SQRT
              | CEILING | FLOOR | ROUND
              | TREE_PARENT | TREE_CHILD
              | KNOMIAL_PARENT | KNOMIAL_CHILD | KNOMIAL_CHILDREN
              | MESH_NEIGHBOR | MESH_COORDINATE
              | TORUS_NEIGHBOR | TORUS_COORDINATE
              | RANDOM_UNIFORM | RANDOM_GAUSSIAN | RANDOM_POISSON

```

All of the above take as an argument one or more integers (which may be the result of an arithmetic expression). The following sections describe each function in turn.

Integer functions

ABS returns the absolute value of its argument. For example, ‘**ABS(99)**’ and ‘**ABS(-99)**’ are both ‘99’.

BITS returns the minimum number of bits needed to store its argument. For example, ‘**BITS(12345)**’ is ‘14’ because 2^{14} is 16,384, which is larger than 12,345, while 2^{13} is 8,192, which is too small. ‘**BITS(0)**’ is defined to be ‘0’. Essentially, ‘**BITS(x)**’ represents $\lceil \log_2 x \rceil$, i.e., the ceiling of the base-2 logarithm of x . Negative numbers are treated as their two’s-complement equivalent. For example, ‘**BITS(-1)**’ returns ‘32’ on a 32-bit system and ‘64’ on a 64-bit system.

CBRT is an integer cube root function. It is essentially just syntactic sugar for the more general **ROOT** function: ‘**CBRT(x)**’ \equiv ‘**ROOT(3, x)**’.

FACTOR10 rounds its argument down (more precisely, towards zero) to the largest single-digit factor of an integral power of 10. ‘**FACTOR10(4975)**’ is therefore ‘4000’. Similarly, ‘**FACTOR10(-4975)**’ is ‘-4000’.

‘**LOG10(x)**’ is $\lfloor \log x \rfloor$, i.e., the floor of the base-10 logarithm of x . For instance, ‘**LOG10(12345)**’ is ‘4’ because 10^4 is the largest integral power of 10 that does not exceed 12,345.

MIN and **MAX** return, respectively, the minimum and maximum value in a list of numbers. Unlike the other built-in functions, **MIN** and **MAX** accept an arbitrary number of arguments (but at least one). For example, ‘**MIN(8,6,7,5,3,0,9)**’ is ‘0’ and ‘**MAX(8,6,7,5,3,0,9)**’ is ‘9’.

‘**ROOT(n, x)**’ returns $\sqrt[n]{x}$, i.e., the n th root of x . More precisely, it returns the largest integer r such that $r^n \leq x$. **ROOT** is not currently defined on negative values of x but this may change in a future release of **CONCEPTUAL**. As an example of **ROOT** usage, ‘**ROOT(5, 245)**’ is ‘3’ because $3^5 = 243 \leq 245$ but $4^5 = 1024 > 245$. Similarly, ‘**ROOT(2, 16)**’ = ‘4’; ‘**ROOT(3, 27)**’ = ‘3’; ‘**ROOT(0, 0)**’ and ‘**ROOT(4, -245)**’ each return a run-time error; and, ‘**ROOT(-3, 8)**’ = ‘0’ (because ‘**ROOT(-3, 8)**’ = ‘1/**ROOT(3, 8)**’ = ‘1/2’ = ‘0’).

SQRT is an integer square root function. It is essentially just syntactic sugar for the more general **ROOT** function: ‘**SQRT(x)**’ \equiv ‘**ROOT(2, x)**’.

Floating-point functions

As stated in [Section 4.2.1 \[Arithmetic expressions\]](#), [page 79](#), there are certain constructs in which expressions are evaluated in floating-point context instead of integer context. In such constructs, all of **CONCEPTUAL**’s built-in functions return floating-point values. Furthermore, the **CBRT**, **LOG10**, **ROOT**, and **SQRT** functions (see [\[Integer functions\]](#), [page 82](#)) compute floating-point results, not integer results which are coerced into floating-point format.

The following functions are not meaningful in integer context but are in floating-point context:

- **CEILING**
- **FLOOR**
- **ROUND**

CEILING returns the smallest integer not less than its argument. For example, ‘**CEILING(-7777/10)**’ is ‘-777’. (-778 is less than -777.7 while -777 is not less than -777.7.)

FLOOR returns the largest integer not greater than its argument. For example, ‘FLOOR(-7777/10)’ is ‘-778’. (-778 is not greater than -777.7 while -777 is greater than -777.7.)

ROUND rounds its argument to the nearest integer. For example, ‘ROUND(-7777/10)’ is ‘-778’.

It is not an error to use CEILING, FLOOR, and ROUND in an integer context; each function merely return its argument unmodified.

n-ary tree functions

n-ary trees are used quite frequently in communication patterns because they require only logarithmic time (in the number of tasks) for a message to propagate from the root to a leaf. CONCEPTUAL supports *n*-ary trees in the form of the TREE_PARENT and TREE_CHILD functions.


TREE_PARENT (*task_ID* [, *fan-out*]) [Function]

TREE_PARENT takes a task number and an optional tree fan-out (*n*) and returns the task’s parent in an *n*-ary tree. *n* defaults to ‘2’, i.e., a binary tree. Taking the TREE_PARENT of any task less than 1 returns the value ‘-1’.

TREE_CHILD (*task_ID*, *child* [, *fan-out*]) [Function]

TREE_CHILD takes a task number, a child number ($0 \leq i < N$), and an optional tree fan-out (*n*), which again defaults to ‘2’. It returns the task number corresponding to the given task’s *child*th child.

The following illustrations show how tasks are numbered in, respectively, a 2-ary and a 3-ary tree:



As shown by the 2-ary tree, task 1’s children are task 3 and task 4. Therefore, ‘TREE_PARENT(3)’ and ‘TREE_PARENT(4)’ are both ‘1’; ‘TREE_CHILD(1, 0)’ is ‘3’; and, ‘TREE_CHILD(1, 1)’ is ‘4’. In a 3-ary tree, each task has three children. Hence, the following expressions hold:

- ‘TREE_PARENT(7, 3)’ \Rightarrow ‘2’
- ‘TREE_PARENT(8, 3)’ \Rightarrow ‘2’
- ‘TREE_PARENT(9, 3)’ \Rightarrow ‘2’
- ‘TREE_CHILD(2, 0, 3)’ \Rightarrow ‘7’
- ‘TREE_CHILD(2, 1, 3)’ \Rightarrow ‘8’
- ‘TREE_CHILD(2, 2, 3)’ \Rightarrow ‘9’

***k*-nomial tree functions**

k-nomial trees are an efficient way to implement collective-communication operations in software. Unlike in an *n*-ary tree, the number of children in a *k*-nomial tree decreases with increasing task depth (i.e., no task has more children than the root). The advantage is that the tasks that start communicating earlier perform more work, which reduces the total latency of the collective operation. In contrast, in an *n*-ary tree, the tasks that start communicating earlier finish earlier, at the expense of increased total latency. CONCEPTUAL supports *k*-nomial trees via the KNOMIAL_PARENT, KNOMIAL_CHILDREN, and KNOMIAL_CHILD functions, as described below.

KNOMIAL_PARENT (*task_ID* [, *fan_out* [, *num_tasks*]]) [Function]

KNOMIAL_PARENT takes a task number, the tree fan-out factor (the “*k*” in “*k*-ary”), and the number of tasks in the tree. It returns the task ID of the given task’s parent. *fan_out* defaults to ‘2’ and the number of tasks defaults to *num_tasks* (see [Section A.2 \[Predeclared variables\]](#), page 184).


KNOMIAL_CHILDREN (*task_ID* [, *fan_out* [, *num_tasks*]]) [Function]

KNOMIAL_CHILDREN takes the same arguments as KNOMIAL_PARENT but returns the number of immediate descendents the given task has.

KNOMIAL_CHILD (*task_ID*, *child* [, *fan_out* [, *num_tasks*]]) [Function]

KNOMIAL_CHILD takes a task number, a child number ($0 \leq i < \text{‘KNOMIAL_CHILDREN(...)’}$), the tree fan-out factor, and the number of tasks in the tree. It returns the task number corresponding to the given task’s *i*th child. As in KNOMIAL_PARENT and KNOMIAL_CHILDREN, *fan_out* defaults to ‘2’ and the number of tasks defaults to *num_tasks* (see [Section A.2 \[Predeclared variables\]](#), page 184).


The following figure shows how CONCEPTUAL numbers tasks in a *k*-nomial tree with *k* = 2 (a.k.a. a 2-nomial or binomial tree).



The figure is structured with time flowing downwards. That is, for a multicast operation expressed over a 2-nomial tree, task 0 sends a message to task 1 in the first time step. Then, task 0 sends to task 2 while task 1 sends to task 3. In the final step, task 0 sends to task 4, task 1 sends to task 5, task 2 sends to task 6, and task 3 sends to task 7—all concurrently. The following expressions also hold, assuming there are a total of eight tasks in the computation:

- ‘KNOMIAL_PARENT(0)’ \Rightarrow ‘-1’
- ‘KNOMIAL_PARENT(1)’ \Rightarrow ‘0’
- ‘KNOMIAL_CHILDREN(1)’ \Rightarrow ‘2’
- ‘KNOMIAL_CHILD(1, 0)’ \Rightarrow ‘3’
- ‘KNOMIAL_CHILD(1, 1)’ \Rightarrow ‘5’
- ‘KNOMIAL_CHILDREN(7)’ \Rightarrow ‘0’
- ‘KNOMIAL_CHILD(7, 0)’ \Rightarrow ‘-1’

k -nomial trees for $k > 2$ are much less common in practice than 2-nomial trees. However, they may perform well when a task has sufficient bandwidth to support multiple, simultaneous, outgoing messages. For example, a trinomial tree (i.e., a k -nomial tree with $k = 3$) should exhibit good performance if there is enough bandwidth to send two messages simultaneously. The following illustration shows how CONCEPTUAL constructs a 27-task trinomial tree:



As before, time flows downward (assuming a multicast operation) and tasks are expected to communicate with their children in order. The following are some CONCEPTUAL k -nomial tree expressions and their evaluations, assuming `num_tasks` is ‘27’:

- ‘KNOMIAL_PARENT(0, 3)’ \Rightarrow ‘-1’
- ‘KNOMIAL_PARENT(2, 3)’ \Rightarrow ‘0’
- ‘KNOMIAL_CHILDREN(2, 3)’ \Rightarrow ‘4’
- ‘KNOMIAL_CHILD(2, 0, 3)’ \Rightarrow ‘5’
- ‘KNOMIAL_CHILD(2, 1, 3)’ \Rightarrow ‘8’
- ‘KNOMIAL_CHILD(2, 2, 3)’ \Rightarrow ‘11’
- ‘KNOMIAL_CHILD(2, 3, 3)’ \Rightarrow ‘20’
- ‘KNOMIAL_CHILD(2, 4, 3)’ \Rightarrow ‘-1’
- ‘KNOMIAL_CHILDREN(8, 3)’ \Rightarrow ‘2’
- ‘KNOMIAL_CHILDREN(8, 3, 26)’ \Rightarrow ‘1’
- ‘KNOMIAL_CHILDREN(8, 3, 10)’ \Rightarrow ‘0’

Mesh functions

CONCEPTUAL provides two functions, `MESH_NEIGHBOR` and `MESH_COORDINATE`, that help treat (linear) task IDs as positions on a multidimensional mesh. Each of these functions takes a variable number of arguments, determined by the dimensionality of the mesh (1-D, 2-D, or 3-D).

`MESH_NEIGHBOR` (*task_ID*, *width*, *x_offset* [, *height*, *y_offset* [, *depth*, *z_offset*]]) [Function]

`MESH_NEIGHBOR` returns a task's neighbor on a 1-D, 2-D, or 3-D mesh. It always takes a task number, the mesh's width, and the desired *x* offset from the given task. For a 2-D or 3-D mesh, the next two arguments are the mesh's height and the desired *y* offset from the given task. For a 3-D mesh only, the next two arguments are the mesh's depth and the desired *z* offset from the given task. Offsets that move off the mesh cause `MESH_NEIGHBOR` to return the value '-1'.

`MESH_COORDINATE` (*task_ID*, *coordinate*, *width* [, *height* [, *depth*]]) [Function]


`MESH_COORDINATE` returns a task's x, y, or z coordinate on a 1-D, 2-D, or 3-D mesh. The first argument to `MESH_COORDINATE` is a task number. The second argument should be '0' to calculate an x coordinate, '1' to calculate a y coordinate, or '2' to calculate a z coordinate. The remaining arguments are the mesh's width, height, and depth, respectively. The height can be omitted for a 1-D mesh and the depth can be omitted for a 2-D mesh. (Both default to '1'.)

`MESH_NEIGHBOR` and `MESH_COORDINATE` number tasks following the right-hand rule: left-to-right, then top-to-bottom, and finally back-to-front, as shown in the following illustrations of a 4-element (1-D) mesh, a 4×3 (2-D) mesh, and a $4 \times 3 \times 2$ (3-D) mesh. Examples of `MESH_NEIGHBOR` and `MESH_COORDINATE` for 1-D, 2-D, and 3-D meshes follow the corresponding illustration.


0 ————— 1 ————— 2 ————— 3

- 'MESH_NEIGHBOR(0, 4, -1)' ⇒ '-1'
 - 'MESH_NEIGHBOR(0, 4, +1)' ⇒ '1'
 - 'MESH_NEIGHBOR(1, 4, -1)' ⇒ '0'
 - 'MESH_NEIGHBOR(1, 4, +1)' ⇒ '2'
 - 'MESH_NEIGHBOR(2, 4, -1)' ⇒ '1'
 - 'MESH_NEIGHBOR(2, 4, +1)' ⇒ '3'
 - 'MESH_NEIGHBOR(3, 4, -1)' ⇒ '2'
 - 'MESH_NEIGHBOR(3, 4, +1)' ⇒ '-1'
-
- 'MESH_COORDINATE(-1, 0, 4)' ⇒ '-1'
 - 'MESH_COORDINATE(0, 0, 4)' ⇒ '0'
 - 'MESH_COORDINATE(1, 0, 4)' ⇒ '1'
 - 'MESH_COORDINATE(2, 0, 4)' ⇒ '2'
 - 'MESH_COORDINATE(3, 0, 4)' ⇒ '3'

- `'MESH_COORDINATE(4, 0, 4)'` \Rightarrow `'-1'`
- `'MESH_COORDINATE(2, 1, 4)'` \Rightarrow `'0'`
- `'MESH_COORDINATE(2, 2, 4)'` \Rightarrow `'0'`



- `'MESH_NEIGHBOR(5, 4, -1, 3, -1)'` \Rightarrow `'0'`
 - `'MESH_NEIGHBOR(5, 4, 0, 3, -1)'` \Rightarrow `'1'`
 - `'MESH_NEIGHBOR(5, 4, +1, 3, -1)'` \Rightarrow `'2'`
 - `'MESH_NEIGHBOR(5, 4, -1, 3, 0)'` \Rightarrow `'4'`
 - `'MESH_NEIGHBOR(5, 4, 0, 3, 0)'` \Rightarrow `'5'`
 - `'MESH_NEIGHBOR(5, 4, +1, 3, 0)'` \Rightarrow `'6'`
 - `'MESH_NEIGHBOR(5, 4, -1, 3, +1)'` \Rightarrow `'8'`
 - `'MESH_NEIGHBOR(5, 4, 0, 3, +1)'` \Rightarrow `'9'`
 - `'MESH_NEIGHBOR(5, 4, +1, 3, +1)'` \Rightarrow `'10'`
-
- `'MESH_COORDINATE(1, 0, 4, 3)'` \Rightarrow `'1'`
 - `'MESH_COORDINATE(6, 0, 4, 3)'` \Rightarrow `'2'`
 - `'MESH_COORDINATE(6, 1, 4, 3)'` \Rightarrow `'1'`
 - `'MESH_COORDINATE(6, 2, 4, 3)'` \Rightarrow `'0'`
 - `'MESH_COORDINATE(8, 0, 4, 3)'` \Rightarrow `'0'`
 - `'MESH_COORDINATE(8, 1, 4, 3)'` \Rightarrow `'2'`
 - `'MESH_COORDINATE(12, 0, 4, 3)'` \Rightarrow `'-1'`



- ‘MESH_NEIGHBOR(0, 4, 0, 3, 0, 2, +1)’ ⇒ ‘12’
- ‘MESH_NEIGHBOR(0, 4, 0, 3, +1, 2, 0)’ ⇒ ‘4’
- ‘MESH_NEIGHBOR(0, 4, +1, 3, 0, 2, 0)’ ⇒ ‘1’
- ‘MESH_NEIGHBOR(0, 4, +1, 3, +1, 2, +1)’ ⇒ ‘17’
- ‘MESH_NEIGHBOR(17, 4, +2, 3, -1, 2, -1)’ ⇒ ‘3’
- ‘MESH_NEIGHBOR(23, 4, +1, 3, +1, 2, +1)’ ⇒ ‘-1’

- ‘MESH_COORDINATE(-5, 0, 4, 3, 2)’ ⇒ ‘-1’
- ‘MESH_COORDINATE(1, 0, 4, 3, 2)’ ⇒ ‘1’
- ‘MESH_COORDINATE(6, 0, 4, 3, 2)’ ⇒ ‘2’
- ‘MESH_COORDINATE(6, 1, 4, 3, 2)’ ⇒ ‘1’
- ‘MESH_COORDINATE(6, 2, 4, 3, 2)’ ⇒ ‘0’
- ‘MESH_COORDINATE(18, 0, 4, 3, 2)’ ⇒ ‘2’
- ‘MESH_COORDINATE(18, 1, 4, 3, 2)’ ⇒ ‘1’
- ‘MESH_COORDINATE(18, 2, 4, 3, 2)’ ⇒ ‘1’
- ‘MESH_COORDINATE(18, 3, 4, 3, 2)’ ⇒ error Invalid coordinate

Torus functions

A torus is a mesh with wraparound edges. CONCEPTUAL provides analogues to MESH_NEIGHBOR and MESH_COORDINATE called TORUS_NEIGHBOR and TORUS_COORDINATE which enable (linear) task IDs to be treated as positions on a 1-D, 2-D, or 3-D torus.

TORUS_NEIGHBOR (*task_ID*, *width*, *x_offset* [, *height*, *y_offset* [Function] [, *depth*, *z_offset*]])

TORUS_NEIGHBOR takes the same arguments as MESH_NEIGHBOR but calculates neighbors on a torus instead of a mesh. See [Mesh functions], page 86, for explanations of the arguments. While task offsets that go out-of-bounds on a mesh return ‘-1’, such offsets merely wrap around a torus.

TORUS_COORDINATE (*task_ID*, *coordinate*, *width* [, *height* [, *depth*]]) [Function]

TORUS_COORDINATE returns a task's x, y, or z coordinate on a 1-D, 2-D, or 3-D torus. It performs exactly the same function as MESH_COORDINATE and is aliased in the language merely for symmetry. See [Mesh functions], page 86, for details.

- 'TORUS_NEIGHBOR(0, 4, -1)' ⇒ '3'
- 'TORUS_NEIGHBOR(0, 4, +1)' ⇒ '1'
- 'TORUS_NEIGHBOR(1, 4, -1)' ⇒ '0'
- 'TORUS_NEIGHBOR(1, 4, +1)' ⇒ '2'
- 'TORUS_NEIGHBOR(2, 4, -1)' ⇒ '1'
- 'TORUS_NEIGHBOR(2, 4, +1)' ⇒ '3'
- 'TORUS_NEIGHBOR(3, 4, -1)' ⇒ '2'
- 'TORUS_NEIGHBOR(3, 4, +1)' ⇒ '0'

- 'TORUS_COORDINATE(-1, 0, 4)' ⇒ '-1'
- 'TORUS_COORDINATE(0, 0, 4)' ⇒ '0'
- 'TORUS_COORDINATE(1, 0, 4)' ⇒ '1'
- 'TORUS_COORDINATE(2, 0, 4)' ⇒ '2'
- 'TORUS_COORDINATE(3, 0, 4)' ⇒ '3'
- 'TORUS_COORDINATE(4, 0, 4)' ⇒ '-1'
- 'TORUS_COORDINATE(2, 1, 4)' ⇒ '0'
- 'TORUS_COORDINATE(2, 2, 4)' ⇒ '0'

- 'TORUS_NEIGHBOR(0, 4, -1, 3, -1)' ⇒ '11'
- 'TORUS_NEIGHBOR(0, 4, 0, 3, -1)' ⇒ '8'
- 'TORUS_NEIGHBOR(0, 4, +1, 3, -1)' ⇒ '9'
- 'TORUS_NEIGHBOR(0, 4, -1, 3, 0)' ⇒ '3'
- 'TORUS_NEIGHBOR(0, 4, 0, 3, 0)' ⇒ '0'
- 'TORUS_NEIGHBOR(0, 4, +1, 3, 0)' ⇒ '1'
- 'TORUS_NEIGHBOR(0, 4, -1, 3, +1)' ⇒ '7'
- 'TORUS_NEIGHBOR(0, 4, 0, 3, +1)' ⇒ '4'
- 'TORUS_NEIGHBOR(0, 4, +1, 3, +1)' ⇒ '5'

- 'TORUS_NEIGHBOR(23, 4, +1, 3, +1, 2, +1)' ⇒ '0'
- 'TORUS_NEIGHBOR(23, 4, +2, 3, +2, 2, +2)' ⇒ '17'
- 'TORUS_NEIGHBOR(23, 4, +3, 3, +3, 2, +3)' ⇒ '10'

- 'TORUS_COORDINATE(1, 0, 4, 3)' ⇒ '1'
- 'TORUS_COORDINATE(6, 0, 4, 3)' ⇒ '2'
- 'TORUS_COORDINATE(6, 1, 4, 3)' ⇒ '1'
- 'TORUS_COORDINATE(6, 2, 4, 3)' ⇒ '0'

- ‘TORUS_COORDINATE(8, 0, 4, 3)’ ⇒ ‘0’
- ‘TORUS_COORDINATE(8, 1, 4, 3)’ ⇒ ‘2’
- ‘TORUS_COORDINATE(12, 0, 4, 3)’ ⇒ ‘-1’

- ‘TORUS_COORDINATE(-5, 0, 4, 3, 2)’ ⇒ ‘-1’
- ‘TORUS_COORDINATE(1, 0, 4, 3, 2)’ ⇒ ‘1’
- ‘TORUS_COORDINATE(6, 0, 4, 3, 2)’ ⇒ ‘2’
- ‘TORUS_COORDINATE(6, 1, 4, 3, 2)’ ⇒ ‘1’
- ‘TORUS_COORDINATE(6, 2, 4, 3, 2)’ ⇒ ‘0’
- ‘TORUS_COORDINATE(18, 0, 4, 3, 2)’ ⇒ ‘2’
- ‘TORUS_COORDINATE(18, 1, 4, 3, 2)’ ⇒ ‘1’
- ‘TORUS_COORDINATE(18, 2, 4, 3, 2)’ ⇒ ‘1’
- ‘TORUS_COORDINATE(18, 3, 4, 3, 2)’ ⇒ error Invalid coordinate

Random-number functions

CONCEPTUAL programs can utilize randomness in one of two ways. The functions described below are *unsynchronized* across tasks. That is, they can—and usually do—return a different value to each task on each invocation. One consequence is that these functions are not permitted within a task expression (see [Section 4.3 \[Task descriptions\]](#), page 92) because randomness would cause the tasks to disagree about who the sources and targets of an operation are. In contrast, the A RANDOM TASK construct described in [Section 4.7.3 \[Binding variables\]](#), page 117 returns a value guaranteed to be synchronized across tasks and thereby enables random-task selection.

RANDOM_UNIFORM (*lower_bound*, *upper_bound*) [Function]

Return a number selected at random from a uniform distribution over the range [*lower_bound*, *upper_bound*).

RANDOM_GAUSSIAN (*mean*, *stddev*) [Function]

Return a number selected at random from a Gaussian distribution with mean *mean* and standard deviation *stddev*.

RANDOM_POISSON (*mean*) [Function]

Return an integer selected at random from a Poisson distribution with mean *mean* and standard deviation $\sqrt{\textit{mean}}$.

4.2.3 Aggregate expressions

Aggregate expressions (*<aggr-expr>*s) are currently used exclusively by the LOGS statement. They represent an expression with a given function applied to the aggregate of all (dynamic) instances of that expression. *<aggr-expr>*s take one of four forms:

```

<aggr-expr> ::= [EACH] <expr>
              | THE <expr>
              | THE <aggr-func> [OF [THE]] <expr>
              | A HISTOGRAM OF [THE] <expr>

```


(In the above, $\langle expr \rangle$ refers to an arithmetic expression defined in [Section 4.2.1 \[Arithmetic expressions\]](#), page 79 and $\langle aggr_func \rangle$ refers to one of the functions defined in [Section 4.2.4 \[Aggregate functions\]](#), page 91.)

The first form does not summarize $\langle expr \rangle$; every individual instance of $\langle expr \rangle$ is utilized. The second form asserts that $\langle expr \rangle$ is a constant (i.e., all values are identical) and utilizes that constant.⁵ The third form applies $\langle aggr_func \rangle$ to the set of all values of $\langle expr \rangle$ and utilizes the result of that function. The fourth form produces a histogram of all values of $\langle expr \rangle$, i.e., a list of $\{unique\ value, tally\}$ pairs, sorted by *unique value*.

4.2.4 Aggregate functions

The following functions, referred to collectively as $\langle aggr_func \rangle$ s, may be used in an aggregate expression (see [Section 4.2.3 \[Aggregate expressions\]](#), page 90):

```
 $\langle aggr\_func \rangle ::=$  [ARITHMETIC] MEAN | HARMONIC MEAN | GEOMETRIC MEAN | MEDIAN |
STANDARD DEVIATION | VARIANCE | SUM | MINIMUM | MAXIMUM | FINAL
```

MEAN and ARITHMETIC MEAN are equivalent. MEDIAN is the value such that there are as many larger as smaller values. If there are an even number of values, MEDIAN is the arithmetic mean of the two medians. FINAL returns only the final value measured. The interpretation of the remaining functions should be unambiguous.

4.2.5 Relational expressions

Relational expressions ($\langle rel_expr \rangle$ s) compare two arithmetic expressions (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79) or test an arithmetic expression for a property. A relational expression can be either TRUE or FALSE.

CONCEPTUAL supports a variety of relational expressions. The following is the language's order of operations from highest to lowest precedence:

```
unary/      IS EVEN, IS ODD
binary/     '=', '<', '>', '<=', '>=', '<>', DIVIDES
ternary     IS IN, IS NOT IN

conjunctive '&'
disjunctive '\'
```

In addition, as in most programming languages, parentheses can be used to group subexpressions.

The unary relation IS EVEN is TRUE if a given arithmetic expression represents an even number and the unary relation IS ODD is TRUE if a given arithmetic expression represents an odd number. For example, '456 IS EVEN' is TRUE and '64 MOD 6 IS ODD' is FALSE.

The CONCEPTUAL operators '=', '<', '>', '<=', '>=', and '<>' represent, respectively, the mathematical relations =, <, >, ≤, ≥, and ≠. These are all binary relations that operate on arithmetic expressions (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79). For example, '2+2 = 4' is TRUE and '2**3 > 2**4' is FALSE. The DIVIDES relation is TRUE if the first expression evenly divides the second, i.e., that $e_2 \equiv 0 \pmod{e_1}$. Hence, '2 DIVIDES 1234' (equivalent to '1234 MOD 2 = 0') is TRUE while '2 DIVIDES 4321' (equivalent to '4321 MOD 2 = 0') is FALSE.

⁵ The program aborts with a run-time error if $\langle expr \rangle$ is not a constant.

The ternary relation IS IN has the form ‘ $\langle expr \rangle$ IS IN [$\langle expr \rangle$, $\langle expr \rangle$]’. The relational expression ‘ x IS IN [a , b]’ is TRUE if x lies within the closed interval [a , b] and FALSE otherwise. The interval bounds a and b can be specified in any order. Hence, ‘ x IS IN [a , b]’ can be more precisely described as “ $(a \leq x \leq b) \vee (b \leq x \leq a)$ ”. To provide a few examples, ‘4 IS IN [3,5]’ and ‘4 IS IN [5,3]’ are both TRUE; ‘5 IS IN [3,5]’ is TRUE; however, ‘6 IS IN [3,5]’ is FALSE. The complementary operation to IS IN is the ternary relation IS NOT IN. Hence, ‘4 IS NOT IN [3,5]’ is FALSE while ‘6 IS NOT IN [3,5]’ is TRUE.

Conjunction (\wedge) and disjunction (\vee) combine multiple relational expressions. $\langle rel_expr \rangle$ ‘ \wedge ’ $\langle rel_expr \rangle$ is TRUE if and only if both $\langle rel_expr \rangle$ s are TRUE, and $\langle rel_expr \rangle$ ‘ \vee ’ $\langle rel_expr \rangle$ is TRUE if and only if either $\langle rel_expr \rangle$ is TRUE. For example, ‘456 IS EVEN \vee 2**3 > 2**4’ is TRUE and ‘456 IS EVEN \wedge 2**3 > 2**4’ is FALSE. Conjunction and disjunction are both short-circuiting operations. Evaluation proceeds left-to-right. Expressions such as ‘ $x <> 0 \wedge 1/x = 1$ ’ will therefore not result in a divide-by-zero error.

CONCEPTUAL does not currently have a logical negation operator.

Formal grammar for relational expressions

For completeness, the following productions formalize the process by which CONCEPTUAL parses relational expressions:

```

 $\langle rel\_expr \rangle$  ::=  $\langle rel\_disj\_expr \rangle$ 
 $\langle rel\_disj\_expr \rangle$  ::= [ $\langle rel\_disj\_expr \rangle$  ‘ $\vee$ ’]  $\langle rel\_conj\_expr \rangle$ 
 $\langle rel\_conj\_expr \rangle$  ::= [ $\langle rel\_conj\_expr \rangle$  ‘ $\wedge$ ’]  $\langle rel\_primary\_expr \rangle$ 
 $\langle rel\_primary\_expr \rangle$  ::=  $\langle eq\_expr \rangle$ 
                       | ‘(’  $\langle rel\_expr \rangle$  ‘)’
 $\langle eq\_expr \rangle$  ::=  $\langle expr \rangle$  ‘=’  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  ‘<’  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  ‘>’  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  ‘<=’  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  ‘>=’  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  ‘<>’  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  DIVIDES  $\langle expr \rangle$ 
                |  $\langle expr \rangle$  IS EVEN
                |  $\langle expr \rangle$  IS ODD
                |  $\langle expr \rangle$  IS IN ‘[’  $\langle expr \rangle$  ‘,’  $\langle expr \rangle$  ‘]’
                |  $\langle expr \rangle$  IS NOT IN ‘[’  $\langle expr \rangle$  ‘,’  $\langle expr \rangle$  ‘]’

```

4.3 Task descriptions

Task descriptions are a powerful way of tersely describing the sources and targets of CONCEPTUAL operations. Task IDs range from 0 to ‘num_tasks-1’ (see [Section A.2 \[Pre-declared variables\]](#), page 184). Operations involving out-of-bound task IDs are silently ignored.

As a side effect, a task description can declare a variable that can be used in subsequent expressions. (See [Section 4.2 \[Expressions\]](#), page 79.) There are two types of task descriptions: one for “source” tasks and one for “target” tasks. The two are syntactically similar but semantically different. Specifically, the scope of a variable declared in a $\langle target_tasks \rangle$ specification is more limited than one declared in a $\langle source_task \rangle$ specification.

Before introducing $\langle source_task \rangle$ and $\langle target_tasks \rangle$ specifications we first introduce the notion of a $\langle restricted_ident \rangle$, which is a variable declaration that can be used to define a set of tasks. We then present CONCEPTUAL’s complete set of mechanisms for describing sets of source and target tasks.

4.3.1 Restricted identifiers

A *restricted identifier* declares a variable, restricting it to the set of tasks that satisfy a given relational expression (see Section 4.2.5 [Relational expressions], page 91). The syntax, shown below, represents the mathematical notion of “ $\forall \langle ident \rangle ((\langle rel_expr \rangle \wedge (0 \leq \langle ident \rangle < \langle \#tasks \rangle))$ ”). That is, “for all $\langle ident \rangle$ such that $\langle rel_expr \rangle$ is TRUE and $\langle ident \rangle$ is between zero and the number of tasks. . .”.

$\langle restricted_ident \rangle ::= \langle ident \rangle \text{ SUCH THAT } \langle rel_expr \rangle$

As an example, ‘`evno SUCH THAT evno IS EVEN`’ describes all even-numbered tasks. On each such task, the variable ‘`evno`’ takes on that task’s ID. Similarly, ‘`thr SUCH THAT 3 DIVIDES thr-1`’ describes tasks 1, 4, 7, 10, 13. . . . On each of those tasks, ‘`thr`’ will be bound to the task ID. On all other tasks, ‘`thr`’ will be undefined. When order matters (as in the cases described in Section 4.4.2 [Sending], page 98 and Section 4.6.5 [Reordering task IDs], page 110), $\langle ident \rangle$ takes on task IDs in increasing order.

4.3.2 Source tasks

A $\langle source_task \rangle$ specification takes one of four forms:

$\langle source_task \rangle ::=$ ALL TASKS
 | ALL TASKS $\langle ident \rangle$
 | TASK $\langle expr \rangle$
 | TASKS $\langle restricted_ident \rangle$

ALL TASKS specifies that each task will perform a given operation. If followed by a variable name ($\langle ident \rangle$), each task will individually bind $\langle ident \rangle$ to its task ID—a number from zero to one less than the total number of tasks. That is, ‘ALL TASKS `me`’ will bind ‘`me`’ to ‘0’ on task 0, ‘1’ on task 1, and so forth.

‘TASK $\langle expr \rangle$ ’ specifies that only the task described by arithmetic expression $\langle expr \rangle$ will perform the given operation. For example, ‘TASK `2*3+1`’ says that only task 7 will act; the other tasks will do nothing.

‘TASKS $\langle restricted_ident \rangle$ ’ describes a set of tasks that will perform a given operation. For instance, ‘TASKS `x SUCH THAT x>0 /\ x<num_tasks-1`’—read as “tasks x such that x is greater than zero and x is less than num_tasks minus one”—expresses that a given operation should be performed on all tasks except the first and last in the computation. On each task that satisfies the relational expression, ‘`x`’ will be bound to the task ID as in ALL TASKS above. Hence, ‘`x`’ will be undefined on task 0, ‘1’ on task 1, ‘2’ on task 2, and so forth up to task ‘`num_tasks-1`’, on which ‘`x`’ will again be undefined.

As per the definitions in Section 4.1 [Primitives], page 77 and Section 4.3.1 [Restricted identifiers], page 93, respectively, $\langle ident \rangle$ s and $\langle restricted_ident \rangle$ s do not accept parentheses. Hence, ‘TASKS (bad SUCH THAT bad IS EVEN)’ and ‘ALL TASKS (no_good)’ result in parse errors while ‘TASKS fine SUCH THAT fine IS EVEN’ and ‘ALL TASKS dandy’ are acceptable constructs. As an analogy, ‘`x = 3`’ is valid in many general-purpose programming languages while ‘`(x) = 3`’ is not.

Variables declared in a ‘`source_task`’ specification are limited in scope to the surrounding statement.

4.3.3 Target tasks

A $\langle target_tasks \rangle$ specification takes one of three forms:

```

 $\langle target\_tasks \rangle$  ::= ALL OTHER TASKS
                    |   TASK  $\langle expr \rangle$ 
                    |   TASKS  $\langle restricted\_ident \rangle$ 

```

ALL OTHER TASKS is just like ALL TASKS in a $\langle source_task \rangle$ specification (see [Section 4.3.2 \[Source tasks\], page 93](#)) but applies to all tasks *except* the source task. Also, unlike ALL TASKS, ALL OTHER TASKS does not accept an $\langle ident \rangle$ term.

‘TASK $\langle expr \rangle$ ’ specifies that only the task described by arithmetic expression $\langle expr \rangle$ is the target of the given operation. $\langle expr \rangle$ can use a variable declared as a $\langle source_task \rangle$. For example, if $\langle source_task \rangle$ is ‘ALL TASKS x’, then a $\langle target_tasks \rangle$ of ‘TASK (x+1) MOD num_tasks’ refers to each task’s right neighbor (with wraparound from num_tasks to ‘0’).

‘TASKS $\langle restricted_ident \rangle$ ’ describes a set of tasks that will perform a given operation. As with its ‘source_task’ counterpart, a $\langle restricted_ident \rangle$ declares a variable. However, in a $\langle target_tasks \rangle$ specification the variable’s scope is limited to the relational expression within the $\langle restricted_ident \rangle$. As an example, ‘TASKS dst SUCH THAT dst>src’ refers to all tasks ‘dst’ with a greater ID than a (previously declared) task ‘src’.

4.4 Communication statements

Communication statements are the core of any CONCEPTUAL program. The CONCEPTUAL language makes it easy to express a variety of communication features:

- synchronous or asynchronous communication
- unaligned, aligned (to arbitrary byte boundaries), or misaligned (from a page boundary) message buffers
- ignored, touched, or verified message contents
- unique or recycled message buffers
- point-to-point or collective operations

Communication statements are performed by an arbitrary $\langle source_task \rangle$ (see [Section 4.3.2 \[Source tasks\], page 93](#)) and may involve arbitrary $\langle target_tasks \rangle$ (see [Section 4.3.3 \[Target tasks\], page 94](#)). After explaining how to describe a message to CONCEPTUAL (see [Section 4.4.1 \[Message specifications\], page 94](#)) this section presents each communication statement in turn and explains its purpose, syntax, and semantics.

4.4.1 Message specifications

A *message specification* describes a set of messages. The following is a formal definition:

```

 $\langle message\_spec \rangle$  ::=  $\langle item\_count \rangle$ 
                    [NONUNIQUE | UNIQUE]
                     $\langle item\_size \rangle$ 
                    [UNALIGNED |
                     $\langle message\_alignment \rangle$  ALIGNED |

```

```

    <message_alignment> MISALIGNED]
MESSAGES
[WITH VERIFICATION | WITH DATA TOUCHING |
WITHOUT VERIFICATION | WITHOUT DATA TOUCHING]
[FROM BUFFER <expr> |
FROM THE DEFAULT BUFFER]

```

Within a RECEIVE statement (see [Section 4.4.3 \[Receiving\]](#), page 100), a *<message_spec>*'s FROM keyword must be replaced with INTO.

A SEND statement's WHO RECEIVES IT clause (see [Section 4.4.2 \[Sending\]](#), page 98) utilizes a slightly different message specification, which is referred to here as a *<recv_message_spec>*:

