



Isabelle/jEdit

Makarius Wenzel

27 August 2014

Abstract

Isabelle/jEdit is a fully-featured Prover IDE, based on Isabelle/Scala and the jEdit text editor. This document provides an overview of general principles and its main IDE functionality.

Isabelle's user interface is no advance over LCF's, which is widely condemned as "user-unfriendly": hard to use, bewildering to beginners. Hence the interest in proof editors, where a proof can be constructed and modified rule-by-rule using windows, mouse, and menus. But Edinburgh LCF was invented because real proofs require millions of inferences. Sophisticated tools — rules, tactics and tacticals, the language ML, the logics themselves — are hard to learn, yet they are essential. We may demand a mouse, but we need better education and training.

Lawrence C. Paulson, "Isabelle: The Next 700 Theorem Provers"

Acknowledgements

Research and implementation of concepts around PIDE and Isabelle/jEdit has started in 2008 and was kindly supported by:

- TU München <http://www.in.tum.de>
- BMBF <http://www.bmbf.de>
- Université Paris-Sud <http://www.u-psud.fr>
- Digiteo <http://www.digiteo.fr>
- ANR <http://www.agence-nationale-recherche.fr>

Contents

1	Introduction	1
1.1	Concepts and terminology	1
1.2	The Isabelle/jEdit Prover IDE	2
1.2.1	Documentation	3
1.2.2	Plugins	4
1.2.3	Options	4
1.2.4	Keymaps	5
1.3	Command-line invocation	5
2	Augmented jEdit functionality	7
2.1	Look-and-feel	7
2.2	Dockable windows	8
2.3	Isabelle symbols	9
2.4	SideKick parsers	12
2.5	Scala console	12
2.6	File-system access	13
3	Prover IDE functionality	15
3.1	Document model	15
3.1.1	Editor buffers and document nodes	15
3.1.2	Theories	16
3.1.3	Auxiliary files	17
3.2	Output	18
3.3	Query	20
3.3.1	Find theorems	22
3.3.2	Find constants	22
3.3.3	Print context	23
3.4	Tooltips and hyperlinks	23
3.5	Completion	24

3.5.1	Varieties of completion	25
3.5.2	Semantic completion context	28
3.5.3	Input events	28
3.5.4	Completion popup	29
3.5.5	Insertion	30
3.5.6	Options	31
3.6	Automatically tried tools	32
3.7	Sledgehammer	34
4	Miscellaneous tools	36
4.1	Timing	36
4.2	Low-level output	37
5	Known problems and workarounds	38
	Bibliography	40
	Index	42

List of Figures

1.1	The Isabelle/jEdit Prover IDE	2
2.1	The Isabelle NEWS file with SideKick tree view	12
3.1	Theories panel with an overview of the document-model, and some jEdit text areas as editable view on some of the document nodes	16
3.2	Multiple views on prover output: gutter area with icon, text area with popup, overview area, Theories panel, Output panel	19
3.3	An instance of the Query panel	21
3.4	Tooltip and hyperlink for some formal entity	23
3.5	Nested tooltips over formal entities	24
3.6	Result of automatically tried tools	33
3.7	An instance of the Sledgehammer panel	35

Introduction

1.1 Concepts and terminology

Isabelle/jEdit is a Prover IDE that integrates *parallel proof checking* [5, 10] with *asynchronous user interaction* [6, 9, 11], based on a document-oriented approach to *continuous proof processing* [7, 8]. Many concepts and system components are fit together in order to make this work. The main building blocks are as follows.

PIDE is a general framework for Prover IDEs based on Isabelle/Scala. It is built around a concept of parallel and asynchronous document processing, which is supported natively by the parallel proof engine that is implemented in Isabelle/ML. The traditional prover command loop is given up; instead there is direct support for editing of source text, with rich formal markup for GUI rendering.

Isabelle/ML is the implementation and extension language of Isabelle, see also [3]. It is integrated into the logical context of Isabelle/Isar and allows to manipulate logical entities directly. Arbitrary add-on tools may be implemented for object-logics such as Isabelle/HOL.

