



The Isabelle System Manual

Markus Wenzel and Stefan Berghofer
TU München

June 8, 2008

Contents

1	The Isabelle system environment	1
1.1	Isabelle settings	1
	Building the environment	2
	Common variables	3
1.2	The Isabelle tools wrapper	5
1.3	The raw Isabelle process	6
1.4	The Isabelle interface wrapper	9
2	Presenting theories	11
2.1	Generating theory browser information	12
2.2	Browsing theory graphs	13
	Invoking the graph browser	13
	Using the graph browser	14
2.3	Creating Isabelle session directories — <code>isatool mkdir</code>	16
2.4	Running Isabelle sessions — <code>isatool usedir</code>	18
2.5	Preparing Isabelle session documents — <code>isatool document</code> .	21
2.6	Running \LaTeX within the Isabelle environment — <code>isatool latex</code>	23
3	Miscellaneous tools	25
3.1	Converting legacy ML scripts — <code>isatool convert</code>	25
3.2	Displaying documents — <code>isatool display</code>	25
3.3	Viewing documentation — <code>isatool doc</code>	26
3.4	Getting logic images — <code>isatool findlogics</code>	26
3.5	Inspecting the settings environment — <code>isatool getenv</code> . . .	27
3.6	Installing standalone Isabelle executables — <code>isatool install</code>	28
3.7	Creating instances of the Isabelle logo — <code>isatool logo</code> . . .	28
3.8	Isabelle’s version of <code>make</code> — <code>isatool make</code>	29
3.9	Make all logics — <code>isatool makeall</code>	29
3.10	Printing documents — <code>isatool print</code>	29
3.11	Run Isabelle with plain tty interaction — <code>isatool tty</code>	30
3.12	Remove awkward symbol names from theory sources — <code>isatool unsymbolize</code>	30

3.13 Output the version identifier of the Isabelle distribution — <code>isatool version</code>	31
3.14 Convert XML to YXML — <code>isatool yxml</code>	31
A Standard Isabelle symbols	32

The Isabelle system environment

This manual describes Isabelle together with related tools and user interfaces as seen from an outside (system oriented) view. See also the *Isabelle/Isar Reference Manual* [3] and the *Isabelle Reference Manual* [2] for the actual Isabelle commands and related functions.

The Isabelle system environment emerges from a few general concepts.

- The *Isabelle settings mechanism* provides environment variables to all Isabelle programs (including tools and user interfaces).
- The *Isabelle tools wrapper* (`isatool`) provides a generic startup platform for Isabelle related utilities. Thus tools automatically benefit from the settings mechanism.
- The raw *Isabelle process* (`isabelle` or `isabelle-process`) runs logic sessions either interactively or in batch mode. In particular, this view abstracts over the invocation of the actual ML system to be used.
- The *Isabelle interface wrapper* (`Isabelle` or `isabelle-interface`) provides some abstraction over the actual user interface to be used. The de-facto standard interface for Isabelle is Proof General [1].

The beginning user would probably just run the default user interface (by invoking the capital `Isabelle`). This assumes that the system has already been installed, of course. In case you have to do this yourself, see the `INSTALL` file in the top-level directory of the distribution of how to proceed; binary packages for various system components are available as well.

1.1 Isabelle settings

The Isabelle system heavily depends on the *settings mechanism*. Essentially, this is a statically scoped collection of environment variables, such as

`ISABELLE_HOME`, `ML_SYSTEM`, `ML_HOME`. These variables are *not* intended to be set directly from the shell, though. Isabelle employs a somewhat more sophisticated scheme of *settings files* — one for site-wide defaults, another for additional user-specific modifications. With all configuration variables in at most two places, this scheme is more maintainable and user-friendly than global shell environment variables.

In particular, we avoid the typical situation where prospective users of a software package are told to put several things into their shell startup scripts, before being able to actually run the program. Isabelle requires none such administrative chores of its end-users — the executables can be invoked straight away.¹

Building the environment

Whenever any of the Isabelle executables is run, their settings environment is put together as follows.

1. The special variable `ISABELLE_HOME` is determined automatically from the location of the binary that has been run.

You should not try to set `ISABELLE_HOME` manually. Also note that the Isabelle executables either have to be run from their original location in the distribution directory, or via the executable objects created by the `install` utility (see §3.6). Just doing a plain copy of the `bin` files will not work!

2. The file `ISABELLE_HOME/etc/settings` is run as a shell script with the auto-export option for variables enabled.

This file holds a rather long list of shell variable assignments, thus providing the site-wide default settings. The Isabelle distribution already contains a global settings file with sensible defaults for most variables. When installing the system, only a few of these may have to be adapted (probably `ML_SYSTEM` etc.).

3. The file `ISABELLE_HOME_USER/etc/settings` (if it exists) is run in the same way as the site default settings. Note that the variable `ISABELLE_HOME_USER` has already been set before — usually to `~/isabelle`.

Thus individual users may override the site-wide defaults. See also file `etc/user-settings.sample` in the distribution. Typically, a user

¹Occasionally, users would still want to put the Isabelle `bin` directory into their shell's search path, but this is not required.

settings file would contain only a few lines, just the assignments that are really changed. One should definitely *not* start with a full copy the basic `$ISABELLE_HOME/etc/settings`. This could cause very annoying maintainance problems later, when the Isabelle installation is updated or changed otherwise.

Note that settings files are actually full GNU bash scripts. So one may use complex shell commands, such as `if` or `case` statements to set variables depending on the system architecture or other environment variables. Such advanced features should be added only with great care, though. In particular, external environment references should be kept at a minimum.

A few variables are somewhat special:

- `ISABELLE` and `ISATOOL` are set automatically to the absolute path names of the `isabelle-process` and `isatool` executables, respectively.
- `ISABELLE_OUTPUT` will have the identifiers of the Isabelle distribution (cf. `ISABELLE_IDENTIFIER`) and the ML system (cf. `ML_IDENTIFIER`) appended automatically to its value.

The Isabelle settings scheme is conceptually simple, but not completely trivial. For debugging purposes, the resulting environment may be inspected with the `getenv` utility, see §3.5.

Common variables

This is a reference of common Isabelle settings variables. Note that the list is somewhat open-ended. Third-party utilities or interfaces may add their own selection. Variables that are special in some sense are marked with `*`.

`ISABELLE_HOME*` is the location of the top-level Isabelle distribution directory. This is automatically determined from the Isabelle executable that has been invoked. Do not attempt to set `ISABELLE_HOME` yourself from the shell.

`ISABELLE_HOME_USER` is the user-specific counterpart of `ISABELLE_HOME`. The default value is `~/isabelle`, under rare circumstances this may be changed in the global setting file. Typically, the `ISABELLE_HOME_USER` directory mimics `ISABELLE_HOME` to some extent. In particular, site-wide defaults may be overridden by a private `etc/settings`.

`ISABELLE*`, `ISATOOL*` are automatically set to the full path names of the `isabelle-process` and `isatool` executables, respectively. Thus other tools and scripts need not assume that the Isabelle `bin` directory is on the current search path of the shell.

`ISABELLE_IDENTIFIER*` refers to the name of this Isabelle distribution, e.g. “`Isabelle2007`”.

`ML_SYSTEM`, `ML_HOME`, `ML_OPTIONS`, `ML_PLATFORM`, `ML_IDENTIFIER*` specify the underlying ML system to be used for Isabelle. There is only a fixed set of admissible `ML_SYSTEM` names (see the `etc/settings` file of the distribution).

The actual compiler binary will be run from the directory `ML_HOME`, with `ML_OPTIONS` as first arguments on the command line. The optional `ML_PLATFORM` may specify the binary format of ML heap images, which is useful for cross-platform installations. The value of `ML_IDENTIFIER` is automatically obtained by composing the values of `ML_SYSTEM`, `ML_PLATFORM` and the Isabelle version values.

`ISABELLE_PATH` is a list of directories (separated by colons) where Isabelle logic images may reside. When looking up heaps files, the value of `ML_IDENTIFIER` is appended to each component internally.

`ISABELLE_OUTPUT*` is a directory where output heap files should be stored by default. The ML system and Isabelle version identifier is appended here, too.

`ISABELLE_BROWSER_INFO` is the directory where theory browser information (HTML text, graph data, and printable documents) is stored (see also §2.1). The default value is `$ISABELLE_HOME_USER/browser_info`.

`ISABELLE_LOGIC` specifies the default logic to load if none is given explicitly by the user. The default value is `HOL`.

