



The Isabelle System Manual

Makarius Wenzel and Stefan Berghofer
TU München

12 February 2013

Contents

1	The Isabelle system environment	1
1.1	Isabelle settings	1
1.1.1	Bootstrapping the environment	2
1.1.2	Common variables	3
1.1.3	Additional components	6
1.2	The raw Isabelle process	7
1.3	The Isabelle tool wrapper	10
2	User interfaces	11
2.1	Isabelle/jEdit Prover IDE	11
2.2	Proof General / Emacs	12
2.3	Plain TTY interaction	12
2.4	Theory graph browser	13
2.4.1	Invoking the graph browser	13
2.4.2	Using the graph browser	14
2.4.3	Syntax of graph definition files	16
3	Isabelle sessions and build management	17
3.1	Session ROOT specifications	17
3.2	System build options	20
3.3	Invoking the build process	22
3.4	Build dialog	25
4	Presenting theories	27
4.1	Generating theory browser information	27
4.2	Preparing session root directories	28
4.3	Preparing Isabelle session documents	29
4.4	Running L ^A T _E X within the Isabelle environment	31

5 Isabelle/Scala development tools	33
5.1 Java Runtime Environment within Isabelle	33
5.2 Scala toplevel	33
5.3 Scala compiler	34
6 Miscellaneous tools	35
6.1 Resolving Isabelle components	35
6.2 Displaying documents	36
6.3 Viewing documentation	36
6.4 Shell commands within the settings environment	37
6.5 Getting logic images	37
6.6 Inspecting the settings environment	37
6.7 Installing standalone Isabelle executables	38
6.8 Creating instances of the Isabelle logo	39
6.9 Isabelle wrapper for make	39
6.10 Creating Isabelle session directories	40
6.11 Printing documents	41
6.12 Remove awkward symbol names from theory sources	41
6.13 Running Isabelle sessions	42
6.14 Output the version identifier of the Isabelle distribution	45
6.15 Convert XML to YXML	45
Bibliography	47
Index	48

The Isabelle system environment

This manual describes Isabelle together with related tools and user interfaces as seen from a system oriented view. See also the *Isabelle/Isar Reference Manual* [3] for the actual Isabelle input language and related concepts, and *The Isabelle/Isar Implementation Manual* [2] for the main concepts of the underlying implementation in Isabelle/ML.

The Isabelle system environment provides the following basic infrastructure to integrate tools smoothly.

1. The *Isabelle settings* mechanism provides process environment variables to all Isabelle executables (including tools and user interfaces).
2. The raw *Isabelle process* (`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. Regular users rarely need to care about the low-level process.
3. The main *Isabelle tool wrapper* (`isabelle`) provides a generic startup environment Isabelle related utilities, user interfaces etc. Such tools automatically benefit from the settings mechanism.

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 clearly defined 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. Occasionally, users would still want to put the `$ISABELLE_HOME/bin` directory into their shell's search path, but this is not required.

1.1.1 Bootstrapping the environment

Isabelle executables need to be run within a proper settings environment. This is bootstrapped as described below, on the first invocation of one of the outer wrapper scripts (such as `isabelle`). This happens only once for each process tree, i.e. the environment is passed to subprocesses according to regular Unix conventions.

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 `isabelle install` tool. Symbolic links are admissible, but a plain copy of the `$ISABELLE_HOME/bin` files will not work!

2. The file `$ISABELLE_HOME/etc/settings` is run as a `bash` 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 `$ISABELLE_HOME_USER/etc/settings` (if it exists) is run in the same way as the site default settings. Note that the variable `ISABELLE_HOME_USER` has already been set before — usually to something like `$USER_HOME/.isabelle/IsabelleXXXX`.

Thus individual users may override the site-wide defaults. See also file `$ISABELLE_HOME/etc/user-settings.sample` in the distribution. Typically, a user 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 maintenance problems later, when the Isabelle installation is updated or changed otherwise.

Since settings files are regular GNU `bash` scripts, 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_PROCESS` and `ISABELLE_TOOL` are set automatically to the absolute path names of the `isabelle-process` and `isabelle` 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.

Note that the settings environment may be inspected with the `isabelle getenv` tool. This might help to figure out the effect of complex settings scripts.

1.1.2 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 `*`.

`USER_HOME*` Is the cross-platform user home directory. On Unix systems this is usually the same as `HOME`, but on Windows it is the regular home directory of the user, not the one of within the Cygwin root file-system.¹

`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!

¹Cygwin itself offers another choice whether its `HOME` should point to the `/home` directory tree or the Windows user home.

`ISABELLE_HOME_USER` is the user-specific counterpart of `ISABELLE_HOME`.

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

`ISABELLE_PLATFORM_FAMILY*` is automatically set to the general platform family: `linux`, `macos`, `windows`. Note that platform-dependent tools usually need to refer to the more specific identification according to `ISABELLE_PLATFORM`, `ISABELLE_PLATFORM32`, `ISABELLE_PLATFORM64`.

`ISABELLE_PLATFORM*` is automatically set to a symbolic identifier for the underlying hardware and operating system. The Isabelle platform identification always refers to the 32 bit variant, even this is a 64 bit machine. Note that the ML or Java runtime may have a different idea, depending on which binaries are actually run.

`ISABELLE_PLATFORM64*` is similar to `ISABELLE_PLATFORM` but refers to the proper 64 bit variant on a platform that supports this; the value is empty for 32 bit. Note that the following bash expression (including the quotes) prefers the 64 bit platform, if that is available:

```
"${ISABELLE_PLATFORM64:-$ISABELLE_PLATFORM}"
```

`ISABELLE_PROCESS*`, `ISABELLE_TOOL*` are automatically set to the full path names of the `isabelle-process` and `isabelle` executables, respectively. Thus other tools and scripts need not assume that the `$ISABELLE_HOME/bin` directory is on the current search path of the shell.