```

<recv_message_spec> ::= [SYNCHRONOUSLY | ASYNCHRONOUSLY]
[AS [A | AN]
[NONUNIQUE | UNIQUE]
[UNALIGNED |
<message_alignment> ALIGNED |
<message_alignment> MISALIGNED]
MESSAGES]
[WITH VERIFICATION | WITH DATA TOUCHING |
WITHOUT VERIFICATION | WITHOUT DATA TOUCHING]
[FROM BUFFER <expr> |
FROM THE DEFAULT BUFFER]

```

Although not indicated by the preceding grammatical rule, a *<recv_message_spec>* is not allowed to be empty. That is, at least one of the optional clauses must be specified.

A REDUCE statement (see [Section 4.4.6 \[Reducing\]](#), page 102) utilizes the following variations of *<message_spec>* and *<recv_message_spec>*, respectively:

```

<reduce_message_spec> ::= <item_count>
[NONUNIQUE | UNIQUE]
[UNALIGNED |
<message_alignment> ALIGNED |
<message_alignment> MISALIGNED]
INTEGERS | DOUBLEWORDS
[WITH DATA TOUCHING | WITHOUT DATA TOUCHING]
[FROM BUFFER <expr> |
FROM THE DEFAULT BUFFER]

<reduce_target_message_spec> ::= [AS <item_count>
[NONUNIQUE | UNIQUE]
[<message_alignment> ALIGNED |
<message_alignment> MISALIGNED]
INTEGERS | DOUBLEWORDS]
[WITH DATA TOUCHING | WITHOUT DATA TOUCHING]
[INTO BUFFER <expr> |
INTO THE DEFAULT BUFFER]

```

We now describe in turn each component of a *<message_spec>*, *<recv_message_spec>*, *<reduce_message_spec>*, and *<reduce_target_message_spec>*.

Item count

The $\langle item_count \rangle$ says how many messages the $\langle message_spec \rangle$ represents:

$\langle item_count \rangle ::= A \mid AN \mid \langle expr \rangle$

A and AN are synonyms for the value ‘1’.

Unique messages

Normally, the CONCEPTUAL backends recycle message memory to reduce the program’s memory requirements and improve performance. By adding the keyword UNIQUE, every message buffer will reside in a unique memory region. NONUNIQUE explicitly specifies the default, buffer-recycling behavior.

Item size

The message size is represented by the $\langle item_size \rangle$ nonterminal. It can be empty or expressed in one of two other ways:

$\langle item_size \rangle ::= \varepsilon$
 $\quad \mid \langle expr \rangle \langle data_multiplier \rangle$
 $\quad \mid \langle data_type \rangle \text{ SIZED}$

A $\langle data_multiplier \rangle$ is a scaling factor that converts a unitless number into a number of bytes. The following are the valid possibilities for $\langle data_multiplier \rangle$ and the number of bytes which which they multiply $\langle expr \rangle$:

$\langle data_multiplier \rangle ::= \text{ BIT } \mid \text{ BYTE } \mid \text{ HALFWORD } \mid \text{ WORD } \mid \text{ INTEGER } \mid \text{ DOUBLEWORD } \mid$
 $\quad \text{ QUADWORD } \mid \text{ PAGE } \mid \text{ KILOBYTE } \mid \text{ MEGABYTE } \mid \text{ GIGABYTE}$

BIT 1/8 bytes, rounded up to the nearest integral number of bytes

BYTE 1 byte

HALFWORD 2 bytes

WORD 4 bytes

INTEGER the number of bytes in the backend’s fundamental integer type

DOUBLEWORD
8 bytes

QUADWORD 16 bytes

PAGE the number of bytes in an operating-system page

KILOBYTE 1,024 bytes

MEGABYTE 1,048,576 bytes

GIGABYTE 1,073,741,824 bytes

A $\langle data_type \rangle$ is an “atomic” unit of data. It can be any of the following:

$\langle data_type \rangle ::= \text{ BYTE } \mid \text{ HALFWORD } \mid \text{ WORD } \mid \text{ INTEGER } \mid \text{ DOUBLEWORD } \mid$
 $\quad \text{ QUADWORD } \mid \text{ PAGE}$

BYTE 1 byte

HALFWORD	2 bytes
WORD	4 bytes
INTEGER	the number of bytes in the backend’s fundamental integer type
DOUBLEWORD	8 bytes
QUADWORD	16 bytes
PAGE	the number of bytes in an operating-system page

Hence, valid $\langle item_size \rangle$ s include, for example, ‘16 MEGABYTE’ or ‘PAGE SIZED’. Empty $\langle item_size \rangle$ s (shown in the grammar as ε) are equivalent to ‘0 BYTE’. Note that `INTEGER` varies in size based on the backend, backend compiler, and CPU architecture but is commonly either 4 or 8 bytes; `PAGE` varies in size from operating system to operating system; each of the other $\langle data_type \rangle$ s has a fixed size, as indicated above.

Message alignment

Messages are normally allocated with arbitrary alignment in memory. However, `CONCEPTUAL` can force a specific alignment relative to the operating-system page size (commonly 4 KB or 8 KB, but significantly larger sizes are gaining popularity). A $\langle message_alignment \rangle$ is represented as follows:

$$\langle message_alignment \rangle ::= \langle data_type \rangle \mid \langle expr \rangle \langle data_multiplier \rangle$$

‘64 BYTE’, ‘3 MEGABYTE’, and ‘QUADWORD’ are therefore all valid examples of $\langle message_alignment \rangle$ s. Bit counts are rounded up to the nearest byte count, so ‘27 BITS’ is in fact equivalent to ‘4 BYTES’.

The `ALIGNED` keyword forces `CONCEPTUAL` to align messages on *exactly* the specified alignment. Hence, a ‘HALFWORD ALIGNED’ message can begin at memory locations 0, 2, 4, 6, 8, . . . , $2k$ ($k \in \mathbb{Z}^+$). In contrast, the `MISALIGNED` keyword forces `CONCEPTUAL` to align messages the given number of bytes (positive or negative) past a page boundary. For example, if pages are 8192 bytes in size then a message described as ‘HALFWORD MISALIGNED’ can begin at memory locations 2, 8194, 16386, 24578, . . . , $8192k + 2$ ($k \in \mathbb{Z}^+$). Unlike `ALIGNED`, `MISALIGNED` supports negative alignments. If the page size is 4096 bytes, then ‘-10 BYTE MISALIGNED’ enables a message to begin at memory locations 4086, 8182, 12278, etc. The `MISALIGNED` alignment is taken modulo the page size. Therefore, with a 4096-byte page size, ‘10000 BYTE MISALIGNED’ is the same as ‘1808 BYTE MISALIGNED’.

The `UNALIGNED` keyword explicitly specifies the default behavior, with messages aligned on arbitrary boundaries.

Data touching

A $\langle message_spec \rangle$ described as being `WITH DATA TOUCHING` will force every word in a message to be both read and written (“touched”). When $\langle message_spec \rangle$ describes an outgoing message, the data will be touched before transmission. When $\langle message_spec \rangle$ describes an incoming message, the data will be touched after reception. In a sense, `WITH DATA TOUCHING` presents a more realistic assessment of network performance, as real applications almost always access the data they send or receive. It also distinguishes between messaging

layers that implicitly touch data and those that can transmit data without having to touch it. One would expect the latter to perform better when the data is not touched, as the former may be paying a penalty for touching the data. However, either could perform better when messages are sent `WITH DATA TOUCHING`, because the latter now has to pay the penalty that the former has already paid.

Another form of data-touching supported by `CONCEPTUAL` is `WITH VERIFICATION`. This causes the source task to write known, but randomly generated, data into the message before transmission and the target task to verify that every bit was correctly received. When a message is received `WITH VERIFICATION`, the `bit_errors` variable (see [Section A.2 \[Predeclared variables\]](#), page 184) is updated appropriately.

`WITHOUT DATA TOUCHING` and `WITHOUT VERIFICATION` are synonymous. Both explicitly specify the default behavior of neither touching nor verifying message contents.

Buffer control

The `CONCEPTUAL` run-time library allocates a unique message buffer for each message sent/received with the `UNIQUE` keyword (see [\[Unique messages\]](#), page 96). The message buffers for `NONUNIQUE` messages are recycled subject to the constraint that no two concurrent transmissions will reference the same buffer. For example, if a task performs a synchronous send followed by a synchronous receive, those operations must be executed serially and will therefore share a message buffer. If, instead, a task performs an asynchronous send followed by an asynchronous receive, those operations may overlap, so `CONCEPTUAL` will use different message buffers for the two operations.

Message specifications enable the programmer to override the default buffer-allocation behavior. If a message is sent `FROM BUFFER <expr>` or received `INTO BUFFER <expr>`, the message is guaranteed to be sent/received using the specified buffer number. For example, `FROM BUFFER` and `INTO BUFFER` can be used to force a synchronous send and synchronous receive to use different buffers or an asynchronous send and asynchronous receive to use the same buffer. If `<expr>` is negative, the behavior is the same as if `FROM BUFFER/INTO BUFFER` was not specified. `FROM THE DEFAULT BUFFER` and `INTO THE DEFAULT BUFFER` also explicitly specify the default buffer-allocation behavior.

Blocking semantics

By default—or if `SYNCHRONOUSLY` is specified—messages are sent synchronously. That is, a sender blocks (i.e., waits) until the message buffer is safe to reuse before it continues and a receiver blocks until it actually receives the message. The `ASYNCHRONOUSLY` keyword specifies that messages should be sent and received asynchronously. That is, the program merely posts the message (i.e., declares that it should eventually be sent and/or received) and immediately continues executing. Asynchronous messages must be *completed* as described in [Section 4.4.4 \[Awaiting completion\]](#), page 101.

4.4.2 Sending

The `SEND` statement is fundamental to `CONCEPTUAL`. It is used to send a multiple messages from multiple source tasks to multiple target tasks. The syntax is formally specified as follows:

```
<send_stmt> ::= <source_task>
```



```

[ASYNCHRONOUSLY] SENDS
  <message_spec>
TO [UNSUSPECTING] <target_tasks>
|
  <source_task>
[ASYNCHRONOUSLY] SENDS
  <message_spec>
TO <target_tasks>
WHO RECEIVE IT
  <recv_message_spec>

```

<source_task> is described in [Section 4.3.2 \[Source tasks\]](#), page 93; *<message_spec>* and *<recv_message_spec>* are described in [Section 4.4.1 \[Message specifications\]](#), page 94; and, *<target_tasks>* is described in [Section 4.3.3 \[Target tasks\]](#), page 94.

The SEND statement's simplest form, “*<source_task>* SENDS *<message_spec>* TO *<target_tasks>*”, is fairly straightforward. The following is an example:

```
TASK 0 SENDS A 0 BYTE MESSAGE TO TASK 1
```

The only subtlety in the preceding statement is that it implicitly causes task 1 to perform a corresponding receive. This receive can be suppressed by adding the keyword UNSUSPECTING before the *<target_tasks>* description:

```
TASK 0 SENDS A 0 BYTE MESSAGE TO UNSUSPECTING TASK 1
```

Here are some further examples of valid *<send_stmt>*s:

- ALL TASKS SEND A 64 KILOBYTE MESSAGE TO TASK 0
- TASK *num_tasks*-1 SENDS 5 53 BYTE PAGE ALIGNED MESSAGES TO ALL OTHER TASKS
- TASKS *upper* SUCH THAT *upper* ≥ *num_tasks*/2 ASYNCHRONOUSLY SEND A 0 BYTE MESSAGE TO TASK *upper*/2
- TASKS *nonzero* SUCH THAT *nonzero* > 0 SEND *nonzero* 1E3 BYTE MESSAGES TO UNSUSPECTING TASK 0

One subtlety of the SEND statement when used without UNSUSPECTING involves the orderings of the sends and receives. The rule is that receives are posted before sends. Furthermore, *<restricted_ident>*s (see [Section 4.3.1 \[Restricted identifiers\]](#), page 93) are evaluated in order from 0 to *num_tasks* - 1. The implication is that a statement such as ‘TASKS *ev* SUCH THAT *ev* IS EVEN /\ *ev* < 6 SEND A 4 WORD MESSAGE TO TASK *ev*+2’ is exactly equivalent to the following ordered sequence of statements (assuming *num_tasks* ≥ 5):

1. TASK 2 RECEIVES A 4 WORD MESSAGE FROM TASK 0
2. TASK 4 RECEIVES A 4 WORD MESSAGE FROM TASK 2
3. TASK 6 RECEIVES A 4 WORD MESSAGE FROM TASK 4
4. TASK 0 SENDS A 4 WORD MESSAGE TO UNSUSPECTING TASK 2
5. TASK 2 SENDS A 4 WORD MESSAGE TO UNSUSPECTING TASK 4
6. TASK 4 SENDS A 4 WORD MESSAGE TO UNSUSPECTING TASK 6

(The RECEIVE statement is described in [Section 4.4.3 \[Receiving\]](#), page 100.)

If the above sequence were executed, tasks 2, 4, and 6 would immediately block on their receives (steps 1–3). Task 0 would awaken task 2 by sending it a message (step 4). Then, task 2 would be able to continue to step 5 at which point it would send a message to task 4.

Task 4 would then finally be able to send a message to task 6 (step 6). Hence, even though the original `CONCEPTUAL` statement encapsulates multiple communication operations, the component communications proceed sequentially because of data dependences and because the operations are blocking.

As should now be apparent, there are a number of attributes associated with every message transmission:

- synchronous vs. asynchronous operation
- unique vs. recycled message buffers
- unaligned vs. aligned vs. misaligned message buffers
- no data touching vs. data touching vs. data verification
- implicit vs. explicit message-buffer selection

When `UNSUSPECTING` is omitted, the implicit `RECEIVE` statement normally inherits all of the attributes of the corresponding `SEND`. However, the second form of a $\langle send_stmt \rangle$, which contains a `WHO RECEIVES IT` (or `WHO RECEIVES THEM`) clause, enables the receiver's attributes to be overridden on a per-attribute basis. For instance, consider the following `SEND` statement:

```
TASK 0 SENDS A 1 MEGABYTE MESSAGE TO TASK 1 WHO RECEIVES IT
ASYNCHRONOUSLY
```

The alternative sequence of statements that does not use `WHO RECEIVES IT` is less straightforward:

1. `TASK 1 ASYNCHRONOUSLY RECEIVES A 1 MEGABYTE MESSAGE FROM TASK 0`
2. `TASK 0 SENDS A 1 MEGABYTE MESSAGE TO UNSUSPECTING TASK 1`

Some further examples of `WHO RECEIVES IT` follow:

```
TASKS left SUCH THAT left IS EVEN SEND 5 2 KILOBYTE 64 BYTE ALIGNED
MESSAGES TO TASKS left+1 WHO RECEIVE THEM AS UNALIGNED MESSAGES WITH
DATA TOUCHING
```

```
TASK num_tasks-1 ASYNCHRONOUSLY SENDS A 1E5 BYTE MESSAGE WITH
VERIFICATION TO TASK 0 WHO RECEIVES IT SYNCHRONOUSLY
```

```
TASK leaf SUCH THAT KNOMIAL_CHILDREN(leaf,2)=0 SENDS A UNIQUE 1536
BYTE MESSAGE WITH DATA TOUCHING TO TASK KNOMIAL_PARENT(leaf,2) WHO
RECEIVES IT ASYNCHRONOUSLY AS A NONUNIQUE QUADWORD ALIGNED MESSAGE
WITHOUT DATA TOUCHING INTO BUFFER KNOMIAL_PARENT(leaf,2)
```

4.4.3 Receiving

Section 4.4.2 [Sending], page 98, mentioned the $\langle send_stmt \rangle$'s `UNSUSPECTING` keyword, which specifies that the targets should not implicitly perform a receive operation. Because every send must have a matching receive, `CONCEPTUAL` offers a `RECEIVE` statement which explicitly receives a set of messages. A $\langle receive_stmt \rangle$ is much like a $\langle send_stmt \rangle$ (see Section 4.4.2 [Sending], page 98) with the $\langle source_task \rangle$ and $\langle target_tasks \rangle$ in the reverse order:

```
 $\langle receive\_stmt \rangle$  ::=  $\langle target\_tasks \rangle$ 
[ASYNCHRONOUSLY] RECEIVE
```

```

    <message_spec>
    FROM <source_task>

```

<target_tasks> is described in [Section 4.3.3 \[Target tasks\], page 94](#); *<message_spec>* is described in [Section 4.4.1 \[Message specifications\], page 94](#); and, *<source_task>* is described in [Section 4.3.2 \[Source tasks\], page 93](#).

For each message sent via a `SEND...TO UNSUSPECTING` statement there must be a `RECEIVE` statement that receives a message of the same size. The *<target_tasks>*'s *<message_spec>* can, however, specify different values for message uniqueness, message alignment, and data touching. In addition, the source and target do not need to agree on the use of the `ASYNCHRONOUSLY` keyword. The only restriction is that `WITH VERIFICATION` will return spurious results if used by the target but not by the source. Hence, the following *<send_stmt>* and *<receive_stmt>* correctly match each other:

```
TASK 0 SENDS 3 4 KILOBYTE MESSAGES TO UNSUSPECTING TASK 1
```

```
TASK 1 ASYNCHRONOUSLY RECEIVES 3 UNIQUE 4 KILOBYTE 48 BYTE ALIGNED
MESSAGES WITH DATA TOUCHING FROM TASK 0.
```

In general, it is better to use a single `SEND` statement with a `WHO RECEIVES IT` clause (see [Section 4.4.2 \[Sending\], page 98](#)) than a `RECEIVE` plus a matching `SEND...TO UNSUSPECTING`; the former is less error-prone than the latter. However, the latter is useful for programs in which a set of receives is posted, then the tasks perform various communication, computation, and synchronization operations, and—towards the end of the program—the matching sends are posted. That sort of split-phase structure requires separate `SEND` and `RECEIVE` statements.

4.4.4 Awaiting completion

When a message is sent or received asynchronously it must eventually be *completed*. In some messaging layers, asynchronous messages are not even sent or received until completion time. `CONCEPTUAL` provides the following statement for completing messages that were send/received asynchronously:

```

<wait_stmt> ::= <source_task>
                AWAITS COMPLETION

```

That is, a *<wait_stmt>* simply specifies the set of tasks that should block until all of their pending communications complete. *<source_task>* is as defined in [Section 4.3.2 \[Source tasks\], page 93](#). Note that a *<wait_stmt>* blocks until *all* pending communications complete. `CONCEPTUAL` does not provide finer-grained control over completions. It is safe, however, for a task to `AWAIT COMPLETION` even if it has no asynchronous messages pending.

4.4.5 Multicasting

Although a single *<send_stmt>* (see [Section 4.4.2 \[Sending\], page 98](#)) can specify multiple messages at once, these messages are sent one at a time. *Multicasting* is a form of collective communication in which a set of tasks collaborates to deliver a message from a source to multiple targets. With many messaging layers, multicasting a message to N tasks is more efficient than sending a sequence of N individual messages. `CONCEPTUAL` supports multicasting as follows:

```

<mcast_stmt> ::= <source_task>

```

```
[ASYNCHRONOUSLY] MULTICASTS
<message_spec>
TO <target_tasks>
```

Unlike $\langle send_stmt \rangle$ s, $\langle mcast_stmt \rangle$ s do not support the `UNSUSPECTING` keyword. This is because `MULTICASTS` is a collective operation: all parties are active participants in delivering messages to the $\langle target_tasks \rangle$.

$\langle source_task \rangle$ (see [Section 4.3.2 \[Source tasks\], page 93](#)) and $\langle target_tasks \rangle$ (see [Section 4.3.3 \[Target tasks\], page 94](#)) can be either disjoint or overlapping sets. That is, either of the following is legal:

```
TASK 0 MULTICASTS A 16 BYTE MESSAGE TO TASKS recip SUCH THAT recip<4
TASK 0 MULTICASTS A 16 BYTE MESSAGE TO TASKS recip SUCH THAT recip>=4
```

Note that in the first $\langle mcast_stmt \rangle$, task 0 both sends and receives a message, while in the second $\langle mcast_stmt \rangle$, task 0 sends but does not receive.

4.4.6 Reducing

A reduction operation is, in a sense, a complementary operation to a multicast (see [Section 4.4.5 \[Multicasting\], page 101](#)). While a multicast delivers a message from one source to multiple targets, a reduction combines messages from multiple sources (by applying a commutative/associative operator to corresponding elements) to a single target. Reduction is a collective operation: all parties collaborate to calculate the reduced value(s). As an example, if tasks 0, 1, and 2 collectively reduce the messages $\{5, 1\}$, $\{2, 7\}$, and $\{3, 4\}$ to task 2 using the “+” operator, then task 2 will receive the message $\{10, 12\}$. In fact, `CONCEPTUAL`’s implementation of reductions also supports reductions to multiple targets with each target receiving a copy of the reduced value.

The following grammatical rules define `CONCEPTUAL`’s many-to-many and many-to-one reduction facilities:

```
<reduce_stmt> ::= <source_task>
                REDUCES
                <reduce_message_spec>
                TO <source_task>
                [WHO RECEIVES THE RESULT <reduce_target_message_spec>]
                |
                <source_task>
                REDUCES
                <reduce_message_spec>
                [TO <reduce_message_spec>]
```

$\langle reduce_message_spec \rangle$ is defined in [Section 4.4.1 \[Message specifications\], page 94](#). Both the data providers and data receivers are specified as $\langle source_task \rangle$ nonterminals (see [Section 4.3.2 \[Source tasks\], page 93](#)). This design enables any set of tasks to provide the data to reduce and any disjoint or overlapping set of tasks to receive the reduced data. As with all communication in `CONCEPTUAL`, message contents are opaque. Furthermore, the grammar does not currently enable the programmer to specify the commutative/associative operator to use.

A simple many-to-one reduction can be expressed in `CONCEPTUAL` with ‘`ALL TASKS REDUCE 5 DOUBLEWORDS TO TASK 0`’. Note that the definition of $\langle reduce_message_spec \rangle$ and $\langle reduce_target_message_spec \rangle$ (see [Section 4.4.1 \[Message specifications\], page 94](#)) supports

reductions only of INTEGERS and DOUBLEWORDS, not arbitrary *<data_type>* values. Omitting the optional ‘TO *<reduce_message_spec>*’, as in ‘TASKS *rt* SUCH THAT 3 DIVIDES *rt* REDUCE AN INTEGER’, specifies that all tasks performing the reduction will receive a copy of the reduced value. The sources and targets can also be designated explicitly as in ‘TASKS *xyz* SUCH THAT *xyz**<num_tasks/2>* REDUCE 100 DOUBLEWORDS TO TASKS *xyz+num_tasks/4*’. When that code is run with 8 tasks, tasks 0–3 reduce 100 doublewords (800 bytes) apiece and tasks 2–5 each receive identical copies of the 100 doublewords of reduced data.

Message data used with REDUCES can be transferred WITH DATA TOUCHING (but not WITH VERIFICATION); data alignment can be specified; and, messages buffers can be named explicitly. The following example represents fairly complex many-to-many usage of REDUCES:

```
TASKS rsrc SUCH THAT rsrc IS EVEN REDUCE 32 64-BYTE-ALIGNED INTEGERS
WITH DATA TOUCHING FROM BUFFER 2 TO TASKS rtarg SUCH THAT
rtarg<num_tasks/4> \/rtarg>(3*num_tasks)/4 WHO RECEIVE THE RESULT AS
32 UNIQUE PAGE-ALIGNED INTEGERS WITHOUT DATA TOUCHING.
```

4.4.7 Synchronizing

CONCEPTUAL enables sets of tasks to perform *barrier synchronization*. The semantics are that no task can finish synchronizing until all tasks have started synchronizing. The syntax is as follows:

```
<sync_stmt> ::= <source_task>
                SYNCHRONIZES
```

A *<sync_stmt>* can be used to ensure that one set of statements has completed before beginning another set. For example, a CONCEPTUAL program might have a set of tasks post a series of asynchronous receives (see [Section 4.4.3 \[Receiving\], page 100](#)), then make ‘ALL TASKS SYNCHRONIZE’ before having another set of tasks perform the corresponding UNSUSPECTING sends (see [Section 4.4.2 \[Sending\], page 98](#)). This procedure ensures that all of the target tasks are ready to receive before the source tasks start sending to them.

4.5 I/O statements

CONCEPTUAL provides two statements for presenting information. One statement writes simple messages to the standard output device and is intended to be used for providing status information during the run of a program. The other statement provides a powerful mechanism for storing performance and correctness data to a log file.

4.5.1 Utilizing log-file comments

*<output_stmt>*s and *<log_stmt>*s have limited access to the *<key:value>* pairs that are written as comments at the top of every log file as shown in [Section 3.5.1 \[Log-file format\], page 44](#). Given a key, *key*, the string expression ‘THE VALUE OF *key*’ represents the value associated with that key or the empty string if *key* does not appear in the log-file comments:

```
<string_or_log_comment> ::= <string>
                            | THE VALUE OF <string>
```

That is, “CPU frequency” means the literal string “CPU frequency” while ‘THE VALUE OF “CPU frequency”’ translates to a string like “1300000000 Hz (1.3 GHz)”. Environment variables are also considered keys and are therefore acceptable input to a THE VALUE OF construct.

4.5.2 Writing to standard output

CONCEPTUAL's `OUTPUT` keyword is used to write a message from one or more source tasks (see [Section 4.3.2 \[Source tasks\]](#), page 93) to the standard output device. This is useful for providing progress reports during the execution of long-running CONCEPTUAL programs. An `<output_stmt>` looks like this:

```
<output_stmt> ::= <source_task>
                OUTPUTS
                <expr> | <string_or_log_comment>
                [AND <expr> | <string_or_log_comment>]*
```

The following are some sample `<output_stmt>`s:

```
TASK 0 OUTPUTS "Hello, world!"
```

```
TASKS nr SUCH THAT nr>0 OUTPUT nr AND "'s parent is " AND nr>>1 AND
" and its children are " AND nr<<1 AND " and " AND nr<<1+1
```

```
ALL TASKS me OUTPUT "Task " AND me AND " is running on host " AND THE
VALUE OF "Host name" AND " and plans to send to task " AND (me+1) MOD
num_tasks
```

`OUTPUT` does not implicitly output spaces between terms. Hence, `'OUTPUT "Yes" AND "No"'` will output `"YesNo"`, not `"Yes No"`. Although it is unlikely that a program would ever need to output two arithmetic expressions with no intervening text, an empty string can be used for this purpose: `'OUTPUT 6 AND "" AND 3'`.

An `<output_stmt>` implicitly outputs a newline character at the end. Additional newline characters can be output by embedding `'\n'` in a string. (see [Section 4.1 \[Primitives\]](#), page 77.) CONCEPTUAL does not provide a means for suppressing the newline, however.

4.5.3 Writing to a log file

After performing a network correctness or performance test it is almost always desirable to store the results in a file. CONCEPTUAL has language support for writing tabular data to a log file. The `<log_stmt>` command does the bulk of the work:

```
<log_stmt> ::= <source_task>
                LOGS
                <aggr_expr> AS <string_or_log_comment>
                [AND <aggr_expr> AS <string_or_log_comment>]*
```

The idea behind a `<log_stmt>` is that a set of source tasks (see [Section 4.3.2 \[Source tasks\]](#), page 93) log an aggregate expression (see [Section 4.2.3 \[Aggregate expressions\]](#), page 90) to a log file under the column heading `<string_or_log_comment>`. Each task individually maintains a separately named log file so there is no ambiguity over which task wrote which entries.

Each (static) `LOGS` statement in a CONCEPTUAL program specifies one or more columns of the log file. Every dynamic execution of a `LOGS` statement writes a single row to the log file. A single `LOGS` statement should suffice for most CONCEPTUAL programs.

The following are some examples of `<log_stmt>`s:

```
ALL TASKS LOG bit_errors AS "Bit errors"
```

```
TASK 0 LOGS THE msgsize AS "Bytes" AND
      THE MEDIAN OF (1E6*bytes_sent)/(1M*elapsed_usecs) AS "MB/s"
```

The first example produces a log file like the following:

```
"Bit errors"
"(all data)"
3
```

The second example produces a log file like this:

```
"Bytes", "MB/s"
"(only value)", "(median)"
65536, 179.9416266
```

Note that in each log file, the CONCEPTUAL run-time system writes two rows of column headers for each column. The first row contains $\langle string_or_log_comment \rangle$ as is. The second row describes the $\langle aggr_func \rangle$ (see [Section 4.2.4 \[Aggregate functions\], page 91](#)) used to aggregate the data. One or more rows of data follow.

Assume that the second $\langle log_stmt \rangle$ presented above appears within a loop (see [Section 4.7.2 \[Iterating\], page 113](#)). It is therefore important to include the **THE** keyword before ‘msgsize’ to assert that the expression ‘msgsize’ is constant across invocations of the $\langle log_stmt \rangle$ and that, consequently, only a single row of data should be written to the log file. Using ‘msgsize’ without the **THE** would produce a column of data with one row per $\langle log_stmt \rangle$ invocation:

```
"Bytes", "MB/s"
"(all data)", "(median)"
65536, 179.9416266
65536,
65536,
65536,
65536,
.
.
.
```

The rules that determine how LOGS statements produce rows and columns of a log file are presented below:

1. Each *static* LOGS statement (and AND clause within a LOGS statement) in a program produces a unique column.
2. Each *dynamic* execution of a LOGS statement appends a row to the column(s) it describes.

- Each top-level complex statement (see [Section 4.9 \[Complete programs\], page 121](#)) produces a new table in the log file.

Note that the choice of column name is inconsequential for determining what columns are written to the log file:

```
TASK 0 LOGS 314/100 AS "Pi" AND 22/7 AS "Pi"
```

```
"Pi","Pi"
"(all data)","(all data)"
3.14,3.142857143
```

Computing aggregates

What if ‘`msgsize`’ takes on a number of values throughout the execution of the program and for each value a number of runs is performed? How would one log the median of each set of data? Using ‘`THE msgsize`’ won’t work because the message size is not constant. Using ‘`msgsize`’ alone won’t work either because `CONCEPTUAL` would then take the median of the times gathered across *all* message sizes, which is undesirable. The solution is for the program to specify explicitly when aggregate functions (`MEDIAN` and all of the other functions listed in [Section 4.2.4 \[Aggregate functions\], page 91](#)) compute a value:

```
<flush_stmt> ::= <source_task>
                COMPUTES AGGREGATES
```

The intention is that an inner loop might `LOG` data after every iteration and an outer loop would ‘`COMPUTE AGGREGATES`’ after each iteration.

4.6 Other statements

`CONCEPTUAL` contains a few more statements than those described in [Section 4.4 \[Communication statements\], page 94](#) and [Section 4.5 \[I/O statements\], page 103](#). As there is no category that clearly describes the remaining statements, they are listed here in this “catch-all” section.

4.6.1 Resetting counters

It is common for benchmarks to perform a few warmup iterations before beginning the actual test. This hides the time needed initially to establish connections, populate the various processor and network caches, etc., making subsequent measurements more uniform. However, these warmup iterations, like everything else in `CONCEPTUAL`, exhibit the side effect of updating `CONCEPTUAL`’s dynamic variables (see [Section A.2 \[Predeclared variables\], page 184](#)) such as `elapsed_usecs` and `msgs_sent`. The solution is to reset all of those dynamic variables to zero before beginning the core part of the program. The `RESET` statement does just that:

```
<reset_stmt> ::= <source_task>
                RESETS ITS COUNTERS
```

Hence, writing ‘`ALL TASKS RESET THEIR COUNTERS`’ causes each task to reset all of the variables listed in [Section A.2 \[Predeclared variables\], page 184](#)—with the exception of `num_tasks`—to zero. Note that `ITS` and `THEIR`, like `RESET` and `RESETS`, are considered synonyms (see [Section 4.1 \[Primitives\], page 77](#)).

4.6.2 Asserting conditions

CONCEPTUAL programs can encode the run-time conditions that must hold in order for the test to run properly. This is achieved through *assertions*, which are expressed as follows:

```
<assert_stmt> ::= ASSERT THAT <string>
                WITH <rel_expr>
```

<string> is a message to be reported to the user if the assertion fails. *<rel_expr>* is a relational expression (as described in [Section 4.2.5 \[Relational expressions\]](#), page 91) that must evaluate to TRUE for the program to continue running. Assertion failures are considered fatal errors. They cause the CONCEPTUAL program to abort immediately.

Here are some sample *<assert_stmt>*s:

```
ASSERT THAT "the bandwidth test requires at least two tasks" WITH
num_tasks >= 2
```

```
ASSERT THAT "pairwise ping-pongs require an even number of task"
WITH num_tasks IS EVEN
```

```
ASSERT THAT "this program requires a square number of tasks" WITH
Sqrt(num_tasks)**2 = num_tasks
```

(For the last example, recall that CONCEPTUAL expressions are of integer type. Hence, the example's *<rel_expr>* is mathematically equivalent to $\lfloor \sqrt{N} \rfloor^2 = N$, which is TRUE if and only if N is a square.)

4.6.3 Delaying execution

It is sometimes interesting to measure the progress of a communication pattern when delays are inserted at various times on various tasks. CONCEPTUAL provides two mechanisms for inserting delays: one that relinquishes the CPU while delaying (SLEEP) and one that hoards it (COMPUTE).

```
<delay_stmt> ::= <source_task>
                SLEEPS | COMPUTES
                FOR <expr> <time_unit>
```

<source_task> (see [Section 4.3.2 \[Source tasks\]](#), page 93) specifies the set of tasks that will stall. *<expr>* (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79) specifies the number of *<time_unit>*s for which to delay and *<time_unit>* represents any of the following measures of time:

```
<time_unit> ::= MICROSECONDS | MILLISECONDS | SECONDS | MINUTES | HOURS | DAYS
```

Delay times are only approximate. SLEEP's accuracy depends upon the operating-system's clock resolution or length of time quantum (commonly measured in milliseconds or tens of milliseconds). COMPUTE, which is implemented by repeatedly reading a variable until the desired amount of time elapses, is calibrated during the CONCEPTUAL run-time system's initialization phase and can be adversely affected by intermittent system load. Both forms of *<delay_stmt>* attempt to measure wall-clock time ("real time"), not just the time the program is running ("virtual time"). Because the delay times are approximate, it is strongly recommended that the `elapsed_usec` variable (see [Section A.2 \[Predeclared variables\]](#), page 184) be employed to determine the actual elapsed time.

4.6.4 Touching memory

“Touching” memory means reading and writing it. The `CONCEPTUAL TOUCHES` statement enables a program to touch memory for one of two purposes: either to simulate computation in an application by thrashing some or all of the memory hierarchy or to preload message buffers into the upper levels of the memory hierarchy in order to better separate communication costs from memory costs.

Simulating computation

While the statements described in [Section 4.6.3 \[Delaying execution\], page 107](#) delay for a specified length of time, it is also possible to delay for the duration of a specified amount of “work”. “Work” is expressed in terms of memory accesses. That is, a `CONCEPTUAL` program can touch (i.e., read plus write) data with a given stride from a memory region of a given size. By varying these parameters, a program can emulate an application’s computation by hoarding the CPU or any level of the memory hierarchy.

```

<touch_stmt> ::= <source_task>
                TOUCHES
                [<expr> <data_type> OF]
                AN <item_size> MEMORY REGION
                [<expr> TIMES]
                [WITH STRIDE <expr> <data_type> | WITH RANDOM STRIDE]

```

`<item_size>` and `<data_type>` are described in [\[Item size\], page 96](#) and `<expr>` is described in [Section 4.2.1 \[Arithmetic expressions\], page 79](#).

As shown by the formal definition of `<touch_stmt>` the required components are a `<source_task>` and the size of the memory region to touch. By default, every `WORD` (see [\[Item size\], page 96](#)) of memory in the region is touched exactly once. The type of data that is touched can be varied with an ‘`<expr> <data_type> OF`’ clause. For instance, ‘`100 BYTES OF`’ of a memory region will touch individual bytes. An optional repeat count enables the memory region (or subset thereof) to be touched multiple times. Hence, if ‘`TASK 0 TOUCHES A 6 MEGABYTE MEMORY REGION 5 TIMES`’, then the touch will be performed as if ‘`TASK 0`’ were told to ‘`TOUCH 5*6M BYTES OF A 6 MEGABYTE MEMORY REGION 1 TIME`’ or simply to ‘`TOUCH 5*6M BYTES OF A 6 MEGABYTE MEMORY REGION`’.

By default, every `<data_type>` of data is touched. However, a `<touch_stmt>` provides for touching only a subset of the `<data_type>`s in the memory region. By writing ‘`WITH STRIDE <expr> <data_type>`’, only the first `<data_type>` out of every `<expr>` will be touched. Instead of specifying an exact stride, the memory region can be accessed in random order using the `WITH RANDOM STRIDE` clause.

Unless the number of touches and data type are specified explicitly, the number of `WORDS` that are touched is equal to the size of the memory region divided by the stride length then multiplied by the repeat count. Therefore, if ‘`TASK 0 TOUCHES AN 8 MEGABYTE MEMORY REGION 2 TIMES WITH STRIDE 8 WORDS`’, then a total of $(2^{23}/(4 \times 8)) \times 2 = 524288$ touches will be performed. For the purpose of the preceding calculation, ‘`WITH RANDOM STRIDE`’ should be treated as if it were ‘`WITH STRIDE 1 WORD`’ (again, unless the number of touches and data type are specified explicitly).

To save memory, all `TOUCH` statements in a `CONCEPTUAL` program access subsets of the same region of memory, whose size is determined by the maximum needed. However, each

dynamic execution of a *<touch_stmt>* starts touching from where the previous execution left off. For example, consider the following statement:

```
TASK 0 TOUCHES 100 WORDS OF A 200 WORD MEMORY REGION
```

The first time that that statement is executed within a loop (see [Section 4.7.2 \[Iterating\]](#), [page 113](#)), the first 200 words are touched. The second time, the second 200 words are touched. The third time, the index into the region wraps around and the first 200 words are touched again.

Each static *<touch_stmt>* maintains its own index into the memory region. Therefore, the first of the following two statements will terminate successfully (assuming it's not executed in the body of a loop) while the second will result in a run-time error because the final byte of the final word does not fit within the given memory region.