`ISABELLE_LINE_EDITOR` specifies the default line editor for `isatool tty` (see also §3.11).

`ISABELLE_USEDIR_OPTIONS` is implicitly prefixed to the command line of any `isatool usedir` invocation (see also §2.4). This typically contains compilation options for object-logics — `usedir` is the basic utility for managing logic sessions (cf. the `IsaMakefiles` in the distribution).

`ISABELLE_FILE_IDENT` specifies a shell command for producing a source file identification, based on the actual content instead of the full physical path and date stamp (which is the default). A typical identification would produce a “digest” of the text, using a cryptographic hash function like SHA-1, for example.

`ISABELLE_LATEX`, `ISABELLE_PDFLATEX`, `ISABELLE_BIBTEX`, `ISABELLE_DVIPS` refer to \LaTeX related tools for Isabelle document preparation (see also §2.6).

`ISABELLE_TOOLS` is a colon separated list of directories that are scanned by `isatool` for external utility programs (see also §1.2).

`ISABELLE_DOCS` is a colon separated list of directories with documentation files.

`ISABELLE_DOC_FORMAT` specifies the preferred document format, typically `dvi` or `pdf`.

`DVI_VIEWER` specifies the command to be used for displaying `dvi` files.

`PDF_VIEWER` specifies the command to be used for displaying `pdf` files.

`PRINT_COMMAND` specifies the standard printer spool command, which is expected to accept `ps` files.

`ISABELLE_TMP_PREFIX*` is the prefix from which any running `isabelle` process derives an individual directory for temporary files. The default is somewhere in `/tmp`.

`ISABELLE_INTERFACE` is an identifier that specifies the actual user interface that the capital `Isabelle` or `isabelle-interface` should invoke. See §1.4 for more details.

1.2 The Isabelle tools wrapper

All Isabelle related utilities are called via a common wrapper — `isatool`:

```
Usage: isatool TOOL [ARGS ...]
```

```
Start Isabelle utility program TOOL with ARGS. Pass "-?" to TOOL
for more specific help.
```

```
Available tools are:
```

```
browser - Isabelle graph browser
...
```

In principle, Isabelle tools are ordinary executable scripts that are run within the Isabelle settings environment, see §1.1. The set of available tools is collected by `isatool` from the directories listed in the `ISABELLE_TOOLS` setting. Do not try to call the scripts directly from the shell. Neither should you add the tool directories to your shell's search path!

Examples

Show the list of available documentation of the current Isabelle installation like this:

```
isatool doc
```

View a certain document as follows:

```
isatool doc isar-ref
```

Create an Isabelle session derived from HOL (see also §2.3 and §3.8):

```
isatool mkdir HOL Test && isatool make
```

Note that `isatool mkdir` is usually only invoked once; existing sessions (including document output etc.) are then updated by `isatool make` alone.

1.3 The raw Isabelle process

The `isabelle` (or `isabelle-process`) executable runs bare-bones Isabelle logic sessions — either interactively or in batch mode. It provides an abstraction over the underlying ML system, and over the actual heap file locations. Its usage is:

Usage: `isabelle [OPTIONS] [INPUT] [OUTPUT]`

Options are:

```

-C          tell ML system to copy output image
-I          startup Isar interaction mode
-P          startup Proof General interaction mode
-S          secure mode -- disallow critical operations
-W          startup process wrapper (interaction via external program)
-X          startup PGIIP interaction mode
-c          tell ML system to compress output image
-e MLTEXT  pass MLTEXT to the ML session
-f          pass 'Session.finish();' to the ML session
-m MODE    add print mode for output
-q          non-interactive session
-r          open heap file read-only
-u          pass 'use"ROOT.ML";' to the ML session
-w          reset write permissions on OUTPUT

```

`INPUT` (default `"$ISABELLE_LOGIC"`) and `OUTPUT` specify in/out heaps. These are either names to be searched in the Isabelle path, or actual file names (containing at least one `/`).

If `INPUT` is `"RAW_ML_SYSTEM"`, just start the bare bones ML system.

Input files without path specifications are looked up in the `ISABELLE_PATH` setting, which may consist of multiple components separated by colons — these are tried in the given order with the value of `ML_IDENTIFIER` appended internally. In a similar way, base names are relative to the directory specified by `ISABELLE_OUTPUT`. In any case, actual file locations may also be given by including at least one slash (`/`) in the name (hint: use `./` to refer to the current directory).

Options

If the input heap file does not have write permission bits set, or the `-r` option is given explicitly, then the session started will be read-only. That is, the ML world cannot be committed back into the image file. Otherwise, a writable session enables commits into either the input file, or into another output heap file (if that is given as the second argument on the command line).

The read-write state of sessions is determined at startup only, it cannot be changed intermediately. Also note that heap images may require considerable amounts of disk space (approximately 40–80 MB). Users are responsible for themselves to dispose their heap files when they are no longer needed.

The `-w` option makes the output heap file read-only after terminating. Thus subsequent invocations cause the logic image to be read-only automatically.

The `-c` option tells the underlying ML system to compress the output heap (fully transparently). On Poly/ML for example, the image is garbage collected and all stored values are maximally shared, resulting in up to 50% less disk space consumption.

The `-C` option tells the ML system to produce a completely self-contained output image, probably including a copy of the ML runtime system itself.

Using the `-e` option, arbitrary ML code may be passed to the Isabelle session from the command line. Multiple `-e`'s are evaluated in the given order. Strange things may happen when erroneous ML code is provided. Also make sure that the ML commands are terminated properly by semicolon.

The `-u` option is a shortcut for `-e` passing `"use"ROOT.ML";"` to the ML session. The `-f` option passes `"Session.finish();"`, which is intended mainly for administrative purposes.

The `-m` option adds identifiers of print modes to be made active for this session. Typically, this is used by some user interface, e.g. to enable output of proper mathematical symbols.

Isabelle normally enters an interactive top-level loop (after processing the `-e` texts). The `-q` option inhibits interaction, thus providing a pure batch mode facility.

The `-I` option makes Isabelle enter Isar interaction mode on startup, instead of the primitive ML top-level. The `-P` option configures the top-level loop for interaction with the Proof General user interface, and the `-X` option enables XML-based PGIP communication. The `-W` option makes Isabelle enter a special process wrapper for interaction via an external program; the protocol is a stripped-down version of Proof General the interaction mode.

The `-S` option makes the Isabelle process more secure by disabling some critical operations, notably runtime compilation and evaluation of ML source code.

Examples

Run an interactive session of the default object-logic (as specified by the `ISABELLE_LOGIC` setting) like this:

```
isabelle
```

Usually `ISABELLE_LOGIC` refers to one of the standard logic images, which are read-only by default. A writable session — based on FOL, but output to Foo (in the directory specified by the `ISABELLE_OUTPUT` setting) — may be invoked as follows:

```
isabelle FOL Foo
```

Ending this session normally (e.g. by typing control-D) dumps the whole ML system state into `Foo`. Be prepared for several tens of megabytes.

The `Foo` session may be continued later (still in writable state) by:

```
isabelle Foo
```

A read-only `Foo` session may be started by:

```
isabelle -r Foo
```

Note that manual session management like this does *not* provide proper setup for theory presentation. This would require the `usedir` utility, see §2.4.

The next example demonstrates batch execution of Isabelle. We print a certain theorem of FOL:

```
isabelle -e "prth allE;" -q -r FOL
```

Note that the output text will be interspersed with additional junk messages by the ML runtime environment.

1.4 The Isabelle interface wrapper

Isabelle is a generic theorem prover, even w.r.t. its user interface. The `Isabelle` (or `isabelle-interface`) executable provides a uniform way for end-users to invoke a certain interface; which one to start is determined by the `ISABELLE_INTERFACE` setting variable, which should give a full path specification to the actual executable. Also note that the `install` utility provides some options to install desktop environment icons (see §3.6).

Presently, the most prominent Isabelle interface is Proof General [1]. There are separate versions for the raw ML top-level and the proper Isabelle/Isar interpreter loop. The Proof General distribution includes separate interface wrapper scripts (in `ProofGeneral/isa` and `ProofGeneral/isar`) for either of these. The canonical settings for Isabelle/Isar are as follows:

```
ISABELLE_INTERFACE=$ISABELLE_HOME/contrib/ProofGeneral/isar/interface
PROOFGENERAL_OPTIONS=""
```

Thus `Isabelle` would automatically invoke Emacs with proper setup of the Proof General Lisp packages. There are some options available, such as `-l` for passing the logic image to be used by default, or `-m` to tune the standard

print mode. The `-I` option allows to switch between the Isar and ML view, independently of the interface script being used.