`ISABELLE_IDENTIFIER*` refers to the name of this Isabelle distribution, e.g. "Isabelle2012".

`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 `$ISABELLE_HOME/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_JDK_HOME` needs to point to a full JDK (Java Development Kit) installation with `javac` and `jar` executables. This is essential for Isabelle/Scala and other JVM-based tools to work properly. Note that conventional `JAVA_HOME` usually points to the JRE (Java Runtime Environment), not JDK.

`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 §4.1). 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 the `isabelle tty` interface.

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

`ISABELLE_LATEX`, `ISABELLE_PDFLATEX`, `ISABELLE_BIBTEX`, `ISABELLE_DVIPS` refer to \LaTeX related tools for Isabelle document preparation (see also §4.4).

`ISABELLE_TOOLS` is a colon separated list of directories that are scanned by `isabelle` for external utility programs (see also §1.3).

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

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

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

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

1.1.3 Additional components

Any directory may be registered as an explicit *Isabelle component*. The general layout conventions are that of the main Isabelle distribution itself, and the following two files (both optional) have a special meaning:

- `etc/settings` holds additional settings that are initialized when bootstrapping the overall Isabelle environment, cf. §1.1.1. As usual, the content is interpreted as a `bash` script. It may refer to the component's enclosing directory via the `COMPONENT` shell variable.

For example, the following setting allows to refer to files within the component later on, without having to hardwire absolute paths:

```
MY_COMPONENT_HOME="$COMPONENT"
```

Components can also add to existing Isabelle settings such as `ISABELLE_TOOLS`, in order to provide component-specific tools that can be invoked by end-users. For example:

```
ISABELLE_TOOLS="$ISABELLE_TOOLS:$COMPONENT/lib/Tools"
```

- `etc/components` holds a list of further sub-components of the same structure. The directory specifications given here can be either absolute (with leading `/`) or relative to the component's main directory.

The root of component initialization is `ISABELLE_HOME` itself. After initializing all of its sub-components recursively, `ISABELLE_HOME_USER` is included in the same manner (if that directory exists). This allows to install private components via `$ISABELLE_HOME_USER/etc/components`, although it is often more convenient to do that programmatically via the `init_component` shell function in the `etc/settings` script of `$ISABELLE_HOME_USER` (or any other component directory). For example:

```
init_component "$HOME/screwdriver-2.0"
```

This is tolerant wrt. missing component directories, but might produce a warning.

More complex situations may be addressed by initializing components listed in a given catalog file, relatively to some base directory:

```
init_components "$HOME/my_component_store" "some_catalog_file"
```

The component directories listed in the catalog file are treated as relative to the given base directory.

See also §6.1 for some tool-support for resolving components that are formally initialized but not installed yet.

1.2 The raw Isabelle process

The `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-process [OPTIONS] [INPUT] [OUTPUT]
```

Options are:

```
-I          startup Isar interaction mode
-P          startup Proof General interaction mode
-S          secure mode -- disallow critical operations
-T ADDR    startup process wrapper, with socket address
-W IN:OUT  startup process wrapper, with input/output fifos
-X          startup PGIP interaction mode
-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 50–200 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.

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 `-T` or `-W` option makes Isabelle enter a special process wrapper for interaction via Isabelle/Scala, see also `~/src/Pure/System/isabelle_process.scala`. The protocol between the ML and JVM process is private to the implementation.

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-process
```

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

```
isabelle-process HOL Test
```

Ending this session normally (e.g. by typing control-D) dumps the whole ML system state into `Test` (be prepared for more than 100 MB):

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

```
isabelle-process Test
```

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

```
isabelle-process -r Test
```

The next example demonstrates batch execution of Isabelle. We retrieve the `Main` theory value from the theory loader within ML (observe the delicate quoting rules for the Bash shell vs. ML):

```
isabelle-process -e 'Thy_Info.get_theory "Main";' -q -r HOL
```

Note that the output text will be interspersed with additional junk messages by the ML runtime environment. The `-W` option allows to communicate with the Isabelle process via an external program in a more robust fashion.

1.3 The Isabelle tool wrapper

All Isabelle related tools and interfaces are called via a common wrapper — `isabelle`:

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

```
Start Isabelle tool NAME with ARGS; pass "-?" for tool specific help.
```

```
Available tools:
```

```
...
```

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 `isabelle` 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 Isabelle distribution:

```
isabelle doc
```

View a certain document as follows:

```
isabelle doc system
```

Query the Isabelle settings environment:

```
isabelle getenv ISABELLE_HOME_USER
```

User interfaces

2.1 Isabelle/jEdit Prover IDE

The `isabelle jedit` tool invokes a version of jEdit¹ that has been augmented with some plugins to provide a fully-featured Prover IDE:

```
Usage: isabelle jedit [OPTIONS]
      [FILES ...]
```

Options are:

```
-J OPTION  add JVM runtime option (default JEDIT_JAVA_OPTIONS)
-b         build only
-d DIR     include session directory
-f         fresh build
-j OPTION  add jEdit runtime option (default JEDIT_OPTIONS)
-l NAME    logic image name (default ISABELLE_LOGIC)
-m MODE    add print mode for output
-n         no build dialog for session image on startup
-s         system build mode for session image
```

Start jEdit with Isabelle plugin setup and open theory FILES (default "\$USER_HOME/Scratch.thy").

The `-l` option specifies the session name of the logic image to be used for proof processing. Additional session root directories may be included via option `-d` to augment that name space (see also §3.3).

By default, the specified image is checked and built on demand, see also `isabelle build_dialog`. The `-s` determines where to store the result session image (see also §3.3). The `-n` option bypasses the session build dialog.

The `-m` option specifies additional print modes for the prover process.

The `-J` and `-j` options allow to pass additional low-level options to the JVM or jEdit, respectively. The defaults are provided by the Isabelle settings environment (§1.1).

The `-b` and `-f` options control the self-build mechanism of Isabelle/jEdit. This is only relevant for building from sources, which also requires an aux-

¹<http://www.jedit.org/>

iliary `jedit_build` component.² Note that official Isabelle releases already include a version of Isabelle/jEdit that is built properly.

2.2 Proof General / Emacs

The `isabelle emacs` tool invokes a version of Emacs and Proof General³ within the regular Isabelle settings environment (§1.1). This is more convenient than starting Emacs separately, loading the Proof General LISP files, and then attempting to start Isabelle with dynamic `PATH` lookup etc., although it might fail if a different version of Proof General is already part of the Emacs installation of the operating system.

The actual interface script is part of the Proof General distribution; its usage depends on the particular version. 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 of the prover process. The following Isabelle settings are particularly important for Proof General:

`PROOFGENERAL_HOME` points to the main installation directory of the Proof General distribution. This is implicitly provided for versions of Proof General that are distributed as Isabelle component, see also §1.1.3; otherwise it needs to be configured manually.

`PROOFGENERAL_OPTIONS` is implicitly prefixed to the command line of any invocation of the Proof General `interface` script. This allows to provide persistent default options for the invocation of `isabelle emacs`.

2.3 Plain TTY interaction

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

²<http://isabelle.in.tum.de/components>

³<http://proofgeneral.inf.ed.ac.uk/>

Usage: `isabelle 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 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. Regular interaction works via the standard Isabelle/Isar toplevel loop. The Isar command `exit` drops out into the bare-bones ML system, which is occasionally useful for debugging of the Isar infrastructure itself. Invoking `Isar.loop ()`; in ML will return to the Isar toplevel.

2.4 Theory graph browser

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.

2.4.1 Invoking the graph browser

The stand-alone version of the graph browser is wrapped up as `isabelle browser`:

Usage: `isabelle browser [OPTIONS] [GRAPHFILE]`

Options are:

```
-b          Admin/build only
-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 `-b` option indicates that this is for administrative build only, i.e. no browser popup if no files are given.

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.

2.4.2 Using the graph browser

The browser’s main window, which is shown in figure 2.1, 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 2.1: 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

Due to Java Applet security restrictions this menu is only available in the full application 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.

2.4.3 Syntax of graph definition files

A graph definition file has the following syntax:

```
graph = { vertex ; }+
vertex = vertex_name vertex_ID dir_name [ + ] path [ < | > ] { vertex_ID }*
```

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

vertex_name The name of the vertex.

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

dir_name The name of the “directory” the vertex should be placed in. A “+” sign after *dir_name* 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.

Isabelle sessions and build management

An Isabelle *session* consists of a collection of related theories that may be associated with formal documents (chapter 4). There is also a notion of *persistent heap* image to capture the state of a session, similar to object-code in compiled programming languages. Thus the concept of session resembles that of a “project” in common IDE environments, but the specific name emphasizes the connection to interactive theorem proving: the session wraps-up the results of user-interaction with the prover in a persistent form.

Application sessions are built on a given parent session, which may be built recursively on other parents. Following this path in the hierarchy eventually leads to some major object-logic session like *HOL*, which itself is based on *Pure* as the common root of all sessions.

Processing sessions may take considerable time. Isabelle build management helps to organize this efficiently. This includes support for parallel build jobs, in addition to the multithreaded theory and proof checking that is already provided by the prover process itself.


3.1 Session ROOT specifications

Session specifications reside in files called `ROOT` within certain directories, such as the home locations of registered Isabelle components or additional project directories given by the user.


The `ROOT` file format follows the lexical conventions of the *outer syntax* of Isabelle/Isar, see also [3]. This defines common forms like identifiers, names, quoted strings, verbatim text, nested comments etc. The grammar for a single *session_entry* is given as syntax diagram below; each `ROOT` file may contain multiple session specifications like this.

Isabelle/jEdit (§2.1) includes a simple editing mode `isabelle-root` for session `ROOT` files, which is enabled by default for any file of that name.


session_entry




body




spec



groups




dir




description




options




opts




value



theories



files



session $A = B + \textit{body}$ defines a new session A based on parent session B , with its content given in \textit{body} (theories and auxiliary source files). Note that a parent (like *HOL*) is mandatory in practical applications: only Isabelle/Pure can bootstrap itself from nothing.

All such session specifications together describe a hierarchy (tree) of sessions, with globally unique names. The new session name A should be sufficiently long and descriptive to stand on its own in a potentially large library.

session A (*groups*) indicates a collection of groups where the new session is a member. Group names are uninterpreted and merely follow certain conventions. For example, the Isabelle distribution tags some important sessions by the group name called “*main*”. Other projects may invent their own conventions, but this requires some care to avoid clashes within this unchecked name space.

session A **in** \textit{dir} specifies an explicit directory for this session; by default this is the current directory of the `ROOT` file.

All theories and auxiliary source files are located relatively to the session directory. The prover process is run within the same as its current working directory.

description \textit{text} is a free-form annotation for this session.

options $[x = a, y = b, z]$ defines separate options (§3.2) that are used when processing this session, but *without* propagation to child sessions. Note that z abbreviates $z = true$ for Boolean options.

theories *options names* specifies a block of theories that are processed within an environment that is augmented by the given options, in addition to the global session options given before. Any number of blocks of **theories** may be given. Options are only active for each **theories** block separately.

files *files* lists additional source files that are involved in the processing of this session. This should cover anything outside the formal content of the theory sources, say some auxiliary T_EX files that are required for document processing. In contrast, files that are loaded formally within a theory, e.g. via **ML_file**, need not be declared again.