```
TASK 0 TOUCHES 100 WORDS OF A 799 BYTE MEMORY REGION THEN
TASK 0 TOUCHES 100 WORDS OF A 799 BYTE MEMORY REGION
```

FOR 2 REPETITIONS TASK 0 TOUCHES 100 WORDS OF A 799 BYTE MEMORY REGION (THEN is described in [Section 4.7.1 \[Combining statements\]](#), [page 112](#), and FOR...REPETITIONS is described in [\[Counted loops\]](#), [page 113](#).) The first statement shown above touches the same 100 words (400 bytes) in each of the two *<touch_stmt>*s. The second statement touches the first 100 words the first time the *<touch_stmt>* is executed and fails when trying to touch the (only partially extant) second 100 words.

Priming message buffers

[\[Data touching\]](#), [page 97](#), describes how a message sent or received WITH DATA TOUCHING will have all of its data touched before or after transmission. Sometimes, however, a program may want to touch message data without actually transmitting a message. For example, a task could touch message data to load it into the cache, then RESET ITS COUNTERS (see [Section 4.6.1 \[Resetting counters\]](#), [page 106](#)), and finally send or receive a message without further data touching. The TOUCHES statement has an alternate form that touches message buffers instead of isolated memory regions:

```
<touch_buffer_stmt> ::= <source_task>
                        TOUCHES
                        ALL MESSAGE BUFFERS
                        | MESSAGE BUFFER <expr>
                        | THE CURRENT MESSAGE BUFFER
```

The first form, 'ALL MESSAGE BUFFERS', touches all message buffers available to the program, even those not yet used at the time TOUCHES is invoked. In the following statement, for example, TOUCHES "knows" that 10 message buffers will be used and touches the data in all 10 of them:

```
ALL TASKS TOUCH ALL MESSAGE BUFFERS THEN
FOR EACH SZ IN {0, ..., 9}
    TASK 0 ASYNCHRONOUSLY SENDS AN SZ MEGABYTE MESSAGE TO TASK 1 THEN
ALL TASKS AWAIT COMPLETION
```

([\[Buffer control\]](#), [page 98](#), explains why the preceding statement requires 10 buffers.) One caveat is that in CONCEPTUAL version 1.0, messages sent or received with the UNIQUE

keyword (see [Unique messages], page 96) are not touched. This limitation may be lifted in a later release of CONCEPTUAL.

The second form of the $\langle touch_buffer_stmt \rangle$ statement, ‘MESSAGE BUFFER $\langle expr \rangle$ ’, touches a specific message buffer. It is expected to be used in programs which send ‘FROM BUFFER $\langle expr \rangle$ ’ or receive ‘INTO BUFFER $\langle expr \rangle$ ’.

‘THE CURRENT MESSAGE BUFFER’, the third and final form of the $\langle touch_buffer_stmt \rangle$ statement, touches the first message buffer that is not currently in use (i.e., is not the source or destination of an asynchronous operation). Usually, this is whichever message buffer will next be sent from or received into. (The exception is when a task AWAITS COMPLETION after touching the current message buffer but before sending or receiving a message—probably a somewhat contrived situation.)

4.6.5 Reordering task IDs

CONCEPTUAL distinguishes between “task IDs”, which are used in task descriptions (see Section 4.3 [Task descriptions], page 92) and “processor IDs”, which are assigned by the underlying communication layer. As stated in Section 3.4 [Running coNcEPTuaL programs], page 41, a CONCEPTUAL program has no control over how processor IDs map to physical processors. It therefore has no way to specify, for instance, that a set of tasks must run on the same multiprocessor node (or on different nodes, for that matter). Initially, every task’s task ID is set equal to its processor ID. However, while processor IDs are immutable, task IDs can be changed dynamically during the execution of a program. Altering task IDs can simplify CONCEPTUAL programs that might otherwise need to evaluate complex expressions to determine peer tasks. CONCEPTUAL enables either a specific or a randomly selected task to be assigned to a given processor:

```
 $\langle processor\_stmt \rangle ::= \langle source\_task \rangle$ 
                     IS ASSIGNED TO
                     PROCESSOR  $\langle expr \rangle$  | A RANDOM PROCESSOR
```

In addition to performing the specified processor assignment, CONCEPTUAL will perform an additional, implicit processor assignment in order to maintain a bijection between task IDs and processor IDs (i.e., every task ID corresponds to exactly one processor ID and every processor ID corresponds to exactly one task ID). Consider the following statement:

```
TASK n SUCH THAT  $n < (\text{num\_tasks} + 1) / 2$  IS ASSIGNED TO PROCESSOR  $n * 2$ 
```

If `num_tasks` is ‘8’ the preceding statement will cause ‘TASK 0’ to refer to processor 0, ‘TASK 1’ to refer to processor processor 2, ‘TASK 2’ to refer to processor 4, and ‘TASK 3’ to refer to processor 6. What may be unintuitive is that the remaining tasks will not map to their original processors, as doing so would violate the bijection invariant. To clarify CONCEPTUAL’s implicit processor assignments the following timeline illustrates the execution of ‘TASK n SUCH THAT $n < (\text{num_tasks} + 1) / 2$ IS ASSIGNED TO PROCESSOR $n * 2$ ’ when `num_tasks` is ‘8’:

‘n’	‘TASK 0’	‘TASK 1’	‘TASK 2’	‘TASK 3’	‘TASK 4’	‘TASK 5’	‘TASK 6’	‘TASK 7’
—	0	1	2	3	4	5	6	7
0	0	1	2	3	4	5	6	7
1	0	2	1	3	4	5	6	7
2	0	2	4	3	1	5	6	7
3	0	2	4	6	1	5	3	7

Initially, task and processor IDs are equal. When `n` takes on the value 0, `CONCEPTUAL` performs the equivalent of ‘TASK 0 IS ASSIGNED TO PROCESSOR 0’, which does not change the task ID to processor ID mapping. When `n` is 1, `CONCEPTUAL` performs the equivalent of ‘TASK 1 IS ASSIGNED TO PROCESSOR 2’, which sets task 1’s processor to 2. However, because task 2 also has processor 2, `CONCEPTUAL` implicitly performs the equivalent of ‘TASK 2 IS ASSIGNED TO PROCESSOR 1’ in order to preserve the unique task ID to processor ID mapping. When `n` is 2, `CONCEPTUAL` performs the equivalent of ‘TASK 2 IS ASSIGNED TO PROCESSOR 4’ and, because task 4 also has processor 4, the equivalent of ‘TASK 4 IS ASSIGNED TO PROCESSOR 1’. Finally, when `n` is 3, `CONCEPTUAL` performs the equivalent of ‘TASK 3 IS ASSIGNED TO PROCESSOR 6’ and, because task 6 also has processor 6, the equivalent of ‘TASK 6 IS ASSIGNED TO PROCESSOR 3’. Thus, `CONCEPTUAL` maintains the invariant that after any processor assignment every task corresponds to a unique processor and every processor corresponds to a unique task.

4.6.6 Injecting arbitrary code

There are some features that are outside the scope of the `CONCEPTUAL` language. However, `CONCEPTUAL` provides a mechanism for inserting backend-specific statements into the control flow of a `CONCEPTUAL` program. This feature is intended for users with specific needs that can’t be satisfied through the conventional `CONCEPTUAL` statements.

```

<backend_stmt> ::= <source_task>
                  BACKEND EXECUTES
                  <expr> | <string>
                  [AND <expr> | <string>]*

```

The following example assumes a C-based backend:

```
ALL TASKS taskID BACKEND EXECUTE "my_c_function(" AND taskID AND ");"
```

The `my_c_function()` function needs to be defined in some object file and linked with the `CONCEPTUAL`-generated code.

Most users will never need to `BACKEND EXECUTE` code. In fact, most users should *not* use `<backend_stmt>`s as they produce nonportable code. One of `CONCEPTUAL`’s goals is for programs to be understandable by people unfamiliar with the language, and `<backend_stmt>`s thwart that goal. However, `<backend_stmt>`s do help ensure that all of the target language/library’s features are available to `CONCEPTUAL`.

All `<expr>`s are passed to the backend in floating-point context (see [\[Evaluation contexts\]](#), page 80). Consequently, all backend code that takes a `CONCEPTUAL` `<expr>` needs to expect a floating-point value (which the backend code can of course cast explicitly to an integer if necessary).

There is an important special case defined for the argument to the `BACKEND EXECUTES` statement: once all of the `<string>` and `<expr>` arguments are concatenated into a single string, all occurrences of the substring ‘`[MESSAGE BUFFER expr]`’ are replaced by a pointer to the `expr`th message buffer created using the `FROM BUFFER` or `INTO BUFFER` keywords (presented in [\[Buffer control\]](#), page 98). This special case makes it easy for a `CONCEPTUAL` program to invoke communication functions provided by the underlying communications library. Note that the notion of a default buffer is not meaningful in the context of `BACKEND EXECUTES`, which is not a communication statement.

4.7 Complex statements

The CONCEPTUAL statements presented in Section 4.4 [Communication statements], page 94, Section 4.5 [I/O statements], page 103, and Section 4.6 [Other statements], page 106 are all known as *simple statements*. This section expands upon the statements already introduced by presenting *complex statements*. In its most basic form, a $\langle complex_stmt \rangle$ is just a $\langle simple_stmt \rangle$. However, the primary purpose of a $\langle complex_stmt \rangle$ is to juxtapose simple statements and other complex statements into more expressive forms.

Complex statements take the following form:

```
 $\langle complex\_stmt \rangle ::= \langle simple\_stmt \rangle [THEN \langle complex\_stmt \rangle]$ 
```

The constituent simple statements include FOR loops, LET bindings, IF conditionals, grouping constructs, and all of the statements introduced in Section 4.4 [Communication statements], page 94, Section 4.5 [I/O statements], page 103, and Section 4.6 [Other statements], page 106:

```
 $\langle simple\_stmt \rangle ::=$ 
  FOR  $\langle expr \rangle$  REPETITIONS [PLUS  $\langle expr \rangle$  WARMUP REPETITIONS
    [AND A SYNCHRONIZATION]]  $\langle simple\_stmt \rangle$ 
  |
  FOR EACH  $\langle ident \rangle$  IN  $\langle range \rangle$  [ $\langle range \rangle$ ]*  $\langle simple\_stmt \rangle$ 
  |
  FOR  $\langle expr \rangle$   $\langle time\_unit \rangle$   $\langle simple\_stmt \rangle$ 
  |
  LET  $\langle let\_binding \rangle$  [AND  $\langle let\_binding \rangle$ ]* WHILE  $\langle simple\_stmt \rangle$ 
  |
  IF  $\langle rel\_expr \rangle$  THEN  $\langle simple\_stmt \rangle$  [OTHERWISE  $\langle simple\_stmt \rangle$ ]
  |
  ‘{’ [ $\langle complex\_stmt \rangle$ ] ‘}’
  |
   $\langle send\_stmt \rangle$ 
  |
   $\langle receive\_stmt \rangle$ 
  |
   $\langle wait\_stmt \rangle$ 
  |
   $\langle mcast\_stmt \rangle$ 
  |
   $\langle reduce\_stmt \rangle$ 
  |
   $\langle sync\_stmt \rangle$ 
  |
   $\langle output\_stmt \rangle$ 
  |
   $\langle log\_stmt \rangle$ 
  |
   $\langle flush\_stmt \rangle$ 
  |
   $\langle reset\_stmt \rangle$ 
  |
   $\langle assert\_stmt \rangle$ 
  |
   $\langle delay\_stmt \rangle$ 
  |
   $\langle touch\_stmt \rangle$ 
  |
   $\langle touch\_buffer\_stmt \rangle$ 
  |
   $\langle processor\_stmt \rangle$ 
  |
   $\langle backend\_stmt \rangle$ 
```

The remainder of this section describes in turn the THEN construct and each of the just-introduced $\langle simple_stmt \rangle$ types.

4.7.1 Combining statements

The THEN keyword separates statements that are to be performed sequentially. For example, a simple ping-pong communication can be expressed as follows:

```
ALL TASKS RESET ALL COUNTERS THEN
TASK 0 SENDS A 0 BYTE MESSAGE TO TASK 1 THEN
```

```
TASK 1 SENDS A 0 BYTE MESSAGE TO TASK 0 THEN
TASK 0 LOGS elapsed_usecs/2 AS "One-way latency"
```

There is no implicit intertask synchronization across THEN statements. Consequently, the two communications specified in the following statement will be performed concurrently:

```
TASK 0 ASYNCHRONOUSLY SENDS AN 8 KILOBYTE MESSAGE TO TASK 1 THEN
TASK 1 ASYNCHRONOUSLY SENDS AN 8 KILOBYTE MESSAGE TO TASK 0 THEN
ALL TASKS AWAIT COMPLETION
```

4.7.2 Iterating

CONCEPTUAL provides a variety of looping constructs designed to repeatedly execute a *<simple_stmt>*.

Counted loops

The simplest form of iteration in CONCEPTUAL repeats a *<simple_stmt>* a given number of times. The syntax is simply “FOR *<expr>* REPETITIONS *<simple_stmt>*”. As could be expected, the *<expr>* term (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79) specifies the number of repetitions to perform. Hence, the following *<simple_stmt>* outputs the phrase “I will not talk in class” 100 times:

```
FOR 100 REPETITIONS ALL TASKS OUTPUT "I will not talk in class."
```

FOR...REPETITIONS can optionally specify a number of “warmup” repetitions to perform in addition to the base number of repetitions. The syntax is “FOR *<expr>* REPETITIONS PLUS *<expr>* WARMUP REPETITIONS *<simple_stmt>*”. During warmup repetitions, the OUTPUTS statement (see [Section 4.5.2 \[Writing to standard output\]](#), page 104), the LOGS statement (see [Section 4.5.3 \[Writing to a log file\]](#), page 104), and the COMPUTES AGGREGATES statement (see [\[Computing aggregates\]](#), page 106) are all suppressed (i.e., they have no effect) and none of the special variables predeclared by CONCEPTUAL (see [Section A.2 \[Predeclared variables\]](#), page 184) are updated. Many benchmarks synchronize all tasks after performing a set of warmup repetitions. This behavior can be expressed conveniently as part of a CONCEPTUAL FOR loop by appending the AND A SYNCHRONIZATION clause:

```
FOR 1000 REPETITIONS PLUS 3 WARMUP REPETITIONS AND A SYNCHRONIZATION
TASK 0 MULTICASTS A 1 MEGABYTE MESSAGE TO ALL OTHER TASKS
```

CONCEPTUAL also provides a separate SYNCHRONIZES statement. This is described in [Section 4.4.7 \[Synchronizing\]](#), page 103.

The importance of performing warmup repetitions is that many communication layers give atypically poor performance on the first few transmissions. This may be because the messages miss in the cache; because the communication layer needs to establish connections between pairs of communicating tasks; or, because the operating system needs to “register” message buffers with the network interface. Regardless of the reason, specifying warmup repetitions helps make performance measurements less variable.

Range loops

Range loops are CONCEPTUAL’s most powerful looping construct. Unlike the counted-loop construct presented in [\[Counted loops\]](#), page 113, a range loop binds a variable to a different value on each iteration. Range loops have the following syntax:

```
FOR EACH <ident>
```

```
IN  $\langle range \rangle$  [,  $\langle range \rangle$ ]*
 $\langle simple\_stmt \rangle$ 
```

A $\langle range \rangle$ is a comma-separated list of $\langle expr \rangle$ s within curly braces. An ellipsis can be used to indicate that CONCEPTUAL should fill in the missing numbers. More formally,

```
 $\langle range \rangle$  ::= '{'
            $\langle expr \rangle$  [' , '  $\langle expr \rangle$ ]*
           [' , ... , '  $\langle expr \rangle$ ]
           '}'
```

For a range loop, CONCEPTUAL successively binds a $\langle restricted_ident \rangle$ (see [Section 4.3.1 \[Restricted identifiers\], page 93](#)) to each $\langle expr \rangle$ in each $\langle range \rangle$ then evaluates the given $\langle simple_stmt \rangle$. Unlike all other uses of a $\langle restricted_ident \rangle$, the variable is not implicitly constrained to being a number from 0 to $num_tasks-1$ within a FOR EACH statement.

The following are some examples of the FOR EACH statement from simplest to most elaborate. Each example uses 'i' as the loop variable and 'TASK 0 OUTPUTS i' as the loop body. The output from each example is shown with a "←" symbol preceding each line.

```
FOR EACH i IN {8, 7, 5, 4, 5} TASK 0 OUTPUTS i
← 8
← 7
← 5
← 4
← 5
```

```
FOR EACH i IN {1, ..., 5} TASK 0 OUTPUTS i
← 1
← 2
← 3
← 4
← 5
```

```
FOR EACH i IN {5, ..., 1} TASK 0 OUTPUTS i
← 5
← 4
← 3
← 2
← 1
```

```
FOR EACH i IN {1, ..., 5}, {8, 7, 5, 4, 5} TASK 0 OUTPUTS i
← 1
← 2
← 3
← 4
← 5
← 8
← 7
← 5
← 4
```



```
→ 5
```

```
FOR EACH i IN {1, 4, 7, ..., 30} TASK 0 OUTPUTS i
```

```
→ 1
→ 4
→ 7
→ 10
→ 13
→ 16
→ 19
→ 22
→ 25
→ 28
```

```
FOR EACH i IN {3**1, 3**2, 3**3, ..., 3**7} TASK 0 OUTPUTS i
```

```
→ 3
→ 9
→ 27
→ 81
→ 243
→ 729
→ 2187
```

```
FOR EACH i IN {0}, {1, 2, 4, ..., 256} TASK 0 OUTPUTS i
```

```
→ 0
→ 1
→ 2
→ 4
→ 8
→ 16
→ 32
→ 64
→ 128
→ 256
```

When `CONCEPTUAL` fills in the numbers summarized by ‘...’ it first looks for an arithmetic progression, then a geometric progression. If `CONCEPTUAL` finds neither an arithmetic nor a geometric progression, it re-evaluates all of the `<expr>`s in floating-point context (see [Section 4.2.1 \[Arithmetic expressions\], page 79](#)) and tries once again to find a geometric progression. If a pattern is still not found, the compiler aborts with a run-time error message. This is why the final example above places the ‘0’ in a separate range; ‘1, 2, 4, ..., 256’ is a geometric pattern in which each number is twice the previous. ‘0, 1, 2, 4, ..., 256’ violates that pattern, because 1 is not twice 0. The number following the ellipsis is not included when calculating the pattern. Hence, the ‘{1, 4, 7, ..., 30}’ range shown above is acceptable to `CONCEPTUAL`. If the number following the ellipsis is less than (respectively, greater than) the first number in an increasing (respectively, decreasing) range, then the loop will silently not execute. Here are some examples to help make those points:

```
FOR EACH i IN {20, 30, 40, ..., 55} TASK 0 OUTPUTS i
  ↪ 20
  ↪ 30
  ↪ 40
  ↪ 50
```

```
FOR EACH i IN {20, 30, 40, ..., 30} TASK 0 OUTPUTS i
  ↪ 20
  ↪ 30
```

```
FOR EACH i IN {20, 30, 40, ..., 20} TASK 0 OUTPUTS i
  ↪ 20
```

```
FOR EACH i IN {20, 30, 40, ..., 10} TASK 0 OUTPUTS i
  ↪ [no output]
```

The examples so far have all used constant expressions. However, as indicated by the definition of *range* above, any arbitrary arithmetic expression (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79) is valid. Assuming that ‘x’ is currently bound to 123, the following example will produce the indicated output:

```
FOR EACH i IN {x*10, x*9, x*8, ..., x*3}, {x+50}, {2*x+3, 3*x+2}
TASK 0 OUTPUTS i
  ↪ 1230
  ↪ 1107
  ↪ 984
  ↪ 861
  ↪ 738
  ↪ 615
  ↪ 492
  ↪ 369
  ↪ 173
  ↪ 249
  ↪ 371
```

Finally, FOR EACH loops with constant progressions are executed exactly once. Note that if the *range* does not contain an ellipsis then all values are used, regardless of order or constancy:

```
FOR EACH i IN {4, 4, 4, ..., 4} TASK 0 OUTPUTS i
  ↪ 4
```

```
FOR EACH i IN {4, 4, 4, 4, 4} TASK 0 OUTPUTS i
  ↪ 4
  ↪ 4
  ↪ 4
  ↪ 4
  ↪ 4
```

Timed loops

A *timed loop* is similar to a counted loop (see [Counted loops], page 113) but instead of running for a given number of iterations it runs for a given length of time. Timed loops are absent from all general-purpose programming languages but can be quite useful in the context of network correctness and performance testing. The syntax of CONCEPTUAL's timed-loop construct is “FOR *<expr>* *<time_unit>* *<simple_stmt>*”. *<time_unit>* is unit of time as listed in Section 4.6.3 [Delaying execution], page 107 and *<expr>* specified the number of *<time_unit>*s for which to execute.

The following example shows how to spend three seconds sending messages from task 0 to task 1:

```
FOR 3 SECONDS TASK 0 SENDS A 1 MEGABYTE MESSAGE TO TASK 1
```

Although CONCEPTUAL tries its best to run for exactly the specified length of time there will invariably be some error in the process. Always use `elapsed_usecs` (see Section A.2 [Predeclared variables], page 184) as the indicator of actual time instead of the time requested in the loop.

4.7.3 Binding variables

There are four ways to bind a value to a variable:

1. as part a source or target task description (see Section 4.3 [Task descriptions], page 92)
2. as part of a range loop (see [Range loops], page 113)
3. via a command-line argument (see Section 4.8.2 [Command-line arguments], page 120)
4. explicitly using the LET keyword (this section)

The LET statement has the following form:

```
LET <let_binding>
[AND <let_binding>]*
WHILE <simple_stmt>
```

where *<let_binding>* is defined as follows:

```
<let_binding> ::= <ident>
                BE
                <expr> | A RANDOM TASK [
                                OTHER THAN <expr>
                                | LESS THAN <expr> [BUT NOT <expr>]
                                | GREATER THAN <expr> [BUT NOT
                                <expr>]
                                | IN '[' <expr> ',' <expr> ']' [BUT NOT
                                <expr>]]
```

Here are some examples of LET:

```
LET reps BE 3 WHILE FOR reps REPETITIONS TASK 0 OUTPUTS "Laissez les
bons temps rouler."
```

```
LET src BE num_tasks-1 AND dest BE num_tasks/2 WHILE TASK src SENDS A
55E6 BIT MESSAGE TO TASK dst
```

```
LET hotspot BE A RANDOM TASK WHILE TASKS other SUCH THAT
```

```

other<>hotspot SEND 1000 1 MEGABYTE MESSAGES TO TASK hotspot

LET target BE A RANDOM TASK OTHER THAN 3 WHILE TASK 3 SENDS A 24 BYTE
MESSAGE TO TASK target

LET x BE A RANDOM TASK AND y be a RANDOM TASK GREATER THAN x WHILE
TASK 0 OUTPUTS "Did you know that " AND x AND " is less than " AND y
AND "?"

LET nonzero BE A RANDOM TASK LESS THAN 10 BUT NOT 0 WHILE TASK nonzero
SLEEPS FOR 2 SECONDS

LET middles BE A RANDOM TASK IN [1, num_tasks-2] WHILE TASK middles
ASYNCHRONOUSLY SENDS A 10 BYTE MESSAGE TO TASK ends SUCH THAT ends IS
NOT IN [1, num_tasks-2]

LET num BE 1000 AND num BE num*2 WHILE TASK 0 OUTPUTS num
+ 2000

```

The last example demonstrates that LET can bind a variable to a function of its previous value. It is important to remember, though, that variables in CONCEPTUAL cannot be assigned, only bound to a value for the duration of the current scope. They can, however, be bound to a different value for the duration of a child scope. The following example is an attempt to clarify the distinction between binding and assignment:

```

LET var BE 123 WHILE FOR 5 REPETITIONS LET var BE var+1 WHILE TASK 0
OUTPUTS var
+ 124
+ 124
+ 124
+ 124
+ 124

```

In that example, 'var' is bound to '123' for the scope containing the FOR statement. Then, within the FOR statement, a new scope begins with 'var' being given one plus the value it had in the outer scope, resulting in '124'. If CONCEPTUAL supported assignment instead of variable-binding, the program would have output '124', '125', '126', '127', and '128'. Note that if A RANDOM TASK were used in the example instead of 'var+1', 'var' would get a different value in each iteration.

When a variable is LET-bound to A RANDOM TASK, all tasks agree on the random number. Otherwise, task A might send a message to task B but task B might be expecting to receive the message from task C, thereby leading to a variety of problems. If there are no valid random tasks, as in the following example, A RANDOM TASK will return '-1':

```

LET invalid_task BE A RANDOM TASK GREATER THAN num_tasks WHILE TASK 0
OUTPUTS invalid_task
+ -1

```

Furthermore, the $\langle expr \rangle$ passed to GREATER THAN is bounded from below by '0' and the $\langle expr \rangle$ passed to LESS THAN is bounded from above by 'num_tasks-1'. Hence, the following

CONCEPTUAL statement will always output values less than or equal to ‘num_tasks-1’ (unless num_tasks is greater than 1×10^6 , of course, in which case it will always output values less than 1×10^6):

```
LET valid_task BE A RANDOM TASK LESS THAN 1E6 WHILE TASK valid_task
  OUTPUTS "Hello from " AND valid_task
```

4.7.4 Conditional execution

Like most programming languages, CONCEPTUAL supports conditional code execution:

```
<if_stmt> ::= IF <rel_expr>
            THEN <simple_stmt>
            [OTHERWISE <simple_stmt>]
```

The semantics of an *<if_stmt>* are that if *<rel_expr>* (see [Section 4.2.5 \[Relational expressions\], page 91](#)) is TRUE then the first *<simple_stmt>* is executed. If *<rel_expr>* is FALSE then the second *<simple_stmt>* is executed. One restriction is that *<rel_expr>* must return the same truth value to every task. Consequently, functions that involve task-specific random numbers (see [\[Random-number functions\], page 90](#)) are forbidden within *<rel_expr>*.

The following is an example of an *<if_stmt>*:

```
IF this>that THEN TASK 0 SENDS A 3 KILOBYTE MESSAGE TO TASK this
  OTHERWISE TASK num_tasks-1 SENDS A 4 KILOBYTE MESSAGE TO TASK that
```

4.7.5 Grouping

FOR loops, LET bindings, and IF statements operate on a single *<simple_stmt>* (or two *<simple_stmt>*s in the case of IF...OTHERWISE). Operating on multiple *<simple_stmt>*s—or, more precisely, operating on a single *<complex_stmt>* that may consist of multiple *<simple_stmt>*s—is a simple matter of placing the *<simple_stmt>*s within curly braces. Contrast the following:

```
FOR 3 REPETITIONS TASK 0 OUTPUTS "She loves me." THEN TASK 0 OUTPUTS
  "She loves me not."
  † She loves me.
  † She loves me.
  † She loves me.
  † She loves me not.
```

```
FOR 3 REPETITIONS {TASK 0 OUTPUTS "She loves me." THEN TASK 0 OUTPUTS
  "She loves me not."}
  † She loves me.
  † She loves me not.
  † She loves me.
  † She loves me not.
  † She loves me.
  † She loves me not.
```

In other words, everything between ‘{’ and ‘}’ is treated as if it were a single statement. Hence, the FOR loop applies only to the “She loves me” output in the first statement above, while the FOR loop applies to both “She loves me” and “She loves me not” in the second statement.

Variable scoping is limited to the *<simple_stmt>* in the body of a LET:

```
LET year BE 1984 WHILE LET year BE 2084 WHILE TASK 0 OUTPUTS year THEN
TASK 0 OUTPUTS year
```

error The second ‘year’ is outside the scope of both ‘LET’ statements.

```
LET year BE 1984 WHILE {{LET year BE 2084 WHILE TASK 0 OUTPUTS year}
THEN TASK 0 OUTPUTS year}
+ 2084
+ 1984
```

As indicated by the grammatical rules presented at the beginning of [Section 4.7 \[Complex statements\]](#), page 112, CONCEPTUAL does support empty pairs of curly braces, which represent a statement that does nothing and takes no time to execute. While never strictly needed, ‘{ }’ may be a convenient mechanism for mechanically produced CONCEPTUAL programs.

4.8 Header declarations

CONCEPTUAL programs may contain a header section that precedes the first statement in a CONCEPTUAL program. The header section contains two types of declarations that affect the remainder of the program: language versioning declarations and declarations of command-line options.

4.8.1 Language versioning

Because the CONCEPTUAL language is still under development, the programmer is forewarned that major changes are likely. Changes may prevent old code from compiling or, even worse, may cause old code to produce incorrect results (e.g., if scoping or block structuring are altered). To mitigate future language changes CONCEPTUAL enables programs to specify which version of the language they were written to. The syntax is straightforward:

```
<version_decl> ::= REQUIRE LANGUAGE VERSION <string>
```

The parser issues a warning message if <string> does not exactly match the language version supported by the compiler. If the program successfully compiles after a version-mismatch warning, the programmer should check the output very carefully to ensure that the program behaved as expected.

The current version of the CONCEPTUAL language is ‘1.0’. Note that the language version does not necessarily correspond to the version of the CONCEPTUAL toolset (see [Chapter 3 \[Usage\]](#), page 14) as a whole.

4.8.2 Command-line arguments

CONCEPTUAL makes it easy to declare command-line parameters, although the syntax is a bit verbose:

```
<param_decl> ::= <ident>
                IS <string>
                AND COMES FROM <string> OR <string>
                WITH DEFAULT <expr>
```

<ident> is the CONCEPTUAL variable being declared. The first <string> is a descriptive string that is provided when the user runs the program with `--help` or `-?`. The ‘<string> OR <string>’ terms list the long name for the command-line option, preceded by ‘--’, and the

short (single-character) name, preceded by '-'. Finally, $\langle expr \rangle$ specifies the value that will be assigned to $\langle ident \rangle$ if the command-line option is not used. $\langle expr \rangle$ must be a constant expression and may not utilize any of the random-number functions listed in [Random-number functions], page 90. Note that short names (also long names) must be unique.

For instance, the declaration 'nummsgs IS "Number of messages to send" AND COMES FROM "--messages" OR "-m" WITH DEFAULT 25*4' declares a new CONCEPTUAL variable called 'nummsgs'. 'nummsgs' is given the value '100' ('25*4') by default. However, if the user running the program specifies, for example, `--messages=55` (or, equivalently, `-m 55`), then 'nummsgs' will be given the value '55'. The following is an example of the output that might be produced if the program is run with `--help` or `-?`:

```
Usage: a.out [OPTION...]
  -m, --messages=<number>    Number of messages to send [default: 100]
  -C, --comment=<string>     Additional commentary to write to the log
                              file, @FILE to import commentary from FILE,
                              or !COMMAND to import commentary from COMMAND
                              (may be specified repeatedly)
  -L, --logfile=<string>     Log-file template [default: "a.out-%p.log"]
  -N, --no-trap=<string>     List of signals that should not be trapped
                              [default: ""]

Help options:
  -?, --help                 Show this help message
  --usage                    Display brief usage message
```

The above is only an example. Depending on what libraries were available when the CONCEPTUAL run-time system was configured, the output could be somewhat different. Also, long options may not be supported if a suitable argument-processing library was not available at configuration time. The above example does indicate one way that help strings could be formatted. It also shows that the CONCEPTUAL run-time system reserves some command-line options for its own purposes. Currently, these all use uppercase letters for their short forms so it should be safe for programs to use any lowercase letter.

4.9 Complete programs

A complete CONCEPTUAL program consists of zero or more header declarations (see Section 4.8 [Header declarations], page 120), each terminated with a '.', followed by one or more complex statements (see Section 4.7 [Complex statements], page 112), each also terminated with a '.'. More formally, CONCEPTUAL's top-level nonterminal is the $\langle program \rangle$:

$$\langle program \rangle ::= ((\langle version_decl \rangle \mid \langle param_decl \rangle) [\cdot \cdot])^* (\langle top_level_complex_stmt \rangle [\cdot \cdot])^+$$

$$\langle top_level_complex_stmt \rangle ::= \langle complex_stmt \rangle$$

Because a $\langle complex_stmt \rangle$ can reduce to a $\langle simple_stmt \rangle$, the most basic, complete CONCEPTUAL program would be a $\langle simple_stmt \rangle$ with a terminating period:

```
ALL TASKS self OUTPUT "Hello from task " AND self AND "!".
```

A fuller example might contain multiple header declarations and multiple $\langle complex_stmt \rangle$ s:

```

# A complete coNcEPTuaL program
# By Scott Pakin <pakin@lanl.gov>

REQUIRE LANGUAGE VERSION "1.0".

maxval IS "Maximum value to loop to" AND COMES FROM "--maxval" OR
"-v" WITH DEFAULT 100.

step IS "Increment after each iteration" AND COMES FROM "--step" OR
"-s" WITH DEFAULT 1.

TASK 0 OUTPUTS "Looping from 0 to " AND maxval AND " by " AND step
AND "...".

FOR EACH loopvar IN {0, step, ..., maxval}
  TASK 0 OUTPUTS "    " AND loopvar.

TASK 0 OUTPUTS "Wasn't that fun?".

```

Technically, the ‘.’ is optional; the language is unambiguous without it. However, for aesthetic purposes it is recommended that you terminate sentences with a period, just like in a natural language. An exception would be when a *<complex_stmt>* ends with a curly brace. The ‘}.’ syntax is unappealing so a simple ‘}’ should be used instead. See [Chapter 5 \[Examples\]](#), [page 129](#), for further examples.

Top-level statements and log files

The reason that CONCEPTUAL distinguishes between *<top_level_complex_stmt>*s and *<complex_stmt>*s is that *<top_level_complex_stmt>*s begin a new table in the log file (see [Section 4.5.3 \[Writing to a log file\]](#), [page 104](#)) while *<complex_stmt>*s add columns to the current table. Consider the following piece of code:

```

TASK 0 LOGS 111 AS "First" AND
          222 AS "Second".

```

Because ‘First’ and ‘Second’ are logged within the same *<simple_stmt>* they appear in the log file within the same table but as separate columns:

```

"First","Second"
"(all data)","(all data)"
111,222

```

The same rule holds when LOGS is used repeatedly across *<simple_stmt>*s but within the same *<complex_stmt>*:

```

TASK 0 LOGS 111 AS "First" THEN
TASK 0 LOGS 222 AS "Second".

```

However, if ‘First’ and ‘Second’ are logged from separate *<top_level_complex_stmt>*s, the CONCEPTUAL run-time library stores them in separate tables:


```
TASK 0 LOGS 111 AS "First".
TASK 0 LOGS 222 AS "Second".
```

```
"First"
"(all data)"
111

"Second"
"(all data)"
222
```

4.10 Summary of the grammar

The following is the complete grammar for the CONCEPTUAL language. The EBNF productions appear here in the order that they were presented in the rest of the chapter.