Note that the world may be also seen the other way round: Emacs may be started first (with proper setup of Proof General mode), and `isabelle` run from within. This requires further Emacs Lisp configuration, see the Proof General documentation [1] for more information.

Presenting theories

Isabelle provides several ways to present the outcome of formal developments, including WWW-based browsable libraries or actual printable documents. Presentation is centered around the concept of *logic sessions*. The global session structure is that of a tree, with Isabelle Pure at its root, further object-logics derived (e.g. HOLCF from HOL, and HOL from Pure), and application sessions in leaf positions (usually without a separate image).

The `mkdir` (see §2.3) and `make` (see §3.8) tools of Isabelle provide the primary means for managing Isabelle sessions, including proper setup for presentation. Here the `usedir` (see §2.4) tool takes care to let the `isabelle` process run any additional stages required for document preparation, notably the tools `document` (see §2.5) and `latex` (see §2.6). The complete tool chain for managing batch-mode Isabelle sessions is illustrated in figure 2.1.

<code>isatool mkdir</code>	invoked once by the user to create the initial source setup (common <code>IsaMakefile</code> plus a single session directory);
<code>isatool make</code>	invoked repeatedly by the user to keep session output up-to-date (HTML, documents etc.);
<code>isatool usedir</code>	part of the standard <code>IsaMakefile</code> entry of a session;
<code>isabelle</code>	run through <code>isatool usedir</code> ;
<code>isatool document</code>	run by the Isabelle process if document preparation is enabled;
<code>isatool latex</code>	universal \LaTeX tool wrapper invoked multiple times by <code>isatool document</code> ; also useful for manual experiments;

Figure 2.1: The tool chain of Isabelle session presentation

2.1 Generating theory browser information

As a side-effect of running a logic sessions, Isabelle is able to generate theory browsing information, including HTML documents that show a theory's definition, the theorems proved in its ML file and the relationship with its ancestors and descendants. Besides the HTML file that is generated for every theory, Isabelle stores links to all theories in an index file. These indexes are linked with other indexes to represent the overall tree structure of logic sessions.

Isabelle also generates graph files that represent the theory hierarchy of a logic. There is a graph browser Java applet embedded in the generated HTML pages, and also a stand-alone application that allows browsing theory graphs without having to start a WWW client first. The latter version also includes features such as generating Postscript files, which are not available in the applet version. See §2.2 for further information.

The easiest way to let Isabelle generate theory browsing information for existing sessions is to append “`-i true`” to the `ISABELLE_USEDIR_OPTIONS` before invoking `isatool make` (or `./build` in the distribution). For example, add something like this to your Isabelle settings file

```
ISABELLE_USEDIR_OPTIONS="-i true"
```

and then change into the `src/FOL` directory of the Isabelle distribution and run `isatool make`, or even `isatool make all`. The presentation output will appear in `$ISABELLE_BROWSER_INFO/FOL`, which usually refers to `~/isabelle/browser_info/FOL`. Note that option `-v true` will make the internal runs of `usedir` more explicit about such details.

Many standard Isabelle sessions (such as `HOL/ex`) also provide actual printable documents. These are prepared automatically as well if enabled like this, using the `-d` option

```
ISABELLE_USEDIR_OPTIONS="-i true -d dvi"
```

Enabling options `-i` and `-d` simultaneously as shown above causes an appropriate “document” link to be included in the HTML index. Documents (or raw document sources) may be generated independently of browser information as well, see §2.5 for further details.

The theory browsing information is stored in a sub-directory directory determined by the `ISABELLE_BROWSER_INFO` setting plus a prefix corresponding to the session identifier (according to the tree structure of sub-sessions by default). A complete WWW view of all standard object-logics and examples of the Isabelle distribution is available at the Cambridge or Munich Isabelle sites:

```
http://www.cl.cam.ac.uk/Research/HVG/Isabelle/library/  
http://isabelle.in.tum.de/library/
```

In order to present your own theories on the web, simply copy the corresponding subdirectory from `ISABELLE_BROWSER_INFO` to your WWW server, having generated browser info like this:

```
isatool usedir -i true HOL Foo
```

This assumes that directory `Foo` contains some `ROOT.ML` file to load all your theories, and `HOL` is your parent logic image (`isatool mkdir` assists in setting up Isabelle session directories, see §2.3). Theory browser information for `HOL` should have been generated already beforehand. Alternatively, one may specify an external link to an existing body of HTML data by giving `usedir` a `-P` option like this:

```
isatool usedir -i true -P http://isabelle.in.tum.de/library/ HOL Foo
```

For production use, the `usedir` tool is usually invoked in an appropriate `IsaMakefile`, via the Isabelle `make` utility. There is a separate `mkdir` tool to provide easy setup of all this, with only minimal manual editing required.

```
isatool mkdir HOL Foo && isatool make
```

See §2.3 for more information on preparing Isabelle session directories, including the setup for documents.

2.2 Browsing theory graphs

The Isabelle graph browser is a general tool for visualizing dependency graphs. Certain nodes of the graph (i.e. theories) can be grouped together in “directories”, whose contents may be hidden, thus enabling the user to collapse irrelevant portions of information. The browser is written in Java, it can be used both as a stand-alone application and as an applet. Note that the option `-g` of `isatool usedir` (see §2.4) creates graph presentations in batch mode for inclusion in session documents.

Invoking the graph browser

The stand-alone version of the graph browser is wrapped up as an Isabelle tool called `browser`:

Usage: browser [OPTIONS] [GRAPHFILE]

Options are:

-c cleanup -- remove GRAPHFILE after use
 -o FILE output to FILE (ps, eps, pdf)

When no filename is specified, the browser automatically changes to the directory `ISABELLE_BROWSER_INFO`.

The `-c` option causes the input file to be removed after use.

The `-o` option indicates batch-mode operation, with the output written to the indicated file; note that `pdf` produces an `eps` copy as well.

The applet version of the browser is part of the standard WWW theory presentation, see the link “theory dependencies” within each session index.

Using the graph browser

The browser’s main window, which is shown in figure 2.2, consists of two sub-windows: In the left sub-window, the directory tree is displayed. The graph itself is displayed in the right sub-window.




Figure 2.2: Browser main window

The directory tree window

We describe the usage of the directory browser and the meaning of the different items in the browser window.

- A red arrow before a directory name indicates that the directory is currently “folded”, i.e. the nodes in this directory are collapsed to one single node. In the right sub-window, the names of nodes corresponding to folded directories are enclosed in square brackets and displayed in red color.
- A green downward arrow before a directory name indicates that the directory is currently “unfolded”. It can be folded by clicking on the directory name. Clicking on the name for a second time unfolds the directory again. Alternatively, a directory can also be unfolded by clicking on the corresponding node in the right sub-window.
- Blue arrows stand before ordinary node names. When clicking on such a name (i.e. that of a theory), the graph display window focuses to the corresponding node. Double clicking invokes a text viewer window in which the contents of the theory file are displayed.

The graph display window

When pointing on an ordinary node, an upward and a downward arrow is shown. Initially, both of these arrows are green. Clicking on the upward or downward arrow collapses all predecessor or successor nodes, respectively. The arrow’s color then changes to red, indicating that the predecessor or successor nodes are currently collapsed. The node corresponding to the collapsed nodes has the name “[...]”. To uncollapse the nodes again, simply click on the red arrow or on the node with the name “[...]”. Similar to the directory browser, the contents of theory files can be displayed by double clicking on the corresponding node.

The “File” menu

Please note that due to Java security restrictions this menu is not available in the applet version. The meaning of the menu items is as follows:

Open ... Open a new graph file.

Export to PostScript Outputs the current graph in Postscript format, appropriately scaled to fit on one single sheet of A4 paper. The resulting file can be printed directly.

Export to EPS Outputs the current graph in Encapsulated Postscript format. The resulting file can be included in other documents.

Quit Quit the graph browser.

*Syntax of graph definition files

A graph definition file has the following syntax:

$$\begin{aligned} \text{graph} &= \{ \text{vertex} ; \}^+ \\ \text{vertex} &= \text{vertexname vertexID dirname} [+]\text{ path} [< | >] \{ \text{vertexID} \}^* \end{aligned}$$

The meaning of the items in a vertex description is as follows:

vertexname The name of the vertex.

vertexID The vertex identifier. Note that there may be two vertices with equal names, whereas identifiers must be unique.

dirname The name of the “directory” the vertex should be placed in. A “+” sign after *dirname* indicates that the nodes in the directory are initially visible. Directories are initially invisible by default.

path The path of the corresponding theory file. This is specified relatively to the path of the graph definition file.