Examples

See `~/src/HOL/ROOT` for a diversity of practically relevant situations, although it uses relatively complex quasi-hierarchic naming conventions like *HOL-SPARK*, *HOL-SPARK-Examples*. An alternative is to use unqualified names that are relatively long and descriptive, as in the Archive of Formal Proofs (<http://afp.sf.net>), for example.

3.2 System build options

See `~/etc/options` for the main defaults provided by the Isabelle distribution. Isabelle/jEdit (§2.1) includes a simple editing mode **isabelle-options** for this file-format.

The following options are particularly relevant to build Isabelle sessions, in particular with document preparation (chapter 4).

- **browser_info** controls output of HTML browser info, see also §4.1.
- **document** specifies the document output format, see **isabelle document** option `-o` in §4.3. In practice, the most relevant values are `document=false` or `document=pdf`.
- **document_output** specifies an alternative directory for generated output of the document preparation system; the default is within the

ISABELLE_BROWSER_INFO hierarchy as explained in §4.1. See also `isabelle mkroot`, which generates a default configuration with output readily available to the author of the document.

- `document_variants` specifies document variants as a colon-separated list of `name=tags` entries, corresponding to `isabelle document` options `-n` and `-t`.

For example, `document_variants=document:outline=/proof,/ML` indicates two documents: the one called `document` with default tags, and the other called `outline` where proofs and ML sections are folded.

Document variant names are just a matter of conventions. It is also possible to use different document variant names (without tags) for different document root entries, see also §4.3.

- `document_graph` tells whether the generated document files should include a theory graph (cf. §2.4 and §4.1). The resulting EPS or PDF file can be included as graphics in \LaTeX .

Note that this option is usually determined as static parameter of some session (e.g. within its `ROOT` file) and *not* given globally or on the command line of `isabelle build`.

- `threads` determines the number of worker threads for parallel checking of theories and proofs. The default 0 means that a sensible maximum value is determined by the underlying hardware. For machines with many cores or with hyperthreading, this is often requires manual adjustment (on the command-line or within personal settings or preferences, not within a session `ROOT`).
- `condition` specifies a comma-separated list of process environment variables (or Isabelle settings) that are required for the subsequent theories to be processed. Conditions are considered “true” if the corresponding environment value is defined and non-empty.

For example, the `condition=ISABELLE_FULL_TEST` may be used to guard extraordinary theories, which are meant to be enabled explicitly via some shell prefix `env ISABELLE_FULL_TEST=true` before invoking `isabelle build`.

- `timeout` specifies a real wall-clock timeout (in seconds) for the session as a whole. The timer is controlled outside the ML process by the JVM that runs Isabelle/Scala. Thus it is relatively reliable in canceling

processes that get out of control, even if there is a deadlock without CPU time usage.

The `isabelle options` tool prints Isabelle system options. Its command-line usage is:

```
Usage: isabelle options [OPTIONS] [MORE_OPTIONS ...]
```

Options are:

```
-b          include $ISABELLE_BUILD_OPTIONS
-l          list options
-x FILE     export to FILE in YXML format
```

Report Isabelle system options, augmented by `MORE_OPTIONS` given as arguments `NAME=VAL` or `NAME`.

The command line arguments provide additional system options of the form `name=value` or `name` for Boolean options.

Option `-b` augments the implicit environment of system options by the ones of `ISABELLE_BUILD_OPTIONS`, cf. §3.3.

Option `-x` specifies a file to export the result in YXML format, instead of printing it in human-readable form.

3.3 Invoking the build process

The `isabelle build` tool invokes the build process for Isabelle sessions. It manages dependencies between sessions, related sources of theories and auxiliary files, and target heap images. Accordingly, it runs instances of the prover process with optional document preparation. Its command-line usage is:¹

¹Isabelle/Scala provides the same functionality via `isabelle.Build.build`.

```
Usage: isabelle build [OPTIONS] [SESSIONS ...]
```

Options are:

```
-D DIR      include session directory and select its sessions
-R         operate on requirements of selected sessions
-a         select all sessions
-b         build heap images
-c         clean build
-d DIR     include session directory
-g NAME    select session group NAME
-j INT     maximum number of parallel jobs (default 1)
-l         list session source files
-n         no build -- test dependencies only
-o OPTION  override session configuration OPTION
           (via NAME=VAL or NAME)
-s         system build mode: produce output in ISABELLE_HOME
-v         verbose
```

```
Build and manage Isabelle sessions, depending on implicit
ISABELLE_BUILD_OPTIONS="..."
```

```
ML_PLATFORM="..."
ML_HOME="..."
ML_SYSTEM="..."
ML_OPTIONS="..."
```

Isabelle sessions are defined via session *ROOT* files as described in (§3.1). The totality of sessions is determined by collecting such specifications from all Isabelle component directories (§1.1.3), augmented by more directories given via options *-d DIR* on the command line. Each such directory may contain a session *ROOT* file with several session specifications.

Any session root directory may refer recursively to further directories of the same kind, by listing them in a catalog file *ROOTS* line-by-line. This helps to organize large collections of session specifications, or to make *-d* command line options persistent (say within $\$ISABELLE_HOME_USER/ROOTS$).

The subset of sessions to be managed is determined via individual *SESSIONS* given as command-line arguments, or session groups that are given via one or more options *-g NAME*. Option *-a* selects all sessions. The build tool takes session dependencies into account: the set of selected sessions is completed by including all ancestors.

Option *-R* reverses the selection in the sense that it refers to its requirements: all ancestor sessions excluding the original selection. This allows to prepare the stage for some build process with different options, before running the main build itself (without option *-R*).