Isabelle/Scala is the system programming language of Isabelle. It extends the pure logical environment of Isabelle/ML towards the “real world” of graphical user interfaces, text editors, IDE frameworks, web services etc. Special infrastructure allows to transfer algebraic datatypes and formatted text easily between ML and Scala, using asynchronous protocol commands.

jEdit is a sophisticated text editor implemented in Java.¹ It is easily extensible by plugins written in languages that work on the JVM, e.g. Scala².

¹<http://www.jedit.org>

²<http://www.scala-lang.org>

Isabelle/jEdit is the main example application of the PIDE framework and the default user-interface for Isabelle. It targets both beginners and experts. Technically, Isabelle/jEdit combines a slightly modified version of the jEdit code base with a special plugin for Isabelle, integrated as standalone application for the main operating system platforms: Linux, Windows, Mac OS X.

The subtle differences of Isabelle/ML versus Standard ML, Isabelle/Scala versus Scala, Isabelle/jEdit versus jEdit need to be taken into account when discussing any of these PIDE building blocks in public forums, mailing lists, or even scientific publications.

1.2 The Isabelle/jEdit Prover IDE

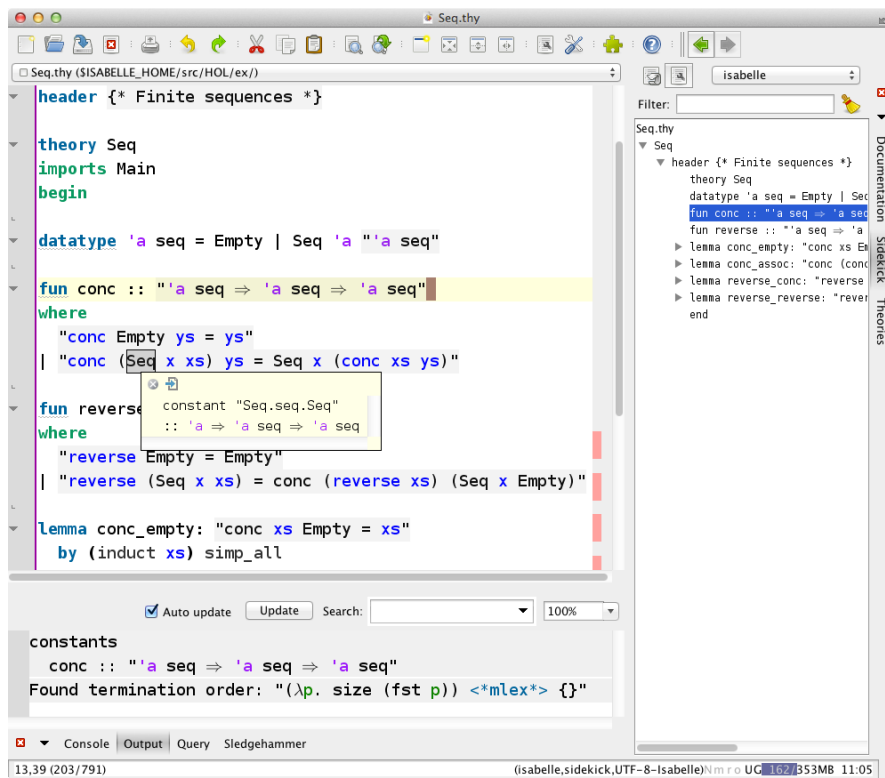


Figure 1.1: The Isabelle/jEdit Prover IDE

Isabelle/jEdit (figure 1.1) consists of some plugins for the jEdit text editor, while preserving its general look-and-feel as far as possible. The main plugin

is called “Isabelle” and has its own menu *Plugins / Isabelle* with access to several panels (see also §2.2), as well as *Plugins / Plugin Options / Isabelle* (see also §1.2.3).

The options allow to specify a logic session name — the same selector is accessible in the *Theories* panel (§3.1.2). On application startup, the selected logic session image is provided automatically by the Isabelle build tool [12]: if it is absent or outdated wrt. its sources, the build process updates it before entering the Prover IDE. Changing the logic session within Isabelle/jEdit requires a restart of the whole application.