List of successor/predecessor nodes A “<” sign before the list means that successor nodes are listed, a “>” sign means that predecessor nodes are listed. If neither “<” nor “>” is found, the browser assumes that successor nodes are listed.

2.3 Creating Isabelle session directories — isatool mkdir

The `mkdir` utility prepares Isabelle session source directories, including a sensible default setup of `IsaMakefile`, `ROOT.ML`, and a `document` directory with a minimal `root.tex` that is sufficient to print all theories of the session (in the order of appearance); see §2.5 for further information on Isabelle document preparation. The usage of `isatool mkdir` is:

Usage: `mkdir [OPTIONS] [LOGIC] NAME`

Options are:

```
-I FILE      alternative IsaMakefile output
-P          include parent logic target
-b          setup build mode (session outputs heap image)
-q          quiet mode
```

Prepare session directory, including `IsaMakefile` and document source, with parent `LOGIC` (default `ISABELLE_LOGIC=$ISABELLE_LOGIC`)

The `mkdir` tool is conservative in the sense that any existing `IsaMakefile` etc. is left unchanged. Thus it is safe to invoke it multiple times, although later runs may not have the desired effect.

Note that `mkdir` is unable to change `IsaMakefile` incrementally — manual changes are required for multiple sub-sessions. In order to get an initial working session, the only editing needed is to add appropriate `use_thy` calls to the generated `ROOT.ML` file.

Options

The `-I` option specifies an alternative to `IsaMakefile` for dependencies. Note that “-” refers to *stdout*, i.e. “`-I-`” provides an easy way to peek at `mkdir`’s idea of `make` setup required for some particular of Isabelle session.

The `-P` option includes a target for the parent `LOGIC` session in the generated `IsaMakefile`. The corresponding sources are assumed to be located within the Isabelle distribution.

The `-b` option sets up the current directory as the base for a new session that provides an actual logic image, as opposed to one that only runs several theories based on an existing image. Note that in the latter case, everything except `IsaMakefile` would be placed into a separate directory `NAME`, rather than the current one. See §2.4 for further information on *build mode* vs. *example mode* of the `usedir` utility.

The `-q` enables quiet mode, suppressing further notes on how to proceed.

Examples

The standard setup of a single “example session” based on the default logic, with proper document generation is generated like this:

```
isatool mkdir Foo && isatool make
```

The theory sources should be put into the `Foo` directory, and its `ROOT.ML` should be edited to load all required theories. Invoking `isatool make` again

would run the whole session, generating browser information and the document automatically. The `IsaMakefile` is typically tuned manually later, e.g. adding source dependencies, or changing the options passed to `usedir`.

Large projects may demand further sessions, potentially with separate logic images being created. This usually requires manual editing of the generated `IsaMakefile`, which is meant to cover all of the sub-session directories at the same time (this is the deeper reason why `IsaMakefile` is not made part of the initial session directory created by `isatool mkdir`). See `src/HOL/IsaMakefile` of the Isabelle distribution for a full-blown example.

2.4 Running Isabelle sessions — `isatool usedir`

The `usedir` utility builds object-logic images, or runs example sessions based on existing logics. Its usage is:

Usage: `usedir [OPTIONS] LOGIC NAME`

Options are:

```

-C BOOL      copy existing document directory to -D PATH (default true)
-D PATH      dump generated document sources into PATH
-M MAX       multithreading: maximum number of worker threads (default 1)
-P PATH      set path for remote theory browsing information
-T LEVEL     multithreading: trace level (default 0)
-V VERSION   declare alternative document VERSION
-b           build mode (output heap image, using current dir)
-c BOOL      tell ML system to compress output image (default true)
-d FORMAT    build document as FORMAT (default false)
-f NAME      use ML file NAME (default ROOT.ML)
-g BOOL      generate session graph image for document (default false)
-i BOOL      generate theory browser information (default false)
-m MODE      add print mode for output
-p LEVEL     set level of detail for proof objects
-r           reset session path
-s NAME      override session NAME
-v BOOL      be verbose (default false)

```

Build object-logic or run examples. Also creates browsing information (HTML etc.) according to settings.

```

ISABELLE_USEDIR_OPTIONS=
HOL_USEDIR_OPTIONS=

```

Note that the value of the `ISABELLE_USEDIR_OPTIONS` setting is implicitly prefixed to *any* `usedir` call. Since the `IsaMakefiles` of all object-logics distributed with Isabelle just invoke `usedir` for the real work, one may control compilation options globally via above variable. In particular, generation of HTML browsing information and document preparation is controlled here.

The `HOL_USEDIR_OPTIONS` setting is specific to the main Isabelle/HOL image; its value is appended to `ISABELLE_USEDIR_OPTIONS` for that particular session only.

Options

Basically, there are two different modes of operation: *build mode* (enabled through the `-b` option) and *example mode* (default).

Calling `usedir` with `-b` runs `isabelle` with input image `LOGIC` and output to `NAME`, as provided on the command line. This will be a batch session, running `ROOT.ML` from the current directory and then quitting. It is assumed that `ROOT.ML` contains all ML commands required to build the logic.

In example mode, `usedir` runs a read-only session of `LOGIC` and automatically runs `ROOT.ML` from within directory `NAME`. It assumes that this file contains appropriate ML commands to run the desired examples.

The `-i` option controls theory browser data generation. It may be explicitly turned on or off — as usual, the last occurrence of `-i` on the command line wins.

The `-P` option specifies a path (or actual URL) to be prefixed to any *non-local* reference of existing theories. Thus user sessions may easily link to existing Isabelle libraries already present on the WWW.

The `-m` options specifies additional print modes to be activated temporarily while the session is processed.

The `-d` option controls document preparation. Valid arguments are `false` (do not prepare any document; this is default), or any of `dvi`, `dvi.gz`, `ps`, `ps.gz`, `pdf`. The logic session has to provide a properly setup document directory. See §2.5 and §2.6 for more details.

The `-V` option declares alternative document versions, consisting of name/tags pairs (cf. options `-n` and `-t` of the `document` tool, §2.5). The standard document is equivalent to “`document=theory,proof,ML`”, which means that all theory begin/end commands, proof body texts, and ML code will be presented faithfully. An alternative version “`outline=/proof/ML`” would fold proof and ML parts, replacing the original text by a short placeholder. The form “`name=-`” means to remove document *name* from the list of versions to be processed. Any number of `-V` options may be given; later declarations have precedence over earlier ones.

The `-g` option produces images of the theory dependency graph (cf. §2.2) for inclusion in the generated document, both as `session_graph.eps` and `session_graph.pdf` at the same time. To include this in the final `LATEX` document one could say `\includegraphics{session_graph}` in `document/root.tex` (omitting the file-name extension enables `LATEX` to select to correct version, either for the DVI or PDF output path).

The `-D` option causes the generated document sources to be dumped at location `PATH`; this path is relative to the session’s main directory. If the `-C` option is true, this will include a copy of an existing `document` directory as provided by the user. For example, `isatool usedir -D generated HOL Foo` produces a complete set of document sources at `Foo/generated`. Subsequent invocation of `isatool document Foo/generated` (see also §2.5) will process the final result independently of an Isabelle job. This decoupled mode of operation facilitates debugging of serious `LATEX` errors, for example.

The `-p` option determines the level of detail for internal proof objects, see also the *Isabelle Reference Manual* [2].

The `-v` option causes additional information to be printed while running the session, notably the location of prepared documents.

The `-M` option specifies the maximum number of parallel threads used for processing independent theory files (multithreading only works on suitable ML platforms). The special value of “0” or “max” refers to the number of actual CPU cores of the underlying machine, which is a good starting point for optimal performance tuning. The `-T` option determines the level of detail in tracing output concerning the internal locking and scheduling in multithreaded operation. This may be helpful in isolating performance bottle-necks, e.g. due to excessive wait states when locking critical code sections.

Any `usedir` session is named by some *session identifier*. These accumulate, documenting the way sessions depend on others. For example, consider `Pure/FOL/ex`, which refers to the examples of FOL, which in turn is built upon `Pure`.

The current session’s identifier is by default just the base name of the `LOGIC` argument (in build mode), or of the `NAME` argument (in example mode). This may be overridden explicitly via the `-s` option.