Option `-D` is similar to `-d`, but selects all sessions that are defined in the given directories.

The build process depends on additional options (§3.2) that are passed to the prover eventually. The settings variable `ISABELLE_BUILD_OPTIONS` allows to provide additional defaults, e.g. `ISABELLE_BUILD_OPTIONS="document=pdf threads=4"`. Moreover, the environment of system build options may be augmented on the command line via `-o name=value` or `-o name`, which abbreviates `-o name=true` for Boolean options. Multiple occurrences of `-o` on the command-line are applied in the given order.

Option `-b` ensures that heap images are produced for all selected sessions. By default, images are only saved for inner nodes of the hierarchy of sessions, as required for other sessions to continue later on.

Option `-c` cleans all descendants of the selected sessions before performing the specified build operation.

Option `-n` omits the actual build process after the preparatory stage (including optional cleanup). Note that the return code always indicates the status of the set of selected sessions.

Option `-j` specifies the maximum number of parallel build jobs (prover processes). Each prover process is subject to a separate limit of parallel worker threads, cf. system option `threads`.

Option `-s` enables *system mode*, which means that resulting heap images and log files are stored in `$ISABELLE_HOME/heaps` instead of the default location `ISABELLE_OUTPUT` (which is normally in `ISABELLE_HOME_USER`, i.e. the user's home directory).

Option `-v` increases the general level of verbosity. Option `-l` lists the source files that contribute to a session.

Examples

Build a specific logic image:

```
isabelle build -b HOLCF
```

Build the main group of logic images:

```
isabelle build -b -g main
```

Provide a general overview of the status of all Isabelle sessions, without building anything:

```
isabelle build -a -n -v
```

Build all sessions with HTML browser info and PDF document preparation:

```
isabelle build -a -o browser_info -o document=pdf
```

Build all sessions with a maximum of 8 parallel prover processes and 4 worker threads each (on a machine with many cores):

```
isabelle build -a -j8 -o threads=4
```

Build some session images with cleanup of their descendants, while retaining their ancestry:

```
isabelle build -b -c HOL-Boogie HOL-SPARK
```

Clean all sessions without building anything:

```
isabelle build -a -n -c
```

Build all sessions from some other directory hierarchy, according to the settings variable `AFP` that happens to be defined inside the Isabelle environment:

```
isabelle build -D '$AFP'
```

Inform about the status of all sessions required for `AFP`, without building anything yet:

```
isabelle build -D '$AFP' -R -v -n
```

3.4 Build dialog

The `isabelle build_dialog` provides a simple GUI wrapper to the tool Isabelle `isabelle build` tool. This enables user interfaces like Isabelle/jEdit §2.1 to provide read-made logic image on startup. Its command-line usage is:

Usage: `isabelle build_dialog [OPTIONS] LOGIC`

Options are:

<code>-L OPTION</code>	default logic via system option
<code>-d DIR</code>	include session directory
<code>-l NAME</code>	logic session name
<code>-s</code>	system build mode: produce output in <code>ISABELLE_HOME</code>

Build Isabelle logic session image via GUI dialog (default: `$ISABELLE_LOGIC`).

Option `-l` specifies an explicit logic session name. Option `-L` specifies a system option name as fall-back to determine the logic session name. If both are omitted or have empty value, `ISABELLE_LOGIC` is used as default.

Options `-d` and `-s` have the same meaning as for the command-line `isabelle build` tool itself.

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 *sessions* (chapter 3). 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 further on in the hierarchy.

The tools `isabelle mkroot` and `isabelle build` provide the primary means for managing Isabelle sessions, including proper setup for presentation; `isabelle build` takes care to have `isabelle-process` run any additional stages required for document preparation, notably the `isabelle document` and `isabelle latex`. The complete tool chain for managing batch-mode Isabelle sessions is illustrated in figure 4.1.

<code>isabelle mkroot</code>	invoked once by the user to initialize the session <code>ROOT</code> with optional <code>document</code> directory;
<code>isabelle build</code>	invoked repeatedly by the user to keep session output up-to-date (HTML, documents etc.);
<code>isabelle-process</code>	run through <code>isabelle build</code> ;
<code>isabelle document</code>	run by the Isabelle process if document preparation is enabled;
<code>isabelle latex</code>	universal \LaTeX tool wrapper invoked multiple times by <code>isabelle document</code> ; also useful for manual experiments;

Figure 4.1: The tool chain of Isabelle session presentation

4.1 Generating theory browser information

As a side-effect of building sessions, Isabelle is able to generate theory browsing information, including HTML documents that show the theory sources

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 the sessions.

Isabelle also generates graph files that represent the theory dependencies within a session. 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 §2.4 for further information.

The easiest way to let Isabelle generate theory browsing information for existing sessions is to invoke `isabelle build` with suitable options:

```
isabelle build -o browser_info -v -c FOL
```

The presentation output will appear in `$ISABELLE_BROWSER_INFO/FOL` as reported by the above verbose invocation of the build process.

Many Isabelle sessions (such as `HOL-Library` in `~/src/HOL/Library`) also provide actual printable documents. These are prepared automatically as well if enabled like this:

```
isabelle build -o browser_info -o document=pdf -v -c HOL-Library
```

Enabling both browser info and document preparation simultaneously causes an appropriate “document” link to be included in the HTML index. Documents may be generated independently of browser information as well, see §4.3 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). In order to present Isabelle applications on the web, the corresponding subdirectory from `ISABELLE_BROWSER_INFO` can be put on a WWW server.