The main job of the Prover IDE is to manage sources and their changes, taking the logical structure as a formal document into account (see also §3.1). The editor and the prover are connected asynchronously in a lock-free manner. The prover is free to organize the checking of the formal text in parallel on multiple cores, and provides feedback via markup, which is rendered in the editor via colors, boxes, squiggly underlines, hyperlinks, popup windows, icons, clickable output etc.

Using the mouse together with the modifier key **CONTROL** (Linux, Windows) or **COMMAND** (Mac OS X) exposes additional formal content via tooltips and/or hyperlinks (see also §3.4). Output (in popups etc.) may be explored recursively, using the same techniques as in the editor source buffer.

Thus the Prover IDE gives an impression of direct access to formal content of the prover within the editor, but in reality only certain aspects are exposed, according to the possibilities of the prover and its many tools.

1.2.1 Documentation

The *Documentation* panel of Isabelle/jEdit provides access to the standard Isabelle documentation: PDF files are opened by regular desktop operations of the underlying platform. The section “Original jEdit Documentation” contains the original *User’s Guide* of this sophisticated text editor. The same is accessible via the **Help** menu or **F1** keyboard shortcut, using the built-in HTML viewer of Java/Swing. The latter also includes *Frequently Asked Questions* and documentation of individual plugins.

Most of the information about generic jEdit is relevant for Isabelle/jEdit as well, but one needs to keep in mind that defaults sometimes differ, and the official jEdit documentation does not know about the Isabelle plugin with its support for continuous checking of formal source text: jEdit is a plain text editor, but Isabelle/jEdit is a Prover IDE.

1.2.2 Plugins

The *Plugin Manager* of jEdit allows to augment editor functionality by JVM modules (jars) that are provided by the central plugin repository, which is accessible via various mirror sites.

Connecting to the plugin server-infrastructure of the jEdit project allows to update bundled plugins or to add further functionality. This needs to be done with the usual care for such an open bazaar of contributions. Arbitrary combinations of add-on features are apt to cause problems. It is advisable to start with the default configuration of Isabelle/jEdit and develop some understanding how it is supposed to work, before loading additional plugins at a grand scale.

The main *Isabelle* plugin is an integral part of Isabelle/jEdit and needs to remain active at all times! A few additional plugins are bundled with Isabelle/jEdit for convenience or out of necessity, notably *Console* with its Isabelle/Scala sub-plugin (§2.5) and *SideKick* with some Isabelle-specific parsers for document tree structure (§2.4). The *Navigator* plugin is particularly important for hyperlinks within the formal document-model (§3.4). Further plugins (e.g. *ErrorList*, *Code2HTML*) are included to saturate the dependencies of bundled plugins, but have no particular use in Isabelle/jEdit.

1.2.3 Options

Both jEdit and Isabelle have distinctive management of persistent options.

Regular jEdit options are accessible via the dialogs *Utilities / Global Options* or *Plugins / Plugin Options*, with a second chance to flip the two within the central options dialog. Changes are stored in `$ISABELLE_HOME_USER/jedit/properties` and `$ISABELLE_HOME_USER/jedit/keymaps`.

Isabelle system options are managed by Isabelle/Scala and changes are stored in `$ISABELLE_HOME_USER/etc/preferences`, independently of other jEdit properties. See also [12], especially the coverage of sessions and command-line tools like `isabelle build` or `isabelle options`.

Those Isabelle options that are declared as **public** are configurable in Isabelle/jEdit via *Plugin Options / Isabelle / General*. Moreover, there are various options for rendering of document content, which are configurable via *Plugin Options / Isabelle / Rendering*. Thus *Plugin Options / Isabelle* in jEdit provides a view on a subset of Isabelle system options. Note that some of these options affect general parameters that are relevant outside Isabelle/jEdit as well, e.g. `threads` or `parallel_proofs` for

the Isabelle build tool [12], but it is possible to use the settings variable `ISABELLE_BUILD_OPTIONS` to change defaults for batch builds without affecting Isabelle/jEdit.

The jEdit action `isabelle.options` opens the options dialog for the Isabelle plugin; it can be mapped to editor GUI elements as usual.