Examples

Refer to the `IsaMakefiles` of the Isabelle distribution’s object-logics as a model for your own developments. For example, see `src/FOL/IsaMakefile`. The Isabelle `mkdir` tool (see §2.3) creates `IsaMakefiles` with proper invocation of `usedir` as well.

2.5 Preparing Isabelle session documents — `isatool` document

The `document` utility prepares logic session documents, processing the sources both as provided by the user and generated by Isabelle. Its usage is:

Usage: `document [OPTIONS] [DIR]`

Options are:

```
-c          cleanup -- be aggressive in removing old stuff
-n NAME     specify document name (default 'document')
-o FORMAT   specify output format: dvi (default), dvi.gz, ps,
            ps.gz, pdf
-t TAGS     specify tagged region markup
```

Prepare the theory session document in `DIR` (default 'document') producing the specified output format.

This tool is usually run automatically as part of the corresponding Isabelle batch process, provided document preparation has been enabled (cf. the `-d` option of the `usedir` utility, §2.4). It may be manually invoked on the generated browser information document output as well, e.g. in case of errors encountered in the batch run.

The `-c` option tells the `document` tool to dispose the document sources after successful operation. This is the right thing to do for sources generated by an Isabelle process, but take care of your files in manual document preparation!

The `-n` and `-o` option specify the final output file name and format, the default is “`document.dvi`”. Note that the result will appear in the parent of the target `DIR`.

The `-t` option tells \LaTeX how to interpret tagged Isabelle command regions. Tags are specified as a comma separated list of modifier/name pairs: “`+foo`” (or just “`foo`”) means to keep, “`-foo`” to drop, and “`/foo`” to fold text tagged as `foo`. The builtin default is equivalent to the tag specification “`/theory,/proof,/ML,+visible,-invisible`”; see also the \LaTeX macros `\isakeeptag`, `\isadroptag`, and `\isafoldtag` in `isabelle.sty`.

Document preparation requires a properly setup “`document`” directory within the logic session sources. This directory is supposed to contain all the files needed to produce the final document — apart from the actual theories which are generated by Isabelle.

For most practical purposes, the `document` tool is smart enough to create any of the specified output formats, taking `root.tex` supplied by the user as a starting point. This even includes multiple runs of \LaTeX to accommodate references and bibliographies (the latter assumes `root.bib` within the same directory).

In more complex situations, a separate `IsaMakefile` for the document sources may be given instead. This should provide targets for any admissible

document format; these have to produce corresponding output files named after `root` as well, e.g. `root.dvi` for target format `dvi`.

When running the session, Isabelle copies the original `document` directory into its proper place within `ISABELLE_BROWSER_INFO` according to the session path. Then, for any processed theory A some \LaTeX source is generated and put there as `A.tex`. Furthermore, a list of all generated theory files is put into `session.tex`. Typically, the root \LaTeX file provided by the user would include `session.tex` to get a document containing all the theories.

The \LaTeX versions of the theories require some macros defined in `isabelle.sty` as distributed with Isabelle. Doing `\usepackage{isabelle}` in `root.tex` should be fine; the underlying Isabelle `latex` utility already includes an appropriate \TeX inputs path.

If the text contains any references to Isabelle symbols (such as `\<forall>`) then `isabellesym.sty` should be included as well. This package contains a standard set of \LaTeX macro definitions `\isasymfoo` corresponding to `\<foo>` (see Appendix A for a complete list of predefined Isabelle symbols). Users may invent further symbols as well, just by providing \LaTeX macros in a similar fashion as in `isabellesym.sty` of the distribution.

For proper setup of PDF documents (with hyperlinks and bookmarks), we recommend to include `pdfsetup.sty` as well. It is safe to do so even without using PDF \LaTeX .

As a final step of document preparation within Isabelle, `isatool document -c` is run on the resulting `document` directory. Thus the actual output document is built and installed in its proper place (as linked by the session's `index.html` if option `-i` of `usedir` has been enabled, cf. §2.1). The generated sources are deleted after successful run of \LaTeX and friends. Note that a separate copy of the sources may be retained by passing an option `-D` to the `usedir` utility when running the session (see also §2.4).

2.6 Running \LaTeX within the Isabelle environment — `isatool latex`

The `latex` utility provides the basic interface for Isabelle document preparation. Its usage is:

Usage: `latex [OPTIONS] [FILE]`

Options are:

`-o FORMAT` specify output format: `dvi` (default), `dvi.gz`, `ps`,
`ps.gz`, `pdf`, `bbl`, `idx`, `sty`, `syms`

Run LaTeX (and related tools) on `FILE` (default `root.tex`),
producing the specified output format.

Appropriate L^AT_EX-related programs are run on the input file, according to the given output format: `latex`, `pdflatex`, `dvips`, `bibtex` (for `bbl`), and `makeindex` (for `idx`). The actual commands are determined from the settings environment (`ISABELLE_LATEX` etc., see §1.1).

The `sty` output format causes the Isabelle style files to be updated from the distribution. This is useful in special situations where the document sources are to be processed another time by separate tools (cf. option `-D` of the `usedir` utility, see §2.4).

The `syms` output is for internal use; it generates lists of symbols that are available without loading additional L^AT_EX packages.

Examples

Invoking `isatool latex` by hand may be occasionally useful when debugging failed attempts of the automatic document preparation stage of batch-mode Isabelle. The abortive process leaves the sources at a certain place within `ISABELLE_BROWSER_INFO`, see the runtime error message for details. This enables users to inspect L^AT_EX runs in further detail, e.g. like this:

```
cd ~/isabelle/browser_info/HOL/Test/document
isatool latex -o pdf
```

Miscellaneous tools

Subsequently we describe various Isabelle related utilities, given in alphabetical order.

3.1 Converting legacy ML scripts — `isatool convert`

The `convert` utility assists in converting legacy ML proof scripts into the new-style format of Isabelle/Isar:

```
Usage: convert [FILES|DIRS...]
```

```
    Recursively find .ML files, converting legacy tactic scripts to  
    Isabelle/Isar tactic emulation.
```

```
    Note: conversion is only approximated, based on some heuristics.
```

```
    Renames old versions of FILES by appending "~0~".
```

```
    Creates new versions of FILES by appending ".thy".
```

The resulting theory text uses the tactic emulation facilities of Isabelle/Isar (see also [2], especially the “Conversion guide” in the appendix). Usually there is some manual tuning required to get an automatically converted script work again — the success rate is around 99% for common ML scripts.

3.2 Displaying documents — `isatool display`

The `display` utility displays documents in DVI format:

```
Usage: display [OPTIONS] FILE

Options are:
  -c          cleanup -- remove FILE after use

Display document FILE (in DVI format).
```

The `-c` option causes the input file to be removed after use. The program for viewing `dvi` files is determined by the `DVI_VIEWER` setting.

3.3 Viewing documentation — `isatool doc`

The `doc` utility displays online documentation:

```
Usage: doc [DOC]

View Isabelle documentation DOC, or show list of available documents.
```

If called without arguments, it lists all available documents. Each line starts with an identifier, followed by a short description. Any of these identifiers may be specified as the first argument in order to have the corresponding document displayed.

The `ISABELLE_DOCS` setting specifies the list of directories (separated by colons) to be scanned for documentations. The program for viewing `dvi` files is determined by the `DVI_VIEWER` setting.

3.4 Getting logic images — `isatool findlogics`

The `findlogics` utility traverses all directories specified in `ISABELLE_PATH`, looking for Isabelle logic images. Its usage is:

```
Usage: findlogics

Collect heap file names from ISABELLE_PATH.
```

The base names of all files found on the path are printed — sorted and with duplicates removed. Also note that lookup in `ISABELLE_PATH` includes the current values of `ML_SYSTEM` and `ML_PLATFORM`. Thus switching to another ML compiler may change the set of logic images available.

3.5 Inspecting the settings environment — isatool getenv

The Isabelle settings environment — as provided by the site-default and user-specific settings files — can be inspected with the `getenv` utility:

```
Usage: getenv [OPTIONS] [VARNAMES ...]
```

Options are:

```
-a          display complete environment
-b          print values only (doesn't work for -a)
```

Get value of VARNAMES from the Isabelle settings.

With the `-a` option, one may inspect the full process environment that Isabelle related programs are run in. This usually contains much more variables than are actually Isabelle settings. Normally, output is a list of lines of the form *name=value*. The `-b` option causes only the values to be printed.