```
<expr> ::= <cond_expr>
<cond_expr> ::= <add_expr> IF <rel_expr> OTHERWISE <add_expr>
<add_expr> ::= <mult_expr>
| <add_expr> '+' <mult_expr>
| <add_expr> '-' <mult_expr>
| <add_expr> '|' <mult_expr>
| <add_expr> XOR <mult_expr>
<mult_expr> ::= <unary_expr>
| <mult_expr> '*' <unary_expr>
| <mult_expr> '/' <unary_expr>
| <mult_expr> MOD <unary_expr>
| <mult_expr> '>>' <unary_expr>
| <mult_expr> '<<' <unary_expr>
| <mult_expr> '&' <unary_expr>
<power_expr> ::= <primary_expr> ['**' <unary_expr>]
<unary_expr> ::= <power_expr>
| <unary_operator> <unary_expr>
<unary_operator> ::= '+' | '-' | NOT
<primary_expr> ::= '(' <expr> ')'
| <ident>
| <integer>
| <func_name> '(' <enumerated_exprs> ')'
| REAL '(' <expr> ')'
<enumerated_exprs> ::= <expr> [',' <expr>]*
<func_name> ::= ABS | BITS | CBRT | FACTOR10 | LOG10 | MAX | MIN | ROOT | SQRT
| CEILING | FLOOR | ROUND
| TREE_PARENT | TREE_CHILD
| KNOMIAL_PARENT | KNOMIAL_CHILD | KNOMIAL_CHILDREN
| MESH_NEIGHBOR | MESH_COORDINATE
```

```

| TORUS_NEIGHBOR | TORUS_COORDINATE
| RANDOM_UNIFORM | RANDOM_GAUSSIAN | RANDOM_POISSON
⟨aggr_expr⟩ ::= [EACH] ⟨expr⟩
| THE ⟨expr⟩
| THE ⟨aggr_func⟩ [OF [THE]] ⟨expr⟩
| A HISTOGRAM OF [THE] ⟨expr⟩
⟨aggr_func⟩ ::= [ARITHMETIC] MEAN | HARMONIC MEAN | GEOMETRIC MEAN | MEDIAN |
STANDARD DEVIATION | VARIANCE | SUM | MINIMUM | MAXIMUM | FINAL
⟨rel_expr⟩ ::= ⟨rel_disj_expr⟩
⟨rel_disj_expr⟩ ::= [⟨rel_disj_expr⟩ '\|'] ⟨rel_conj_expr⟩
⟨rel_conj_expr⟩ ::= [⟨rel_conj_expr⟩ '/\'] ⟨rel_primary_expr⟩
⟨rel_primary_expr⟩ ::= ⟨eq_expr⟩
| '(' ⟨rel_expr⟩ ')'
⟨eq_expr⟩ ::= ⟨expr⟩ '=' ⟨expr⟩
| ⟨expr⟩ '<' ⟨expr⟩
| ⟨expr⟩ '>' ⟨expr⟩
| ⟨expr⟩ '<=' ⟨expr⟩
| ⟨expr⟩ '>=' ⟨expr⟩
| ⟨expr⟩ '<>' ⟨expr⟩
| ⟨expr⟩ DIVIDES ⟨expr⟩
| ⟨expr⟩ IS EVEN
| ⟨expr⟩ IS ODD
| ⟨expr⟩ IS IN '[' ⟨expr⟩ ',' ⟨expr⟩ ']
| ⟨expr⟩ IS NOT IN '[' ⟨expr⟩ ',' ⟨expr⟩ ']
⟨restricted_ident⟩ ::= ⟨ident⟩ SUCH THAT ⟨rel_expr⟩
⟨source_task⟩ ::= ALL TASKS
| ALL TASKS ⟨ident⟩
| TASK ⟨expr⟩
| TASKS ⟨restricted_ident⟩
⟨target_tasks⟩ ::= ALL OTHER TASKS
| TASK ⟨expr⟩
| TASKS ⟨restricted_ident⟩
⟨message_spec⟩ ::= ⟨item_count⟩
[NONUNIQUE | UNIQUE]
⟨item_size⟩
[UNALIGNED |
⟨message_alignment⟩ ALIGNED |
⟨message_alignment⟩ MISALIGNED]
MESSAGES
[WITH VERIFICATION | WITH DATA TOUCHING |
WITHOUT VERIFICATION | WITHOUT DATA TOUCHING]
[FROM BUFFER ⟨expr⟩ |
FROM THE DEFAULT BUFFER]
⟨recv_message_spec⟩ ::= [SYNCHRONOUSLY | ASYNCHRONOUSLY]

```

```

[AS [A | AN]
[NONUNIQUE | UNIQUE]
[UNALIGNED |
<message_alignment> ALIGNED |
<message_alignment> MISALIGNED]
MESSAGES]
[WITH VERIFICATION | WITH DATA TOUCHING |
WITHOUT VERIFICATION | WITHOUT DATA TOUCHING]
[FROM BUFFER <expr> |
FROM THE DEFAULT BUFFER]
<reduce_message_spec> ::= <item_count>
[NONUNIQUE | UNIQUE]
[UNALIGNED |
<message_alignment> ALIGNED |
<message_alignment> MISALIGNED]
INTEGERS | DOUBLEWORDS
[WITH DATA TOUCHING | WITHOUT DATA TOUCHING]
[FROM BUFFER <expr> |
FROM THE DEFAULT BUFFER]
<reduce_target_message_spec> ::= [AS <item_count>
[NONUNIQUE | UNIQUE]
[<message_alignment> ALIGNED |
<message_alignment> MISALIGNED]
INTEGERS | DOUBLEWORDS]
[WITH DATA TOUCHING | WITHOUT DATA TOUCHING]
[INTO BUFFER <expr> |
INTO THE DEFAULT BUFFER]
<item_count> ::= A | AN | <expr>
<item_size> ::= ε
| <expr> <data_multiplier>
| <data_type> SIZED
<data_multiplier> ::= BIT | BYTE | HALFWORD | WORD | INTEGER | DOUBLEWORD |
QUADWORD | PAGE | KILOBYTE | MEGABYTE | GIGABYTE
<data_type> ::= BYTE | HALFWORD | WORD | INTEGER | DOUBLEWORD |
QUADWORD | PAGE
<message_alignment> ::= <data_type>
| <expr> <data_multiplier>
<send_stmt> ::= <source_task>
[ASYNCHRONOUSLY] SENDS
<message_spec>
TO [UNSUSPECTING] <target_tasks>
| <source_task>
[ASYNCHRONOUSLY] SENDS
<message_spec>
TO <target_tasks>

```

```

                                WHO RECEIVE IT
                                <recv_message_spec>
<receive_stmt> ::= <target_tasks>
                    [ASYNCHRONOUSLY] RECEIVE
                    <message_spec>
                    FROM <source_task>

<wait_stmt> ::= <source_task>
                AWAITS COMPLETION

<mcast_stmt> ::= <source_task>
                [ASYNCHRONOUSLY] MULTICASTS
                <message_spec>
                TO <target_tasks>

<reduce_stmt> ::= <source_task>
                REDUCES
                <reduce_message_spec>
                TO <source_task>
                [WHO RECEIVES THE RESULT <reduce_target_message_spec>]
                |
                <source_task>
                REDUCES
                <reduce_message_spec>
                [TO <reduce_message_spec>]

<sync_stmt> ::= <source_task>
                SYNCHRONIZES

<string_or_log_comment> ::= <string>
                            | THE VALUE OF <string>

<output_stmt> ::= <source_task>
                OUTPUTS
                <expr> | <string_or_log_comment>
                [AND <expr> | <string_or_log_comment>]*

<log_stmt> ::= <source_task>
                LOGS
                <aggr_expr> AS <string_or_log_comment>
                [AND <aggr_expr> AS <string_or_log_comment>]*

<flush_stmt> ::= <source_task>
                COMPUTES AGGREGATES

<reset_stmt> ::= <source_task>
                RESETS ITS COUNTERS

<assert_stmt> ::= ASSERT THAT <string>
                WITH <rel_expr>

<delay_stmt> ::= <source_task>
                SLEEPS | COMPUTES
                FOR <expr> <time_unit>

<time_unit> ::= MICROSECONDS | MILLISECONDS | SECONDS | MINUTES | HOURS | DAYS

<touch_stmt> ::= <source_task>

```

```

    TOUCHES
    [⟨expr⟩ ⟨data-type⟩ OF]
    AN ⟨item-size⟩ MEMORY REGION
    [⟨expr⟩ TIMES]
    [WITH STRIDE ⟨expr⟩ ⟨data-type⟩ | WITH RANDOM STRIDE]
⟨touch_buffer_stmt⟩ ::=  ⟨source_task⟩
    TOUCHES
        ALL MESSAGE BUFFERS
        | MESSAGE BUFFER ⟨expr⟩
        | THE CURRENT MESSAGE BUFFER

⟨processor_stmt⟩ ::=  ⟨source_task⟩
    IS ASSIGNED TO
    PROCESSOR ⟨expr⟩ | A RANDOM PROCESSOR

⟨backend_stmt⟩ ::=  ⟨source_task⟩
    BACKEND EXECUTES
    ⟨expr⟩ | ⟨string⟩
    [AND ⟨expr⟩ | ⟨string⟩]*

⟨complex_stmt⟩ ::=  ⟨simple_stmt⟩ [THEN ⟨complex_stmt⟩]
⟨simple_stmt⟩ ::=  FOR ⟨expr⟩ REPETITIONS [PLUS ⟨expr⟩ WARMUP REPETITIONS
    [AND A SYNCHRONIZATION]] ⟨simple_stmt⟩
    | FOR EACH ⟨ident⟩ IN ⟨range⟩ [‘,’ ⟨range⟩]* ⟨simple_stmt⟩
    | FOR ⟨expr⟩ ⟨time_unit⟩ ⟨simple_stmt⟩
    | LET ⟨let_binding⟩ [AND ⟨let_binding⟩]* WHILE ⟨simple_stmt⟩
    | IF ⟨rel_expr⟩ THEN ⟨simple_stmt⟩ [OTHERWISE ⟨simple_stmt⟩]
    | ‘{’ [(⟨complex_stmt⟩)] ‘}’
    | ⟨send_stmt⟩
    | ⟨receive_stmt⟩
    | ⟨wait_stmt⟩
    | ⟨mcast_stmt⟩
    | ⟨reduce_stmt⟩
    | ⟨sync_stmt⟩
    | ⟨output_stmt⟩
    | ⟨log_stmt⟩
    | ⟨flush_stmt⟩
    | ⟨reset_stmt⟩
    | ⟨assert_stmt⟩
    | ⟨delay_stmt⟩
    | ⟨touch_stmt⟩
    | ⟨touch_buffer_stmt⟩
    | ⟨processor_stmt⟩
    | ⟨backend_stmt⟩

⟨range⟩ ::=  ‘{’
    ⟨expr⟩ [‘,’ ⟨expr⟩]*
    [‘,’ ... ,’ ⟨expr⟩]
    ‘}’

```

```

<let_binding> ::= <ident>
                BE
                <expr> | A RANDOM TASK [
                                OTHER THAN <expr>
                                | LESS THAN <expr> [BUT NOT <expr>]
                                | GREATER THAN <expr> [BUT NOT
                                <expr>]
                                | IN '[' <expr> ',' <expr> ']' [BUT NOT
                                <expr>]]

<if_stmt> ::= IF <rel_expr>
              THEN <simple_stmt>
              [OTHERWISE <simple_stmt>]

<version_decl> ::= REQUIRE LANGUAGE VERSION <string>

<param_decl> ::= <ident>
                IS <string>
                AND COMES FROM <string> OR <string>
                WITH DEFAULT <expr>

<program> ::= (<version_decl> | <param_decl> ['.'])*
             (<top_level_complex_stmt> ['.'])+

<top_level_complex_stmt> ::= <complex_stmt>

```

The primitives *<ident>*, *<string>*, and *<integer>* are described in [Section 4.1 \[Primitives\]](#), page 77.

5 Examples

This chapter presents a variety of examples of complete `CONCEPTUAL` programs. The purpose is to put in context the grammatical elements described in [Chapter 4 \[Grammar\]](#), [page 77](#) and also to illustrate `CONCEPTUAL`'s power and expressiveness.

5.1 Latency

One of the most common network performance benchmarks is a ping-pong latency test. Not surprisingly, such a test is straightforward to implement in `CONCEPTUAL`:

```
# A ping-pong latency test written in coNcEPTuaL

Require language version "1.0".

# Parse the command line.
reps is "Number of repetitions of each message size" and comes from
"--reps" or "-r" with default 1000.
maxbytes is "Maximum number of bytes to transmit" and comes from
"--maxbytes" or "-m" with default 1M.

# Ensure the we have a peer with whom to communicate.
Assert that "the latency test requires at least two tasks" with
  num_tasks>=2.

# Perform the benchmark.
For each msgsize in {0}, {1, 2, 4, ..., maxbytes} {
  for reps repetitions {
    task 0 resets its counters then
    task 0 sends a msgsize byte message to task 1 then
    task 1 sends a msgsize byte message to task 0 then
    task 0 logs the msgsize as "Bytes" and
      the median of elapsed_usecs/2 as "1/2 RTT (usecs)"
  } then
  task 0 computes aggregates
}
```

Note that the outer `FOR` loop specifies two *range*s (see [\[Range loops\]](#), [page 113](#)). This is because `{0, 1, 2, 4, ..., maxbytes}` is not a geometric progression. Hence, that incorrect *range* is split into the singleton `{0}` and the geometric progression `{1, 2, 4, ..., maxbytes}`.

5.2 Hot potato

One way to measure performance variance on a parallel system is with a “hot potato” test. The idea is that the tasks send a message in a ring pattern, then the first task logs the minimum, mean, and variance of the per-hop latency. Ideally, the minimum should

equal the mean and these should both maintain a constant value as the number of tasks increases. Also, the variance should be small and constant. The following CONCEPTUAL code implements a hot-potato test.

```
# Virtual ring "hot potato" test

Require language version "1.0".

trials is "Number of trials to perform" and comes from "--trials" or
"-t" with default 100000.

Assert that "the hot-potato test requires at least two tasks" with
  num_tasks>=2.

Let len be 0 while {
  for each task_count in {2, ..., num_tasks} {
    task 0 outputs "Performing " and trials and " " and
      task_count and "-task runs..." then
    for trials repetitions plus 5 warmup repetitions {
      task 0 resets its counters then
      task 0 sends a len byte message to unsuspecting task 1 then
      task (n+1) mod task_count receives a len byte message from task
        n such that n<task_count then
      task n such that n>0 /\ n<task_count sends a len byte message
        to unsuspecting task (n+1) mod task_count then
      task 0 logs the task_count as "# of tasks" and
        the minimum of elapsed_usecs/task_count as
          "Latency (usecs)" and
        the mean of elapsed_usecs/task_count as
          "Latency (usecs)" and
        the variance of elapsed_usecs/task_count as
          "Latency (usecs)"
    } then
    task 0 computes aggregates
  }
}
```

All tasks receive from their left neighbor and send to their right neighbor. However, in order to avoid a deadlock situation, task 0 sends then receives while all of the other tasks receive then send.

5.3 Hot spot

Different systems react differently to resource contention. A hot-spot test attempts to measure the performance degradation that occurs when a task is flooded with data. That is, all tasks except 0 concurrently send a batch of messages to task 0. Task 0 reports the incoming bandwidth, i.e., the number of bytes it received divided by the time it took to

receive that many bytes. The two independent variables are the message size and the number of tasks.

```
# Hot-spot bandwidth

Require language version "1.0".

maxbytes is "Maximum message size in bytes" and comes from
  "--maxbytes" or "-x" with default 1024.
numtrials is "Number of bursts of each size" and comes from "--trials"
  or "-t" with default 100.
burst is "Number of messages in each burst" and comes from
  "--burstsize" or "-b" with default 1000.

Assert that "the hot-spot test requires at least two tasks" with
  num_tasks>=2.

For each maxtask in {2, ..., num_tasks}
  for each msgsize in {1, 2, 4, ..., maxbytes} {
    task 0 outputs "Performing " and numtrials and " " and
      maxtask and "-task trials with " and
      msgsize and "-byte messages" then
      for numtrials repetitions plus 3 warmup repetitions {
        task 0 resets its counters then
        task sender such that sender>0 /\ sender<maxtask asynchronously
          sends burst msgsize byte messages to task 0 then
        all tasks await completion then
        task 0 logs the maxtask as "Tasks" and
          the msgsize as "Message size (B)" and
          the mean of (1E6*bytes_received)/(1M*elapsed_usec)
          as "Incoming BW (MB/s)"
      } then
    task 0 computes aggregates
  }
}
```

5.4 Multicast trees

It may be worth comparing the performance of a native multicast operation to the performance achieved by multicasting over a k -nomial tree to gauge how well the underlying communication layer implements multicasts. The following code records a wealth of data, varying the tree arity (i.e., k), the number of tasks receiving the multicast, and the message size. It provides a good demonstration of how to use the `KNOMIAL_CHILDREN` and `KNOMIAL_CHILD` functions.

```

# Test the performance of multicasting over various k-nomial trees
# By Scott Pakin <pakin@lanl.gov>

Require language version "1.0".

# Parse the command line.
minsize is "Min. message size (bytes)" and comes from "--minbytes" or "-n"
  with default 1.
maxsize is "Max. message size (bytes)" and comes from "--maxbytes" or "-x"
  with default 1M.
reps is "Repetitions to perform" and comes from "--reps" or "-r" with
  default 100.
maxarity is "Max. arity of the tree" and comes from "--maxarity" or "-a"
  with default 2.

Assert that "this program requires at least two processors" with
  num_tasks>=2.

# Send messages from task 0 to 1, 2, 3, ... other tasks in a k-nomial tree.
For each arity in {2, ..., maxarity} {
  for each num_targets in {1, ..., num_tasks-1} {
    for each msgsize in {minsize, minsize*2, minsize*4, ..., maxsize} {
      task 0 outputs "Multicasting a " and msgsize and "-byte message to "
        and num_targets and " target(s) over a " and arity and
        "-nomial tree ..." then
      for reps repetitions {
        task 0 resets its counters then
        for each src in {0, ..., num_tasks}
          for each dstnum in {0, ..., knomial_children(src, arity,
            num_targets+1)}
            task src sends a msgsize byte message to task
              knomial_child(src, dstnum, arity) then
        all tasks synchronize then
        task 0 logs the arity as "k-nomial arity" and
          the num_targets as "# of recipients" and
          the msgsize as "Message size (bytes)" and
          the median of (1E6/1M)*(msgsize/elapsed_usecs) as
            "Incoming bandwidth (MB/s)" and
          the median of (num_targets*msgsize/elapsed_usecs)*
            (1E6/1M) as "Outgoing bandwidth (MB/s)"
      } then
      task 0 computes aggregates
    }
  }
}

```

5.5 Calling MPI functions

The CONCEPTUAL language is designed to be highly portable. Any CONCEPTUAL program can be compiled using any of the backends listed in [Section 3.3 \[Supplied backends\]](#), [page 20](#). A consequence of this portability is that CONCEPTUAL does not include primitives that are specific to any particular target language or communication library.

The BACKEND EXECUTES statement (see [Section 4.6.6 \[Injecting arbitrary code\]](#), [page 111](#)) gives a programmer the ability to sacrifice portability for the ability to measure the performance of features provided by a specific target language or communication library. Hence, it is possible to write the core parts of a benchmark in a lower-level language while letting CONCEPTUAL handle the setup, measurement, logging, and other mundane operations.

The following program uses BACKEND EXECUTES to measure the performance of the `MPI_Allgather()` function provided by an MPI library. Because it utilizes C code to call an MPI function, the code builds only with the `c_mpi` backend.

```
# Measure the performance of MPI_Allgather()
# By Scott Pakin <pakin@lanl.gov>
#
# N.B. Requires the c_mpi backend.

Require language version "1.0".

# Parse the command line.
numwords is "Message size (words)" and comes from "--msgsize" or "-s" with
default 1.

# Allocate a send buffer and a receive buffer.
Task 0 multicasts a numwords-word message from buffer 0 to all other tasks.
Task 0 multicasts a numwords*num_tasks word message from buffer 1 to all
other tasks.

# Measure the performance of MPI_Allgather().
Task 0 resets its counters then
for 100 repetitions plus 3 warmup repetitions
  all tasks backend execute "
    MPI_Allgather([MESSAGE BUFFER 0], (int)" and numwords and ", MPI_INT,
                  [MESSAGE BUFFER 1], (int)" and numwords and ", MPI_INT,
                  MPI_COMM_WORLD);
  " then
task 0 logs elapsed_usecs/100 as "Gather time (us)".
```

The preceding code demonstrates a few useful techniques:

- A pair of MULTICASTS statements (see [Section 4.4.5 \[Multicasting\]](#), [page 101](#)) are used to allocate the two message buffers.
- The BACKEND EXECUTE statement uses a value provided on the command line (`numwords`, via `--msgsize`) in the call to `MPI_Allgather()`.

- Pointers to the two message buffers are passed to `MPI_Allgather()` using the '[MESSAGE BUFFER *expr*]' substitution described in [Section 4.6.6 \[Injecting arbitrary code\]](#), page 111.

6 Implementation

CONCEPTUAL could have been implemented as a benchmarking library instead of as a special-purpose language. In addition to improved readability and the practicality of including entire source programs in every log file, one advantage of the language approach is that the same CONCEPTUAL source code can be used to compare the performance of multiple communication libraries. A compiler command-line option selects a particular backend module to use to generate code. Each backend outputs code for a different combination of low-level language and communication library.

The CONCEPTUAL compiler is structured into a pipeline of modules. Thus, the backend can be replaced without altering the front end, lexer, or parser modules. CONCEPTUAL ensures consistency across backends by providing a run-time library that generated code can link to. The run-time library encapsulates many of the mundane tasks a network correctness or performance test needs to perform.

6.1 Overview

Compiler

The CONCEPTUAL compiler is written in Python and is based on the PLY (Python Lex-Yacc) compiler framework. Compiler execution follows a basic pipeline structure. Compilation starts with the top-level file (`ncpt1.py`), which processes the command line then transfers control to the lexer (`ncpt1_lexer.py`). The lexer inputs CONCEPTUAL source code and outputs a stream of tokens (`ncpt1_token.py`). Next, the parser (`ncpt1_parser.py`) finds structure in those tokens based on CONCEPTUAL's grammatical rules and outputs an abstract syntax tree (`ncpt1_ast.py`). Finally, the code generator (`codegen_language_library.py`) that was designated on the command line walks the abstract syntax tree, converting it to code in the target language and for the target communication library.

Run-time library

CONCEPTUAL makes a large run-time library (`runtime.lib.c`) available to generated programs. The CONCEPTUAL run-time library, which is written in C, provides consistent functionality across target languages and communication layers as well as across hardware architectures and operating systems. The library also simplifies code generation by implementing functions for such tasks as memory allocation, queue management, and data logging. The functions in this library are described in [Section 6.3 \[Run-time library functions\]](#), page 145.

Build process

CONCEPTUAL is built using the GNU Autotools (Autoconf, Automake, and Libtool). Consequently, changes should be made to original files, not generated files. Specifically, `configure.ac` and `acinclude.m4` should be edited in place of `configure`; `ncpt1.h.in` should be edited in place of `ncpt1.h`; and, the various `Makefile.am` files should be edited in place of the corresponding `Makefile`'s. See the *Autoconf* documentation, the *Automake* documentation, and the *Libtool* documentation for information about how these various tools operate.

If `configure` is given the `--enable-maintainer-mode` option, `make` will automatically re-run `aclocal`, `autoheader`, `automake`, `autoconf`, and/or `./configure` as needed. Developers who plan to modify any of the “maintainer” files (`acinclude.m4`, `configure.ac`, and the various `Makefile.am` files) are strongly encouraged to configure `CONCEPTUAL` with `--enable-maintainer-mode` in order to ensure that the build process is kept current with any changes.

6.2 Backend creation

The `CONCEPTUAL` compiler’s backend generates code from an abstract syntax tree (AST). The compiler was designed to support a variety of code generators, each targeting a particular programming language and communication library. There are two ways to create a new `CONCEPTUAL` backend. Either a `codegen_language_library.py` backend supporting an arbitrary language and communication library can be written from scratch or a C-based `codegen_c_library.py` backend can be derived from `codegen_c_generic.py`.

In the former case, the backend must define an `NCPTL_CodeGen` class. `NCPTL_CodeGen` class must contain a `generate` method with the following signature:

```
def generate(self, ast, filesource='<stdin>', filetarget='',
            sourcecode=None):
```

That is, `generate` takes as arguments a class object, the root of an abstract-syntax tree (as defined in `ncptl_ast.py`), the name of the input file containing `CONCEPTUAL` code (to be used for outputting error messages), the name of the output file to produce, and the complete `CONCEPTUAL` source code (which is both stored in prologue comments and passed to the run-time library). `generate` should invoke `self.postorder_traversal` to traverse the AST in a post-order fashion, calling various code-generating methods as it proceeds. The `NCPTL_CodeGen` must implement all of the methods listed in [Section B.1 \[Method calls\]](#), [page 185](#), each of which corresponds to some component of the `CONCEPTUAL` grammar. Each method takes a “self” class object and a node of the AST (of type `AST`).

The compiler front-end, `ncptl`, invokes the following two methods, which must be defined by the backend’s `NCPTL_CodeGen` class:

```
def compile_only(self, progfilename, codelines, outfilename,
                verbose, keepints):
```

```
def compile_and_link(self, progfilename, codelines, outfilename,
                    verbose, keepints):
```

The `compile_only` method compiles the backend-specific code into an object file. The `compile_and_link` method compiles the backend-specific code into an object file and links it into an executable file. For some backends, the notions of “compile” and “link” may not be appropriate. In that situation, the backend should perform the closest meaningful operations. For example, the `dot_ast` backend (see [Section 3.3.10 \[The dot_ast backend\]](#), [page 38](#)) compiles to a `.dot` file and links into the target graphics format (`.ps` by default).

For both the `compile_only` and `compile_and_link` methods, `progfilename` is the name of the CONCEPTUAL input file specified on the ‘`ncptl`’ command line or the string ‘`<command line>`’ if a program was specified literally with `--program`. `codelines` is the output from the `generate` method, i.e., a list of lines of backend-specific code. `outfile` is the name of the target file specified on the ‘`ncptl`’ command line with `--output` or the string ‘`-`’ if `--output` was not used. If `verbose` is ‘`1`’, the method should write each operation it plans to perform to the standard-error device. For consistency, comment lines should begin with ‘`#`’; shell commands should be output verbatim. If `verbose` is ‘`0`’, corresponding to the ‘`ncptl`’ `--quiet` option, the method should output nothing but error messages. Finally, `keepints` corresponds to the `--keep-ints` option to ‘`ncptl`’. If equal to ‘`0`’, all intermediate files should be deleted before returning; if equal to ‘`1`’, intermediate files should be preserved. See [Section 3.2 \[Compiling coNCePTuaL programs\]](#), page 18, for a description of the various command-line options to ‘`ncptl`’.

As long as NCPTL_CodeGen implements all of the required functions it is free to generate code in any way that it sees fit. However, [Section B.1 \[Method calls\]](#), page 185, lists a large number of methods, many of which will be identical across multiple code generators for the same language but different communication libraries. To simplify a common case, C plus some messaging library, CONCEPTUAL provides ‘`codegen_c_generic.py`’, to which the remainder of the Implementation chapter is devoted to explaining.

6.2.1 Hook methods

Multiple code generators for the same language but different communication libraries are apt to contain much code in common. Because C is a popular language, CONCEPTUAL provides a ‘`codegen_c_generic.py`’ module which implements a virtual NCPTL_CodeGen base class. This base class implements all of the methods listed in [Section B.1 \[Method calls\]](#), page 185. However, rather than support a particular communication library, the ‘`codegen_c_generic.py`’ implementation of NCPTL_CodeGen merely provides a number of calls to “hook” methods—placeholders that implement library-specific functionality. See [Section B.2 \[C hooks\]](#), page 187, for a list of all of the hooks that ‘`codegen_c_generic.py`’ defines. For clarity, hooks are named after the method from which they’re called but with an all-uppercase tag appended. Hook methods take a single parameter, a read-only dictionary (the result of invoking Python’s `locals()` function) of all of the local variables in the caller’s scope. They return C code in the form of a list with one line of C per element. A hook method is invoked only if it exists, which gives the backend developer some flexibility in selecting places at which to insert code. Of course, for coarser-grained control, a backend developer can override complete methods in ‘`codegen_c_generic.py`’ if desired. Generally, this will not be necessary as hook invocations are scattered liberally throughout the file.

An example

‘`codegen_c_generic.py`’ defines a method named `code_specify_include_files`. (‘`codegen_c_generic.py`’ names all of its code-generating helper methods ‘`code_something`’.) `code_specify_include_files` pushes a sequence of `#include` directives onto a queue of lines of C code. The method is shown below in its entirety:

```

def code_specify_include_files(self, node):
    "Load all of the C header files the generated code may need."

    # Output a section comment.
    self.pushmany([
        "/*****",
        " * Include files *",
        " *****/",
        ""])

    # Enable hooks both before and after the common includes.
    self.pushmany(self.invoke_hook("code_specify_include_files_PRE",
                                   locals(),
                                   before=[
                                       "/* Header files specific to the %s backend */" %
                                       self.backend_name],
                                   after=[""]))

    self.pushmany([
        "/* Header files needed by all C-based backends */",
        "#include <stdio.h>",
        "#include <string.h>",
        "#include <ncptl/ncptl.h>"])
    self.pushmany(self.invoke_hook("code_specify_include_files_POST",
                                   locals(),
                                   before=[
                                       "",
                                       "/* Header files specific to the %s backend */" %
                                       self.backend_name]))

```

`code_specify_include_files` uses the `pushmany` method (see [Section 6.2.4 \[Internals\]](#), page 142) to push each element in a list of lines of C code onto the output queue. It starts by pushing a section comment—‘`codegen_c_generic.py`’ outputs fully commented C code. Next, it invokes the `code_specify_include_files_PRE` hook if it exists and pushes that method’s return value onto the queue. Then, it pushes all of the `#includes` needed by the generated C code. Finally, it invokes the `code_specify_include_files_POST` hook if it exists and pushes that method’s return value onto the queue.

A backend that requires additional header files from those included by `code_specify_include_files` need only define `code_specify_include_files_PRE` to add extra header files before the standard ones or `code_specify_include_files_POST` to add extra header files after them. The following is a sample (hypothetical) hook definition:

```

def code_specify_include_files_POST(self, localvars):
    "Specify extra header files needed by the c_threads backend."
    return [
        "#include <errno.h>",
        "#include <pthread.h>"]

```


Although the top-level structure of `codegen_c_generic.py` is described in [Section 6.2.4 \[Internals\]](#), [page 142](#), a backend developer will normally need to study the `codegen_c_generic.py` source code to discern the purpose of each hook method and its relation to the surrounding code.

6.2.2 A minimal C-based backend

A backend derived from `codegen_c_generic.py` starts by defining an `NCPTL_CodeGen` child class that inherits much of its functionality from the parent `NCPTL_CodeGen` class. There are only two items that a C-based backend *must* define: `backend_name`, the name of the backend in the form `'c_library'`; and, `backend_desc`, a brief phrase describing the backend. (These are used for error messages and file comments.) Also, a backend's `__init__` method must accept an `options` parameter, which is given a list of command-line parameters not recognized by `ncptl.py`. After `NCPTL_CodeGen`'s `__init__` method processes the entries in `options` that it recognizes, it should pass the remaining options to its parent class's `__init__` method for further processing. (For proper initialization, the parent class's `__init__` method must be called, even if there are no remaining options to process.)

The following is the complete source code to a minimal `CONCEPTUAL` backend. This backend, `codegen_c_seq.py`, supports only sequential `CONCEPTUAL` programs (e.g., `TASK 0 OUTPUTS "Hello, world!"`); any attempt to use communication statements (see [Section 4.4 \[Communication statements\]](#), [page 94](#)) will result in a compile-time error.

```

#!/usr/bin/env python

#####
# Code generation module for the coNCePTuaL language: #
# Minimal C-based backend -- all communication      #
# operations result in a compiler error            #
#                                                    #
# By Scott Pakin <pakin@lanl.gov>                  #
#####

import codegen_c_generic

class NCPTL_CodeGen(codegen_c_generic.NCPTL_CodeGen):
    def __init__(self, options):
        "Initialize the sequential C code generation module."
        self.backend_name = "c_seq"
        self.backend_desc = "C, sequential code only"
        codegen_c_generic.NCPTL_CodeGen.__init__(self, options)

        # We don't have our own command-line options but we handle
        # --help, nevertheless.
        for arg in range(0, len(options)):
            if options[arg] == "--help":
                # Output a help message.
                self.show_help()
                raise SystemExit, 0

```

The `c_seq` backend can be used like any other:

```

ncptl --backend=c_seq \
  --program='For each i in {10, ..., 1} task 0 outputs i.' | \
  indent > myprogram.c

```

(‘`codegen_c_generic.py`’ outputs unindented code, deferring attractive formatting to the Unix ‘`indent`’ utility.)

One sequential construct the `c_seq` backend does not support is randomness, as needed by `A RANDOM PROCESSOR` (see [Section 4.6.5 \[Reordering task IDs\]](#), page 110) and `A RANDOM TASK` (see [Section 4.7.3 \[Binding variables\]](#), page 117). ‘`codegen_c_generic.py`’ cannot support randomness itself because doing so requires broadcasting the seed for the random-number generator to all tasks. Broadcasting requires messaging-layer support, which a derived backend provides through the `code_def_init_reseed_BCAST` hook (see [Section 6.2.1 \[Hook methods\]](#), page 137). For the sequential backend presented above, a broadcast can be implemented as a no-op:

```
def code_def_init_reseed_BCAST(self, localvars):
    'Broadcast" a random-number seed to all tasks.'
    return []
```

In fact, that same do-nothing hook method is used by the `c_udgram` backend. `c_udgram` seeds the random-number generator before calling `fork()`, thereby ensuring that all tasks have the same seed without requiring an explicit broadcast.

6.2.3 Generated code

‘`codegen_c_generic.py`’ generates thoroughly commented C code. However, the overall structure of the generated code may be somewhat unintuitive, as it does not resemble the code that a human would write to accomplish a similar task. The basic idea behind the generated C code is that it expands the entire program into a list of “events”, then starts the clock, then executes all of the events in a single loop. Regardless of the `CONCEPTUAL` program being compiled, the body of the generated C code will look like this:

```
for (i=0; i<numevents; i++) {
    CONC_EVENT *thisev = &eventlist[i];

    switch (thisev->type) {
        case event_1:
            :
        case event_2:
            :
    }
}
```

Programs generated by ‘`codegen_c_generic.py`’ define a number of event types which are summarized in [Section B.3 \[Event types\]](#), page 190. The `EV_CODE` event is used, for example, by the `BACKEND EXECUTES` (see [Section 4.6.6 \[Injecting arbitrary code\]](#), page 111), `LOGS` (see [Section 4.5.3 \[Writing to a log file\]](#), page 104), and `OUTPUTS` (see [Section 4.5.2 \[Writing to standard output\]](#), page 104) statements. Note that there are no loop events—in fact, there are no complex statements (see [Section 4.7 \[Complex statements\]](#), page 112) whatsoever. Complex statements are expanded into multiple simple statements at initialization time.

The advantage of completely expanding a `CONCEPTUAL` program during the initialization phase—essentially, “pre-executing” the entire program—is that that enables all of the expensive, non-communication-related setup to be hoisted out of the timing loop, which is how a human would normally express a network benchmark. Pre-execution is possible because the `CONCEPTUAL` language is not a Turing machine; infinite loops are not expressible by the language and message contents and timings cannot affect program behavior, for instance. During its initialization phase, the generated C code allocates memory for message

buffers, evaluates numerical expressions, verifies program assertions, unrolls loops, and does everything else that's not directly relevant to communication performance. For instance, the CONCEPTUAL program 'TASK tx SUCH THAT $tx > 4$ SENDS 10 1 MEGABYTE MESSAGES TO TASK $tx/2$ ' would cause each task to perform the following steps during initialization:

- determine if its task ID is greater than 4, making the task a sender
- determine if its task ID is equal to ' $tx/4$ ' (rounded down to the nearest integer) for some task ' tx ' in the program, making the task a receiver
- allocate 1 MB for a message buffer
- allocate and initialize a repeat event, specifying that the subsequent event should repeat 10 times
- allocate a send or receive event

The final two of those steps repeat as necessary. For example, task 3 receives 10 messages from each of task 6 and task 7.

Note that each task's receive events (if any) are allocated before its send events (if any), as described [Section 4.4.2 \[Sending\], page 98](#). Also, note that only a single message buffer is allocated because the CONCEPTUAL source did not specify the UNIQUE keyword (see [\[Unique messages\], page 96](#)).

An event is implemented as a C struct that contains all of the state needed to perform a particular operation. For example, an event corresponding to a synchronous or asynchronous send operation (CONC_SEND_EVENT) stores the destination task ID, the number of bytes to send, the message alignment, the number of outstanding asynchronous sends and receives, a flag indicating whether the data is to be touched, and a flag indicating that the message should be filled with data the receiver can verify. In addition, the `code_declare_datatypes_SEND_STATE` hook (see [Section 6.2.1 \[Hook methods\], page 137](#)) enables a backend to include additional, backend-specific state in the (CONC_SEND_EVENT) data structure.

6.2.4 Internals

'codegen_c_generic.py' is a fairly substantial piece of code. It is divided into ten sections:

1. methods exported to the compiler front end
2. utility functions that do not generate code
3. utility functions that do generate code
4. methods for outputting language atoms (see [Section 4.1 \[Primitives\], page 77](#))
5. methods for outputting miscellaneous language constructs (e.g., restricted identifiers; see [Section 4.3.1 \[Restricted identifiers\], page 93](#))
6. methods for outputting expressions (see [Section 4.2 \[Expressions\], page 79](#))
7. methods for outputting complete programs (see [Section 4.9 \[Complete programs\], page 121](#))
8. methods for outputting complex statements (see [Section 4.7 \[Complex statements\], page 112](#))
9. methods for outputting simple statements (e.g., communication statements; see [Section 4.4 \[Communication statements\], page 94](#))

10. methods for outputting nodes with non-textual names (e.g., ‘...’ and various operators)

The `NCPTL_CodeGen` class defined in ‘`codegen_c_generic.py`’ generates code as follows. The `generate` method, which is invoked by ‘`ncptl.py`’, calls upon PLY to process the abstract-syntax tree (AST) in postorder fashion. `NCPTL_CodeGen` maintains a stack (`codestack`) on which code fragments are pushed and popped but that ends up containing a complete line of C code in each element. For example, in the `CONCEPTUAL` program ‘`TASK 0 OUTPUTS 1+2*3`’, the `n_outputs` method will pop ‘`[('expr', '(1)+((2)*(3)))]`’ (a list containing the single expression ‘`1+2*3`’) and ‘`(('task_expr', '0'))`’ (a tuple designating a task by the expression ‘`0`’) and push multiple lines of code that prepare task 0 to evaluate and output the given expression.

The utility functions are the most useful for backend developers to understand, as they are frequently called from hook methods (see [Section 6.2.1 \[Hook methods\]](#), page 137). The following should be of particular importance:

`push`

`pushmany` Push a single value (typically a string of C code) or each value in a list of values onto a stack.

`error_fatal`

`error_internal`

Output a generic error message or an “internal error” error message and abort the program.

`code_declare_var`

Push (using the `push` method) a line of C code that declares a variable with an optionally specified type, name, initial value, and comment. Return the variable name actually used.

See the definitions in ‘`codegen_c_generic.py`’ of each of the above to determine required and optional parameters. The following, adapted from ‘`codegen_c_udgram.py`’ demonstrates some of the preceding methods:

```

def n_for_count_SYNC_ALL(self, localvars):
    "Synchronize all of the tasks in the job."
    synccode = []
    self.push("{", synccode)
    loopvar = self.code_declare_var(suffix="task",
        comment="Loop variable that iterates over all (physical) ranks",
        stack=synccode)
    self.pushmany([
        "thisev_sync->s.sync.peerqueue = ncptl_queue_init (sizeof(int));",
        "for (%s=0; %s<var_num_tasks; %s++)" %
        (loopvar, loopvar, loopvar),
        "*(int *)ncptl_queue_allocate(thisev_sync->s.sync.peerqueue) = %s;" %
        loopvar,
        "thisev_sync->s.sync.syncrank = physrank;",
        "}],
        stack=synccode)
    return synccode

```

That definition of the `n_for_count_SYNC_ALL` hook method defines a new stack (`synccode`) and pushes a `{` onto it. It then declares a loop variable, letting `code_declare_var` select a name but dictating that it end in `_task`. The hook method then pushes some additional C code onto the `synccode` stack and finally returns the stack (which is really just a list of lines of C code).

Some useful variables defined by `NCPTL_CodeGen` include the following:

base_global_parameters

a list of 6-ary tuples defining extra command-line parameters to parse (format: `{type, variable, long_name, short_name, description, default_value}`)

events_used

a dictionary containing the names of events actually used by the program being compiled

Some methods in `codegen_c_generic.py` that are worth understanding but are unlikely to be used directly in a derived backend include the following:

pop Pop a value from a stack.

push_marker

Push a specially designated “marker” value onto a stack.

combine_to_marker

Pop all items off a stack up to the first marker value found; discard the marker; then, push the popped items as a single list of items. This is used, for example, by a complex statement (see [Section 4.7 \[Complex statements\]](#), page 112) that applies to a list of statements, which can be popped as a unit using `combine_to_marker`.

invoke_hook

Call a hook method, specifying code to be pushed before/after the hook-produced code and alternative text (or Python code) to be pushed (or executed) in the case that a hook method is not provided.

6.3 Run-time library functions

To simplify the backend developer’s task and to provide consistent functionality across backends, CONCEPTUAL provides a run-time library that encapsulates many of the common operations needed for network-correctness and performance-testing programs. This section describes all of the functions that the library exports (plus a few important types and variables). The library is written in C, so all of the type/variable/function prototypes are expressed with C syntax. The library includes, among others, functions that manage heap-allocated memory, accurately read the time, write results to log files, control queues of arbitrary data, and implement various arithmetic operations. All of these functions should be considered “slow” and should therefore generally not be invoked while execution is being timed.¹

6.3.1 Constants, variables, and data types

The following constants, variables, and data types are used by various run-time library functions and directly by backends.

`ncptl_int` [Data type]

The internal data type of the CONCEPTUAL run-time library is `ncptl_int`. This is normally a 64-bit signed integer type selected automatically by ‘`configure`’ (see [Section 2.1 \[configure\], page 5](#)) but can be overridden with the `--with-datatype` option to ‘`configure`’. ‘`ncptl.h`’ defines a string macro called `NICS` that can be used to output an `ncptl_int` regardless of how the `ncptl_int` type is declared:

```
ncptl_fatal ("My variable contains a negative value (%" NICS ")",
            my_ncptl_int_var);
```

`ncptl_int` constants declared by backends derived from ‘`codegen_c_generic.py`’ are given an explicit suffix which defaults to ‘`LL`’ but can be overridden at configuration time using the `--with-const-suffix` option.

`NCPTL_INT_MIN` [Constant]

`NCPTL_INT_MIN` is a C preprocessor macro which represents the smallest (i.e., most negative) number that can be assigned to an `ncptl_int`. For example, if `ncptl_int` is a 64-bit signed integer type, then `NCPTL_INT_MIN` will be the value ‘`-9223372036854775808`’.

`NCPTL_CMDLINE` [Data type]

The `NCPTL_CMDLINE` structure describes an acceptable command-line option. It contains a `type`, which is either `NCPTL_TYPE_INT` for an `ncptl_int` or `NCPTL_TYPE_STRING` for a `char *`, a pointer to a variable that will receive the value specified on the command line, the long name of the argument (without the ‘`--`’), the one-letter short name of the argument (without the ‘`-`’), a textual description of what the argument represents, and a default value to use if the option is not specified on the command line.

¹ Some notable exceptions are the functions described in [Section 6.3.4 \[Message-buffer manipulation functions\], page 148](#), which implement CONCEPTUAL’s `WITH DATA TOUCHING` and `WITH VERIFICATION` constructs.

`NCPTL_QUEUE` [Data type]

An `NCPTL_QUEUE` is an opaque data type that represents a dynamically growing queue that can be flattened to an array for more convenient access. `NCPTL_QUEUES` have proved to be quite useful when implementing `CONCEPTUAL` backends.

`NCPTL_LOG_FILE_STATE` [Data type]

Every `CONCEPTUAL` log file is backed by a unique `NCPTL_LOG_FILE_STATE` opaque data type. An `NCPTL_LOG_FILE_STATE` data type represents all of the state needed to maintain that file, such as file descriptors, prologue comments, and data which has not yet been aggregated.

`int ncptl_pagesize` [Variable]

This variable is initialized by `ncptl_init()` to the number of bytes in an operating-system memory page. `ncptl_pagesize` can be used by backends to implement `CONCEPTUAL`'s `PAGE SIZED` (see [Item size], page 96) and `PAGE ALIGNED` (see [Message alignment], page 97) keywords.

`int ncptl_fast_init` [Variable]

The `ncptl_init()` function (see Section 6.3.2 [Initialization functions], page 146) can take many seconds to complete. Much of this time is spent calibrating and measuring the quality of the various timers the run-time library uses. For backends such as `picl` (see Section 3.3.8 [The `picl` backend], page 33) which do not measure real time there is little need to have an accurate timer. Setting `ncptl_fast_init` to '1' before invoking `ncptl_init()` skips the timer calibration and measurement steps, thereby leading to faster initialization times. A user can also override the setting of `ncptl_fast_init` at run time by setting the `NCPTL_FAST_INIT` environment variable to either '0' or '1', as appropriate.

6.3.2 Initialization functions

The following functions are intended to be called fairly early in the generated code.

`void ncptl_init (int version, char *program_name)` [Function]

Initialize the `CONCEPTUAL` run-time library. `version` is used to verify that '`runtime.lib.c`' corresponds to the version of '`ncptl.h`' used by the generated code. The caller must pass in `NCPTL_RUN_TIME_VERSION` for `version`. `program_name` is the name of the executable program and is used for outputting error messages. The caller should pass in `argv[0]` for `program_name`. `ncptl_init()` must be the first library function called by the generated code (with a few exceptions, as indicated below).

`void ncptl_permit_signal (int signalnum)` [Function]

Indicate that the backend relies on signal `signalnum` for correct operation. Because signal handling has performance implications, the `CONCEPTUAL` run-time library normally terminates the program upon receiving a signal. Hence, the user can be assured that if a program runs to completion then no signals have affected its performance. See Section 3.4 [Running `coNCePTuaL` programs], page 41, for a description of the `--no-trap` command-line option, which enables a user to permit additional signals to be delivered to the program (e.g., when linking with a particular implementation of a communication library that relies on signals). `ncptl_permit_signal()` must be invoked before `ncptl_parse_command_line()` to have any effect.