Options are usually loaded on startup and saved on shutdown of Isabelle/jEdit. Editing the machine-generated `$ISABELLE_HOME_USER/jedit/properties` or `$ISABELLE_HOME_USER/etc/preferences` manually while the application is running is likely to cause surprise due to lost update!

1.2.4 Keymaps

Keyboard shortcuts used to be managed as jEdit properties in the past, but recent versions (2013) have a separate concept of *keymap* that is configurable via *Global Options / Shortcuts*. The `imported` keymap is derived from the initial environment of properties that is available at the first start of the editor; afterwards the keymap file takes precedence.

This is relevant for Isabelle/jEdit due to various fine-tuning of default properties, and additional keyboard shortcuts for Isabelle-specific functionality. Users may change their keymap later, but need to copy some keyboard shortcuts manually (see also `$ISABELLE_HOME_USER/jedit/keymaps` versus `shortcut` properties in `$ISABELLE_HOME/src/Tools/jEdit/src/jEdit.props`).

1.3 Command-line invocation

Isabelle/jEdit is normally invoked as standalone application, with platform-specific executable wrappers for Linux, Windows, Mac OS X. Nonetheless it is occasionally useful to invoke the Prover IDE on the command-line, with some extra options and environment settings as explained below. The command-line usage of `isabelle jedit` is as follows:

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

Options are:

<code>-J OPTION</code>	add JVM runtime option (default <code>JEDIT_JAVA_OPTIONS</code>)
<code>-b</code>	build only
<code>-d DIR</code>	include session directory
<code>-f</code>	fresh build
<code>-j OPTION</code>	add jEdit runtime option (default <code>JEDIT_OPTIONS</code>)
<code>-l NAME</code>	logic image name (default <code>ISABELLE_LOGIC</code>)
<code>-m MODE</code>	add print mode for output
<code>-n</code>	no build of session image on startup
<code>-s</code>	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 of `isabelle build` [12].

By default, the specified image is checked and built on demand. The `-s` option determines where to store the result session image of `isabelle build`. The `-n` option bypasses the implicit build process for the selected session image.

The `-m` option specifies additional print modes for the prover process. Note that the system option `jedit_print_mode` allows to do the same persistently (e.g. via the *Plugin Options* dialog of Isabelle/jEdit), without requiring command-line invocation.

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 [12], but note that these only work for the command-line tool described here, and not the regular application.

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 auxiliary `jedit_build` component from <http://isabelle.in.tum.de/components>. The official Isabelle release already includes a pre-built version of Isabelle/jEdit.

Augmented jEdit functionality

2.1 Look-and-feel

jEdit is a Java/AWT/Swing application with some ambition to support “native” look-and-feel on all platforms, within the limits of what Oracle as Java provider and major operating system distributors allow (see also §5).

Isabelle/jEdit enables platform-specific look-and-feel by default as follows.

Linux: The platform-independent *Nimbus* is used by default.

GTK+ works under the side-condition that the overall GTK theme is selected in a Swing-friendly way.¹

Windows: Regular *Windows* is used by default, but *Windows Classic* also works.

Mac OS X: Regular *Mac OS X* is used by default.

The bundled *MacOSX* plugin provides various functions that are expected from applications on that particular platform: quit from menu or dock, preferences menu, drag-and-drop of text files on the application, full-screen mode for main editor windows. It is advisable to have the *MacOSX* plugin enabled all the time on that platform.

Users may experiment with different look-and-feels, but need to keep in mind that this extra variance of GUI functionality is unlikely to work in arbitrary combinations. The platform-independent *Nimbus* and *Metal* should always work. The historic *CDE/Motif* is better avoided.

After changing the look-and-feel in *Global Options / Appearance*, it is advisable to restart Isabelle/jEdit in order to take full effect.

¹GTK support in Java/Swing was once marketed aggressively by Sun, but never quite finished. Today (2013) it is lagging behind further development of Swing and GTK. The graphics rendering performance can be worse than for other Swing look-and-feels.