Examples

Get the ML system name and the location where the compiler binaries are supposed to reside as follows:

```
isatool getenv ML_SYSTEM ML_HOME
ML_SYSTEM=polym1
ML_HOME=/usr/share/polym1/x86-linux
```

The next one peeks at the output directory for `isabelle` logic images:

```
isatool getenv -b ISABELLE_OUTPUT
/home/me/isabelle/heapspolym1_x86-linux
```

Here we have used the `-b` option to suppress the `ISABELLE_OUTPUT=` prefix. The value above is what became of the following assignment in the default settings file:

```
ISABELLE_OUTPUT="$ISABELLE_HOME_USER/heapspolym1_x86-linux"
```

Note how the `ML_IDENTIFIER` value got appended automatically to each path component. This is a special feature of `ISABELLE_OUTPUT`.

3.6 Installing standalone Isabelle executables — `isatool install`

By default, the Isabelle binaries (`isabelle`, `isatool` etc.) are just run from their location within the distribution directory, probably indirectly by the shell through its `PATH`. Other schemes of installation are supported by the `install` utility:

Usage: `install [OPTIONS]`

Options are:

`-d DISTDIR` use `DISTDIR` as Isabelle distribution
(default `ISABELLE_HOME`)
`-p DIR` install standalone binaries in `DIR`

Install Isabelle executables with absolute references to the current distribution directory.

The `-d` option overrides the current Isabelle distribution directory as determined by `ISABELLE_HOME`.

The `-p` option installs executable wrapper scripts for `isabelle`, `isatool`, `Isabelle`, containing proper absolute references to the Isabelle distribution directory. A typical `DIR` specification would be some directory expected to be in the shell's `PATH`, such as `/usr/local/bin`. It is important to note that a plain manual copy of the original Isabelle executables just would not work!

3.7 Creating instances of the Isabelle logo — `isatool logo`

The `logo` utility creates any instance of the generic Isabelle logo as an Encapsulated Postscript file (EPS):

Usage: `logo [OPTIONS] NAME`

Create instance `NAME` of the Isabelle logo (as EPS).

Options are:

`-o OUTFILE` set output file (default determined from `NAME`)
`-q` quiet mode

You are encouraged to use this to create a derived logo for your Isabelle project. For example, `isatool logo Bali` creates `isabelle_bali.eps`.

3.8 Isabelle's version of make — isatool make

The Isabelle `make` utility is a very simple wrapper for ordinary Unix `make`:

```
Usage: make [ARGS ...]
```

```
Compile the logic in current directory using IsaMakefile.  
ARGS are directly passed to the system make program.
```

Note that the Isabelle settings environment is also active. Thus one may refer to its values within the `IsaMakefile`, e.g. `$(ISABELLE_OUTPUT)`. Furthermore, programs started from the `make` file also inherit this environment. Typically, `IsaMakefiles` defer the real work to the `usedir` utility, see §2.4.

The basic `IsaMakefile` convention is that the default target builds the actual logic, including its parents if appropriate. The `images` target is intended to build all local logic images, while the `test` target shall build all related examples. The `all` target shall do `images` and `test`.

Examples

Refer to the `IsaMakefiles` of the Isabelle distribution's object-logics as a model for your own developments. For example, see `src/FOL/IsaMakefile`.

3.9 Make all logics — isatool makeall

The `makeall` utility applies Isabelle `make` to all logic directories of the distribution:

```
Usage: makeall [ARGS ...]
```

```
Apply isatool make to all logics (passing ARGS).
```

The arguments `ARGS` are just passed verbatim to each `make` invocation.

3.10 Printing documents — isatool print

The `print` utility prints documents:

Usage: `print [OPTIONS] FILE`

Options are:

`-c` cleanup -- remove FILE after use

Print document FILE.

The `-c` option causes the input file to be removed after use. The printer spool command is determined by the `PRINT_COMMAND` setting.

3.11 Run Isabelle with plain tty interaction — `isatool tty`

The `tty` utility runs the Isabelle process interactively within a plain terminal session:

Usage: `tty [OPTIONS]`

Options are:

`-l NAME` logic image name (default `ISABELLE_LOGIC`)
`-m MODE` add print mode for output
`-p NAME` line editor program name (default `ISABELLE_LINE_EDITOR`)

Run Isabelle process with plain tty interaction, and optional line editor.

The `-l` option specifies the logic image. The `-m` option specifies additional print modes. The `-p` option specifies an alternative line editor (such as the `rlwrap` wrapper for GNU `readline`); the fall-back is to use raw standard input.

3.12 Remove awkward symbol names from theory sources — `isatool unsymbolize`

The `unsymbolize` utility tunes Isabelle theory sources to improve readability for plain ASCII output (e.g. in email communication). Most notably, `unsymbolize` replaces awkward arrow symbols such as `\<Longrightarrow>` by `==>`.

Usage: `unsymbolize [FILES|DIRS...]`

Recursively find `.thy/.ML` files, removing unreadable symbol names.

Note: this is an ad-hoc script; there is no systematic way to replace symbols independently of the inner syntax of a theory!

Renames old versions of FILES by appending `"~~"`.

3.13 Output the version identifier of the Isabelle distribution — `isatool version`

The `version` utility outputs the full version string of the Isabelle distribution being used, e.g. “Isabelle2007: November 2007”. There are no options nor arguments.

3.14 Convert XML to YXML — `isatool yxml`

The `yxml` utility converts a standard XML document (`stdin`) to the much simpler and more efficient YXML format of Isabelle (`stdout`). The YXML format is defined as follows.

1. The encoding is always UTF-8.
2. Body text is represented verbatim (no escaping, no named entities, no CDATA chunks, no comments).
3. Markup elements are represented via ASCII control characters $X = 5$ and $Y = 6$ as follows:

XML	YXML
<code><name attribute=value ... ></code>	<code>X Y name Y attribute=value... X</code>
<code></name ></code>	<code>X Y X</code>

There is no special case for empty body text, i.e. `<foo/>` is treated like `<foo></foo>`. Also note that X and Y may never occur in well-formed XML documents.

Parsing YXML is pretty straight-forward: split the text into chunks separated by X , then split each chunk into sub-chunks separated by Y . Markup chunks start with an empty sub-chunk, and a second empty sub-chunk indicates close of an element. Any other non-empty chunk consists of plain text.

YXML documents may be detected quickly by checking that the first two characters are $X Y$.

Standard Isabelle symbols

Isabelle supports an infinite number of non-ASCII symbols, which are represented in source text as $\langle name \rangle$ (where *name* may be any identifier). It is left to front-end tools how to present these symbols to the user. The collection of predefined standard symbols given below is available by default for Isabelle document output, due to appropriate definitions of $\backslash isasymname$ for each $\langle name \rangle$ in the `isabellesym.sty` file. Most of these symbols are displayed properly in Proof General if used with the X-Symbol package.

Moreover, any single symbol (or ASCII character) may be prefixed by $\langle \wedge sup \rangle$ for superscript and $\langle \wedge sub \rangle$ for subscript, such as $A \langle \wedge sup \rangle \langle star \rangle$ for A^* ; the alternative versions $\langle \wedge isub \rangle$ and $\langle \wedge isup \rangle$ are considered as quasi letters and may be used within identifiers. Sub- and superscripts that span a region of text are marked up with $\langle \wedge bsub \rangle \dots \langle \wedge esub \rangle$ and $\langle \wedge bsup \rangle \dots \langle \wedge esup \rangle$, respectively. Furthermore, all ASCII characters and most other symbols may be printed in bold by prefixing $\langle \wedge bold \rangle$, such as $\langle \wedge bold \rangle \langle alpha \rangle$ for α . Note that $\langle \wedge bold \rangle$ may *not* be combined with sub- or superscripts for single symbols.

Further details of Isabelle document preparation are covered in chapter 2.