```
void ncptl_parse_command_line (int argc, char *argv[],           [Function]
                             NCPTL_CMDLINE *arglist, int numargs)
```

Parse the command line. *argc* and *argv* should be the argument count and argument vector passed to the generated code by the operating system. *arglist* is a list of descriptions of acceptable command-line arguments and *numargs* is the length of that list.

Because `ncptl_init()` takes many seconds to run, it is common for generated code to scan the command line for `--help` or `-?` and, if found, skip `ncptl_init()` and immediately invoke `ncptl_parse_command_line()`. Doing so gives the user immediate feedback when requesting program usage information. Skipping `ncptl_init()` is safe in this context because `ncptl_parse_command_line()` terminates the program after displaying usage information; it does not require any information discovered by `ncptl_init()`.

Most generated programs have a `--seed/-S` option that enables the user to specify explicitly a seed for the random-number generator with `--help/-?` showing the default seed. `ncptl_seed_random_task()` must therefore be called before `ncptl_parse_command_line()` which, as stated in the previous paragraph, can be invoked without a prior invocation of `ncptl_init()`. Consequently, it can be considered safe also to invoke `ncptl_seed_random_task()` before `ncptl_init()`.

A generated program's initialization routine will generally exhibit a structure based on the following pseudocode:

```
if "--help" or "-?" in command-line options then
    only_help := TRUE
else
    only_help := FALSE
    ncptl_init(...)
end if
random_seed := ncptl_seed_random_task(0)
ncptl_parse_command_line(...)
if only_help = TRUE then
    ncptl_error("Internal error; should have exited")
end if
ncptl_seed_random_task(random_seed)
```

6.3.3 Memory-allocation functions

The CONCEPTUAL run-time library provides its own wrappers for `malloc()`, `free()`, `realloc()`, and `strdup()` as well as a specialized `malloc()` designed specifically for allocating message buffers. The wrappers' "value added" is that they support the explicit data alignments needed by `ALIGNED` messages (see [\[Message alignment\]](#), page 97) and that they automatically call `ncptl_fatal()` on failure, so the return value does not need to be checked for `NULL`.

```
void * ncptl_malloc (ncptl.int numbytes, ncptl.int alignment) [Function]
```

Allocate *numbytes* bytes of memory aligned to an *alignment*-byte boundary. If *alignment* is '0', `ncptl_malloc()` will use whatever alignment is "natural" for the underlying architecture. `ncptl_malloc()` will automatically call `ncptl_fatal()` if memory allocation fails. Therefore, unlike `malloc()`, there is no need to check the return value for `NULL`.

`void * ncptl_malloc_misaligned (ncptl_int numbytes, ncptl_int misalignment)` [Function]

Allocate *numbytes* bytes of memory from the heap aligned *misalignment* bytes past a page boundary. If *alignment* is ‘0’, `ncptl_malloc_misaligned()` will return page-aligned memory. `ncptl_malloc_misaligned()` will automatically call `ncptl_fatal()` if memory allocation fails. Therefore, unlike `malloc()`, there is no need to check the return value for NULL.

`void ncptl_free (void *pointer)` [Function]

Free memory previously allocated by `ncptl_malloc()`. It is an error to pass `ncptl_free()` memory not allocated by `ncptl_malloc()`.

`void * ncptl_realloc (void *pointer, ncptl_int numbytes, ncptl_int alignment)` [Function]

Given a pointer returned by `ncptl_malloc()`, change its size to *numbytes* and byte-alignment to *alignment* without altering the contents (except for truncation in the case of a smaller target size). If *alignment* is ‘0’, `ncptl_realloc()` will use whatever alignment is “natural” for the underlying architecture. `ncptl_realloc()` will automatically call `ncptl_fatal()` if memory allocation fails. Therefore, unlike `realloc()`, there is no need to check the return value for NULL.

`char * ncptl_strdup (const char *string)` [Function]

`ncptl_strdup()` copies a string as does the standard C `strdup()` function. However, `ncptl_strdup()` uses `ncptl_malloc()` instead of `malloc()` to allocate memory for the copy, which must therefore be deallocated using `ncptl_free()`.

`void * ncptl_malloc_message (ncptl_int numbytes, ncptl_int alignment, ncptl_int outstanding, int misaligned)` [Function]

Allocate *numbytes* bytes of memory from the heap either aligned on an *alignment*-byte boundary (if *misaligned* is ‘0’) or *alignment* bytes past a page boundary (if *misaligned* is ‘1’). All calls with the same value of *outstanding* will share a buffer. `ncptl_malloc_message()` is intended to be used in two passes. The first time the function is called on a set of messages it merely determines how much memory to allocate. The second time, it returns valid memory buffers. Note that the returned pointer can be neither `free()`d nor `ncptl_free()`d.

`void * ncptl_get_message_buffer (ncptl_int buffernum)` [Function]

Return a pointer to a message buffer previously allocated (and finalized) by `ncptl_malloc_message()`. The *buffernum* argument to `ncptl_get_message_buffer()`, which corresponds to the *outstanding* argument to `ncptl_malloc_message()`, specifies the number of the buffer to return. `ncptl_get_message_buffer()` returns NULL if buffer *buffernum* is either unallocated or uninitialized.

6.3.4 Message-buffer manipulation functions

The CONCEPTUAL language facilitates verifying message contents and touching every word in a message (see [Data touching], page 97). The following functions implement those features.

`void ncptl_fill_buffer (void *buffer, ncptl_int numbytes, int validity)` [Function]

Fill a region of memory with known values. If *validity* is '+1', `ncptl_fill_buffer()` will fill the first *numbytes* bytes of *buffer* with a verifiable sequence of integers (see [Data touching], page 97). If *validity* is '-1', `ncptl_fill_buffer()` will pollute the first *numbytes* bytes of *buffer*. Receive buffers should be polluted before reception to avoid false negatives caused, for example, by an inadvertently dropped message destined for a previously validated buffer.

`ncptl_int ncptl_verify (void *buffer, ncptl_int numbytes)` [Function]

Verify the contents of memory filled by `ncptl_fill_buffer()`. The function returns the number of erroneous bits. `ncptl_verify()` is used to implement CONCEPTUAL's WITH VERIFICATION construct (see [Data touching], page 97).

`void ncptl_touch_data (void *buffer, ncptl_int numbytes)` [Function]

Touch every byte in a given buffer. `ncptl_touch_data()` is used to implement the WITH DATA TOUCHING construct described in [Data touching], page 97.

The following function, `ncptl_touch_memory()`, does not actually manipulate message buffers. It is included in this section because of its similarity to `ncptl_touch_data()`. `ncptl_touch_data()` touches message buffers to implement the WITH DATA TOUCHING construct; `ncptl_touch_memory()` touches a hidden memory region to implement the TOUCHES statement.

`void ncptl_touch_memory (void *buffer, ncptl_int bufferbytes, ncptl_int wordbytes, ncptl_int firstbyte, ncptl_int numaccesses, ncptl_int bytestride)` [Function]

Walk a memory region *buffer* of size *bufferbytes* bytes. Each “word” to touch contains *wordbytes* bytes. *firstbyte* indicates the byte index into *buffer* from which to start touching. The function will read and write *numaccesses* “words” with stride *bytestride* bytes. For example, '`ncptl_touch_memory (mybuffer, 1048576, 64, 192, 10000, 4096)`' will read and write (but otherwise do nothing with) the 64 bytes that lie 192 bytes into a 1 MB memory region, then the 64 bytes starting at offset 4288, then the 64 bytes starting at $192 + 4096 \times 2$, then at $192 + 4096 \times 3$, then at $192 + 4096 \times 4$, and so forth up to $192 + 4096 \times 10000$, wrapping around the 1 MB region as necessary.

A *bytestride* of NCPTL_INT_MIN implies a random stride.

As an important special case, if *firstbyte* is '-1', then `ncptl_touch_memory()` will touch one or more message buffers (cf. `ncptl_malloc_message()` in Section 6.3.3 [Memory-allocation functions], page 147) instead of the given *buffer*. In that case, *bufferbytes* stores the buffer number (a nonnegative number). If *bufferbytes* is '-1', however, then *all* message buffers are touched. Note that when *firstbyte* is '-1', all parameters to `ncptl_touch_memory()` other than *bufferbytes* are ignored.

`ncptl_touch_memory()` is intended to to be used to implement the TOUCHES statement (see Section 4.6.4 [Touching memory], page 108).

6.3.5 Time-related functions

An essential component of any benchmarking system is an accurate timer. `CONCEPTUAL`'s `ncptl_time()` function selects from a variety of timers at configuration time, first favoring lower-overhead cycle-accurate timers, then higher-overhead cycle-accurate, and finally non-cycle-accurate timers. `ncptl_init()` measures the actual timer overhead and resolution and `ncptl_log_write_prologue()` writes this information to the log file. Furthermore, the `'validatetimer'` program (see [\[Validating the coNcEPTuaL timer\]](#), page 10) can be used to verify that the timer used by `ncptl_init()` truly does correspond to wall-clock time.

The `CONCEPTUAL` language provides a few time-related functions. These are supported by the functions described below.

`uint64_t ncptl_time (void)` [Function]

Return the time in microseconds. The timer ticks even when the program is not currently scheduled. No assumptions can be made about the relation of the value returned to the time of day; `ncptl_time()` is intended to be used strictly for computing elapsed time. The timer's resolution and accuracy are logged to the log file by `ncptl_log_write_prologue()` (more precisely, by the internal `log_write_prologue_timer()` function, which is called by `ncptl_log_write_prologue()`). Note that `ncptl_time()` always returns a 64-bit unsigned value, regardless of how `ncptl_int` is declared.

The `CONCEPTUAL` `'configure'` script (see [Section 2.1 \[configure\]](#), page 5) searches for a number of high-resolution timers and selects the best timer mechanism from among the ones available. The selection criteria is as follows:

1. If `./configure` was passed `--with-gettimeofday` (see [Section 2.1 \[configure\]](#), page 5) then `ncptl_time()` uses `gettimeofday()` as its timer mechanism.
2. If `./configure` was passed `--with-mpi-wtime` (see [Section 2.1 \[configure\]](#), page 5) then `ncptl_time()` uses `MPI_Wtime()` as its timer mechanism.
3. If `./configure` recognizes the CPU architecture, knows how to instruct the C compiler to insert inline assembly code, and can determine the number of clock cycles per second, then `ncptl_time()` reads the timer using inline assembly code. If the cycle counter is likely to wrap around during a moderately long benchmark (i.e., because the cycle counter is a 32-bit register), `ncptl_time()` augments the inline assembly code with calls to `gettimeofday()` in an attempt to produce accurate timings that don't suffer from clock wraparound.
4. If the Linux `get_cycles()` function is available and `./configure` can determine the number of clock cycles per second then `ncptl_time()` uses `get_cycles()` to measure execution time.
5. If the PAPI library is available, `ncptl_time()` becomes a call to PAPI's `PAPI_get_real_usec()` function.
6. If the System V `clock_gettime()` function is available and the `CLOCK_SGI_CYCLE` macro is defined, `ncptl_time()` invokes `clock_gettime()` with the `CLOCK_SGI_CYCLE` argument. If `CLOCK_SGI_CYCLE` is not defined but `CLOCK_REALTIME` is, then `ncptl_time()` invokes `clock_gettime()` with the `CLOCK_REALTIME` argument.

7. Intel’s (now obsolete) supercomputers provide a `dclock()` function for reading the time. `ncptl_time()` makes use of `dclock()` if it’s available.
8. Microsoft Windows provides functions for reading a high-resolution timer (`QueryPerformanceCounter()`) and for determining the number of ticks per second that the timer measures (`QueryPerformanceFrequency()`). If those functions are available, `ncptl_time()` uses them.
9. As a last resort, `ncptl_time()` uses `gettimeofday()` to measure execution time.

Furthermore, CONCEPTUAL makes use of the High-Precision Event Timers (HPET) device if and only if all of the following conditions hold at run time:

- neither `--with-gettimeofday`, `--with-mpi-wtime`, nor `--disable-hpet` was specified to `./configure`
- `/dev/hpet`—or an alternative device specified at configuration time with `--enable-hpet`—exists, is readable, and can be memory-mapped into user space
- the HPET device’s main counter is a 64-bit value, not a 32-bit value
- the period of the HPET counter is within the ranges required by the HPET specification, i.e., between 0 and 100 nanoseconds

Failing any of those conditions, CONCEPTUAL falls back to the timer selected at configuration time. See the [HPET specification](#) for more information on HPET.

```
void ncptl_set_flag_after_usecs (volatile int *flag, uint64_t          [Function]
    delay)
```

`ncptl_set_flag_after_usecs()` uses the operating system’s interval timer to asynchronously set a variable to ‘1’ after a given number of microseconds. This function is intended to be used to support the ‘FOR *time*’ construct (see [\[Timed loops\]](#), page 117). Note that *delay* is a 64-bit unsigned value, regardless of how `ncptl_int` is declared.

`ncptl_set_flag_after_usecs()` is implemented in terms of the `setitimer()` function and issues a run-time error if the `setitimer()` function is not available.

```
void ncptl_udelay (uint64_t delay, int spin0block1)                    [Function]
```

If *spin0block1* is ‘0’, `ncptl_udelay()` spins for *delay* microseconds (i.e., using the CPU). If *spin0block1* is ‘1’, `ncptl_udelay()` sleeps for *delay* microseconds (i.e., relinquishing the CPU). Note that *delay* is a 64-bit unsigned value, regardless of how `ncptl_int` is declared. `ncptl_udelay()` is intended to be used to support the CONCEPTUAL language’s SLEEPS and COMPUTES statements (see [Section 4.6.3 \[Delaying execution\]](#), page 107).

When *spin0block1* is ‘0’, `ncptl_udelay()` uses `ncptl_time()` to determine when *delay* microseconds have elapsed. Unless `ncptl_time()` is known to utilize an extremely low-overhead timer, `ncptl_udelay()` intersperses calls to `ncptl_time()` with writes to a dummy variable. When *spin0block1* is ‘0’, `ncptl_udelay()` invokes `nanosleep()` to introduce delays. `ncptl_udelay()` issues a run-time error if the `nanosleep()` function is not available.

6.3.6 Log-file functions

Benchmarking has limited value without a proper record of the performance results. The CONCEPTUAL run-time library provides functions for writing data to log files. It takes care of much of the work needed to calculate statistics on data columns and to log a thorough experimental setup to every log file.

The library treats a log file as a collection of tables of data. Each table contains a number of rows, one per dynamic invocation of the LOGS statement (see [Section 4.5.3 \[Writing to a log file\]](#), page 104). Each row contains a number of columns, one per aggregate expression (see [Section 4.2.3 \[Aggregate expressions\]](#), page 90) expressed statically in a CONCEPTUAL program.² Log-file functions should be called only if the CONCEPTUAL source code accesses a log file (see [Section 4.5.3 \[Writing to a log file\]](#), page 104).

`void ncptl_log_add_comment (const char *key, const char *value)` [Function]
`ncptl_log_add_comment()` makes it possible for a backend to add backend-specific *<key:value>* pairs to the set of prologue or epilogue comments that get written to a log file (see [Section 3.5.1 \[Log-file format\]](#), page 44). `ncptl_log_add_comment()` can be called repeatedly. All calls which precede `ncptl_log_open()` are included in the log-file prologue. All calls which follow `ncptl_log_open()` are included in the log-file epilogue. Note that `ncptl_log_add_comment()` makes a copy of *key* and *value*, so these need not be heap-allocated.

`NCPTL_LOG_FILE_STATE * ncptl_log_open (char *template,` [Function]
`ncptl_int processor)`

Given a filename template containing a ‘%d’ placeholder and a processor number (i.e., the process’s physical rank in the computation), `ncptl_log_open()` creates and opens a log file named by the template with ‘%d’ replaced by *processor*. For example, if *template* is ‘/home/me/myprog-%d.log’ and *processor* is ‘3’, the resulting filename will be ‘/home/me/myprog-3.log’. `ncptl_log_open()` must be called before any of the other `ncptl_log_something()` functions—except for `ncptl_log_add_comment()`, which should be called before `ncptl_log_open()`. `ncptl_log_open()` returns a pointer to an opaque `NCPTL_LOG_FILE_STATE` value; the backend will need to pass this pointer to nearly all of the other log-file functions described in this section.

There are two special cases of *template*. First, if *template* points to an empty string, all log-file output is sent to the null device (i.e., ‘/dev/null’ on Unix and Unix-like operating systems). Second, if *template* is a null pointer, all log-file output is sent to the standard output device.

`char * ncptl_log_generate_uuid (void)` [Function]
 Return a random string of hexadecimal digits formatted as “xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx” (36 bytes plus a NULL byte) to pass to `ncptl_log_write_prologue()`. The caller should `ncptl_free()` the string when it is no longer needed (generally, as soon as `ncptl_log_write_prologue()` returns).

² Writing A HISTOGRAM OF THE *<expr>* produces two columns, one for values and one for tallies.

```
void ncptl_log_write_prologue (NCPTL_LOG_FILE_STATE      [Function]
    *logstate, char *prognam, char *uuid, char *backend_name, char
    *backend_desc, ncptl_int numtasks, NCPTL_CMDLINE *arglist, int
    numargs, char **sourcecode)
```

`ncptl_log_write_prologue()` standardizes the prologue with which all log files begin. *prognam* is the name of the program executable (`argv[0]` in C). *uuid* is a value returned by `ncptl_log_generate_uuid()`. Note that every process in a program must pass the same value of *uuid* to `ncptl_log_write_prologue()`. *backend_name* is the name of the backend in ‘*language_library*’ format (e.g., ‘*java_rmi*’). *backend_desc* is a brief description of the backend (e.g., ‘*Java + RMI*’). *numtasks* is the total number of tasks in the program. *arglist* is the list of arguments passed to `ncptl_parse_command_line()` and *numargs* is the number of entries in that list. *sourcecode* is the complete CONCEPTUAL source code stored as a NULL-terminated list of NULL-terminated strings.

```
char * ncptl_log_lookup_string (NCPTL_LOG_FILE_STATE      [Function]
    *logstate, char *key)
```

`ncptl_log_write_prologue()` stores every *<key:value>* comment it writes into an in-memory database. `ncptl_log_lookup_string()` searches the comment database for a key and returns the corresponding value. The function returns the empty string if the key is not found in the database. In either case, the caller should not deallocate the result. `ncptl_log_lookup_string()` is intended to be used to implement CONCEPTUAL’s THE VALUE OF construct (see [Section 4.5.1 \[Utilizing log-file comments\]](#), [page 103](#)).

```
void ncptl_log_write (NCPTL_LOG_FILE_STATE *logstate, int      [Function]
    logcolumn, char *description, LOG_AGGREGATE aggregate, double
    value)
```

Push value *value* onto column *logcolumn* of the current table. Gaps between columns are automatically elided. *description* is used as the column header for column *logcolumn*. Acceptable values for *aggregate* are defined in [Section B.4 \[Representing aggregate functions\]](#), [page 191](#).

```
void ncptl_log_compute_aggregates (NCPTL_LOG_FILE_STATE      [Function]
    *logstate)
```

`ncptl_log_compute_aggregates()` implements the COMPUTES AGGREGATES construct described in [\[Computing aggregates\]](#), [page 106](#). When `ncptl_log_compute_aggregates()` is invoked, the CONCEPTUAL run-time library uses the aggregate function specified by `ncptl_log_write()` to aggregate all of the data that accumulated in each column since the last invocation of `ncptl_log_compute_aggregates()`. Note that `ncptl_log_compute_aggregates()` is called implicitly by `ncptl_log_commit_data()`.

```
void ncptl_log_commit_data (NCPTL_LOG_FILE_STATE      [Function]
    *logstate)
```

The CONCEPTUAL run-time library keeps the current data table in memory and doesn’t write anything to the log file until `ncptl_log_commit_data()` is called, at

which point the run-time library writes all accumulated data to the log file and begins a new data table. Note that `ncptl_log_commit_data()` is called implicitly by `ncptl_log_close()`. Furthermore, a backend should call `ncptl_log_commit_data()` when beginning execution of a new statement in a `CONCEPTUAL` program. For instance, the `'codegen_c_generic.py'` backend invokes `ncptl_log_commit_data()` from `code_def_main_newstmt`.

```
void ncptl_log_write_epilogue (NCPTL_LOG_FILE_STATE      [Function]
                             *logstate)
```

Write a stock set of comments as an epilogue to the log file.

```
void ncptl_log_close (NCPTL_LOG_FILE_STATE *logstate)    [Function]
    Close the log file. No ncptl_log_something() function should be called after
    ncptl_log_close() is invoked.
```

6.3.7 Random-task functions

Randomness appears in various forms in the `CONCEPTUAL` language, such as when assigning a task to `A RANDOM PROCESSOR` (see [Section 4.6.5 \[Reordering task IDs\]](#), page 110) or when let-binding `A RANDOM TASK` or `A RANDOM TASK OTHER THAN` a given task ID to a variable (see [Section 4.7.3 \[Binding variables\]](#), page 117). The following functions are used to select tasks at random. `CONCEPTUAL` currently uses the Mersenne Twister as its random-number generator. Hence, given the same seed, a `CONCEPTUAL` program will see the same random-number sequence on every platform.

```
int ncptl_seed_random_task (int seed, ncptl_int procID)  [Function]
    Initialize the random-number generator needed by ncptl_random_task().
    If seed is zero, ncptl_seed_random_task() selects an arbitrary seed value.
    ncptl_seed_random_task() returns the seed that was used. procID specifies the
    (physical) processor ID of the calling task and is needed to seed the task-local
    random-number generators used by some of the functions in Section 6.3.11
    [Language-visible functions], page 157.
```

```
ncptl_int ncptl_random_task (ncptl_int lower_bound, ncptl_int      [Function]
                             upper_bound, ncptl_int excluded)
```

Return a randomly selected task number from `lower_bound` to `upper_bound` (both inclusive). If `excluded` is nonnegative then that task number will never be selected, even if it's within range.

6.3.8 Task-mapping functions

```
NCPTL_VIRT_PHYS_MAP * ncptl_allocate_task_map (ncptl_int      [Function]
                                                numtasks)
```

Allocate and initialize an opaque `NCPTL_VIRT_PHYS_MAP` object and return a pointer to it.

```
ncptl_int ncptl_virtual_to_physical (NCPTL_VIRT_PHYS_MAP      [Function]
                                     *procmap, ncptl_int virtID)
```

Given a process map allocated by `ncptl_allocate_task_map()` and a (virtual) task ID, return the corresponding (physical) process ID.

`ncptl_int ncptl_physical_to_virtual (NCPTL_VIRT_PHYS_MAP [Function]
*procmap, ncptl_int physID)`

Given a process map allocated by `ncptl_allocate_task_map()` and a (physical) process ID, return the corresponding (virtual) task ID.

`ncptl_int ncptl_assign_processor (ncptl_int virtID, ncptl_int [Function]
physID, NCPTL_VIRT_PHYS_MAP *procmap, ncptl_int physrank)`

Assign a (physical) processor ID, *physID*, to a (virtual) task ID, *virtID* given a virtual-to-physical mapping table, *procmap*. Return a new task ID for the caller's processor given its processor ID, *physrank*. `ncptl_assign_processor()` is intended to implement the IS ASSIGNED TO construct (see [Section 4.6.5 \[Reordering task IDs\]](#), page 110).

6.3.9 Queue functions

Because queues are a widely applicable construct, the run-time library provides support for queues of arbitrary datatypes. In the current implementation, these can more precisely be termed “dynamically growing lists” than “queues”. However, they may be extended in a future version of the library to support more queue-like functionality.

`NCPTL_QUEUE * ncptl_queue_init (ncptl_int eltbytes) [Function]`

`ncptl_queue_init()` creates and initializes a dynamically growing queue in which each element occupies *eltbytes* bytes of memory.

`void * ncptl_queue_allocate (NCPTL_QUEUE *queue) [Function]`

Allocate a new data element at the end of queue *queue*. The queue passed to `ncptl_queue_allocate()` must be one returned by `ncptl_queue_init()`. `ncptl_queue_allocate()` returns a pointer to the data element allocated.

`void * ncptl_queue_push (NCPTL_QUEUE *queue, void *element) [Function]`

Push (via a memory copy) the element pointed to by *element* onto the end of queue *queue* and return a pointer to the copy in the queue. The queue passed to `ncptl_queue_allocate()` must be one returned by `ncptl_queue_init()`. (`ncptl_queue_push()` is actually implemented in terms of `ncptl_queue_allocate()`.)

`void * ncptl_queue_pop (NCPTL_QUEUE *queue) [Function]`

Pop a pointer to the element at the head of queue *queue*. If *queue* is empty, return NULL. The pointer returned by `ncptl_queue_pop()` is guaranteed to be valid until the next invocation of `ncptl_queue_empty()`.

`void * ncptl_queue_contents (NCPTL_QUEUE *queue, int [Function]
copyelts)`

Return queue *queue* as an array of elements. If `ncptl_queue_contents()` is passed ‘1’ for *copyelts*, a new array is allocated using `ncptl_malloc()`; the queue's internal array is copied to the newly allocated array; and, this new array is returned to the caller. It is the caller's responsibility to pass the result to `ncptl_free()` when the array is no longer needed. If `ncptl_queue_contents()` is passed ‘0’ for *copyelts*, a pointer to the queue's internal array is returned without first copying it. This pointer should not be passed to `ncptl_free()` as it is still needed by *queue*.

`ncptl_int ncptl_queue_length (NCPTL_QUEUE *queue)` [Function]
Return the number of elements in queue *queue*.

`void ncptl_queue_empty (NCPTL_QUEUE *queue)` [Function]
Empty a queue, freeing the memory that had been allocated for its elements. Queue contents returned by `ncptl_queue_contents()` with *copyelts* set to '0' are also invalidated. The queue itself can continue to be used and should be deallocated with `ncptl_free()` (see [Section 6.3.3 \[Memory-allocation functions\], page 147](#)) when no longer needed.

6.3.10 Unordered-set functions

Because unordered collections of data are a widely applicable construct, the run-time library provides support for sets. A set contains zero or more keys, each of which must be unique within the set. Furthermore, each key can be associated with a data value. Sets are currently implemented in the CONCEPTUAL run-time library as hash tables.

`NCPTL_SET * ncptl_set_init (ncptl_int numelts, ncptl_int keybytes, ncptl_int valuebytes)` [Function]

Allocate and initialize a set object and return a pointer to it. Each element in the set maps a *keybytes*-byte key to a *valuebytes*-byte value. The *numelts* parameter is an estimate of the maximum number of elements in the set. `ncptl_set_init()` returns a pointer to the set.

`void * ncptl_set_find (NCPTL_SET *set, void *key)` [Function]
Given a set and a pointer to a key, return a pointer to the associated value or NULL if the key is not found in the set.

`void ncptl_set_insert (NCPTL_SET *set, void *key, void *value)` [Function]
Insert a key into a set and associate a value with it. `ncptl_set_insert()` copies both *key* and *value* so stack-allocated keys and values are acceptable inputs. The run-time library aborts with an error message if the key is already in the set.

`void ncptl_set_walk (NCPTL_SET *set, void (*userfunc)(void *, void *))` [Function]

Execute function *userfunc* for every $\langle key:value \rangle$ pair in a set. *userfunc* must take two `void *` values as input: a pointer to a key and a pointer to a value. The order in which keys and values are passed to the function is unspecified.

`void ncptl_set_remove (NCPTL_SET *set, void *key)` [Function]
Remove a key and its associated value from a set. The run-time library aborts with an error message if the key is not found in the set.

`void ncptl_set_empty (NCPTL_SET *set)` [Function]
Empty a set, freeing the memory that had been allocated for its contents. The set itself can continue to be used and should be deallocated with `ncptl_free()` (see [Section 6.3.3 \[Memory-allocation functions\], page 147](#)) when no longer needed.

`ncptl_int ncptl_set_length (NCPTL_SET *set)` [Function]
Return the number of $\langle key:value \rangle$ pairs in set *set*.

6.3.11 Language-visible functions

The CONCEPTUAL language contains a number of built-in functions that perform various operations on floating-point numbers (used when writing to a log file or the standard output device) and integers (used at all other times) and that determine the IDs of neighboring tasks on a variety of topologies. Each function occurs in two forms: `ncptl_func_function`, which maps `ncptl_ints` to `ncptl_ints`, and `ncptl_dfunc_function`, which maps `doubles` to `doubles`. See [Section 4.2.2 \[Built-in functions\]](#), page 81, for additional details about each function’s semantics.

Although some of the functions described in this section are fairly simple, including them in the run-time library ensures that each function returns the same value across different backends and across different platforms.

Integer functions

`ncptl_int ncptl_func_sqrt (ncptl_int num)` [Function]

`double ncptl_dfunc_sqrt (double num)` [Function]

`ncptl_func_sqrt()` returns the unique integer x such that $x^2 \leq num \wedge (x + 1)^2 > num$. `ncptl_dfunc_sqrt()` returns \sqrt{num} in double-precision arithmetic.

`ncptl_int ncptl_func_cbrt (ncptl_int num)` [Function]

`double ncptl_dfunc_cbrt (double num)` [Function]

`ncptl_func_cbrt()` returns the unique integer x such that $x^3 \leq num \wedge (x + 1)^3 > num$. `ncptl_dfunc_cbrt()` returns $\sqrt[3]{num}$ in double-precision arithmetic.

`ncptl_int ncptl_func_root (ncptl_int root, ncptl_int num)` [Function]

`double ncptl_dfunc_root (double root, double num)` [Function]

Return the *root*th root of a number *num*. More precisely, `ncptl_dfunc_root()` returns $\sqrt[{\text{root}}]{num}$ while `ncptl_func_root()` returns the largest-in-magnitude integer x with the same sign as *num* such that $|x^{\text{root}}| \leq |num|$. Currently, *num* must be nonnegative but this may change in a future release of CONCEPTUAL.

`ncptl_int ncptl_func_bits (ncptl_int num)` [Function]

`double ncptl_dfunc_bits (double num)` [Function]

Return the minimum number of bits needed to represent a given integer. (*num* is rounded up to the nearest integer in the case of `ncptl_dfunc_bits()`.)

`ncptl_int ncptl_func_shift_left (ncptl_int num, ncptl_int bits)` [Function]

`double ncptl_dfunc_shift_left (double num, double bits)` [Function]

Shift a number to the left by *bits* bits. This corresponds to multiplying *num* by 2^{bits} . In the case of `ncptl_dfunc_shift_left()`, *num* and *bits* are first converted to `ncptl_int` values. There are no corresponding `ncptl_func_shift_right()` or `ncptl_dfunc_shift_right()` functions because shifting right by x is defined to be equivalent to shifting left by $-x$.

`ncptl_int ncptl_func_log10 (ncptl_int num)` [Function]

`double ncptl_dfunc_log10 (double num)` [Function]

Return $\log_{10}(num)$ In the case of `ncptl_func_log10()`, this value is rounded down to the nearest integer.

`ncptl_int ncptl_func_factor10 (ncptl_int num)` [Function]

`double ncptl_dfunc_factor10 (double num)` [Function]

Return *num* rounded down to the nearest single-digit factor of a power of 10.

`ncptl_int ncptl_func_abs (ncptl_int num)` [Function]

`double ncptl_dfunc_abs (double num)` [Function]

Return $|num|$. In the case of `ncptl_func_log10()`, this value is rounded down to the nearest integer.

`ncptl_int ncptl_func_power (ncptl_int base, ncptl_int exponent)` [Function]

`double ncptl_dfunc_power (double base, double exponent)` [Function]

Return *base* raised to the power of *exponent*.

`ncptl_int ncptl_func_modulo (ncptl_int numerator, ncptl_int denominator)` [Function]

`double ncptl_dfunc_modulo (double numerator, double denominator)` [Function]

Return the remainder of dividing *numerator* by *denominator*. The result is guaranteed to be a nonnegative integer. `ncptl_dfunc_modulo()` rounds each of *numerator* and *denominator* down to the nearest integer before dividing and taking the remainder.

`ncptl_int ncptl_func_min (ncptl_int count, ...)` [Function]

`double ncptl_dfunc_min (double count, ...)` [Function]

Return the minimum of a list of numbers. The first argument specifies the number of remaining arguments and must be a positive integer.

`ncptl_int ncptl_func_max (ncptl_int count, ...)` [Function]

`double ncptl_dfunc_max (double count, ...)` [Function]

Return the maximum of a list of numbers. The first argument specifies the number of remaining arguments and must be a positive integer.

Floating-point functions

`ncptl_int ncptl_func_floor (ncptl_int num)` [Function]

`double ncptl_dfunc_floor (double num)` [Function]

Return $\lfloor num \rfloor$. (This is the identity function in the case of `ncptl_func_floor()`.)

`ncptl_int ncptl_func_ceiling (ncptl_int num)` [Function]

`double ncptl_dfunc_ceiling (double num)` [Function]

Return $\lceil num \rceil$. (This is the identity function in the case of `ncptl_func_ceiling()`.)

`ncptl_int ncptl_func_round (ncptl_int num)` [Function]

`double ncptl_dfunc_round (double num)` [Function]

Return *num* rounded to the nearest integer. (This is the identity function in the case of `ncptl_func_round()`.)

Topology functions

In the following functions, the ‘dfunc’ versions merely cast their arguments to `ncptl_ints` and call the corresponding ‘func’ versions.

`ncptl_int ncptl_func_tree_parent (ncptl_int task, ncptl_int` [Function]
`arity)`

`double ncptl_dfunc_tree_parent (double task, double arity)` [Function]
 Return task *task*’s parent in an *arity*-ary tree.

`ncptl_int ncptl_func_tree_child (ncptl_int task, ncptl_int child,` [Function]
`ncptl_int arity)`

`double ncptl_dfunc_tree_child (double task, double child, double` [Function]
`arity)`
 Return child *child* of task *task* in an *arity*-ary tree.

`ncptl_int ncptl_func_grid_coord (ncptl_int vartask, ncptl_int` [Function]
`coord, ncptl_int width, ncptl_int height, ncptl_int depth)`

`double ncptl_dfunc_grid_coord (double vartask, double coord,` [Function]
`double width, double height, double depth)`
 Return task *task*’s *x* coordinate (*coord* = 0), *y* coordinate (*coord* = 1), or *z* coordinate (*coord* = 2) on a *width* × *height* × *depth* mesh (or torus).

`ncptl_int ncptl_func_grid_neighbor (ncptl_int task, ncptl_int` [Function]
`torus, ncptl_int width, ncptl_int height, ncptl_int depth, ncptl_int xdelta,`
`ncptl_int ydelta, ncptl_int zdelta)`

`double ncptl_dfunc_grid_neighbor (double task, double torus,` [Function]
`double width, double height, double depth, double xdelta, double ydelta,`
`double zdelta)`

Return one of task *task*’s neighbors—not necessarily an immediate neighbor—on a 3-D mesh or torus. For the following explanation, assume that task *task* lies at coordinates (*x*, *y*, *z*) on a *width* × *height* × *depth* mesh or torus. In the mesh case (*torus* = 0), the value returned is the task ID corresponding to coordinates (*x* + *xdelta*, *y* + *ydelta*, *z* + *zdelta*). In the torus case (*torus* = 1), the value returned is the task ID corresponding to coordinates ((*x* + *xdelta*) mod *width*, (*y* + *ydelta*) mod *height*, (*z* + *zdelta*) mod *depth*).

Note that there are no 1-D or 2-D grid functions. Instead, the appropriate 3-D function should be used with *depth* and—in the 1-D case—*height* set to ‘1’.

`ncptl_int ncptl_func_knomial_parent (ncptl_int task, ncptl_int` [Function]
`arity, ncptl_int numtasks)`

`double ncptl_dfunc_knomial_parent (double task, double arity,` [Function]
`double numtasks)`

Return task *task*’s parent in an *arity*-nomial tree of *numtasks* tasks.

`ncptl_int ncptl_func_knomial_child (ncptl_int task, ncptl_int` [Function]
`child, ncptl_int arity, ncptl_int numtasks, ncptl_int count_only)`

`double ncptl_dfunc_knomial_child (double task, double child, [Function]
 double arity, double numtasks, double count_only)`

If *count_only* is '0', return task *task*'s *child*th child in an *arity*-nomial tree of *numtasks* tasks. If *count_only* is '1', return the number of children task *task* has in an *arity*-nomial tree of *numtasks* tasks.

Random-number functions

`ncptl_int ncptl_func_random_uniform (ncptl_int lower_bound, [Function]
 ncptl_int upper_bound)`

`double ncptl_dfunc_random_uniform (double lower_bound, double [Function]
 upper_bound)`

Return a number in the interval [*lower_bound*, *upper_bound*) selected at random with a uniform distribution.

`ncptl_int ncptl_func_random_gaussian (ncptl_int mean, ncptl_int [Function]
 stddev)`

`double ncptl_dfunc_random_gaussian (double mean, double [Function]
 stddev)`

Return a number selected at random from a Gaussian distribution with mean *mean* and standard deviation *stddev*.

`ncptl_int ncptl_func_random_poisson (ncptl_int mean) [Function]`

`double ncptl_dfunc_random_poisson (double mean) [Function]`

Return an integer selected at random from a Poisson distribution with mean *mean* and standard deviation $\sqrt{\text{mean}}$.

6.3.12 Finalization functions

The following function should be called towards the end of the generated code's execution.

`void ncptl_finalize (void) [Function]`

Shut down the CONCEPTUAL run-time library. No run-time library functions should be invoked after `ncptl_finalize()`.

`void ncptl_fatal (char *format, ...) [Function]`

Output an error message and abort the program. `ncptl_fatal()` takes the same types of arguments as C's `printf()` routine.

7 Tips and Tricks

The following sections present some ways to make better use of `CONCEPTUAL` in terms of producing simpler, more efficient programs or being able to run on complex computer systems.

7.1 Using out-of-bound task IDs to simplify code

See [Section 4.3 \[Task descriptions\], page 92](#), mentions a language feature that can substantially simplify `CONCEPTUAL` programs: Operations involving out-of-bound task IDs are silently ignored. The beauty of this feature is that it reduces the need for special cases at network boundaries. Consider, for example, a simple pipeline pattern in which each task in turn sends a message to the subsequent task:

```
ALL TASKS t SEND A 64 DOUBLEWORD MESSAGE TO TASK t+1.
```

Because implicit receives are posted before the corresponding sends (see [Section 4.4.2 \[Sending\], page 98](#)), all tasks except task 0 start by posting a blocking receive. (No task is sending to task 0.) Task 0 is therefore free to send a message to task 1. Receipt of that message unblocks task 1, who then sends a message to task 2, thereby unblocking task 3, and so forth. Without needing an explicit special case in the program, task `'num_tasks-1'` receives a message from task `'num_tasks-2'` but does not attempt to send a message to nonexistent task `'num_tasks'`, thanks to the rule that communication with nonexistent tasks turns into a no-op (i.e., is elided from the program).

As a more complex variation of the same program, consider a wavefront communication pattern that progresses from the upper-left corner of a mesh to the lower-right corner. Such a pattern can be expressed in just four lines of `CONCEPTUAL` (receive left, receive up, send right, send down) by relying on the property that communication with a nonexistent task is simply not executed:

```
TASK MESH_NEIGHBOR(src, xsize, +1, ysize, 0) RECEIVES A
  64 DOUBLEWORD MESSAGE FROM ALL TASKS src THEN
TASK MESH_NEIGHBOR(src, xsize, 0, ysize, +1) RECEIVES A
  64 DOUBLEWORD MESSAGE FROM ALL TASKS src THEN
ALL TASKS src SEND A 64 DOUBLEWORD MESSAGE TO
  UNSUSPECTING TASK MESH_NEIGHBOR(src, xsize, +1, ysize, 0) THEN
ALL TASKS src SEND A 64 DOUBLEWORD MESSAGE TO
  UNSUSPECTING TASK MESH_NEIGHBOR(src, xsize, 0, ysize, +1).
```

To understand the preceding program recall that `MESH_NEIGHBOR` returns `'-1'` for nonexistent neighbors. Because `'-1'` is outside of the range `[0, num_tasks)` communication with a nonexistent neighbor is ignored. To help the reader understand the preceding program, we present a trace of the events it posts as it runs with a 2×2 arrangement of tasks:

```
[TRACE] phys: 0 | virt: 0 | action: SEND | event: 1 / 2 | lines: 3 - 3
[TRACE] phys: 1 | virt: 1 | action: RECV | event: 1 / 2 | lines: 1 - 1
[TRACE] phys: 2 | virt: 2 | action: RECV | event: 1 / 2 | lines: 2 - 2
[TRACE] phys: 3 | virt: 3 | action: RECV | event: 1 / 2 | lines: 1 - 1

[TRACE] phys: 0 | virt: 0 | action: SEND | event: 2 / 2 | lines: 4 - 4
[TRACE] phys: 1 | virt: 1 | action: SEND | event: 2 / 2 | lines: 4 - 4
```

```
[TRACE] phys: 2 | virt: 2 | action: SEND | event: 2 / 2 | lines: 3 - 3
```

```
[TRACE] phys: 3 | virt: 3 | action: RECV | event: 2 / 2 | lines: 2 - 2
```

The `c_trace` backend (see [Section 3.3.4 \[The `c_trace` backend\]](#), page 23) was used to produce that trace. To increase clarity, we manually added blank lines to group concurrent events (i.e., there is no significance to the order of the `TRACE` lines within each group). The important thing to notice is that there are exactly four receives and exactly four sends:

- Although all tasks are instructed to receive a message from the left, only tasks 1 and 3 actually do so;
- although all tasks are instructed to receive a message from above, only tasks 2 and 3 actually do so;
- although all tasks are instructed to send a message to the right, only tasks 0 and 2 actually do so; and,
- although all tasks are instructed to send a message downwards, only tasks 0 and 1 actually do so.

Because communication with nonexistent tasks is elided at program initialization time there is no run-time cost for such operations—as evidenced by the `c_trace` output presented above. Furthermore, there is no reliance on the backend to drop messages from nonexistent senders or to nonexistent receivers; it is perfectly safe to utilize no-op’ed communication in any `CONCEPTUAL` program and when using any backend.

7.2 Proper use of conditionals

`CONCEPTUAL` supports two forms of conditional execution: conditional expressions (see [Section 4.2.1 \[Arithmetic expressions\]](#), page 79) and conditional statements (see [Section 4.7.4 \[Conditional execution\]](#), page 119). From the perspective of code readability and “thinking in `CONCEPTUAL`” it is generally preferable to use restricted identifiers (see [Section 4.3.1 \[Restricted identifiers\]](#), page 93) to select groups of tasks rather than a loop with a conditional as would be typical in other programming languages. For example, consider the following code in which certain even-numbered tasks each send a message to the right:

```
FOR EACH evtask IN {0, ..., num_tasks-1}
  IF evtask IS EVEN /\ evtask MOD 3 <> 2 THEN
    TASK evtask SENDS A 64 BYTE MESSAGE TO TASK evtask+1
```

While the preceding control flow is representative of that in other programming languages, `CONCEPTUAL` can express the same communication pattern without needing either a loop or an explicit conditional statement:

```
TASK evtask SUCH THAT evtask IS EVEN /\ evtask MOD 3 <> 2 SENDS A 64
  BYTE MESSAGE TO TASK evtask+1
```

One situation in which conditional statements do not have a convenient analogue is when a program selects among multiple disparate subprograms based on a command-line parameter:


```

func IS "Operation to perform (1=op1; 2=op2; 3=op3)" AND COMES FROM
"--function" OR "-f" WITH DEFAULT 1.

ASSERT THAT "the function must be 1, 2, or 3" WITH func>=1 /\ func<=3.

IF func = 1 THEN {
  Perform operation op1.
}
OTHERWISE IF func = 2 THEN {
  Perform operation op2.
}
OTHERWISE IF func = 3 THEN {
  Perform operation op3.
}

```

7.3 Memory efficiency

As described in [Section 6.2.3 \[Generated code\]](#), page 141, the `c_generic` backend (and therefore all derived backends) generates programs that run by executing a sequence of events in an event list. While this form of program execution makes it possible to hoist a significant amount of computation out of the timing loop, it does imply that a program's memory requirements are proportional to the number of statements that the program executes.

CONCEPTUAL's memory usage can be reduced by taking advantage of repeat counts within statements that support such a construct. The language's `<send_stmt>` (see [Section 4.4.2 \[Sending\]](#), page 98), `<receive_stmt>` (see [Section 4.4.3 \[Receiving\]](#), page 100), and `<touch_stmt>` (see [Section 4.6.4 \[Touching memory\]](#), page 108) are all examples of statements that accept repeat counts. For other statements and for groups of statements that repeat, the `FOR...REPETITIONS` statement produces a single `EV_REPEAT` event followed by a single instance of the events in the loop body. This technique is valid because CONCEPTUAL knows *a priori* that every iteration is identical to every other iteration. In contrast, the more general `FOR EACH` statement can induce different behavior each iteration based on the value of the loop variable so programs must conservatively instantiate the events in the loop body for every iteration. Consider the following examples:

Least efficient:

```

'FOR EACH i IN {1, ..., 1000} TASK 0 TOUCHES A 1 WORD MEMORY REGION'
(1000 EV_TOUCH events on task 0)

```

More efficient:

```

'FOR 1000 REPETITIONS TASK 0 TOUCHES A 1 WORD MEMORY REGION'
(an EV_REPEAT event and an EV_TOUCH event on task 0)

```

Most efficient:

```

'TASK 0 TOUCHES A 1 WORD MEMORY REGION 1000 TIMES'
(one EV_TOUCH event on task 0)

```

Least efficient:

```
'FOR EACH i IN {1, . . . , 1000} TASK 0 SENDS A 32 KILOBYTE MESSAGE TO
TASK 1'
```

(1000 `EV_SEND` events on task 0 and 1000 `EV_RECV` events on task 1)

More efficient:

```
'FOR 1000 REPETITIONS TASK 0 SENDS A 32 KILOBYTE MESSAGE TO TASK 1'
```

(an `EV_REPEAT` event and an `EV_SEND` event on task 0 plus an `EV_REPEAT` event and an `EV_RECV` event on task 1)

Most efficient:

```
'TASK 0 SENDS 1000 32 KILOBYTE MESSAGES TO TASK 1'
```

(currently the same as the above although a future release of `CONCEPTUAL` may reduce this to a single `EV_SEND` event on task 0 and a single `EV_RECV` event on task 1)

7.4 Cross-compilation

Some experimental architectures lack native compilers and therefore require programs to be compiled on an architecturally distinct front-end machine. `CONCEPTUAL` can be configured to support such an arrangement. As an example, we follow how `CONCEPTUAL` was recently configured to run on a Cray/SNL Red Storm prototype. The Red Storm prototype requires programs to be cross-compiled from a compile server then launched from an execution server to run on the compute nodes. The CPUs in all three node types use the x86-64 architecture but the compile server and execution server run Linux while the compute nodes run the lightweight Catamount kernel. Cross-compilation is necessary to prevent the `./configure` script which runs on the compile server from attempting to execute small programs to test run-time features—these would be guaranteed to fail.

`./configure` (see [Section 2.1 \[configure\], page 5](#)) enters cross-compilation mode if the build system—specified with `--build`—has a different CPU, comes from a different vendor, or runs a different operating system from the execution system—specified with `--host`. The `'config.guess'` script outputs the build-system type in the form *CPU-vendor-operating system*:

```
% ./config.guess
x86_64-unknown-linux-gnu
```

(That is, the CPU architecture is `'x86_64'`; the computer vendor is `'unknown'`; and, the operating system is `'linux-gnu'`.)

As a special case of cross-compilation, if the build system and execution system utilize the same CPU and operating system¹ (but presumably differ in terms of other feature), `./configure` provides a `--with-cross-compilation` option to force a cross-compile. The alternative is to modify one of the build-system name components; best is to modify the vendor component as that's used solely for informational purposes. For our Red Storm

¹ As long as the host and build operating systems have moderately similar interfaces (e.g., both are Unix-like) and the same CPU architecture (e.g., both are x86-64 variants), `--with-cross-compilation` should be applicable.

configuration (which predated the `--with-cross-compilation` option) we renamed `vendor` from ‘unknown’ to ‘cray’.

Because ‘configure’ assumes it cannot execute small test programs on the execution system, it is unable to determine valid memory-buffer alignments. Consequently, the `--with-alignment` option must also be passed to ‘configure’ to specify explicitly the minimum number of bytes at which data must be aligned. (If not specified, the minimum alignment defaults to ‘8’.) We know *a priori* that the x86-64 architecture can support byte-aligned data. Hence, we specify `--with-alignment=1`.

A command line like the following was used to configure CONCEPTUAL for the Red Storm prototype:

```
% ./configure --build=x86_64-unknown-linux-gnu --host=x86_64-cray-linux-gnu
  --with-alignment=1 CC=mpicc
```

Once CONCEPTUAL is configured to cross-compile, there is nothing special about performing the compilation itself. The `make` command runs unmodified from its description in [Section 2.2 \[make\], page 7](#).

Running `make check` can be tricky because it involves both compilation and execution. As stated previously, these cannot be performed on the same servers in the Red Storm prototype. The solution is first to run `make check` on the compile server. This compiles all of the CONCEPTUAL regression tests—and unsuccessfully attempts to run them. Then, when `make check` is run on the execution server it does not need to compile any of the tests (which it can’t do successfully, anyway) but can simply run each of them.

One catch is that the Red Storm execution server cannot directly run compute-node programs. Rather, it needs to spawn a job launcher (`yod`) for each test program. The mechanism for doing this is the `TESTS_ENVIRONMENT` variable, which `make check` prepends to every command it executes. We were therefore able to regression-test the CONCEPTUAL run-time library on the Red Storm prototype with the following command:

```
% make TESTS_ENVIRONMENT="yod -list 4" check
```

If the system for which you’re cross-compiling is unable to run Python or unable to build the interpreter-based backends, you might consider building CONCEPTUAL twice—cross-compiled for the target system and compiled regularly for the front-end system. The `--prefix` option to ‘configure’ (see [Section 2.1 \[configure\], page 5](#)) specifies the top-level directory for the CONCEPTUAL installation.

7.5 Implicit dynamic-library search paths

By default, CONCEPTUAL installs both a static and dynamic library with the dynamic library taking precedence (at least on most Unix and Unix-like operating systems). While dynamic libraries offer a number of benefits—such as not requiring applications to be relinked every time a library is upgraded—one inconvenience is the need to set the `LD_LIBRARY_PATH` environment variable to point to the CONCEPTUAL library directory if the CONCEPTUAL run-time library is not installed in a standard location (e.g., ‘`/usr/lib`’).

An alternative to setting `LD_LIBRARY_PATH` is to specify the target library directory at configuration time via the `LDFLAGS` variable and let the linker embed that directory into the CONCEPTUAL run-time library's dynamic search path. Many linkers accept a `-rpath` option for exactly that purpose. With the GNU C compiler and linker the appropriate option is `LDFLAGS="-Wl,-rpath,directory"` (e.g., `'LDFLAGS="-Wl,-rpath,/usr/local/ncpt1/lib"'`). Other compilers may have analogous mechanisms for passing flags directly to the linker.

7.6 Running without installing

You can compile CONCEPTUAL programs without having first to do a *make install* by using the same mechanism as was discussed in [Section 7.5 \[Implicit dynamic-library search paths\]](#), page 165.

First, Libtool builds the CONCEPTUAL run-time library in the `.libs` subdirectory. Hence, you should add `'LDFLAGS="-L'pwd'/.libs -Wl,-rpath,'pwd'/.libs"'` to the `./configure` command line (see [Section 2.1 \[configure\]](#), page 5) to point both the static and dynamic linkers to the CONCEPTUAL build directory.

Second, CONCEPTUAL-generated C code specifies `#include <ncpt1/ncpt1.h>`. To ensure that the generated code can find `ncpt1.h` you should add `'CPPFLAGS="-I'pwd' "'` to the `./configure` command line to point the C compiler to the CONCEPTUAL build directory. Then, create a symbolic link from the CONCEPTUAL build directory to an `ncpt1` directory by running `ln -s 'pwd' ncpt1` from the build directory.

After running `./configure` and *make*, you can compile CONCEPTUAL programs by invoking the compiler as `./ncpt1.py` instead of the usual `ncpt1`.

7.7 Reporting configuration information

On platforms that support it, the CONCEPTUAL run-time library is also an executable program that can be run from the command line. Executing the library outputs to the standard output device a complete log-file prologue and epilogue but no data.² This feature makes it quite convenient to determine all of the configuration options, compiler features, etc. that were used to build the run-time library.

² This is the same information produced by *make empty.log* (see [Section 2.2 \[make\]](#), page 7).

8 Troubleshooting

In any complex system, things are bound to go wrong. The following sections present solutions to various problems that have been encountered when building `CONCEPTUAL` and running `CONCEPTUAL` programs.

8.1 Problems with `configure`

For typical `CONCEPTUAL` usage, the most important function of the `'configure'` script is to prepare the system to build the `CONCEPTUAL` run-time library. `CONCEPTUAL`'s functionality is severely restricted without that library. This section provides information to help ensure that configuration succeeds.

8.1.1 Interpreting `configure` warnings

The `'configure'` script performs a large number of tests to ensure that `CONCEPTUAL` will compile properly and function as expected. In particular, any missing or improperly functioning feature upon which the C run-time library relies causes `./configure` to issue a `'not building the C run-time library'` warning. Without its run-time library, `CONCEPTUAL`'s functionality is severely limited so it's worth every effort to get `./configure` to build that.

Like all Autoconf scripts, `'configure'` logs detailed information to a `'config.log'` file. As a general diagnostic technique one should search for puzzling output in `'config.log'` and examine the surrounding context. For instance, on one particular system, `./configure` output `'no'` following `'checking if we can run a trivial program linked with "-lrt -lm -lpopt "'` and then refused to build the run-time library. The following relevant lines appeared in `'config.log'`:

```

configure:12845: checking if we can run a trivial program linked with "-lrt
-lm -lpopt "
configure:12862: /usr/local/bin/gcc -o conftest -g -O2 conftest.c -lrt
-lm -lpopt >&5
configure:12865: $? = 0
configure:12867: ./conftest
ld.so.1: ./conftest: fatal: libpopt.so.0: open failed: No such file or
directory
./configure: line 1: 5264 Killed                  ./conftest$ac_exeext
configure:12870: $? = 137
configure: program exited with status 137
configure: failed program was:
#line 12851 "configure"
#include "confdefs.h"

int
main (int argc, char *argv[])
{
    return 0;
}
configure:12879: result: no

```

Note the error message from ‘ld.so.1’ about ‘libpopt.so.0’ not being found. Further investigation revealed that although ‘/usr/local/bin/gcc’ knew to look in ‘/usr/local/lib/’ for shared libraries, that directory was not in the search path utilized by ‘ld.so.1’. Consequently, it couldn’t find ‘/usr/local/lib/libpopt.so.0’. The solution in this case was to add ‘/usr/local/lib/’ to the `LD_LIBRARY_PATH` environment variable before running `./configure`.

In general, ‘`config.log`’ should be the first place to look when trying to interpret warnings issued by `./configure`. Furthermore, note that certain command-line options to `./configure` (see [Section 2.1 \[configure\], page 5](#)) may help bypass problematic operations that the script stumbles over.

8.1.2 ‘PRIId64 is not a valid printf conversion specifier’

The following configuration warning was encountered on various BSD Unix systems (OpenBSD 3.4 and FreeBSD 4.10-BETA, both IA-32) and inhibited the building of the CONCEPTUAL run-time library:

```
PRIId64 is not a valid printf conversion specifier for values of type int64_t
```

The ‘`config.log`’ file indicated the source of the problem was a ‘syntax error before ‘PRIId64’ which was reported when compiling a sample program. A brief search revealed that the PRIId64 macro is not defined in any of the standard C header files on the systems in question. The solution turned out to be to configure with `./configure --with-printf-format="%lld"` to instruct the C compiler to use `printf()` format strings such as ‘%lld’ when outputting 64-bit signed integers. The extra pair of double quotes is required because the conversion specifier is used in constructs like the following:

```
printf ("The number is %10" conversion-specifier "!\\n", num);
```

Most—but apparently not all—C compilers define the `PRId64` macro (“PRInt signed decimal number of length 64 bits”) in one of the standard header files. Typical values of `PRId64` include the strings `"lld"` and `"ld"`.

8.1.3 Header is ‘present but cannot be compiled’

On some platforms `./configure` may output one or more warnings of the following form:

```
WARNING: filename.h: present but cannot be compiled
WARNING: filename.h:      check for missing prerequisite headers?
WARNING: filename.h:      see the Autoconf documentation
WARNING: filename.h:      section "Present But Cannot Be Compiled"
WARNING: filename.h: proceeding with the preprocessor's result
WARNING: filename.h: in the future, the compiler will take precedence
WARNING:      ## ----- ##
WARNING:      ## Report this to pakin@lanl.gov ##
WARNING:      ## ----- ##
```

Typically, such warnings indicate that the header file in question was written for one compiler but a different compiler is being used to build `CONCEPTUAL`. Unless the header file in question poses problems during the run of `make`, the ‘present but cannot be compiled’ warnings can be ignored and there is no need to report the compiler-preprocessor mismatch to the e-mail address shown.

8.1.4 Slow ‘checking the maximum length of command line arguments...’

Recent versions of Libtool need to know the maximum supported command-line length. Normally, ‘`configure`’ determines this value by invoking a test script with successively longer command lines until such an invocation fails. On most platforms, the maximum command-line length is determined almost instantaneously. However, on one test system, a 600 MHz *x86* running the original—not GNU—Bourne shell under Solaris, `./configure` was stuck ‘checking the maximum length of command line arguments...’ for approximately 45 *minutes*.

Inspecting the ‘`configure`’ script revealed that the result of executing the length-checking code in ‘`configure`’ is an assignment to the `lt_cv_sys_max_cmd_len` shell variable. Consequently, running `./configure lt_cv_sys_max_cmd_len=8192` (or some other conservative estimate of the maximum command-line length) caused `./configure` to skip the unduly slow length-checking test and use the given value instead.

8.1.5 ‘`configure`’ is slow

On a few systems, `./configure` has been observed to take an extremely long time to run. A common source of the problem is a slow filesystem. During the course of its execution the ‘`configure`’ script creates and compiles a large number of small files. If the current directory resides in a filesystem which exhibits poor small-file performance, then this may explain why `./configure` runs slowly. Try to configure `CONCEPTUAL` from a local filesystem (e.g., `/tmp`) and see if it runs faster.

A second source of poor `CONCEPTUAL` configuration speed regards poor implementations of the Bourne shell. One test system, a 1.9 GHz PowerPC running the original—not

GNU—Bourne shell under AIX, took several *hours* to run `./configure` to completion. Fortunately, the system administrator had installed the GNU Bourne Again shell (`'bash'`), which does not exhibit the same poor performance as the default Bourne shell. `./configure` can be instructed to use `'bash'`—or any Bourne-compatible shell—by setting the `CONFIG_SHELL` environment variable:

```
env CONFIG_SHELL=/usr/local/bin/bash ./configure
```

The preceding command reduced configuration time from over two hours to under two minutes.

8.1.6 Problems with ‘C compiler used for Python extension modules’

The following message from `'configure'` is not uncommon:

- * Not building the Python interface to the coNcEPTuaL run-time library because the C compiler used for Python extension modules (`gcc -pthread`) can't link against the output of `gcc`

Without the Python interface, none of the Python-based backends can be built, either:

- * Not installing the coNcEPTuaL interpreter because it depends upon the Python interface to the coNcEPTuaL run-time library
- * Not installing the statistics backend because it depends upon the coNcEPTuaL interpreter
- * Not installing the PICL backend because it depends upon the coNcEPTuaL interpreter
- * Not installing the LaTeX visualization backend because it depends upon the coNcEPTuaL interpreter

Note that `CONCEPTUAL` is still quite usable without the Python-based backends; the C+MPI backend (see [Section 3.3.2 \[The c_mpi backend\], page 22](#)), for example, does not rely on Python.

A typical cause of the `'can't link'` message is that the `'configure'` script is unable to locate the Python development files that are needed to build a Python module. Most Python installations exclude the development files by default, requiring that they be installed separately. In some Linux distributions these files are provided by a package with a name like `'python-dev'`. A key file to look for is `'Python.h'`. If `./configure` cannot find `'Python.h'`, the `'can't link'` message will almost certainly appear.

8.1.7 Manual configuration

One of the most important by-products of running `./configure` is a `'config.h'` file which is used to build the `CONCEPTUAL` run-time library. If the approach outlined in [Section 8.1.1 \[Interpreting configure warnings\], page 167](#) proves unable to convince `CONCEPTUAL` to build the library or other component, it may be possible to address the failed tests manually by editing any incorrect definitions in `'config.h'`. In addition, running `./configure` with the `--enable-broken-components` option will force `make` to attempt to build and install *everything*, no matter how unlikely the prospects of success are. You may need to `make` individual rules from the makefile in order to skip over unsalvageable parts of the build process.

The `--with-header-code` option to `'configure'` lets you specify a single line of C code to insert into every test file which `'configure'` generates and also into the header files used to build the run-time library. This extra line of code can be used to make up for missing functionality or to load nonstandard header files. If you need to inject more than one line of code use `./configure --with-header=code='#include "myheaders.h"'` and create an appropriate `'myheaders.h'` file.

It is highly unlikely that `./configure` will fail completely. If it does, a working CONCEPTUAL system may still be possible. Manually edit `'config.h'` for your system. (If `'config.h'` could not be created, copy `'config.h.in'` to `'config.h'`.) Do likewise for `'ncpt1.h'`. Finally, copy `'Makefile.simple.in'` to `'Makefile.simple'` if necessary and replace all text bracketed by `@` signs with appropriate values or blanks. In particular, `'@DEFS@'` should be replaced by the string `'-DHAVE_CONFIG_H'` as this instructs the various C files to include `'config.h'`. See [Section 8.2.8 \[Building on problematic platforms\], page 173](#), for more information about `'Makefile.simple'`. Although `'Makefile.in'` can be copied to `'Makefile'` and edited, doing so requires many more string replacements. The advantage is that the result will support all of the options described in [Section 2.2 \[make\], page 7](#).

8.2 Problems with make

Once CONCEPTUAL is configured, the next step is to build the various components. This section explains what to do if a compile fails or `make` is otherwise unable to perform all of its operations.

8.2.1 'Too many columns in multitable item'

Very old versions of `makeinfo` are unable to process `'conceptual.texi'` into Emacs `'info'`-format documentation; instead, they fail with a large number of `'Too many columns in multitable item (max 1)'` errors. Simply re-running `make` should bypass the failing documentation-building step.

8.2.2 Can't find 'compiler_version.h'

The CONCEPTUAL build process currently requires `make` to be executed before `make install`. Skipping the `make` step results in a `'compiler_version.h'` error resembling the following:

```
logfilefuncs.c:18:30: compiler_version.h: No such file or directory
```

Ergo, make sure always to run `make` before running `make install`.

8.2.3 'could not read symbols'

An attempt to link the CONCEPTUAL run-time library on one x86-64 Linux platform aborted with the following error:

```
/usr/lib/libpopt.so: could not read symbols: File in wrong format
```

An attempt on another x86-64 Linux platform aborted with a similar error:

```
/usr/lib/libpopt.so: could not read symbols: Invalid operation
```

The problem on both platforms turned out to be that the `/usr/lib/libpopt.so` is a 32-bit binary and could not be linked with a 64-bit library. The solution was to configure with `./configure LDFLAGS=-L/usr/lib64` to indicate that libraries should be read from `/usr/lib64` instead of the default of `/usr/lib64`.

8.2.4 Incorrect tools/flags are utilized

You might find that although you specified a particular tool or flag at configuration time (e.g., with `MPICC=mypicompiler`; see [Section 2.1 \[configure\], page 5](#)), the `ncptl` compiler seems to ignore it. A likely culprit is that an environment variable with the same name as the configuration parameter is set (e.g., the `MPICC` environment variable may be set to `othermpicompiler`) and therefore overrides all prior settings. This situation sometimes arises when a compiler or MPI library is made available using the Environment Modules system (e.g., with the `module load` command).

The solution is simply to undefine or properly redefine the offending environment variable. In most Unix environments the `env` command can be used to redefine an environment variable for the duration of a single command:

```
env MPICC=mypicompiler ncptl --backend=c_mpi myprogram.ncptl
```

8.2.5 Compaq compilers on Alpha CPUs

Although `CONCEPTUAL` builds fine on Alpha-based systems when using a `gcc` compiler, Compaq's C compilers are sometimes problematic. For instance, the `libncptl_wrap.c` source file fails to compile on a system with the following versions of the operating system, C compiler, and Python interpreter:

```
% uname -a
OSF1 qsc14 V5.1 2650 alpha
% cc -V
Compaq C V6.5-011 on Compaq Tru64 UNIX V5.1B (Rev. 2650)
Compiler Driver V6.5-003 (sys) cc Driver
% python -V
Python 2.3
```

On the system that was tested, `cc` aborts with a set of `'Missing type specifier or type qualifier'` messages. The problem appears to be that some of Compaq's standard C header files fail to `#include` various header files they depend upon. A workaround is to insert the following C preprocessor directives in `libncptl_wrap.c` before the line reading `'#include "Python.h"'`:

```
#include <sys/types.h>
#include <sys/time.h>
#include <sys/stat.h>
```

A second problem with Compaq compilers on Alpha-based systems occurs under Linux when using Compaq's `ccc` compiler:

```
% uname -a
Linux wi 2.4.21-3.7qsnet #2 SMP Fri Oct 17 14:08:00 MDT 2003 alpha unknown
% ccc -V
Compaq C T6.5-002 on Linux 2.4.21-3.7qsnet alpha
Compiler Driver T6.5-001 (Linux) cc Driver
Installed as ccc-6.5.6.002-1
Red Hat Linux release 7.2 (Enigma)
Using /usr/lib/gcc-lib/alpha-redhat-linux/2.96 (4).
```

When linking files into a shared object, ‘ccc’ aborts with a ‘[...]/.libs: file not recognized: Is a directory’ error message. The problem appears to be that ‘libtool’ is confused about the arguments it’s supposed to pass to the linker; ‘libtool’ uses ‘--rpath’ (two hyphens) where the Compaq linker expects ‘-rpath’ (one hyphen). As a workaround, you can edit the ‘libtool’ file after running `./configure` but before running `make`. Simply replace ‘--rpath’ with ‘-rpath’ in the following ‘libtool’ line and the problem should go away:

```
hardcode_libdir_flag_spec="\${wl}--rpath \${wl}\$libdir"
```

8.2.6 ‘undefined type, found ‘DEFINE_____’

The GCC 2.95.2 compiler on at least one MacOS 10.1 system complains repeatedly about an ‘undefined type, found ‘DEFINE_____’ while trying to preprocess a data file generated by ‘makehelper.py’ and based upon the contents of ‘config.h’ ‘substitutions.dat’. Fortunately, the complaint is only a warning and can be ignored. The problem is that the C preprocessor performs “smart preprocessing” if given a filename on the command line (as is the case here) but “basic preprocessing” when reading from the standard input device. Because the input does not represent syntactically correct C code—the C preprocessor is used only as a convenient device for performing macro substitutions—the syntax-aware smart preprocessing fails. However, the C preprocessor then reprocesses the file with basic preprocessing (as indicated by the message ‘cpp-precomp: warning: errors during smart preprocessing, retrying in basic mode’) and succeeds.

8.2.7 *makehelper.py* config fails

When ‘makehelper.py’ is run with the ‘config’ option, it generates a temporary ‘.c’ file which it runs through the C preprocessor. Problems (e.g., a ‘gcc.exe: no input files’ error message) may occur when running under Cygwin (a Unix-like user environment for Microsoft Windows) with a Cygwin Python interpreter but a non-Cygwin C preprocessor (e.g., MinGW’s) because ‘makehelper.py’ may try to pass a Unix-style filename to the C preprocessor, which expects to receive a Windows-style filename.

A workaround is to specify explicitly a temporary directory for ‘makehelper.py’ to use. Use forward slashes and filenames without spaces (e.g., DOS-style 8.3 filenames) as in the following example:

```
make TEMP=C:/DOCUME~1/user/LOCALS~1/Temp
```

8.2.8 Building on problematic platforms

Some experimental systems require rather specialized build procedures that thwart CONCEPTUAL’s standard makefile. Unfortunately, ‘Makefile’ is complex and difficult to edit

by hand. Users comfortable with Automake should edit ‘`Makefile.am`’—which is used to generate ‘`Makefile`’—and re-run `automake`, `autoconf`, and `./configure` as described in Section 2.2 [make], page 7.

CONCEPTUAL includes an alternate Makefile called ‘`Makefile.simple`’ (generated at configuration time from ‘`Makefile.simple.in`’). ‘`Makefile.simple`’ is a stripped-down version of ‘`Makefile`’ which is designed to be easy to edit by hand. `make -f Makefile.simple` builds a static version of the CONCEPTUAL run-time library in the current directory. `make -f Makefile.simple clean` deletes the run-time library and all of the object files used to build it. ‘`Makefile.simple`’ supports no other features. The intention is to provide the bare minimum needed to get backends such as `c_mpi` to produce executable programs even when running in unusual environments.

8.3 Problems running

After CONCEPTUAL is configured, compiled, built, and installed, there is still the chance that CONCEPTUAL-generated executables fail to run. This section addresses some common problems and presents their solutions.

As a quick tip, a very conservative way to run a CONCEPTUAL program is with the `NCPTL_NOFORK` environment variable set to ‘1’, the `NCPTL_CHECKPOINT` environment variable set to ‘0’, and with `--no-trap=1-63` on the command line. Such usage should work around some of the most common—and some of the hardest to diagnose—problems that may impact a CONCEPTUAL program. Read the corresponding sections (Section 8.3.2 [Miscellaneous mysterious hangs or crashes], page 175; Section 8.3.3 [Extremely noisy measurements], page 176; and, Section 8.3.4 [Keeping programs from dying on a signal], page 176) for information about the drawbacks of each of the preceding settings.

8.3.1 ‘cannot open shared object file’

A common problem on many workstation clusters is that the head node (on which `./configure` is run) has more libraries installed than do the compute nodes (on which CONCEPTUAL programs themselves are run). Consequently, CONCEPTUAL programs that compile and link properly on the head node will fail with a ‘cannot open shared object file’ error if they try to dynamically link a shared object that is absent on the compute nodes.

Ideally, the missing shared objects should be installed on each of the compute nodes. If they cannot be installed in their standard locations (e.g., ‘`/usr/lib`’), they can be installed elsewhere (e.g., in the user’s home directory) and that location can be pointed to using `-rpath` or `LD_LIBRARY_PATH`, as described below. Alternatively, ‘`configure`’ can be instructed to disregard certain libraries—even if they exist on the head node—by passing the `--with-ignored-libs` option to ‘`configure`’ (see Section 2.1 [configure], page 5). For example, if ‘`libpopt.so`’ is not installed on the compute nodes one can run `./configure --with-ignored-libs=popt` to prevent ‘`libpopt.so`’ from being used.

The `--with-ignored-libs` option is of little use if the CONCEPTUAL library itself cannot be found at run time. By default, `make` will build and `make install` will install both a static and a dynamic version of the CONCEPTUAL run-time library. Most linkers give precedence to the dynamic library over the static library unless the static library is requested explicitly. As a consequence, the dynamic version of the CONCEPTUAL run-time

library needs to be available at program-load time in order to avoid error messages like the following:

```
a.out: error while loading shared libraries: libncptl.so.0: cannot open
shared object file: No such file or directory
```

There are a few alternatives for pointing the dynamic loader to the CONCEPTUAL run-time library. On systems that support it, the best option is the `-rpath` approach described in [Section 7.5 \[Implicit dynamic-library search paths\]](#), page 165. The second-best option is to add the directory in which ‘libncptl.so’ was installed (by default, ‘/usr/local/lib’) to your `LD_LIBRARY_PATH` environment variable. Finally, as a last resort, you can use the `--disable-shared` configuration option (see [Section 2.1 \[configure\]](#), page 5) to prevent CONCEPTUAL from building the dynamic version of the run-time library altogether, thereby forcing the linker to use the static version:

```
make uninstall
make clean
./configure --disable-shared ...
make
make install
```

As mentioned in [Section 2.1 \[configure\]](#), page 5, however, ‘libncptlmodule.so’ can’t be built when `--disable-shared` is in effect.

8.3.2 Miscellaneous mysterious hangs or crashes

In some implementations of the software stack for InfiniBand—and possibly for some other networks as well—invocations of the `fork()` system call made while the network device is open can corrupt process memory and hang or crash the corresponding process. (On Unix and Unix-like operating systems this is typically via a segmentation fault.)

The workaround is to build CONCEPTUAL using the `--without-fork` option (see [Section 2.1 \[configure\]](#), page 5). Alternatively, you can set the `NCPTL_NOFORK` environment variable when running a CONCEPTUAL program. Either option suppresses the CONCEPTUAL run-time library’s use of all process-spawning functions. See [Appendix C \[Environment Variables\]](#), page 192, for a description of `NCPTL_NOFORK`’s side effects. The same description also applies to `--without-fork`.

The C library that comes bundled with a lightweight run-time kernel or other custom operating system may provide broken versions of some of the functions on which the CONCEPTUAL run-time library relies. If a debugger or crash-analysis tool indicates that a particular function is problematic it may be possible to disable that function at configuration time. For example, the current (at the time of this writing) release of Red Storm’s software stack includes a faulty `getrusage()` function that crashes consistently. Observing that ‘configure’ checks for `getrusage()` we searched the generated ‘config.log’ file and encountered the corresponding shell variable, ‘ac_cv_func_getrusage’, which was set to ‘yes’. Rerunning `./configure` with ‘ac_cv_func_getrusage=no’ on the command line

disabled the `CONCEPTUAL` run-time library's use of `getrusage()` and thereby resulted in crash-free execution of `CONCEPTUAL` programs.

8.3.3 Extremely noisy measurements

Sometimes a `CONCEPTUAL` program runs to completion but the data written to the log files exhibit high levels of variability across runs or even across trials within a single run. A possible source of this variability—especially for long-running programs running on a large number of processors—is the run-time library's log-file checkpointing mechanism. Because each process in a `CONCEPTUAL` programs writes its own log file, poorly scalable shared filesystems, limited spare network bandwidth, and asynchronous operating-system buffer flushes may each impact program performance in an unpredictable manner and at unpredictable times.

Setting the `NCPTL_CHECKPOINT` environment variable to '0' disables log-file checkpointing and may thereby reduce some of the data variability. The caveat is that a program that aborts abnormally will leave behind *no* data in its log files. See [Appendix C \[Environment Variables\]](#), page 192, for more information about `NCPTL_CHECKPOINT`.

8.3.4 Keeping programs from dying on a signal

By default, `CONCEPTUAL` programs terminate upon receiving *any* unexpected signal. The error message list the signal number and, if available, a human-readable signal name:

```
myprogram: Received signal 28 (Window changed); specify --no-trap=28 to ignore
```

The motivation behind this decision to abort on unexpected signals is that signal-handling adversely affects a program's performance. Hence, by running to completion, a program indicates that it did not receive any unexpected signals. However, some messaging layers use signals internally (most commonly `SIGUSR1` and `SIGUSR2`) to coordinate helper processes. To permit a program to deliver such signals to the messaging layer a user should use the program's `--no-trap` command-line option as described in [Section 3.4 \[Running coNCePTuaL programs\]](#), page 41.

8.3.5 'Unaligned access' warnings

On some platforms you may encounter messages like the following written to the console and/or various system log files (e.g., `/var/log/messages`):

```
myprog(25044): unaligned access to 0x6000000000001022,
ip=0x40000000000009e1
```

Alternatively:

```
Unaligned access pid=7890104 <myprog> va=0x140004221 pc=0x1200012b4
ra=0x1200012a4 inst=0xb449fff8
```

What's happening is that some CPUs require *n*-byte-wide data to be aligned on an *n*-byte boundary. For example, a 64-bit datatype can be accessed properly only from memory locations whose address is a multiple of 64 bits (8 bytes). On some platforms, misaligned accesses abnormally terminate the program, typically with a `SIGBUS` signal. On other platforms, misaligned accesses interrupt the operating system. The operating system fixes

the access by splitting it into multiple aligned accesses plus some bit masking and shifting and then notifying the user and/or system administrator that a fixup occurred.

CONCEPTUAL's `configure` script automatically determines what data alignments are allowed by the architecture but it has no way to determine if fixups occurred as these are transparent to programs. The result is annoying “unaligned access” messages such as those quoted above. One solution is to use the `--with-alignment` option to `configure` to specify explicitly the minimum data alignment that CONCEPTUAL should be permitted to use. Alternatively, some operating systems provide a mechanism to cause misaligned accesses to result in a SIGBUS signal instead of a fixup and notification message. On Linux/IA-64 this is achieved with the command `prctl --unaligned=signal`. On OSF1/Alpha the equivalent command is `uac p sigbus`. Be sure to rerun `configure` after issuing those commands to make it reexamine the set of valid data alignments.

8.3.6 ‘Unable to determine the OS page size’

To implement PAGE ALIGNED messages (see [Message alignment], page 97) a program needs to be able to query the operating system's page size. The CONCEPTUAL run-time library performs this query using one of the `getpagesize()` or `sysconf()` operating-system calls. If neither call is available or functional the run-time library aborts with an ‘Unable to determine the OS page size’ error.

As a workaround, `configure` provides a `--with-page-size` option which enables the user to manually specify the page size. Because any manually specified value is prone to error, CONCEPTUAL log files include a ‘Page size was specified manually at configuration time’ warning if the run-time library was configured using `--with-page-size`.

8.3.7 Invalid timing measurements

Although the CONCEPTUAL `configure` script is usually good at selecting a mechanism for measuring elapsed time, there are a few systems that confuse the script. For example, different Intel processors use different mechanisms for mapping cycle counts to time (cf. <http://en.wikipedia.org/wiki/RDTSC>); the correct mapping cannot always be determined at configuration time and may require administrator privileges to calculate. Consequently, if `configure` determines that the CONCEPTUAL timer should read the cycle counter directly, incorrect times may be reported.

The best way to test the timer quality is to run the `validatetimer` program (see [Validating the coNcePTuaL timer], page 10). If the difference between wall-clock time and CONCEPTUAL-reported time is great, the `--with-gettimeofday` configuration option (see Section 2.1 [configure], page 5) is usually a safe bet for improving accuracy, albeit at a slight cost in measurement overhead.

8.3.8 ‘TeX capacity exceeded’

TeX—and by consequence, L^ATeX—does not use dynamically allocated memory. Therefore, attempting to produce a very large diagram with the `latex_vis` backend (see Section 3.3.9 [The latex_vis backend], page 34) will likely exceed TeX's hardwired memory capacity:

```
! TeX capacity exceeded, sorry [main memory size=350001].
\psm@endnode@i ... \endgroup \psm@endmath \egroup
                                                    \use@par \@psttrue
1.489 \task
          {0} & \task{1} \\
No pages of output.
```

TeX/LaTeX’s memory capacity can be increased but the mechanism for doing so varies from one TeX distribution to another and is rarely straightforward. See <http://www.tex.ac.uk/cgi-bin/texfaq2html?label=enlarge> for a few terse pointers. In general, it is best to try to minimize the number of loop repetitions when running programs through the `latex_vis` backend.

An alternative is to use a prebuilt large-memory ‘`latex`’. Some TeX distributions come with a ‘`hugelatex`’ executable, which is just like ‘`latex`’ but compiled with larger memory limits. ‘`ncptl`’ can be told to use ‘`hugelatex`’ by setting the `LATEX` environment variable (e.g., with ‘`env LATEX=hugelatex ncptl ...`’). This may be the most convenient way to produce complex diagrams with `latex_vis` when simplifying the run is not an option.

8.3.9 Bad bounding boxes from `latex_vis`

Very tall program visualizations are susceptible to Ghostscript bug #202735, “bbox device doesn’t allow min coords < 0”. `latex_vis` attempts to work around that bug by defining a large PostScript ImagingBBox as described in http://bugs.ghostscript.com/show_bug.cgi?id=202735. While no problems with the `latex_vis` workaround have yet been reported, if problems do occur it should be sufficient to set the `GS` environment variable to ‘.’ or the name of a program known not to exist. `latex_vis` will then issue a warning message and generate an EPS file with an acceptable but slightly loose bounding box, as mentioned in Section 3.3.9 [The `latex_vis` backend], page 34.

8.4 When all else fails

The CONCEPTUAL project pages on [SourceForge.net](http://sourceforge.net) provide a variety of mechanisms for providing feedback to the CONCEPTUAL developers:

Mailing lists

Read and post messages on the CONCEPTUAL mailing lists at http://sourceforge.net/mail/?group_id=117615.

Trackers

Search for old or post new [bug reports](#), [feature requests](#), or [code patches](#) at http://sourceforge.net/tracker/?group_id=117615.

Appendix A Reserved Words

As mentioned in [Section 4.1 \[Primitives\]](#), page 77, not all identifiers can be used as variables. The following sections provide a complete list of identifiers that are forbidden as variable names. These identifiers fall into two categories: keywords, which are never allowed as variable names, and predeclared variables, which are “read-only” variables; they can be utilized just like any other variables but cannot be redeclared.

A.1 Keywords

The following is a list of all currently defined keywords in the `CONCEPTUAL` language. It is an error to try to use any of these as identifiers.

- A
- ABS
- AGGREGATES
- ALIGNED
- ALL
- AN
- AND
- ARE
- ARITHMETIC
- AS
- ASSERT
- ASSIGNED
- ASYNCHRONOUSLY
- AWAIT
- AWAITS
- BACKEND
- BE
- BIT
- BITS
- BUFFER
- BUFFERS
- BUT
- BYTE
- BYTES
- CBRT
- CEILING
- COMES
- COMPLETION
- COMPLETIONS

- COMPUTE
- COMPUTES
- COUNTERS
- CURRENT
- DATA
- DAY
- DAYS
- DEFAULT
- DEVIATION
- DIVIDES
- DOUBLEWORD
- DOUBLEWORDS
- EACH
- EVEN
- EXECUTE
- EXECUTES
- FACTOR10
- FINAL
- FLOOR
- FOR
- FROM
- GEOMETRIC
- GIGABYTE
- GREATER
- HALFWORD
- HALFWORDS
- HARMONIC
- HISTOGRAM
- HOUR
- HOURS
- IF
- IN
- INTEGER
- INTEGERS
- INTO
- IS
- IT
- ITS

- KILOBYTE
- KNOMIAL_CHILD
- KNOMIAL_CHILDREN
- KNOMIAL_PARENT
- LANGUAGE
- LESS
- LET
- LOG
- LOG10
- LOGS
- MAX
- MAXIMUM
- MEAN
- MEDIAN
- MEGABYTE
- MEMORY
- MESH_NEIGHBOR
- MESH_COORDINATE
- MESSAGE
- MESSAGES
- MICROSECOND
- MICROSECONDS
- MILLISECOND
- MILLISECONDS
- MIN
- MINIMUM
- MINUTE
- MINUTES
- MISALIGNED
- MOD
- MULTICAST
- MULTICASTS
- NONUNIQUE
- NOT
- ODD
- OF
- OR
- OTHER

- OTHERWISE
- OUTPUT
- OUTPUTS
- PAGE
- PAGES
- PLUS
- PROCESSOR
- PROCESSORS
- QUADWORD
- QUADWORDS
- RANDOM
- RANDOM_GAUSSIAN
- RANDOM_POISSON
- RANDOM_UNIFORM
- REAL
- RECEIVE
- RECEIVES
- REDUCE
- REDUCES
- REGION
- REPETITION
- REPETITIONS
- REQUIRE
- RESET
- RESETS
- RESULT
- RESULTS
- ROOT
- ROUND
- SECOND
- SECONDS
- SEND
- SENDS
- SIZED
- SLEEP
- SLEEPS
- SQRT
- STANDARD

- STRIDE
- SUCH
- SUM
- SYNCHRONIZATION
- SYNCHRONIZE
- SYNCHRONIZES
- SYNCHRONOUSLY
- TASK
- TASKS
- THAN
- THAT
- THE
- THEIR
- THEM
- THEN
- TIME
- TIMES
- TO
- TORUS_NEIGHBOR
- TORUS_COORDINATE
- TOUCH
- TOUCHES
- TOUCHING
- TREE_PARENT
- TREE_CHILD
- UNALIGNED
- UNIQUE
- UNSUSPECTING
- VALUE
- VARIANCE
- VERIFICATION
- VERSION
- WARMUP
- WHILE
- WHO
- WITH
- WITHOUT
- WORD
- WORDS
- XOR

A.2 Predeclared variables

CONCEPTUAL predeclares a set of variables that programs can use but not redeclare. These variables and their descriptions are listed below.

<code>bit_errors</code>	Total number of bit errors observed
<code>bytes_received</code>	Total number of bytes received
<code>bytes_sent</code>	Total number of bytes sent
<code>elapsed_usecs</code>	Elapsed time in microseconds
<code>msgs_received</code>	Total number of messages received
<code>msgs_sent</code>	Total number of messages sent
<code>num_tasks</code>	Number of tasks running the program
<code>total_bytes</code>	Sum of bytes sent and bytes received
<code>total_msgs</code>	Sum of messages sent and messages received

As should be evident from their descriptions, CONCEPTUAL's predeclared variables are updated dynamically. Each access can potentially return a different result. Furthermore, unlike user-declared variables, predeclared variables in can have a different value on each task.

The following notes clarify the semantics that relate to the updating of some of the preceding variables:

ASYNCHRONOUSLY (see [\[Blocking semantics\]](#), page 98)

A message that is *sent* ASYNCHRONOUSLY immediately increments each of `msgs_sent` and `total_msgs` by 1 and each of `bytes_sent` and `total_bytes` by the message size. A message that is *received* ASYNCHRONOUSLY increments none of the predeclared variables. However, after the receiving task **AWAITS COMPLETION** (see [Section 4.4.4 \[Awaiting completion\]](#), page 101) it increments each of `msgs_received` and `total_msgs` by 1 and each of `bytes_received` and `total_bytes` by the message size.

MULTICASTS (see [Section 4.4.5 \[Multicasting\]](#), page 101)

None of the byte or message variables are updated as part of a multicast operation, regardless of how the underlying messaging layer implements multicasts.

SYNCHRONIZES (see [Section 4.4.7 \[Synchronizing\]](#), page 103)

None of the byte or message variables are updated as part of a barrier synchronization, regardless of how the underlying messaging layer implements barriers.

Appendix B Backend Developer's Reference

Programmers wishing to develop their own `CONCEPTUAL` backends can refer to the information presented in this appendix as a complement to the more tutorial-in-nature [Chapter 6 \[Implementation\]](#), page 135.

B.1 Method calls

The following method calls must be defined when writing a `CONCEPTUAL` backend from scratch. They are invoked indirectly as part of PLY's AST traversal. See [Section 6.2 \[Backend creation\]](#), page 136, for more information.

- `n_add_expr`
- `n_aggregate_expr`
- `n_aggregate_func`
- `n_an`
- `n_assert_stmt`
- `n_await_completion`
- `n_backend_stmt`
- `n_buffer_number`
- `n_computes_for`
- `n_data_multiplier`
- `n_data_type`
- `n_empty_stmt`
- `n_eq_expr`
- `n_expr_list`
- `n_for_count`
- `n_for_each`
- `n_for_time`
- `n_func_call`
- `n_ident`
- `n_if_stmt`
- `n_ifelse_expr`
- `n_integer`
- `n_item_size`
- `n_let_binding`
- `n_let_binding_list`
- `n_let_stmt`
- `n_log_expr_list`
- `n_log_expr_list_elt`
- `n_log_flush_stmt`
- `n_log_stmt`

- `n_mcast_stmt`
- `n_message_alignment`
- `n_message_spec`
- `n_mult_expr`
- `n_no_touching`
- `n_output_stmt`
- `n_param_decl`
- `n_power_expr`
- `n_processor_stmt`
- `n_program`
- `n_range`
- `n_range_list`
- `n_real`
- `n_receive_attrs`
- `n_receive_stmt`
- `n_recv_buffer_number`
- `n_reduce_message_spec`
- `n_reduce_stmt`
- `n_rel_conj_expr`
- `n_rel_disj_expr`
- `n_reset_stmt`
- `n_send_attrs`
- `n_send_stmt`
- `n_simple_stmt_list`
- `n_sleeps_for`
- `n_stride`
- `n_string`
- `n_string_or_expr_list`
- `n_string_or_log_comment`
- `n_sync_stmt`
- `n_task_expr`
- `n_time_unit`
- `n_top_level_stmt`
- `n_touch_buffer_stmt`
- `n_touch_repeat_count`
- `n_touch_stmt`
- `n_touching`
- `n_trivial_node`

- `n_unary_expr`
- `n_unique`
- `n_verification`
- `n_version_decl`

B.2 C hooks

To save the backend developer from having to implement `CONCEPTUAL` backends entirely from scratch, `CONCEPTUAL` provides a `'codegen_c_generic.py'` module which defines a base class for code generators that output C code. The base class handles the features that are specific to C but independent of any messaging library. Derived classes need only define those “hook” functions that are needed to implement library-specific functionality.

Hooks are named after the method from which they're called but with an all-uppercase tag appended. The following list shows each hook-calling method in `'codegen_c_generic.py'` and the set of hooks it calls. See [Section 6.2.1 \[Hook methods\]](#), [page 137](#), for more information.

- `code_declare_datatypes`
 - `code_declare_datatypes_EXTRA_EVENTS`
 - `code_declare_datatypes_EXTRA_EVENT_STATE`
 - `code_declare_datatypes_EXTRA_EVS`
 - `code_declare_datatypes_MCAST_STATE`
 - `code_declare_datatypes_POST`
 - `code_declare_datatypes_PRE`
 - `code_declare_datatypes_RECV_STATE`
 - `code_declare_datatypes_REDUCE_STATE`
 - `code_declare_datatypes_SEND_STATE`
 - `code_declare_datatypes_SYNC_STATE`
 - `code_declare_datatypes_WAIT_STATE`
- `code_declare_globals`
 - `code_declare_globals_DUMMY_VAR`
 - `code_declare_globals_EXTRA`
- `code_def_alloc_event`
 - `code_def_alloc_event_DECLS`
 - `code_def_alloc_event_POST`
 - `code_def_alloc_event_PRE`
- `code_def_exit_handler`
 - `code_def_exit_handler_BODY`
- `code_def_finalize`
 - `code_def_finalize_DECL`
 - `code_def_finalize_POST`
 - `code_def_finalize_PRE`

- `code_def_init_check_pending`
 - `code_def_init_check_pending_POST`
 - `code_def_init_check_pending_PRE`
- `code_def_init_cmd_line`
 - `code_def_init_cmd_line_POST_ARGS`
 - `code_def_init_cmd_line_POST_PARSE`
 - `code_def_init_cmd_line_PRE_ARGS`
 - `code_def_init_cmd_line_PRE_PARSE`
- `code_def_init_decls`
 - `code_def_init_decls_POST`
 - `code_def_init_decls_PRE`
- `code_def_init_init`
 - `code_def_init_init_POST`
 - `code_def_init_init_PRE`
- `code_def_init_misc`
 - `code_def_init_misc_EXTRA`
 - `code_def_init_misc_PRE_LOG_OPEN`
- `code_def_init_msg_mem`
 - `code_def_init_msg_mem_EACH_TAG`
 - `code_def_init_msg_mem_POST`
 - `code_def_init_msg_mem_PRE`
- `code_def_init_reseed`
 - `code_def_init_reseed_BCAST`
- `code_def_init_seed`
 - `code_def_init_seed_POST`
 - `code_def_init_seed_PRE`
- `code_def_init_uuid`
 - `code_def_init_uuid_BCAST`
- `code_def_mark_used`
 - `code_def_mark_used_POST`
 - `code_def_mark_used_PRE`
- `code_def_procev`
 - `code_def_procev_DECL`
 - `code_def_procev_EVENTS_DECL`
 - `code_def_procev_EXTRA_EVENTS`
 - `code_def_procev_POST`
 - `code_def_procev_POST_SWITCH`
 - `code_def_procev_PRE`

- code_def_procev_PRE_SWITCH
- code_def_procev_arecv
 - code_def_procev_arecv_BODY
- code_def_procev_asend
 - code_def_procev_asend_BODY
- code_def_procev_etime
 - code_def_procev_etime_REDUCE_MIN
- code_def_procev_mcast
 - code_def_procev_mcast_BODY
- code_def_procev_newstmt
 - code_def_procev_newstmt_BODY
- code_def_procev_recv
 - code_def_procev_recv_BODY
- code_def_procev_reduce
 - code_def_procev_reduce_BODY
- code_def_procev_repeat
 - code_def_procev_repeat_BODY
- code_def_procev_send
 - code_def_procev_send_BODY
- code_def_procev_sync
 - code_def_procev_sync_BODY
- code_def_procev_wait
 - code_def_procev_wait_BODY_RECVS
 - code_def_procev_wait_BODY_SENDS
- code_def_small_funcs
 - code_def_small_funcs_POST
 - code_def_small_funcs_PRE
- code_define_functions
 - code_define_functions_INIT_COMM_1
 - code_define_functions_INIT_COMM_2
 - code_define_functions_INIT_COMM_3
 - code_define_functions_POST
 - code_define_functions_PRE
- code_define_macros
 - code_define_macros_POST
 - code_define_macros_PRE
- code_define_main
 - code_define_main_DECL

- `code_define_main_POST_EVENTS`
- `code_define_main_POST_INIT`
- `code_define_main_PRE_EVENTS`
- `code_define_main_PRE_INIT`
- `code_output_header_comments`
 - `code_output_header_comments_EXTRA`
- `code_specify_include_files`
 - `code_specify_include_files_POST`
 - `code_specify_include_files_PRE`

B.3 Event types

Programs generated by `'codegen_c_generic.py'` define the following event types:

<code>EV_ARECV</code>	Asynchronous receive
<code>EV_ASEND</code>	Asynchronous send
<code>EV_BTIME</code>	Beginning of a timed loop
<code>EV_DELAY</code>	Spin or sleep
<code>EV_ETIME</code>	Ending of a timed loop
<code>EV_FLUSH</code>	Compute aggregate functions for log-file columns
<code>EV_MCAST</code>	Synchronous multicast
<code>EV_NEWSTMT</code>	Beginning of a new top-level statement
<code>EV_RECV</code>	Synchronous receive
<code>EV_REDUCE</code>	Reduction with or without a subsequent multicast
<code>EV_REPEAT</code>	Repeatedly process the next N events
<code>EV_RESET</code>	Reset counters
<code>EV_SEND</code>	Synchronous send
<code>EV_SUPPRESS</code>	Suppress writing to the log and standard output
<code>EV_SYNC</code>	Barrier synchronization
<code>EV_TOUCH</code>	Touch a region of memory
<code>EV_WAIT</code>	Wait for all asynchronous sends/receives to complete
<code>EV_CODE</code>	None of the above

[Section 6.2.3 \[Generated code\], page 141](#), motivates the use of event-based execution for `CONCEPTUAL` programs.

B.4 Representing aggregate functions

The `LOG_AGGREGATE` enumerated type, defined in `'ncptl.h'`, accepts the following values:

`NCPTL_FUNC_NO_AGGREGATE`

Log all data points.

`NCPTL_FUNC_MEAN`

Log only the arithmetic mean.

`NCPTL_FUNC_HARMONIC_MEAN`

Log only the harmonic mean.

`NCPTL_FUNC_GEOMETRIC_MEAN`

Log only the geometric mean.

`NCPTL_FUNC_MEDIAN`

Log only the median.

`NCPTL_FUNC_STDEV`

Log only the standard deviation.

`NCPTL_FUNC_VARIANCE`

Log only the variance.

`NCPTL_FUNC_SUM`

Log only the sum.

`NCPTL_FUNC_MINIMUM`

Log only the minimum.

`NCPTL_FUNC_MAXIMUM`

Log only the maximum.

`NCPTL_FUNC_FINAL`

Log only the final measurement.

`NCPTL_FUNC_ONLY`

Log any data point, aborting if they're not all identical.

`NCPTL_FUNC_HISTOGRAM`

Log a histogram of the data points

Appendix C Environment Variables

The CONCEPTUAL compiler (`ncptl`) honors the following environment variables:

NCPTL_BACKEND

Name a default backend for the compiler to use. For example, setting *NCPTL_BACKEND* to `c_udgram` tells `ncptl` to use the `c_udgram` backend unless the `--backend` compiler option designates a different backend. See [Section 3.2 \[Compiling coNCePTuaL programs\]](#), page 18, for more information.

NCPTL_PATH

Specify a colon-separated list of directories in which to search for compiler backends. See [Section 3.2 \[Compiling coNCePTuaL programs\]](#), page 18, for more information.

The following environment variables are honored when running a CONCEPTUAL program (any backend):

NCPTL_CHECKPOINT

Specify the minimum number of seconds between log-file checkpoints (default: 60). The CONCEPTUAL run-time library buffers logged data in memory because CONCEPTUAL programs are not restricted to writing data in a top-to-bottom, left-to-right format. The following program, for example, writes row 9, column 2 before writing row 5, column 1:

```
FOR EACH i IN {1, ..., 10} {IF i IS EVEN THEN TASK 0 LOGS i
  AS "Even numbers" THEN TASK 0 LOGS i AS "All numbers"}
```

Buffering data in memory enables that sort of “two-dimensional” logging. However, it is not robust to computer crashes or uncatchable signals (e.g., SIGKILL). Consequently, every LOGS statement (see [Section 4.5.3 \[Writing to a log file\]](#), page 104) which is executed at least *NCPTL_CHECKPOINT* seconds after the previous one forces the CONCEPTUAL run-time library to write its partial data to the log file. Specifically, the library flushes its partial-data buffers then rewinds the write pointer to the beginning of that partial data. Hence, the partial data is overwritten by a later checkpoint or by the complete data set.

Smaller values of *NCPTL_CHECKPOINT* provide more robustness to crashes and uncatchable signals. Larger values put less stress on the filesystem. As a special case, if *NCPTL_CHECKPOINT* is set to `0` then log-file checkpointing is disabled altogether.

NCPTL_FAST_INIT

If set to `1`, more quickly initialize the run-time library by skipping the timer calibration and measurement steps. As a consequence, all timing measurements will be meaningless. *NCPTL_FAST_INIT* may be useful during the development of a CONCEPTUAL program or compiler backend to enable shorter turnaround times. If set to `0`, *NCPTL_FAST_INIT* forces a thorough initialization even for backends which do not rely on timing measurements, for instance the `picl` backend (see [Section 3.3.8 \[The picl backend\]](#), page 33). See [Section 6.3.1 \[Variables and data types\]](#), page 145, for more information.

NCPTL_LOG_DELAY

Artificially delay each log-file open and flush operation by a random number of milliseconds in the range $0 \dots NCPTL_LOG_DELAY$ (default: '0'). For example, if 1000 CPUs share a single filesystem, an *NCPTL_LOG_DELAY* of '2000' (i.e., 2 seconds) will probabilistically ensure that a log-file creation request will be issued only once every $2000 \div 1000 = 2$ milliseconds instead of all at once. *NCPTL_LOG_DELAY* is intended to help CONCEPTUAL programs run atop broken filesystems which are unable to handle large numbers of concurrent accesses—an all-too-common problem on large-scale workstation clusters and parallel computers in which hundreds or thousands of diskless compute nodes compete for access to the same filesystem.

NCPTL_LOG_ONLY

Limit the set of processes which produce log files. *NCPTL_LOG_ONLY* accepts a comma-separated list of dash-separated process ranges such as '0-3,12-16,24,25,32-48'. Only processes included in the list produce log files. See [Section 3.4 \[Running coNcEPTuaL programs\], page 41](#), for more information.

NCPTL_NOFORK

If set, inhibit the use of `fork()`, `system()`, `popen()`, and other process-spawning functions. The result is that some information will be omitted from the log-file prologue. *NCPTL_NOFORK* is intended to be used on systems in which such functions corrupt messaging-layer state, hang or crash processes, or wreak other such havoc.

Referenced Applications

A number of third-party applications are mentioned throughout this document. For your convenience, the following list states the URL of each application's home page.

a2ps	http://www.gnu.org/software/a2ps/
dot	See Graphviz.
Emacs	http://www.gnu.org/software/emacs/
Environment Modules	http://modules.sourceforge.net/
Graphviz	http://www.graphviz.org/
L ^A T _E X	http://www.latex-project.org/
MPICL	http://www.csm.ornl.gov/picl/
ParaGraph	http://www.csar.uiuc.edu/software/paragraph/
PICL	See MPICL.
pkg-config	http://pkg-config.freedesktop.org/wiki/
PSTricks	http://www.tug.org/applications/PSTricks/
Python	http://www.python.org/
SLOCCount	http://www.dwheeler.com/sloccount/
Source-highlight	http://www.gnu.org/software/src-highlite/
T _E X	See L ^A T _E X.
Vim	http://www.vim.org/

License

Copyright © 2009, Los Alamos National Security, LLC
All rights reserved.

Copyright (2009). Los Alamos National Security, LLC. This software was produced under U.S. Government contract DE-AC52-06NA25396 for Los Alamos National Laboratory (LANL), which is operated by Los Alamos National Security, LLC (LANS) for the U.S. Department of Energy. The U.S. Government has rights to use, reproduce, and distribute this software. NEITHER THE GOVERNMENT NOR LANS MAKES ANY WARRANTY, EXPRESS OR IMPLIED, OR ASSUMES ANY LIABILITY FOR THE USE OF THIS SOFTWARE. If software is modified to produce derivative works, such modified software should be clearly marked, so as not to confuse it with the version available from LANL.

Additionally, redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Los Alamos National Security, LLC, Los Alamos National Laboratory, the U.S. Government, nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY LANS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL LANS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Summary

This is a BSD license with the additional proviso that modified versions of `CONCEPTUAL` must indicate that they are, in fact, modified.

Index

#

#include 137, 138

-

-- 20, 21

--after 49, 50

--all-events 34

--annotate 36, 37

--arrow-width 37

--backend 18, 22, 23, 192

--before 49, 50

--binary-tasks 36

--booktabs 60, 62

--breakpoint 24, 25, 26

--build 164

--chdir 62, 63

--colbegin 51, 53, 54, 56, 57, 59

--colend 51, 53, 54, 56, 57, 59

--colsep 51, 53, 54, 56, 57, 59, 61

--columns 61, 65

--combegin 64

--comend 64

--comment 41, 42

--curses 24, 25, 26

--dcolumn 60

--delay 24, 26

--disable-hpet 151

--disable-proc-interrupts 48

--disable-shared 5, 175

--dumpkeys 66

--enable-broken-components 170

--enable-hpet 151

--enable-maintainer-mode 136

--envformat 61, 62

--every-event 36

--excel 52, 54, 60

--exclude 33, 60, 61, 66

--expand-lists 32

--extra-dot 40

--extract 49, 50, 51, 60, 62, 63, 64, 67

--filter 19, 20, 37, 38

--force-merge 49, 50, 66

--format 32, 33, 40, 49, 50, 51, 53, 54, 56, 57,

58, 60, 61, 62, 63, 64, 65, 67, 68

--frequency 34

--hcolbegin 52, 53, 55, 56, 58, 59

--hcolend 52, 53, 55, 56, 58, 59

--hcolsep 52, 53, 55, 56, 58, 59

--help... 2, 5, 7, 18, 20, 21, 41, 48, 49, 50, 70, 71,

72, 74, 75, 120, 121, 147

--hierarchy 28, 30

--host 164

--hrowbegin 52, 53, 55, 56, 58, 59

--hrowend 52, 53, 55, 56, 58, 59

--hrowsep 52, 53, 55, 56, 58, 59

--include 60, 61, 66

--indent 64, 65

--itembegin 64, 65

--itemend 64, 65

--keep-ints 19, 36, 38, 137

--kwbegin 64

--kwend 64

--lenient 19, 20

--linebegin 63

--lineend 63

--listbegin 64, 65

--listend 64, 65

--logfile 32, 34, 35, 41, 43, 75, 76

--longtable 60, 62, 67

--man 48, 49, 50, 70, 71, 72, 74, 75

--mcastsync 28, 32, 34, 35

--memcache 75, 76

--merge 52, 54, 55, 57, 58, 60, 66, 67

--messages 121

--monitor 24, 26

--newlines 62, 63

--no-attrs 40

--no-compile 19, 22, 23, 35, 40

--no-lines 40

--no-link 19, 22, 23, 40

--no-source 40

--no-trap 41, 43, 146, 174, 176

--node-code 40

--noenv 61, 62

--noheaders 51, 53, 54, 56, 57, 58

--noparams 61, 62

--output 20, 28, 49, 51, 71, 72, 137

--prefix 5, 11, 165

--procs 49, 50, 75, 76

--profile 26

--program 19, 20, 137

--quiet 19, 49, 50, 75, 76, 137

--quote 52, 54, 55, 57, 58, 60

--reduce 22

--rowbegin 51, 53, 54, 56, 57, 59, 61

--rowend 52, 53, 55, 56, 57, 59, 61

--rowsep 51, 53, 54, 56, 59

--seed 41, 43, 147

--showfnames 52, 54, 55, 57, 58, 60, 66

--simplify 71, 72, 73

--sort 61, 62

--source-lines 37

--ssend 22

--stagger 37

--strbegin 64

--strend 64

--tablebegin 52, 54, 55, 57, 58, 59

--tableend 52, 54, 55, 57, 58, 60

--tablesep 52, 54, 55, 57, 58, 59

--tabularx 62, 67
 --tasks 23, 28, 32, 34, 35
 --this 67
 --trace 23, 26
 --unquote 52, 54, 55, 57, 58, 60
 --unset 62, 63
 --usage 41, 48, 49, 70, 71, 72, 74, 75
 --verbose 49, 50, 51
 --with-alignment 165, 177
 --with-const-suffix 145
 --with-cross-compilation 164, 165
 --with-datatype 145
 --with-gettimeofday 6, 11, 150, 151, 177
 --with-header-code 171
 --with-ignored-libs 5, 174
 --with-mpi-wtime 6, 22, 150, 151
 --with-page-size 177
 --without-fork 5, 175
 --wrap 64, 65
 --zero-latency 37
 -? 41, 120, 121, 147
 -a 50
 -b 18, 50
 -B 24, 25
 -c 19
 -C 2, 41
 -D 24
 -e 50
 -E 19
 -f 19, 50
 -F 50
 -h 50, 72, 75
 -K 19
 -L 19, 41, 75
 -m 50, 72, 75, 121
 -M 24, 76
 -o 20, 51, 72
 -p 20, 50, 76
 -q 19, 50, 76
 -rpath 166, 174, 175
 -s 72, 73
 -S 41, 147
 -u 72, 75
 -v 50
 -W 41

 .
 .libs 166

 /
 /proc/interrupts 48
 /var/log/messages 176

 -
 __init__ 139

A

A 78, 95, 96, 125, 179
 A HISTOGRAM OF 90, 124
 A HISTOGRAM OF THE 152
 A RANDOM PROCESSOR 110, 127, 140, 154
 A RANDOM TASK 90, 117, 118, 128, 140, 154
 A RANDOM TASK OTHER THAN 154
 a.out 20
 a2ps 9
 a2ps 12
 a2ps 194
 ABS 81, 82, 123, 179
 abstract-syntax tree (AST) .. 39, 40, 136, 143, 185
 acinclude.m4 135, 136
 aclocal 8, 136
 Add conditional (GUI menu item) 17
 Add Row (GUI button) 17
 add_expr 81, 123
 Advanced (GUI menu) 17
 aggr_expr 90, 104, 124, 126
 aggr_func 90, 91, 105, 124
 aggregate expressions 90
 aggregate functions 91, 191
 AGGREGATES 179
 ALIGNED 94, 95, 97, 124, 125, 147, 179
 alignment, message 97
 ALL 179
 ALL MESSAGE BUFFERS 109, 127
 ALL OTHER TASKS 94, 124
 ALL TASKS 93, 94, 124
 AN 78, 95, 96, 108, 125, 127, 179
 AND 104, 105, 111, 112, 117, 126, 127, 179
 AND A SYNCHRONIZATION 112, 113, 127
 AND COMES FROM 120, 128
 ANNOTATIONS 37
 ARE 179
 ARITHMETIC 91, 124, 179
 arithmetic expressions 79
 ARITHMETIC MEAN 91
 AS 95, 104, 125, 126, 179
 ASSERT 179
 ASSERT THAT 107, 126
 assert_stmt 107, 112, 126, 127
 ASSIGNED 179
 assigning values to variables 117
 AST 136
 AST (abstract-syntax tree) .. 39, 40, 136, 143, 185
 ASYNCHRONOUSLY 30, 95, 98, 99, 100, 101, 102,
 124, 125, 126, 179, 184
 autoconf 8, 136, 174
 Autoconf 3, 8, 135, 167
 autoheader 8, 136
 automake 8, 136, 174
 Automake 3, 8, 135, 174
 average *see* MEAN
 AWAIT 179
 AWAIT COMPLETION 16, 101
 AWAITs 179

- AWAITS COMPLETION 30, 101, 110, 126, 184
awk 27, 33
- ## B
- BACKEND 179
backend creation 136
BACKEND EXECUTE 111, 133
BACKEND EXECUTES 80, 111, 127, 133, 141
backend_desc 139
backend_name 139
backend_stmt 111, 112, 127
backends, supplied 20
backquoted commands 42
backslash 78
base_global_parameters 144
bash 42, 170
BE 117, 128, 179
binding variables 117
BIT 96, 125, 179
bit_errors 184
BITS 81, 82, 123, 179
blocking 98
bounding box 35, 178
Bourne shell 5, 169, 170
breakpoints 24, 25, 26
BUFFER 95, 124, 125, 179
BUFFERS 179
buffers, message 98, 109, 148
build directory 166
building 135, 173
built-in functions 81
BUT 179
BUT NOT 117, 128
buttons, GUI 17, 18
BYTE 96, 125, 179
BYTES 179
bytes_received 184
bytes_sent 184
- ## C
- C code 111
C preprocessor 6, 7, 145, 172, 173
c_generic ('codegen_c_generic.py') ... 21, 23, 26, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 154, 163, 187, 190
c_mpi ('codegen_c_mpi.py') .. 18, 20, 21, 22, 23, 33, 133, 174
c_profile ('codegen_c_profile.py') ... 20, 21, 26
c_seq ('codegen_c_seq.py') 20, 21, 23, 139, 140
c_trace ('codegen_c_trace.py') 18, 20, 21, 23, 24, 25, 26, 162
c_udgram ('codegen_c_udgram.py') .. 19, 20, 21, 23, 27, 141, 143, 192
can't link against the output of gcc 170
case sensitivity 78
Catamount 164
CBRT 81, 82, 123, 179
cc 4, 172
CC 6, 7, 22, 23
ccc 4, 172, 173
CEILING 81, 82, 83, 123, 179
CFLAGS 6, 22
char * 145
clock, incorrect readings 177
clock_gettime 150
CLOCK_REALTIME 150
CLOCK_SGI_CYCLE 150
code examples 129
code, generated 141
code, injecting arbitrary 111
code_declare_datatypes 187
code_declare_datatypes_EXTRA_EVENT_STATE 187
code_declare_datatypes_EXTRA_EVENTS 187
code_declare_datatypes_EXTRA_EVS 187
code_declare_datatypes_MCAST_STATE 187
code_declare_datatypes_POST 187
code_declare_datatypes_PRE 187
code_declare_datatypes_RECV_STATE 187
code_declare_datatypes_REDUCE_STATE 187
code_declare_datatypes_SEND_STATE ... 142, 187
code_declare_datatypes_SYNC_STATE 187
code_declare_datatypes_WAIT_STATE 187
code_declare_globals 187
code_declare_globals_DUMMY_VAR 187
code_declare_globals_EXTRA 187
code_declare_var 143, 144
code_def_alloc_event 187
code_def_alloc_event_DECLS 187
code_def_alloc_event_POST 187
code_def_alloc_event_PRE 187
code_def_exit_handler 187
code_def_exit_handler_BODY 187
code_def_finalize 187
code_def_finalize_DECL 187
code_def_finalize_POST 187
code_def_finalize_PRE 187
code_def_init_check_pending 188
code_def_init_check_pending_POST 188
code_def_init_check_pending_PRE 188
code_def_init_cmd_line 188
code_def_init_cmd_line_POST_ARGS 188
code_def_init_cmd_line_POST_PARSE 188
code_def_init_cmd_line_PRE_ARGS 188
code_def_init_cmd_line_PRE_PARSE 188
code_def_init_decls 188
code_def_init_decls_POST 188
code_def_init_decls_PRE 188
code_def_init_init 188
code_def_init_init_POST 188
code_def_init_init_PRE 188
code_def_init_misc 188

- code_def_init_misc_EXTRA 188
- code_def_init_misc_PRE_LOG_OPEN 188
- code_def_init_msg_mem 188
- code_def_init_msg_mem_EACH_TAG 188
- code_def_init_msg_mem_POST 188
- code_def_init_msg_mem_PRE 188
- code_def_init_reseed 188
- code_def_init_reseed_BCAST 140, 188
- code_def_init_seed 188
- code_def_init_seed_POST 188
- code_def_init_seed_PRE 188
- code_def_init_uuid 188
- code_def_init_uuid_BCAST 188
- code_def_main_newstmt 154
- code_def_mark_used 188
- code_def_mark_used_POST 188
- code_def_mark_used_PRE 188
- code_def_procev 188
- code_def_procev_arecv 189
- code_def_procev_arecv_BODY 189
- code_def_procev_asend 189
- code_def_procev_asend_BODY 189
- code_def_procev_DECL 188
- code_def_procev_etime 189
- code_def_procev_etime_REDUCE_MIN 189
- code_def_procev_EVENTS_DECL 188
- code_def_procev_EXTRA_EVENTS 188
- code_def_procev_mcast 189
- code_def_procev_mcast_BODY 189
- code_def_procev_newstmt 189
- code_def_procev_newstmt_BODY 189
- code_def_procev_POST 188
- code_def_procev_POST_SWITCH 188
- code_def_procev_PRE 188
- code_def_procev_PRE_SWITCH 189
- code_def_procev_recv 189
- code_def_procev_recv_BODY 189
- code_def_procev_reduce 189
- code_def_procev_reduce_BODY 189
- code_def_procev_repeat 189
- code_def_procev_repeat_BODY 189
- code_def_procev_send 189
- code_def_procev_send_BODY 189
- code_def_procev_sync 189
- code_def_procev_sync_BODY 189
- code_def_procev_wait 189
- code_def_procev_wait_BODY_RECVS 189
- code_def_procev_wait_BODY_SENDS 189
- code_def_small_funcs 189
- code_def_small_funcs_POST 189
- code_def_small_funcs_PRE 189
- code_define_functions 189
- code_define_functions_INIT_COMM_1 189
- code_define_functions_INIT_COMM_2 189
- code_define_functions_INIT_COMM_3 189
- code_define_functions_POST 189
- code_define_functions_PRE 189
- code_define_macros 189
- code_define_macros_POST 189
- code_define_macros_PRE 189
- code_define_main 189
- code_define_main_DECL 189
- code_define_main_POST_EVENTS 190
- code_define_main_POST_INIT 190
- code_define_main_PRE_EVENTS 190
- code_define_main_PRE_INIT 190
- code_output_header_comments 190
- code_output_header_comments_EXTRA 190
- code_specify_include_files 137, 138, 190
- code_specify_include_files_POST 138, 190
- code_specify_include_files_PRE 138, 190
- codegen_c_generic.py .. 21, 23, 26, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 154, 163, 187, 190
- codegen_c_mpi.py 18, 20, 21, 22, 23, 33, 133, 174
- codegen_c_profile.py 20, 21, 26
- codegen_c_seq.py 20, 21, 23, 139, 140
- codegen_c_trace.py 18, 20, 21, 23, 24, 25, 26, 162
- codegen_c_udgram.py 19, 20, 21, 23, 27, 141, 143, 192
- codegen_dot_ast.py 21, 38, 39, 40, 78, 136
- codegen_interpret.py.. 20, 27, 28, 29, 30, 32, 34, 35
- codegen_language_library.py 135, 136
- codegen_latex_vis.py.. 19, 21, 34, 35, 36, 37, 38, 177, 178
- codegen_picl.py 21, 33, 34, 146, 192
- codegen_stats.py 21, 31, 32, 33
- codestack 143
- COLORS 37
- combine_to_marker 144
- COMES 179
- Command line options (GUI menu item) 17
- command-line arguments 120
- Communicate (GUI button) 17
- COMMUNICATION 37
- communication diagrams 34
- communication statements 94
- compile_and_link 136, 137
- compile_only 136, 137
- compiler, incorrect 172
- compiler_version.h 171
- compiling 18, 135, 172
- COMPLETION 179
- COMPLETIONS 179
- complex_stmt 112, 119, 121, 122, 127, 128
- components, GUI 15
- computation, simulating 107, 108
- COMPUTE 16, 78, 107, 180
- Compute (GUI button) 17
- COMPUTES 27, 33, 78, 107, 126, 151, 180
- COMPUTES AGGREGATES 106, 113, 126, 153
- CONC_SEND_EVENT 142
- coNcEPTUAL GUI 10, 14, 15, 16, 17

conceptual-1.0.tar.gz 10
 conceptual.info* 9
 conceptual.pdf 2, 9
 conceptual.texi 171
 conceptual.xml 9
 conceptual_1.0 9, 10
 cond_expr 80, 81, 123
 conditionals 119, 162
 config.guess 164
 config.h 170, 171, 173
 config.h.in 171
 config.log 7, 167, 168, 175
 config.status 7
 CONFIG_SHELL 170
 configuration information 166
 configuration, manual 170
 configure 4, 5, 6, 7, 8, 9, 11, 13, 22, 41, 135,
 136, 145, 150, 151, 164, 165, 166, 167, 168,
 169, 170, 171, 172, 173, 174, 175, 177
 configure.ac 135, 136
 constants, language 77
 constants, run-time library 145
 Copy (GUI menu item) 17
 counters 106
 COUNTERS 180
 cpp-precomp 173
 CPPFLAGS 6, 7, 22, 23
 crashed processes 175
 cross-compilation 164
 csh 42
 CURRENT 180
 curses 20, 24, 25
 Cut (GUI menu item) 17
 Cygwin 4, 173

D

DATA 180
 data touching 97, 108
 data types, run-time library 145
 data_multiplier 96, 97, 125
 data_type 96, 97, 103, 108, 125, 127
 DAY 180
 DAYS 107, 126, 180
 dcllock 151
 DEADLOCK 37
 DEBUG 8
 debugging 23
 DEFAULT 180
 delay_stmt 107, 112, 126, 127
 delaying execution 107
 Delete (GUI button) 17
 DEVIATION 180
 diagrams, communication 34
 DIVIDES 91, 92, 124, 180
 DOCUMENT 37
 dot 21, 38
 dot 40

dot 194
 DOT 40
 dot_ast ('codegen_dot_ast.py') 21, 38, 39,
 40, 78, 136
 double 157
 DOUBLEWORD 96, 97, 103, 125, 180
 DOUBLEWORDS 95, 125, 180
 dvips 35
 DVIPS 35
 dynamic libraries 165, 174

E

EACH 90, 124, 180
 ecc 4, 6
 Edit (GUI menu) 17
 elapsed_usecs 106, 107, 117, 184
 Emacs 9, 10, 12, 171, 194
 empty.log 9
 END 37
 enumerated_exprs 81, 123
 env 172
 Environment Modules 9, 172
 environment variables .. 2, 6, 7, 8, 9, 10, 18, 22, 23,
 28, 35, 40, 43, 44, 146, 165, 166, 168, 170, 172,
 174, 175, 176, 178, 192, 193
 eq_expr 92, 124
 error_fatal 143
 error_internal 143
 escape character 78
 EV_ARECV 190
 EV_ASEND 190
 EV_BTIME 190
 EV_CODE 141, 190
 EV_DELAY 190
 EV_ETIME 190
 EV_FLUSH 190
 EV_MCAST 190
 EV_NEWSTMT 190
 EV_RECV 164, 190
 EV_REDUCE 190
 EV_REPEAT 163, 164, 190
 EV_RESET 190
 EV_SEND 164, 190
 EV_SUPPRESS 190
 EV_SYNC 190
 EV_TOUCH 163, 190
 EV_WAIT 190
 evaluation contexts 80
 EVEN 180
 event types, defined by 'codegen_c_generic.py'
 190
 events_used 144
 example programs 129
 EXECUTE 180
 EXECUTES 180
 expr .. 79, 80, 81, 90, 91, 92, 93, 94, 95, 96, 97, 98,
 104, 107, 108, 109, 110, 111, 112, 113, 114,

115, 117, 118, 120, 121, 123, 124, 125, 126,
127, 128, 152

expr1 80

expr2 80

expressions, aggregate 90

expressions, arithmetic 79

expressions, relational 91

Extend (GUI button) 18

F

FACTOR10 81, 82, 123, 180

File (GUI menu) 17

filetype.vim 12

FINAL 91, 124, 180

finalization functions 160

flags, incorrect 172

floating-point functions 82, 158

FLOOR 81, 82, 83, 123, 180

flush_stmt 106, 112, 126, 127

FOR 16, 107, 109, 112, 113, 117, 118, 119, 126,
127, 129, 151, 163, 180

FOR EACH 80, 112, 113, 114, 116, 127, 163

fork 5, 6, 141, 175, 193

fprintf 20, 24

FROM 95, 101, 124, 125, 126, 180

FROM BUFFER 23, 98, 111

FROM THE DEFAULT BUFFER 98

func_name 81, 123

function 79

functions, aggregate 91, 191

functions, built-in 81

functions, finalization 160

functions, floating-point 82, 158

functions, initialization 146

functions, integer 82, 157

functions, language-visible 157

functions, log-file 152

functions, memory-allocation 147

functions, mesh 86

functions, message-buffer 148

functions, queue 155

functions, random-number 90, 160

functions, random-task 154

functions, set 156

functions, task-mapping 154

functions, time-related 150

functions, topology 159

functions, torus 88

functions, tree 83, 84

G

gcc 4, 172

generate 136, 137, 143

generated code 141

GEOMETRIC 180

GEOMETRIC MEAN 91, 124

get_cycles 150

getpagesize 177

getrusage 175, 176

gettimeofday 6, 150, 151

Ghostscript 35, 178

GIGABYTE 96, 125, 180

GNU 3, 4, 8, 12, 135, 166, 169, 170

grammar 77

grammar, summary 123

graphical user interface 14

graphical view of communication 34

Graphviz 21, 38, 40, 194

GREATER 180

GREATER THAN 117, 118, 128

grouping 119

gs 35

GS 35, 178

GUI buttons 17, 18

GUI menus 17

H

HALFWORD 96, 97, 125, 180

HALFWORDS 180

HARMONIC 180

HARMONIC MEAN 91, 124

High-Precision Event Timers (HPET) 151

HISTOGRAM 180

hooks 137

hot-potato benchmark 129

hot-spot benchmark 130

HOUR 180

HOURS 107, 126, 180

HPET *see* High-Precision Event Timers

hugelatex 178

hung processes 175

I

icc 4

ident 2, 78, 81, 93, 94, 112, 113, 117, 120, 121,
123, 124, 127, 128

identifiers, restricted 93

IF 79, 80, 81, 112, 119, 123, 127, 128, 180

if_stmt 119, 128

illustrations of communication patterns 34

implementation 135

IN 112, 114, 117, 127, 128, 180

indent 140

InfiniBand 5, 175

initialization functions 146

injecting arbitrary code 111

installation 3, 166
 integer 81, 123, 128
 INTEGER 96, 97, 103, 125, 180
 integer functions 82, 157
 integers 78
 INTEGERS 95, 125, 180
 internals 142
 interpret ('codegen_interpret.py') ... 20, 27,
 28, 29, 30, 32, 34, 35
 INTO 95, 125, 180
 INTO BUFFER 23, 98, 111
 INTO THE DEFAULT BUFFER 98
 invoke_hook 144
 IS 120, 128, 180
 IS ASSIGNED TO 110, 127, 155
 IS EVEN 91, 92, 124
 IS IN 91, 92, 124
 IS NOT IN 91, 92, 124
 IS ODD 91, 92, 124
 IT 180
 item_count 94, 95, 96, 124, 125
 item_size 94, 96, 97, 108, 124, 125, 127
 iteration 113, 117
 ITS 180

J

Java 10, 14
 javac 10
 Jython 10
 jythonc 10

K

KEYWORD 77
 keywords 78, 179
 KILOBYTE 96, 125, 181
 KNOMIAL_CHILD 81, 84, 123, 131, 181
 KNOMIAL_CHILDREN 81, 84, 123, 131, 181
 KNOMIAL_PARENT 81, 84, 123, 181
 kpsewhich 11

L

lang.map 12
 LANGUAGE 181
 latency benchmark 129
 latex 34, 35, 178
 L^AT_EX 9, 21, 34, 35, 37, 38, 60, 61, 65, 67, 68,
 177, 178, 194
 L^AT_EX 35, 178
 latex_vis ('codegen_latex_vis.py') ... 19, 21,
 34, 35, 36, 37, 38, 177, 178
 LD_LIBRARY_PATH 9, 165, 166, 168, 174, 175
 LDFLAGS 6, 7, 22, 23, 166
 LESS 181
 LESS THAN 117, 118, 128
 LET 112, 117, 118, 119, 127, 181

let_binding 112, 117, 127, 128
 libncptl 13
 libncptl.so 175
 libncptl_wrap.c 172
 libncptlmodule.so 5, 175
 libpopt.so 168, 171, 172, 174
 libraries, dynamic 174
 library, run-time 135, 145
 LIBS 6, 7, 22, 23
 libtool 173
 Libtool 3, 8, 135, 166, 169
 libtoolize 8
 Linux 4, 7, 150, 164, 170, 171, 172, 177
 listings 12
 locals() 137
 LOG 106, 181
 log files 44, 48, 70, 74, 103, 104, 152
 LOG_AGGREGATE 191
 log_stmt 103, 104, 105, 112, 126, 127
 log_write_prologue_timer 150
 LOG10 81, 82, 123, 181
 LOGS 16, 26, 43, 44, 80, 90, 104, 105, 113, 122,
 126, 141, 152, 181, 192
 Loop (GUI button) 17
 loops 113, 117
 lspci 42
 lt_cv_sys_max_cmd_len 169

M

M-x font-lock-mode 12
 make .. 4, 7, 8, 10, 11, 136, 165, 166, 169, 170, 171,
 173, 174, 175
 make all 10
 make check 3, 7, 8, 165
 make clean 8, 175
 make dist 10
 make distclean 8
 make docbook 9
 make empty.log 9, 166
 make gui 10
 make info 9
 make install 2, 4, 5, 8, 11, 166, 171, 174, 175
 make maintainer-clean 8
 make modulefile 9, 10
 make ncptl-logextract.html 9
 make pdf 9
 make stylesheets 9, 11
 make tags 10
 make uninstall 8, 175
 Makefile 5, 7, 9, 11, 135, 171, 173, 174
 Makefile.am 135, 136, 174
 Makefile.in 171
 Makefile.simple 171, 174
 Makefile.simple.in 171, 174
 makehelper.py 173
 makeinfo 171
 malloc 148

- MANPATH 9
 - manual configuration 170
 - math 37
 - MAX 81, 82, 123, 181
 - MAXIMUM 91, 124, 181
 - mcast_stmt 101, 102, 112, 126, 127
 - MEAN 91, 124, 181
 - Measure (GUI button) 17
 - MEDIAN 91, 106, 124, 181
 - MEGABYTE 96, 125, 181
 - MEMORY 181
 - memory efficiency 163
 - MEMORY REGION 108, 127
 - memory-allocation functions 147
 - menus, GUI 17
 - mesh functions 86
 - MESH_COORDINATE 81, 86, 88, 89, 123, 181
 - MESH_NEIGHBOR 81, 86, 88, 123, 161, 181
 - MESSAGE 181
 - message alignment 97
 - MESSAGE BUFFER 109, 127
 - message buffers 98, 109, 148
 - Message Passing Interface (MPI) . . . 1, 6, 7, 18, 20, 22, 23, 33, 41, 133, 170, 172
 - message_alignment 94, 95, 97, 124, 125
 - message_spec 94, 95, 96, 97, 99, 101, 102, 124, 125, 126
 - messages 94
 - MESSAGES 95, 124, 125, 181
 - MICROSECOND 181
 - MICROSECONDS 107, 126, 181
 - MILLISECOND 181
 - MILLISECONDS 107, 126, 181
 - MIN 81, 82, 123, 181
 - MinGW 173
 - MINIMUM 91, 124, 181
 - MINUTE 181
 - MINUTES 107, 126, 181
 - MISALIGNED 95, 97, 124, 125, 181
 - MOD 79, 81, 123, 181
 - module 10, 172
 - MODULEPATH 10
 - MPI *see* Message Passing Interface
 - MPI_Allgather 133, 134
 - MPI_Allreduce 22, 23
 - MPI_Barrier 22
 - MPI_Bcast 22
 - MPI_Comm_rank 22
 - MPI_Comm_size 22
 - MPI_Comm_split 22
 - MPI_Errhandler_create 22
 - MPI_Errhandler_set 22
 - MPI_Finalize 22
 - MPI_Init 6, 22
 - MPI_Irecv 22
 - MPI_Isend 22
 - MPI_Recv 22
 - MPI_Reduce 22, 23
 - MPI_Send 22
 - MPI_Ssend 22
 - MPI_SUM 22
 - MPI_Waitall 22
 - MPI_Wtime 6, 22, 150
 - mpicc 7
 - MPICC 7, 22, 172
 - MPICFLAGS 22
 - MPICL 33, 194
 - MPICPPFLAGS 7, 22
 - MPILDFLAGS 7, 22
 - MPILIBS 7, 22
 - mpirun 22
 - msgs_received 184
 - msgs_sent 106, 184
 - mult_expr 81, 123
 - MULTICAST 16, 181
 - Multicast (GUI button) 18
 - multicast-tree benchmark 131
 - MULTICASTS 102, 126, 133, 181, 184
 - multipliers 78, 79
- ## N
- n_add_expr 185
 - n_aggregate_expr 185
 - n_aggregate_func 185
 - n_an 185
 - n_assert_stmt 185
 - n_awaits_completion 185
 - n_backend_stmt 185
 - n_buffer_number 185
 - n_computes_for 185
 - n_data_multiplier 185
 - n_data_type 185
 - n_empty_stmt 185
 - n_eq_expr 185
 - n_expr_list 185
 - n_for_count 185
 - n_for_count_SYNC_ALL 144
 - n_for_each 185
 - n_for_time 185
 - n_func_call 185
 - n_ident 185
 - n_if_stmt 185
 - n_ifelse_expr 185
 - n_integer 185
 - n_item_size 185
 - n_let_binding 185
 - n_let_binding_list 185
 - n_let_stmt 185
 - n_log_expr_list 185
 - n_log_expr_list_elt 185
 - n_log_flush_stmt 185
 - n_log_stmt 185
 - n_mcast_stmt 186
 - n_message_alignment 186
 - n_message_spec 186

<code>n_mult_expr</code>	186	<code>ncptl_allocate_task_map</code>	154, 155
<code>n_no_touching</code>	186	<code>ncptl_assign_processor</code>	155
<code>n_output_stmt</code>	186	<code>ncptl_ast.py</code>	135, 136
<code>n_outputs</code>	143	<code>NCPTL_BACKEND</code>	18, 192
<code>n_param_decl</code>	186	<code>NCPTL_CHECKPOINT</code>	174, 176, 192
<code>n_power_expr</code>	186	<code>NCPTL_CMDLINE</code>	145
<code>n_processor_stmt</code>	186	<code>NCPTL_CodeGen</code>	136, 137, 139, 143, 144
<code>n_program</code>	186	<code>ncptl_dfunc_abs</code>	158
<code>n_range</code>	186	<code>ncptl_dfunc_bits</code>	157
<code>n_range_list</code>	186	<code>ncptl_dfunc_cbrt</code>	157
<code>n_real</code>	186	<code>ncptl_dfunc_ceiling</code>	158
<code>n_receive_attrs</code>	186	<code>ncptl_dfunc_factor10</code>	158
<code>n_receive_stmt</code>	186	<code>ncptl_dfunc_floor</code>	158
<code>n_rcv_buffer_number</code>	186	<code>ncptl_dfunc_grid_coord</code>	159
<code>n_reduce_message_spec</code>	186	<code>ncptl_dfunc_grid_neighbor</code>	159
<code>n_reduce_stmt</code>	186	<code>ncptl_dfunc_knomial_child</code>	160
<code>n_rel_conj_expr</code>	186	<code>ncptl_dfunc_knomial_parent</code>	159
<code>n_rel_disj_expr</code>	186	<code>ncptl_dfunc_log10</code>	157
<code>n_reset_stmt</code>	186	<code>ncptl_dfunc_max</code>	158
<code>n_send_attrs</code>	186	<code>ncptl_dfunc_min</code>	158
<code>n_send_stmt</code>	186	<code>ncptl_dfunc_modulo</code>	158
<code>n_simple_stmt_list</code>	186	<code>ncptl_dfunc_power</code>	158
<code>n_sleeps_for</code>	186	<code>ncptl_dfunc_random_gaussian</code>	160
<code>n_stride</code>	186	<code>ncptl_dfunc_random_poisson</code>	160
<code>n_string</code>	186	<code>ncptl_dfunc_random_uniform</code>	160
<code>n_string_or_expr_list</code>	186	<code>ncptl_dfunc_root</code>	157
<code>n_string_or_log_comment</code>	186	<code>ncptl_dfunc_round</code>	158
<code>n_sync_stmt</code>	186	<code>ncptl_dfunc_shift_left</code>	157
<code>n_task_expr</code>	186	<code>ncptl_dfunc_shift_right</code>	157
<code>n_time_unit</code>	186	<code>ncptl_dfunc_sqrt</code>	157
<code>n_top_level_stmt</code>	186	<code>ncptl_dfunc_tree_child</code>	159
<code>n_touch_buffer_stmt</code>	186	<code>ncptl_dfunc_tree_parent</code>	159
<code>n_touch_repeat_count</code>	186	<code>ncptl_fast_init</code>	146
<code>n_touch_stmt</code>	186	<code>NCPTL_FAST_INIT</code>	146, 192
<code>n_touching</code>	186	<code>ncptl_fatal</code>	147, 148, 160
<code>n_trivial_node</code>	186	<code>ncptl_fill_buffer</code>	149
<code>n_unary_expr</code>	187	<code>ncptl_finalize</code>	160
<code>n_unique</code>	187	<code>ncptl_free</code>	148, 152, 155, 156
<code>n_verification</code>	187	<code>ncptl_func_abs</code>	158
<code>n_version_decl</code>	187	<code>ncptl_func_bits</code>	157
<code>nanosleep</code>	151	<code>ncptl_func_cbrt</code>	157
<code>ncptl</code>	18, 19, 20, 22, 23, 39, 40, 136, 137, 166, 172, 178, 192	<code>ncptl_func_ceiling</code>	158
<code>ncptl-logextract</code>	9, 44, 48	<code>ncptl_func_factor10</code>	158
<code>ncptl-logextract.html</code>	9	<code>NCPTL_FUNC_FINAL</code>	191
<code>ncptl-logmerge</code>	43, 70, 74	<code>ncptl_func_floor</code>	158
<code>ncptl-logunmerge</code>	74	<code>NCPTL_FUNC_GEOMETRIC_MEAN</code>	191
<code>ncptl-mode.el</code>	9, 12	<code>ncptl_func_grid_coord</code>	159
<code>ncptl-mode.elc</code>	9, 12	<code>ncptl_func_grid_neighbor</code>	159
<code>ncptl-replaytrace</code>	26	<code>NCPTL_FUNC_HARMONIC_MEAN</code>	191
<code>ncptl.h</code>	13, 135, 145, 146, 166, 171, 191	<code>NCPTL_FUNC_HISTOGRAM</code>	191
<code>ncptl.h.in</code>	135	<code>ncptl_func_knomial_child</code>	160
<code>ncptl.lang</code>	9, 12	<code>ncptl_func_knomial_parent</code>	159
<code>ncptl.pc</code>	13	<code>ncptl_func_log10</code>	157, 158
<code>ncptl.py</code>	135, 139, 143	<code>ncptl_func_max</code>	158
<code>ncptl.ssh</code>	9, 12	<code>NCPTL_FUNC_MAXIMUM</code>	191
<code>ncptl.sty</code>	9, 11, 12	<code>NCPTL_FUNC_MEAN</code>	191
<code>ncptl.vim</code>	9, 12	<code>NCPTL_FUNC_MEDIAN</code>	191
		<code>ncptl_func_min</code>	158

- NCPTL_FUNC_MINIMUM..... 191
 ncptl_func_modulo..... 158
 NCPTL_FUNC_NO_AGGREGATE..... 191
 NCPTL_FUNC_ONLY..... 191
 ncptl_func_power..... 158
 ncptl_func_random_gaussian..... 160
 ncptl_func_random_poisson..... 160
 ncptl_func_random_uniform..... 160
 ncptl_func_root..... 157
 ncptl_func_round..... 158
 ncptl_func_shift_left..... 157
 ncptl_func_shift_right..... 157
 ncptl_func_sqrt..... 157
 NCPTL_FUNC_STDEV..... 191
 NCPTL_FUNC_SUM..... 191
 ncptl_func_tree_child..... 159
 ncptl_func_tree_parent..... 159
 NCPTL_FUNC_VARIANCE..... 191
 ncptl_get_message_buffer..... 148
 ncptl_init..... 146, 147, 150
 ncptl_int..... 145, 150, 151, 157, 159
 NCPTL_INT_MIN..... 145, 149
 ncptl_lexer.py..... 135
 ncptl_log_add_comment..... 152
 ncptl_log_close..... 154
 ncptl_log_commit_data..... 153, 154
 ncptl_log_compute_aggregates..... 153
 NCPTL_LOG_DELAY..... 193
 NCPTL_LOG_FILE_STATE..... 146, 152
 ncptl_log_generate_uuid..... 152, 153
 ncptl_log_lookup_string..... 153
 NCPTL_LOG_ONLY..... 28, 43, 44, 193
 ncptl_log_open..... 152
 ncptl_log_write..... 153
 ncptl_log_write_epilogue..... 154
 ncptl_log_write_prologue..... 150, 152, 153
 ncptl_malloc..... 147, 148, 155
 ncptl_malloc_message..... 148, 149
 ncptl_malloc_misaligned..... 148
 NCPTL_NOFORK..... 174, 175, 193
 ncptl_pagesize..... 146
 ncptl_parse_command_line..... 146, 147, 153
 ncptl_parser.py..... 135
 NCPTL_PATH..... 18, 192
 ncptl_permit_signal..... 146
 NCPTL_QUEUE..... 146
 ncptl_queue_allocate..... 155
 ncptl_queue_contents..... 155, 156
 ncptl_queue_empty..... 155, 156
 ncptl_queue_init..... 155
 ncptl_queue_length..... 156
 ncptl_queue_pop..... 155
 ncptl_queue_push..... 155
 ncptl_random_task..... 154
 ncptl_realloc..... 148
 NCPTL_RUN_TIME_VERSION..... 146
 ncptl_seed_random_task..... 147, 154
 ncptl_set_empty..... 156
 ncptl_set_find..... 156
 ncptl_set_flag_after_usecs..... 151
 ncptl_set_init..... 156
 ncptl_set_insert..... 156
 ncptl_set_length..... 156
 ncptl_set_remove..... 156
 ncptl_set_walk..... 156
 ncptl_strdup..... 148
 ncptl_time..... 47, 150, 151
 ncptl_token.py..... 135
 ncptl_touch_data..... 149
 ncptl_touch_memory..... 149
 NCPTL_TYPE_INT..... 145
 NCPTL_TYPE_STRING..... 145
 ncptl_udelay..... 151
 ncptl_verify..... 149
 NCPTL_VIRT_PHYS_MAP..... 154
 ncptl_virtual_to_physical..... 154
 ncptlGUI-1.0.jar..... 10, 14
 ncurses..... 24
 New (GUI menu item)..... 17
 newline..... 78
 newlines in backquoted commands..... 42
 NICS..... 145
 NODESHAPE..... 37
 nonblocking..... 98
 nonterminal..... 77
 NONUNIQUE..... 94, 95, 96, 98, 124, 125, 181
 Normalize (GUI button)..... 18
 NOT..... 79, 81, 123, 181
 num_tasks..... 84, 85, 94, 106, 110, 119, 161, 184
- ## O
- ODD..... 181
 OF..... 90, 108, 124, 127, 181
 Open (GUI menu item)..... 17
 openc.c..... 4
 options..... 139
 Options (GUI menu)..... 17
 OR..... 120, 128, 181
 OTHER..... 181
 OTHER THAN..... 117, 128
 OTHERWISE..... 79, 80, 81, 112, 119, 123, 127, 128, 182
 OUTPUT..... 78, 104, 182
 output_stmt..... 103, 104, 112, 126, 127
 OUTPUTS..... 26, 34, 78, 80, 104, 113, 126, 141, 182

P

PACKAGES 37
 PAGE 96, 97, 125, 182
 PAGE ALIGNED 146, 177
 PAGE SIZED 146
 PAGES 182
 PAPI 150
 PAPI_get_real_usec 150
 ParaGraph 33, 34, 194
 param_decl 120, 121, 128
 Paste (GUI menu item) 17
 PATH 2, 9
 pathcc 4
 pdsh 22
 pgcc 4
 PICL 21, 33, 34, 194
 picl ('codegen_picl.py')... 21, 33, 34, 146, 192
 pictures of communication patterns 34
 pkg-config 13, 194
 PLUS 112, 113, 127, 182
 PLY *see* Python Lex-Yacc
 pop 144
 popen 6, 193
 PostScript 21, 34, 36, 37, 40, 178
 power_expr 81, 123
 prctl 177
 predeclared variables 184
 PRId64 168, 169
 primary_expr 81, 123
 primitives, language 77
 Print (GUI menu item) 17
 printf 43, 160, 168
 PROCESSOR 110, 127, 182
 processor_stmt 110, 112, 127
 PROCESSORS 182
 program 121, 128
 prun 22
 PSMATRIX 37
 PSTricks 34, 35, 37, 38, 194
 push 143
 push_marker 144
 pushmany 138, 143
 Python... 5, 10, 18, 19, 37, 79, 135, 137, 144, 165,
 170, 172, 173, 194
 Python interface, failure to build 170
 Python Lex-Yacc (PLY) 135, 143, 185
 Python.h 170

Q

QUADWORD 96, 97, 125, 182
 QUADWORDS 182
 QueryPerformanceCounter 151
 QueryPerformanceFrequency 151
 queue functions 155
 Quit (GUI menu item) 17
 quotes 42, 44, 78, 168

R

RANDOM 182
 random numbers 90, 160
 RANDOM TASK 43
 random tasks 154
 RANDOM_GAUSSIAN 81, 90, 124, 182
 RANDOM_POISSON 81, 90, 124, 182
 RANDOM_UNIFORM 81, 90, 124, 182
 range 112, 114, 116, 127, 129
 REAL 79, 80, 81, 123, 182
 RECEIVE 2, 16, 78, 95, 99, 100, 101, 126, 182
 receive_stmt 100, 101, 112, 126, 127, 163
 Received signal 176
 RECEIVES 182
 recv_message_spec 95, 99, 124, 126
 Red Storm 164, 165, 175
 REDUCE 17, 22, 95, 182
 Reduce (GUI button) 18
 reduce_message_spec 95, 102, 103, 125, 126
 reduce_stmt 102, 112, 126, 127
 reduce_target_message_spec 95, 102, 125, 126
 REDUCES 102, 103, 126, 182
 REGION 182
 rel_conj_expr 92, 124
 rel_disj_expr 92, 124
 rel_expr... 79, 80, 81, 91, 92, 93, 107, 112, 119, 123,
 124, 126, 127, 128
 rel_primary_expr 92, 124
 relational expressions 91
 reordering task IDs 110
 REPETITION 182
 REPETITIONS 109, 112, 113, 127, 163, 182
 REQUIRE 182
 REQUIRE LANGUAGE VERSION 120, 128
 reserved words 179
 RESET 106, 182
 RESET ITS COUNTERS 109
 reset_stmt 106, 112, 126, 127
 RESETS 182
 RESETS ITS COUNTERS 106, 126
 restricted identifiers 93
 restricted_ident 93, 94, 99, 114, 124
 RESULT 182
 RESULTS 182
 ROOT 81, 82, 123, 182
 ROUND 81, 82, 83, 123, 182
 run-time library 135, 145
 running programs 41, 174
 runtime_random 8
 runtime_lib.c 135, 146

S

sample programs 129
 Save (GUI menu item) 17
 Save As (GUI menu item) 17
 SECOND 182
 SECONDS 107, 126, 182
 sed 19, 20
 segmentation fault 175
 SEND 16, 78, 95, 98, 99, 100, 101, 182
 send_stmt 98, 99, 100, 101, 102, 112, 125, 127, 163
 SENDS 99, 125, 182
 set functions 156
 setitimer 151
 Settings (GUI menu item) 17
 shared objects 174
 sheets 12
 shuffling task IDs 110
 signals 176
 simple_stmt 24, 112, 113, 114, 117, 119, 121, 122, 127, 128
 single-stepping 24, 25, 26
 SIZED 96, 125, 182
 SLEEP 16, 107, 182
 SLEEPS 27, 33, 107, 126, 151, 182
 SLOccount 13, 194
 smart preprocessing 173
 Source-highlight 9, 12, 194
 source_task 92, 93, 94, 98, 99, 100, 101, 102, 103, 104, 106, 107, 108, 109, 110, 111, 124, 125, 126, 127
 space 78
 spinning 107
 SQRT 81, 82, 123, 182
 STANDARD 182
 STANDARD DEVIATION 91, 124
 standard output 104
 statements, communication 94
 stats ('codegen_stats.py') 21, 31, 32, 33
 strdup 148
 STRIDE 183
 string 103, 107, 111, 120, 126, 127, 128
 string_or_log_comment 103, 104, 105, 126
 strings 78
 stylesheets 11
 substitutions.dat 173
 SUCH 183
 SUCH THAT 93, 124
 SUM 91, 124, 183
 sync_stmt 103, 112, 126, 127
 SYNCHRONIZATION 183
 SYNCHRONIZE 17, 183
 Synchronize (GUI button) 18
 SYNCHRONIZES 103, 113, 126, 183, 184
 SYNCHRONOUSLY 95, 98, 124, 183
 sysconf 177
 system 6, 193

T

tag 23
 TAGS 10
 target_tasks 92, 93, 94, 99, 100, 101, 102, 124, 125, 126
 TASK 78, 93, 94, 124, 183
 task-mapping functions 154
 tasks 92, 93, 94, 161
 TASKS 78, 93, 94, 124, 183
 tcsh 42
 TESTS_ENVIRONMENT 165
 TeX 11, 12, 36, 177, 178, 194
 TEXTOEPS 37
 THAN 183
 THAT 183
 THE 90, 105, 124, 183
 THE CURRENT MESSAGE BUFFER 109, 127
 THE DEFAULT BUFFER 95, 124, 125
 THE VALUE OF 103, 126, 153
 THEIR 183
 THEM 183
 THEN 109, 112, 113, 119, 127, 128, 183
 TIME 183
 time-related functions 150
 time_unit 107, 112, 117, 126, 127
 TIMELINE 37
 timer selection 150
 timer validation 10
 TIMES 108, 127, 183
 timings, invalid 177
 tips and Tricks 161
 TO 99, 102, 103, 125, 126, 183
 TO UNSUSPECTING 101
 top_level_complex_stmt 121, 122, 128
 topology functions 159
 torus functions 88
 TORUS_COORDINATE 81, 88, 89, 124, 183
 TORUS_NEIGHBOR 81, 88, 124, 183
 total_bytes 184
 total_msgs 184
 TOUCH 108, 183
 touch_buffer_stmt 109, 110, 112, 127
 touch_stmt 108, 109, 112, 126, 127, 163
 TOUCHES 27, 33, 108, 109, 127, 149, 183
 TOUCHING 183
 touching data 97, 108
 tracing 23
 tree functions 83, 84
 TREE_CHILD 81, 83, 123, 183
 TREE_PARENT 81, 83, 123, 183
 trees 131
 troubleshooting 167
 typesetting conventions 2

U

uac 177
 UNALIGNED 94, 95, 97, 124, 125, 183
 unaligned accesses 176
 unary_expr 81, 123
 unary_operator 81, 123
 undefined type 173
 Undo (GUI menu item) 17
 UNIQUE .. 23, 94, 95, 96, 98, 109, 124, 125, 142, 183
 UNSUSPECTING 99, 100, 102, 103, 125, 183

V

validatetimer 6, 10, 11, 150, 177
 VALUE 183
 variability in data values 176
 variable assignment 117
 variables, binding 117
 variables, environment 2, 6, 7, 8, 9, 10, 18, 22,
 23, 28, 35, 40, 43, 44, 146, 165, 166, 168, 170,
 172, 174, 175, 176, 178, 192, 193
 variables, predeclared 184
 variables, run-time library 145
 VARIANCE 91, 124, 183
 VERIFICATION 183
 VERSION 183
 version_decl 120, 121, 128
 versioning 120
 Vim 9, 12, 194

W

Wait (GUI button) 17
 wait_stmt 101, 112, 126, 127
 WARMUP 112, 127, 183
 WARMUP REPETITIONS 113

warnings 7
 WHILE 112, 117, 127, 183
 whitespace 78
 WHO 183
 WHO RECEIVES IT 99, 126
 WHO RECEIVES IT 95, 100, 101
 WHO RECEIVES THE RESULT 102, 126
 WHO RECEIVES THEM 100
 WITH 107, 126, 183
 WITH DATA TOUCHING 95, 97, 98, 103, 109, 124,
 125, 145, 149
 WITH DEFAULT 120, 128
 WITH RANDOM STRIDE 108, 127
 WITH STRIDE 108, 127
 WITH VERIFICATION 43, 95, 98, 101, 103, 124,
 125, 145, 149
 WITHOUT 183
 WITHOUT DATA TOUCHING 95, 98, 124, 125
 WITHOUT VERIFICATION 95, 98, 124, 125
 WORD 96, 97, 108, 125, 183
 WORDS 183

X

x86-64 4, 164, 165, 171
 XEmacs 12
 xlc 4, 7
 XOR 79, 81, 123, 183

Y

yod 165