2.2 Dockable windows

In jEdit terminology, a *view* is an editor window with one or more *text areas* that show the content of one or more *buffers*. A regular view may be surrounded by *dockable windows* that show additional information in arbitrary format, not just text; a *plain view* does not allow dockables. The *dockable window manager* of jEdit organizes these dockable windows, either as *floating* windows, or *docked* panels within one of the four margins of the view. There may be any number of floating instances of some dockable window, but at most one docked instance; jEdit actions that address *the* dockable window of a particular kind refer to the unique docked instance.

Dockables are used routinely in jEdit for important functionality like *Hyper-Search Results* or the *File System Browser*. Plugins often provide a central dockable to access their key functionality, which may be opened by the user on demand. The Isabelle/jEdit plugin takes this approach to the extreme: its plugin menu merely provides entry-points to panels that are managed as dockable windows. Some important panels are docked by default, e.g. *Documentation*, *Output*, *Query*, but the user can change this arrangement easily.

Compared to plain jEdit, dockable window management in Isabelle/jEdit is slightly augmented according to the the following principles:

- Floating windows are dependent on the main window as *dialog* in the sense of Java/AWT/Swing. Dialog windows always stay on top of the view, which is particularly important in full-screen mode. The desktop environment of the underlying platform may impose further policies on such dependent dialogs, in contrast to fully independent windows, e.g. some window management functions may be missing.
- Keyboard focus of the main view vs. a dockable window is carefully managed according to the intended semantics, as a panel mainly for output or input. For example, activating the *Output* (§3.2) panel via the dockable window manager returns keyboard focus to the main text area, but for *Query* (§3.3) the focus is given to the main input field of that panel.
- Panels that provide their own text area for output have an additional dockable menu item *Detach*. This produces an independent copy of the current output as a floating *Info* window, which displays that content independently of ongoing changes of the PIDE document-model. Note that Isabelle/jEdit popup windows (§3.4) provide a similar *Detach* operation as an icon.

2.3 Isabelle symbols

Isabelle sources consist of *symbols* that extend plain ASCII to allow infinitely many mathematical symbols within the formal sources. This works without depending on particular encodings and varying Unicode standards.² See also [7].

For the prover back-end, formal text consists of ASCII characters that are grouped according to some simple rules, e.g. as plain “a” or symbolic “\<alpha>”. For the editor front-end, a certain subset of symbols is rendered physically via Unicode glyphs, in order to show “\<alpha>” as “α”, for example. This symbol interpretation is specified by the Isabelle system distribution in \$ISABELLE_HOME/etc/symbols and may be augmented by the user in \$ISABELLE_HOME_USER/etc/symbols.

The appendix of [4] gives an overview of the standard interpretation of finitely many symbols from the infinite collection. Uninterpreted symbols are displayed literally, e.g. “\<foobar>”. Overlap of Unicode characters used in symbol interpretation with informal ones (which might appear e.g. in comments) needs to be avoided. Raw Unicode characters within prover source files should be restricted to informal parts, e.g. to write text in non-latin alphabets in comments.

Encoding. Technically, the Unicode view on Isabelle symbols is an *encoding* in jEdit (not in the underlying JVM) that is called **UTF-8-Isabelle**. It is provided by the Isabelle/jEdit plugin and enabled by default for all source files. Sometimes such defaults are reset accidentally, or malformed UTF-8 sequences in the text force jEdit to fall back on a different encoding like ISO-8859-15. In that case, verbatim “\<alpha>” will be shown in the text buffer instead of its Unicode rendering “α”. The jEdit menu operation *File / Reload with Encoding / UTF-8-Isabelle* helps to resolve such problems (after repairing malformed parts of the text).

Font. Correct rendering via Unicode requires a font that contains glyphs for the corresponding codepoints. Most system fonts lack that, so Isabelle/jEdit prefers its own application font **IsabelleText**, which ensures that standard collection of Isabelle symbols are actually seen on the screen (or printer).

²Raw Unicode characters within formal sources would compromise portability and reliability in the face of changing interpretation of special features of Unicode, such as Combining Characters or Bi-directional Text.