4.2 Preparing session root directories

The `isabelle mkroot` tool configures a given directory as session root, with some `ROOT` file and optional document source directory. Its usage is:

```
Usage: isabelle mkroot [OPTIONS] [DIR]
```

Options are:

```
-d          enable document preparation
-n NAME     alternative session name (default: DIR base name)
```

```
Prepare session root DIR (default: current directory).
```

The results are placed in the given directory *dir*, which refers to the current directory by default. The `isabelle mkroot` tool is conservative in the sense that it does not overwrite existing files or directories. Earlier attempts to generate a session root need to be deleted manually.

Option `-d` indicates that the session shall be accompanied by a formal document, with `DIR/document/root.tex` as its \LaTeX entry point (see also chapter 4).

Option `-n` allows to specify an alternative session name; otherwise the base name of the given directory is used.

The implicit Isabelle settings variable `ISABELLE_LOGIC` specifies the parent session, and `ISABELLE_DOCUMENT_FORMAT` the document format to be filled into the generated `ROOT` file.

Examples

Produce session `Test` (with document preparation) within a separate directory of the same name:

```
isabelle mkroot -d Test && isabelle build -D Test
```

Upgrade the current directory into a session `ROOT` with document preparation, and build it:

```
isabelle mkroot -d && isabelle build -D .
```

4.3 Preparing Isabelle session documents

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

Usage: `isabelle 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 Isabelle build process, provided document preparation has been enabled via suitable options. 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 `isabelle document` 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 `\isakepttag`, `\isadroptag`, and `\isafoldtag`, in `~/lib/texinputs/isabelle.sty`.

Document preparation requires a `document` directory within the 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, `isabelle document` 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 `build` script for the document sources may be given. It is invoked with command-line arguments for the document format and the document variant name. The script needs to produce corresponding output files, e.g. `root.pdf` for target format `pdf` (and default variants). The main work can be again delegated to `isabelle latex`, but it is also possible to harvest generated \LaTeX sources and copy them elsewhere.

When running the session, Isabelle copies the content of the original `document` directory into its proper place within `ISABELLE_BROWSER_INFO`, according to the session path and document variant. 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 `~/lib/texinputs/isabelle.sty`. Doing `\usepackage{isabelle}` in `root.tex` should be fine; the underlying `isabelle latex` already includes an appropriate path specification for \TeX inputs.

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 [2] 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 `~/lib/texinputs/isabellesym.sty` of the Isabelle distribution.

For proper setup of DVI and PDF documents (with hyperlinks and bookmarks), we recommend to include `~/lib/texinputs/pdfsetup.sty` as well.

As a final step of Isabelle document preparation, `isabelle document -c` is run on the resulting copy of the `document` directory. Thus the actual output document is built and installed in its proper place. The generated sources are deleted after successful run of \LaTeX and friends.

Some care is needed if the document output location is configured differently, say within a directory whose content is still required afterwards!

4.4 Running \LaTeX within the Isabelle environment

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

```
Usage: isabelle 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.).

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.

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

Examples

Invoking `isabelle 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/IsabelleXXXX/browser_info/HOL/Test/document  
isabelle latex -o pdf
```

Isabelle/Scala development tools

Isabelle/ML and Isabelle/Scala are the two main language environments for Isabelle tool implementations. There are some basic command-line tools to work with the underlying Java Virtual Machine, the Scala toplevel and compiler. Note that Isabelle/jEdit (§2.3) provides a Scala Console for interactive experimentation within the running application.

5.1 Java Runtime Environment within Isabelle

The `isabelle java` tool is a direct wrapper for the Java Runtime Environment, within the regular Isabelle settings environment (§1.1). The command line arguments are that of the underlying Java version. It is run in `-server` mode if possible, to improve performance (at the cost of extra startup time). The `java` executable is the one within `ISABELLE_JDK_HOME`, according to the standard directory layout for official JDK distributions. The class loader is augmented such that the name space of `Isabelle/Pure.jar` is available, which is the main Isabelle/Scala module.

For example, the following command-line invokes the main method of class `isabelle.GUI_Setup`, which opens a windows with some diagnostic information about the Isabelle environment:

```
isabelle java isabelle.GUI_Setup
```

5.2 Scala toplevel

The `isabelle scala` tool is a direct wrapper for the Scala toplevel; see also `isabelle java` above. The command line arguments are that of the underlying Scala version.

This allows to interact with Isabelle/Scala in TTY mode like this:

```
isabelle scala
scala> isabelle.Isabelle_System.getenv("ISABELLE_HOME")
scala> val options = isabelle.Options.init()
scala> options.bool("browser_info")
```

5.3 Scala compiler

The `isabelle scalac` tool is a direct wrapper for the Scala compiler; see also `isabelle scala` above. The command line arguments are that of the underlying Scala version.

This allows to compile further Scala modules, depending on existing Isabelle/Scala functionality. The resulting class or jar files can be added to the `CLASSPATH` via the `classpath` Bash function that is provided by the Isabelle process environment. Thus add-on components can register themselves in a modular manner, see also §1.1.3.

Note that `jEdit` (§2.1) has its own mechanisms for adding plugin components, which needs special attention since it overrides the standard Java class loader.

Miscellaneous tools

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

6.1 Resolving Isabelle components

The `isabelle components` tool resolves Isabelle components:

```
Usage: isabelle components [OPTIONS] [COMPONENTS ...]
```

Options are:

```
-I          init user settings
-R URL      component repository
             (default $ISABELLE_COMPONENT_REPOSITORY)
-a          all missing components
-l          list status
```

```
Resolve Isabelle components via download and installation.
COMPONENTS are identified via base name.
```

```
ISABELLE_COMPONENT_REPOSITORY="http://isabelle.in.tum.de/components"
```

Components are initialized as described in §1.1.3 in a permissive manner, which can mark components as “missing”. This state is amended by letting `isabelle components` download and unpack components that are published on the default component repository `http://isabelle.in.tum.de/components/` in particular.