$\langle zero \rangle$	0	$\langle one \rangle$	1
$\langle two \rangle$	2	$\langle three \rangle$	3
$\langle four \rangle$	4	$\langle five \rangle$	5
$\langle six \rangle$	6	$\langle seven \rangle$	7
$\langle eight \rangle$	8	$\langle nine \rangle$	9
$\langle A \rangle$	<i>A</i>	$\langle B \rangle$	<i>B</i>
$\langle C \rangle$	<i>C</i>	$\langle D \rangle$	<i>D</i>
$\langle E \rangle$	<i>E</i>	$\langle F \rangle$	<i>F</i>
$\langle G \rangle$	<i>G</i>	$\langle H \rangle$	<i>H</i>
$\langle I \rangle$	<i>I</i>	$\langle J \rangle$	<i>J</i>
$\langle K \rangle$	<i>K</i>	$\langle L \rangle$	<i>L</i>
$\langle M \rangle$	<i>M</i>	$\langle N \rangle$	<i>N</i>
$\langle O \rangle$	<i>O</i>	$\langle P \rangle$	<i>P</i>
$\langle Q \rangle$	<i>Q</i>	$\langle R \rangle$	<i>R</i>

<code>\<S></code>	\mathcal{S}	<code>\<T></code>	\mathcal{T}
<code>\<U></code>	\mathcal{U}	<code>\<V></code>	\mathcal{V}
<code>\<W></code>	\mathcal{W}	<code>\<X></code>	\mathcal{X}
<code>\<Y></code>	\mathcal{Y}	<code>\<Z></code>	\mathcal{Z}
<code>\<a></code>	a	<code>\</code>	b
<code>\<c></code>	c	<code>\<d></code>	d
<code>\<e></code>	e	<code>\<f></code>	f
<code>\<g></code>	g	<code>\<h></code>	h
<code>\<i></code>	i	<code>\<j></code>	j
<code>\<k></code>	k	<code>\<l></code>	l
<code>\<m></code>	m	<code>\<n></code>	n
<code>\<o></code>	o	<code>\<p></code>	p
<code>\<q></code>	q	<code>\<r></code>	r
<code>\<s></code>	s	<code>\<t></code>	t
<code>\<u></code>	u	<code>\<v></code>	v
<code>\<w></code>	w	<code>\<x></code>	x
<code>\<y></code>	y	<code>\<z></code>	z
<code>\<AA></code>	\mathfrak{A}	<code>\<BB></code>	\mathfrak{B}
<code>\<CC></code>	\mathfrak{C}	<code>\<DD></code>	\mathfrak{D}
<code>\<EE></code>	\mathfrak{E}	<code>\<FF></code>	\mathfrak{F}
<code>\<GG></code>	\mathfrak{G}	<code>\<HH></code>	\mathfrak{H}
<code>\<II></code>	\mathfrak{I}	<code>\<JJ></code>	\mathfrak{J}
<code>\<KK></code>	\mathfrak{K}	<code>\<LL></code>	\mathfrak{L}
<code>\<MM></code>	\mathfrak{M}	<code>\<NN></code>	\mathfrak{N}
<code>\<OO></code>	\mathfrak{O}	<code>\<PP></code>	\mathfrak{P}
<code>\<QQ></code>	\mathfrak{Q}	<code>\<RR></code>	\mathfrak{R}
<code>\<SS></code>	\mathfrak{S}	<code>\<TT></code>	\mathfrak{T}
<code>\<UU></code>	\mathfrak{U}	<code>\<VV></code>	\mathfrak{V}
<code>\<WW></code>	\mathfrak{W}	<code>\<XX></code>	\mathfrak{X}
<code>\<YY></code>	\mathfrak{Y}	<code>\<ZZ></code>	\mathfrak{Z}
<code>\<aa></code>	a	<code>\<bb></code>	b
<code>\<cc></code>	c	<code>\<dd></code>	d
<code>\<ee></code>	e	<code>\<ff></code>	f
<code>\<gg></code>	g	<code>\<hh></code>	h
<code>\<ii></code>	i	<code>\<jj></code>	j
<code>\<kk></code>	k	<code>\<ll></code>	l
<code>\<mm></code>	m	<code>\<nn></code>	n
<code>\<oo></code>	o	<code>\<pp></code>	p
<code>\<qq></code>	q	<code>\<rr></code>	r

<code>\<ss></code>	s	<code>\<tt></code>	t
<code>\<uu></code>	u	<code>\<vv></code>	v
<code>\<ww></code>	w	<code>\<xx></code>	x
<code>\<yy></code>	y	<code>\<zz></code>	z
<code>\<alpha></code>	α	<code>\<beta></code>	β
<code>\<gamma></code>	γ	<code>\<delta></code>	δ
<code>\<epsilon></code>	ϵ	<code>\<zeta></code>	ζ
<code>\<eta></code>	η	<code>\<theta></code>	θ
<code>\<iota></code>	i	<code>\<kappa></code>	κ
<code>\<lambda></code>	λ	<code>\<mu></code>	μ
<code>\<nu></code>	ν	<code>\<xi></code>	ξ
<code>\<pi></code>	π	<code>\<rho></code>	ρ
<code>\<sigma></code>	σ	<code>\<tau></code>	τ
<code>\<upsilon></code>	υ	<code>\<phi></code>	φ
<code>\<chi></code>	χ	<code>\<psi></code>	ψ
<code>\<omega></code>	ω	<code>\<Gamma></code>	Γ
<code>\<Delta></code>	Δ	<code>\<Theta></code>	Θ
<code>\<Lambda></code>	Λ	<code>\<Xi></code>	Ξ
<code>\<Pi></code>	Π	<code>\<Sigma></code>	Σ
<code>\<Upsilon></code>	Υ	<code>\<Phi></code>	Φ
<code>\<Psi></code>	Ψ	<code>\<Omega></code>	Ω
<code>\<bool></code>	\mathbb{B}	<code>\<complex></code>	\mathbb{C}
<code>\<nat></code>	\mathbb{N}	<code>\<rat></code>	\mathbb{Q}
<code>\<real></code>	\mathbb{R}	<code>\<int></code>	\mathbb{Z}
<code>\<leftarrow></code>	\leftarrow	<code>\<longleftarrow></code>	\longleftarrow
<code>\<rightarrow></code>	\rightarrow	<code>\<longrightarrow></code>	\longrightarrow
<code>\<Leftarrow></code>	\Lleftarrow	<code>\<Longleftarrow></code>	\Lleftarrow
<code>\<Rightarrow></code>	\Rightarrow	<code>\<Longrightarrow></code>	\Longrightarrow
<code>\<leftrightarrow></code>	\leftrightarrow	<code>\<longleftrightarrow></code>	\longleftrightarrow
<code>\<Leftrightarrow></code>	\Leftrightarrow	<code>\<Longleftrightarrow></code>	\Leftrightarrow
<code>\<mapsto></code>	\mapsto	<code>\<longmapsto></code>	\longmapsto
<code>\<midarrow></code>	\mid	<code>\<Midarrow></code>	\equiv
<code>\<hookleftarrow></code>	\hookleftarrow	<code>\<hookrightarrow></code>	\hookrightarrow
<code>\<leftharpoondown></code>	\leftharpoondown	<code>\<rightharpoondown></code>	\rightharpoondown
<code>\<leftharpoonup></code>	\leftharpoonup	<code>\<rightharpoonup></code>	\rightharpoonup
<code>\<rightleftharpoons></code>	\rightleftharpoons	<code>\<leadsto></code>	\rightsquigarrow
<code>\<downharpoonleft></code>	\downdownarrows	<code>\<downharpoonright></code>	\downdownarrows
<code>\<upharpoonleft></code>	\upharpoonleft	<code>\<upharpoonright></code>	\upharpoonright
<code>\<restriction></code>	\upharpoonright	<code>\<Colon></code>	$::$