Note that a Java/AWT/Swing application can load additional fonts only if they are not installed on the operating system already! Some outdated version of `IsabelleText` that happens to be provided by the operating system would prevent Isabelle/jEdit to use its bundled version. This could lead to missing glyphs (black rectangles), when the system version of `IsabelleText` is older than the application version. This problem can be avoided by refraining to “install” any version of `IsabelleText` in the first place, although it is occasionally tempting to use the same font in other applications.

Input methods. In principle, Isabelle/jEdit could delegate the problem to produce Isabelle symbols in their Unicode rendering to the underlying operating system and its *input methods*. Regular jEdit also provides various ways to work with *abbreviations* to produce certain non-ASCII characters. Since none of these standard input methods work satisfactorily for the mathematical characters required for Isabelle, various specific Isabelle/jEdit mechanisms are provided.

This is a summary for practically relevant input methods for Isabelle symbols.

1. The *Symbols* panel: some GUI buttons allow to insert certain symbols in the text buffer. There are also tooltips to reveal the official Isabelle representation with some additional information about *symbol abbreviations* (see below).
2. Copy/paste from decoded source files: text that is rendered as Unicode already can be re-used to produce further text. This also works between different applications, e.g. Isabelle/jEdit and some web browser or mail client, as long as the same Unicode view on Isabelle symbols is used.
3. Copy/paste from prover output within Isabelle/jEdit. The same principles as for text buffers apply, but note that *copy* in secondary Isabelle/jEdit windows works via the keyboard shortcut `C+c`, while jEdit menu actions always refer to the primary text area!
4. Completion provided by Isabelle plugin (see §3.5). Isabelle symbols have a canonical name and optional abbreviations. This can be used with the text completion mechanism of Isabelle/jEdit, to replace a prefix of the actual symbol like `\<lambda>`, or its name preceded by backslash `\lambda`, or its ASCII abbreviation `%` by the Unicode rendering.

The following table is an extract of the information provided by the standard `$ISABELLE_HOME/etc/symbols` file:

symbol	name with backslash	abbreviation
λ	<code>\lambda</code>	<code>%</code>
\Rightarrow	<code>\Rightarrow</code>	<code>=></code>
\Longrightarrow	<code>\Longrightarrow</code>	<code>==></code>
\wedge	<code>\And</code>	<code>!!</code>
\equiv	<code>\equiv</code>	<code>==</code>
\forall	<code>\forall</code>	<code>!</code>
\exists	<code>\exists</code>	<code>?</code>
\longrightarrow	<code>\longrightarrow</code>	<code>--></code>
\wedge	<code>\and</code>	<code>&</code>
\vee	<code>\or</code>	<code> </code>
\neg	<code>\not</code>	<code>~</code>
\neq	<code>\noteq</code>	<code>~=</code>
\in	<code>\in</code>	<code>:</code>
\notin	<code>\notin</code>	<code>~:</code>

Note that the above abbreviations refer to the input method. The logical notation provides ASCII alternatives that often coincide, but sometimes deviate. This occasionally causes user confusion with very old-fashioned Isabelle source that use ASCII replacement notation like `!` or `ALL` directly in the text.

On the other hand, coincidence of symbol abbreviations with ASCII replacement syntax helps to update old theory sources via explicit completion (see also `C+b` explained in §3.5).

Control symbols. There are some special control symbols to modify the display style of a single symbol (without nesting). Control symbols may be applied to a region of selected text, either using the *Symbols* panel or keyboard shortcuts or jEdit actions. These editor operations produce a separate control symbol for each symbol in the text, in order to make the whole text appear in a certain style.

style	symbol	shortcut	action
superscript	<code>\<^sup></code>	C+e UP	<code>isabelle.control-sup</code>
subscript	<code>\<^sub></code>	C+e DOWN	<code>isabelle.control-sub</code>
bold face	<code>\<^bold></code>	C+e RIGHT	<code>isabelle.control-bold</code>
reset		C+e LEFT	<code>isabelle.control-reset</code>

To produce a single control symbol, it is also possible to complete on `\sup`, `\sub`, `\bold` as for regular symbols.

2.4 SideKick parsers

The *SideKick* plugin provides some general services to display buffer structure in a tree view.