Option `-R` specifies an alternative component repository. Note that `file:///` URLs can be used for local directories.

Option `-a` selects all missing components to be installed. Explicit components may be named as command line-arguments as well. Note that components are uniquely identified by their base name, while the installation takes place in the location that was specified in the attempt to initialize the component before.

Option `-l` lists the current state of available and missing components with their location (full name) within the file-system.

Option `-I` initializes the user settings file to subscribe to the standard components specified in the Isabelle repository clone — this does not make any sense for regular Isabelle releases. If the file already exists, it needs to be edited manually according to the printed explanation.

6.2 Displaying documents

The `isabelle display` tool displays documents in DVI or PDF format:

```
Usage: isabelle 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.

6.3 Viewing documentation

The `isabelle doc` tool displays online documentation:

```
Usage: isabelle 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.

6.4 Shell commands within the settings environment

The `isabelle env` tool is a direct wrapper for the standard `/usr/bin/env` command on POSIX systems, running within the Isabelle settings environment (§1.1).

The command-line arguments are that of the underlying version of `env`. For example, the following invokes an instance of the GNU Bash shell within the Isabelle environment:

```
isabelle env bash
```

6.5 Getting logic images

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

```
Usage: isabelle 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.

6.6 Inspecting the settings environment

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

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

Options are:

```
-a          display complete environment
-b          print values only (doesn't work for -a)
-d FILE    dump complete environment to FILE
           (null terminated entries)
```

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.

Option `-d` produces a dump of the complete environment to the specified file. Entries are terminated by the ASCII null character, i.e. the C string terminator.

Examples

Get the location of `ISABELLE_HOME_USER` where user-specific information is stored:

```
isabelle getenv ISABELLE_HOME_USER
```

Get the value only of the same settings variable, which is particularly useful in shell scripts:

```
isabelle getenv -b ISABELLE_OUTPUT
```

6.7 Installing standalone Isabelle executables

By default, the main Isabelle binaries (`isabelle` 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 `isabelle install` tool:

```
Usage: isabelle install [OPTIONS] BINDIR
```

Options are:

```
-d DISTDIR    refer to DISTDIR as Isabelle distribution
              (default ISABELLE_HOME)
```

Install Isabelle executables with absolute references to the distribution directory.

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

The `BINDIR` argument tells where executable wrapper scripts for `isabelle-process` and `isabelle` should be placed, which is typically a directory in the shell's `PATH`, such as `$HOME/bin`.

It is also possible to make symbolic links of the main Isabelle executables manually, but making separate copies outside the Isabelle distribution directory will not work!

6.8 Creating instances of the Isabelle logo

The `isabelle logo` tool creates instances of the generic Isabelle logo as EPS and PDF, for inclusion in \LaTeX documents.

```
Usage: isabelle logo [OPTIONS] XYZ
```

Create instance XYZ of the Isabelle logo (as EPS and PDF).

Options are:

```
-n NAME      alternative output base name (default "isabelle_xyz")
-q          quiet mode
```

Option `-n` specifies an alternative (base) name for the generated files. The default is `isabelle_xyz` in lower-case.

Option `-q` omits printing of the result file name.

Implementors of Isabelle tools and applications are encouraged to make derived Isabelle logos for their own projects using this template.

6.9 Isabelle wrapper for make

The old `isabelle make` tool is a very simple wrapper for ordinary Unix `make`:

Usage: `isabelle 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_HOME)`. Furthermore, programs started from the make file also inherit this environment.

6.10 Creating Isabelle session directories

The old `isabelle mkdir` tool 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 §4.3 for further information on Isabelle document preparation. The usage of `isabelle mkdir` is:

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

Options are:

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

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

The `isabelle 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 `isabelle mkdir` is unable to change `IsaMakefile` incrementally — manual changes are required for multiple sub-sessions. On 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 `isabelle mkdir`’s idea of `isabelle 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 §6.13 for further information on *build mode* vs. *example mode* of `isabelle usedir`.

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

6.11 Printing documents

The `isabelle print` tool prints documents:

```
Usage: isabelle 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.

6.12 Remove awkward symbol names from theory sources

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

```
Usage: isabelle 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 `"~"`.

6.13 Running Isabelle sessions

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

```
Usage: isabelle 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
-Q INT       set threshold for sub-proof parallelization (default 50)
-T LEVEL     multithreading: trace level (default 0)
-V VARIANT   declare alternative document VARIANT
-b           build mode (output heap image, using current dir)
-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 (default 0)
-q LEVEL     set level of parallel proof checking (default 1)
-r           reset session path
-s NAME      override session NAME
-t BOOL      internal session timing (default false)
-v BOOL      be verbose (default false)
```

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