<code>\<up></code>	\uparrow	<code>\<Up></code>	\Uparrow
<code>\<down></code>	\downarrow	<code>\<Down></code>	\Downarrow
<code>\<updown></code>	\updownarrow	<code>\<Updown></code>	\Updownarrow
<code>\<langle></code>	\langle	<code>\<rangle></code>	\rangle
<code>\<lceil></code>	\lceil	<code>\<rceil></code>	\rceil
<code>\<lfloor></code>	\lfloor	<code>\<rfloor></code>	\rfloor
<code>\<lparr></code>	$\langle\! $	<code>\<rparr></code>	$\! \rangle$
<code>\<lbrakk></code>	\llbracket	<code>\<rbrakk></code>	\rrbracket
<code>\<lbrace></code>	$\{$	<code>\<rbrace></code>	$\}$
<code>\<guillemotleft></code>	\ll	<code>\<guillemotright></code>	\gg
<code>\<bottom></code>	\perp	<code>\<top></code>	\top
<code>\<and></code>	\wedge	<code>\<And></code>	\bigwedge
<code>\<or></code>	\vee	<code>\<Or></code>	\bigvee
<code>\<forall></code>	\forall	<code>\<exists></code>	\exists
<code>\<nexists></code>	\nexists	<code>\<not></code>	\neg
<code>\<box></code>	\square	<code>\<diamond></code>	\diamond
<code>\<turnstile></code>	\vdash	<code>\<Turnstile></code>	\Vdash
<code>\<tturnstile></code>	\Vdash	<code>\<TTurnstile></code>	\Vdash
<code>\<stileturn></code>	\dashv	<code>\<surd></code>	\surd
<code>\<le></code>	\leq	<code>\<ge></code>	\geq
<code>\<lless></code>	\ll	<code>\<ggreater></code>	\gg
<code>\<lesssim></code>	\lesssim	<code>\<greatersim></code>	\gtrsim
<code>\<lessapprox></code>	\approx	<code>\<greaterapprox></code>	\gtrapprox
<code>\<in></code>	\in	<code>\<notin></code>	\notin
<code>\<subset></code>	\subset	<code>\<supset></code>	\supset
<code>\<subseteq></code>	\subseteq	<code>\<supseteq></code>	\supseteq
<code>\<sqsubset></code>	\sqsubset	<code>\<sqsupset></code>	\sqsupset
<code>\<sqsubsetq></code>	\sqsubseteq	<code>\<sqsupseteq></code>	\sqsupseteq
<code>\<inter></code>	\cap	<code>\<Inter></code>	\bigcap
<code>\<union></code>	\cup	<code>\<Union></code>	\bigcup
<code>\<squnion></code>	\sqcup	<code>\<Squnion></code>	\bigsqcup
<code>\<sqinter></code>	\sqcap	<code>\<Sqinter></code>	\bigsqcap
<code>\<setminus></code>	\setminus	<code>\<propto></code>	\propto
<code>\<uplus></code>	\oplus	<code>\<Uplus></code>	\bigoplus
<code>\<noteq></code>	\neq	<code>\<sim></code>	\sim
<code>\<doteq></code>	\doteq	<code>\<simeq></code>	\simeq
<code>\<approx></code>	\approx	<code>\<asymp></code>	\asymp
<code>\<cong></code>	\cong	<code>\<smile></code>	\smile
<code>\<equiv></code>	\equiv	<code>\<frown></code>	\frown

<code>\<Join></code>	\bowtie	<code>\<bowtie></code>	\bowtie
<code>\<prec></code>	\prec	<code>\<succ></code>	\succ
<code>\<preceq></code>	\preceq	<code>\<succeq></code>	\succeq
<code>\<parallel></code>	\parallel	<code>\<bar></code>	$ $
<code>\<plusminus></code>	\pm	<code>\<minusplus></code>	\mp
<code>\<times></code>	\times	<code>\<div></code>	\div
<code>\<cdot></code>	\cdot	<code>\<star></code>	\star
<code>\<bullet></code>	\bullet	<code>\<circ></code>	\circ
<code>\<dagger></code>	\dagger	<code>\<ddagger></code>	\ddagger
<code>\<lhd></code>	\triangleleft	<code>\<rh></code>	\triangleright
<code>\<unlhd></code>	\triangleleft	<code>\<unrhd></code>	\triangleright
<code>\<triangleleft></code>	\triangleleft	<code>\<triangleright></code>	\triangleright
<code>\<triangle></code>	\triangle	<code>\<triangleq></code>	\triangleq
<code>\<oplus></code>	\oplus	<code>\<Oplus></code>	\oplus
<code>\<otimes></code>	\otimes	<code>\<Otimes></code>	\otimes
<code>\<odot></code>	\odot	<code>\<Odot></code>	\odot
<code>\<ominus></code>	\ominus	<code>\<oslash></code>	\oslash
<code>\<dots></code>	\dots	<code>\<cdots></code>	\dots
<code>\<Sum></code>	\sum	<code>\<Prod></code>	\prod
<code>\<Coproduct></code>	\coprod	<code>\<infinity></code>	∞
<code>\<integral></code>	\int	<code>\<ointegral></code>	\oint
<code>\<clubsuit></code>	\clubsuit	<code>\<diamondsuit></code>	\diamondsuit
<code>\<heartsuit></code>	\heartsuit	<code>\<spadesuit></code>	\spadesuit
<code>\<aleph></code>	\aleph	<code>\<emptyset></code>	\emptyset
<code>\<nabla></code>	∇	<code>\<partial></code>	∂
<code>\<Re></code>	\Re	<code>\<Im></code>	\Im
<code>\<flat></code>	\flat	<code>\<natural></code>	\natural
<code>\<sharp></code>	\sharp	<code>\<angle></code>	\angle
<code>\<copyright></code>	\copyright	<code>\<registered></code>	$\text{\textcircled{R}}$
<code>\<hyphen></code>	$-$	<code>\<inverse></code>	$^{-1}$
<code>\<onesuperior></code>	1	<code>\<onequarter></code>	$\frac{1}{4}$
<code>\<twosuperior></code>	2	<code>\<onehalf></code>	$\frac{1}{2}$
<code>\<threesuperior></code>	3	<code>\<threequarters></code>	$\frac{3}{4}$
<code>\<ordfeminine></code>	a	<code>\<ordmasculine></code>	o
<code>\<section></code>	\S	<code>\<paragraph></code>	\P
<code>\<exclamdown></code>	$¡$	<code>\<questiondown></code>	$¿$
<code>\<euro></code>	€	<code>\<pounds></code>	£
<code>\<yen></code>	¥	<code>\<cent></code>	¢
<code>\<currency></code>	¤	<code>\<degree></code>	$^\circ$

<code>\<amalg></code>	II	<code>\<mho></code>	Ω
<code>\<lozenge></code>	◇	<code>\<wp></code>	∅
<code>\<wrong></code>	∫	<code>\<struct></code>	◇
<code>\<acute></code>	´	<code>\<index></code>	1
<code>\<dieresis></code>	¨	<code>\<cedilla></code>	¸
<code>\<hungarumlaut></code>	˘	<code>\<spacespace></code>	
<code>\<module></code>	⟨module⟩	<code>\<some></code>	ε

Bibliography

- [1] David Aspinall. Proof General. <http://proofgeneral.inf.ed.ac.uk/>.
- [2] Lawrence C. Paulson. *The Isabelle Reference Manual*.
<http://isabelle.in.tum.de/doc/ref.pdf>.
- [3] Markus Wenzel. *The Isabelle/Isar Reference Manual*.
<http://isabelle.in.tum.de/doc/isar-ref.pdf>.

Index

browser tool, 13
convert tool, 25
display tool, 25
doc tool, 26
document tool, 21
DVI_VIEWER setting, 5
findlogics tool, 26
getenv tool, 27
HOL_USEDIR_OPTIONS setting, 19
HTML, 19
INSTALL, 1
install tool, 28
ISABELLE setting, 3, 4
Isabelle, 1, 9
isabelle, 1, 6
isabelle-interface, 9
isabelle-process, 6
ISABELLE_BIBTEX setting, 5
ISABELLE_BROWSER_INFO setting, 4, 12
ISABELLE_DOC_FORMAT setting, 5
ISABELLE_DOCS setting, 5
ISABELLE_DVIPS setting, 5
ISABELLE_FILE_IDENT setting, 5
ISABELLE_HOME setting, 2, 3
ISABELLE_HOME_USER setting, 3
ISABELLE_IDENTIFIER setting, 4
ISABELLE_INTERFACE setting, 5, 9
ISABELLE_LATEX setting, 5
ISABELLE_LINE_EDITOR setting, 4
ISABELLE_LOGIC setting, 4
ISABELLE_OUTPUT setting, 3, 4
ISABELLE_PATH setting, 4
ISABELLE_PDFLATEX setting, 5
ISABELLE_TMP_PREFIX setting, 5
ISABELLE_TOOLS setting, 5
ISABELLE_USEDIR_OPTIONS setting, 4, 12, 19
IsaMakefile, 19, 29
ISATOOL setting, 3, 4
isatool, 1, 5
latex tool, 23
logo tool, 28
make tool, 29
makeall tool, 29
mkdir tool, 16
ML_HOME setting, 4
ML_IDENTIFIER setting, 4
ML_OPTIONS setting, 4
ML_PLATFORM setting, 4
ML_SYSTEM setting, 4
PDF_VIEWER setting, 5
print tool, 29
PRINT_COMMAND setting, 5
settings, **1**
theory browsing information, **12**
theory graph browser, **13**
tty tool, 30
unsymbolize tool, 30
usedir tool, 18
user interface
 Proof General, 9

version tool, 31

yxml tool, 31