```
ISABELLE_USEDIR_OPTIONS=...
```

```
ML_PLATFORM=...
```

```
ML_HOME=...
```

```
ML_SYSTEM=...
```

```
ML_OPTIONS=...
```

Note that the value of the `ISABELLE_USEDIR_OPTIONS` setting is implicitly prefixed to *any* `isabelle usedir` call. Since the `IsaMakefiles` of all object-logics distributed with Isabelle just invoke `isabelle 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.

Options

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

Calling `isabelle usedir` with `-b` runs `isabelle-process` 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, `isabelle 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 §4.3 and §4.4 for more details.

The `-V` option declares alternative document variants, consisting of name/tags pairs (cf. options `-n` and `-t` of `isabelle document`). 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 variant “`outline=/proof/ML`” would fold proof and ML parts, replacing the original text by a short place-holder. The form “`name=-,`” means to remove document *name* from the list of variants to be processed. Any number of `-V` options may be given; later declarations have precedence over earlier ones.

Some document variant *name* may use an alternative L^AT_EX entry point called `document/root_name.tex` if that file exists; otherwise the common `document/root.tex` is used.

The `-g` option produces images of the theory dependency graph (cf. §2.4)

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 L^AT_EX document one could say `\includegraphics{session_graph}` in `document/root.tex` (omitting the file-name extension enables L^AT_EX 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, `isabelle usedir -D generated HOL Foo` produces a complete set of document sources at `Foo/generated`. Subsequent invocation of `isabelle document Foo/generated` (see also §4.3) will process the final result independently of an Isabelle job. This decoupled mode of operation facilitates debugging of serious L^AT_EX errors, for example.

The `-p` option determines the level of detail for internal proof objects, see also the *Isabelle Reference Manual* [1].

The `-q` option specifies the level of parallel proof checking: 0 no proofs, 1 toplevel proofs (default), 2 toplevel and nested Isar proofs. The option `-Q` specifies a threshold for `-q2`: nested proofs are only parallelized when the current number of forked proofs falls below the given value (default 50), multiplied by the number of worker threads (see option `-M`).

The `-t` option produces a more detailed internal timing report of the session.

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 worker threads used for processing independent tasks when checking theory sources (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 `isabelle 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.

6.14 Output the version identifier of the Isabelle distribution

The `isabelle version` tool displays Isabelle version information:

```
Usage: isabelle version [OPTIONS]

Options are:
  -i          short identification (derived from Mercurial id)

Display Isabelle version information.
```

The default is to output the full version string of the Isabelle distribution, e.g. "Isabelle2012: May 2012."

The `-i` option produces a short identification derived from the Mercurial id of the `ISABELLE_HOME` directory.

6.15 Convert XML to YXML

The `isabelle yxml` tool 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 special treatment of white space, 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>XYnameYattribute=value...X</code>
<code></name></code>	<code>XYX</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. For example, see `~/src/Pure/PIDE/yxml.ML` or `~/src/Pure/PIDE/yxml.scala`. YXML documents may be detected quickly by checking that the first two characters are **XY**.

Bibliography

- [1] L. C. Paulson. *The Old Isabelle Reference Manual*.
<http://isabelle.in.tum.de/doc/ref.pdf>.
- [2] M. Wenzel. *The Isabelle/Isar Implementation*.
<http://isabelle.in.tum.de/doc/implementation.pdf>.
- [3] M. Wenzel. *The Isabelle/Isar Reference Manual*.
<http://isabelle.in.tum.de/doc/isar-ref.pdf>.

Index

- bash (executable), 2, **3**
- browser (tool), **13**
- browser_info (system option), **20**
- build (tool), **22**, 27
- build_dialog (tool), **25**

- components (tool), **35**
- condition (system option), **21**

- display (tool), **36**
- doc (tool), **36**
- document (system option), **20**
- document (tool), 27, **29**, 43
- document_graph (system option), **21**
- document_output (system option), **20**
- document_variants (system option), **21**
- DVI_VIEWER (setting), **6**

- emacs (tool), **12**
- env (tool), **37**

- findlogics (tool), **37**

- getenv (tool), **37**

- HTML, 42

- install (tool), **38**
- isabelle (executable), 1, 2
- isabelle-process (executable), 1, **7**, 27
- ISABELLE_BIBTEX (setting), **5**
- ISABELLE_BROWSER_INFO (setting), **5**, 28
- ISABELLE_BUILD_OPTIONS (setting), 24
- ISABELLE_DOC_FORMAT (setting), **5**
- ISABELLE_DOCS (setting), **5**
- ISABELLE_DVIPS (setting), **5**
- ISABELLE_HOME (setting), **2**, **3**
- ISABELLE_HOME_USER (setting), **4**
- ISABELLE_IDENTIFIER (setting), **4**
- ISABELLE_JDK_HOME (setting), **5**
- ISABELLE_LATEX (setting), **5**
- ISABELLE_LINE_EDITOR (setting), **5**
- ISABELLE_LOGIC (setting), **5**
- ISABELLE_OUTPUT (setting), 3, **5**
- ISABELLE_PATH (setting), **5**
- ISABELLE_PDFLATEX (setting), **5**
- ISABELLE_PLATFORM (setting), **4**
- ISABELLE_PLATFORM64 (setting), **4**
- ISABELLE_PLATFORM_FAMILY (setting), **4**
- ISABELLE_PROCESS (setting), **3**, **4**
- ISABELLE_TMP_PREFIX (setting), **6**
- ISABELLE_TOOL (setting), **3**
- ISABELLE_TOOLS (setting), **5**, **6**
- ISABELLE_USEDIR_OPTIONS (setting), **5**, 42

- java (tool), **33**
- jedit (tool), **11**

- latex (tool), 27, **31**

- logo (tool), **39**

- make (tool), **39**
- mkdir (tool), **40**
- mkroot (tool), 27, **28**
- ML_HOME (setting), **4**
- ML_IDENTIFIER (setting), **4**
- ML_OPTIONS (setting), **4**
- ML_PLATFORM (setting), **4**
- ML_SYSTEM (setting), **4**

- options (tool), **22**

- PDF_VIEWER (setting), **6**
- print (tool), **41**
- PRINT_COMMAND (setting), **6**
- PROOFGENERAL_HOME (setting), **12**
- PROOFGENERAL_OPTIONS (setting), **12**

- rlwrap (executable), **13**

- scala (tool), **33**
- scalac (tool), **34**
- session_entry (syntax), **18**
- settings, **1**

- theory browsing information, **27**
- threads (system option), **21**, 24
- timeout (system option), **21**
- tty (tool), 5, **12**

- unsymbolize (tool), **41**
- usedir (tool), 5, **42**
- USER_HOME (setting), **3**

- version (tool), **45**

- yxml (tool), **45**