Isabelle/jEdit provides SideKick parsers for its main mode for theory files, as well as some minor modes for the NEWS file (see figure 2.1), session ROOT files, and system options.

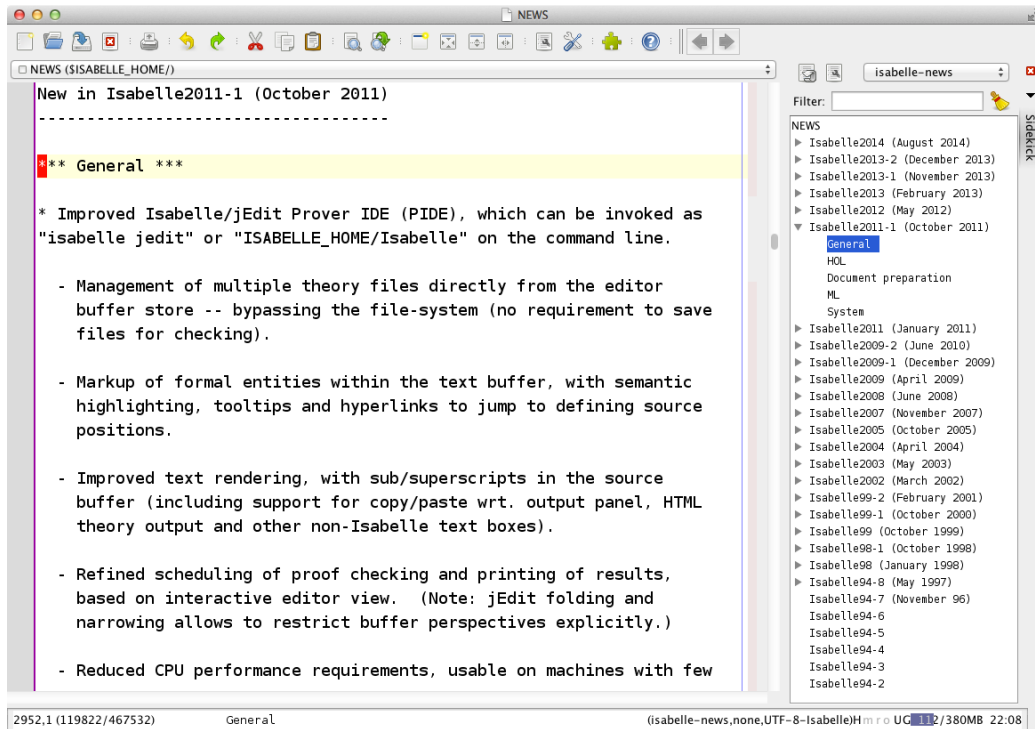


Figure 2.1: The Isabelle NEWS file with SideKick tree view

Moreover, the special SideKick parser *isabelle-markup* provides access to the full (uninterpreted) markup tree of the PIDE document model of the current buffer. This is occasionally useful for informative purposes, but the amount of displayed information might cause problems for large buffers, both for the human and the machine.

2.5 Scala console

The *Console* plugin manages various shells (command interpreters), e.g. *BeanShell*, which is the official jEdit scripting language, and the cross-

platform *System* shell. Thus the console provides similar functionality than the Emacs buffers `*scratch*` and `*shell*`.

Isabelle/jEdit extends the repertoire of the console by *Scala*, which is the regular Scala toplevel loop running inside the same JVM process as Isabelle/jEdit itself. This means the Scala command interpreter has access to the JVM name space and state of the running Prover IDE application. The default environment imports the full content of packages `isabelle` and `isabelle.jedit`.

For example, `PIDE` refers to the Isabelle/jEdit plugin object, and `view` to the current editor view of jEdit. The Scala expression `PIDE.snapshot(view)` makes a PIDE document snapshot of the current buffer within the current editor view.

This helps to explore Isabelle/Scala functionality interactively. Some care is required to avoid interference with the internals of the running application, especially in production use.

2.6 File-system access

File specifications in jEdit follow various formats and conventions according to *Virtual File Systems*, which may be also provided by additional plugins. This allows to access remote files via the `http:` protocol prefix, for example. Isabelle/jEdit attempts to work with the file-system model of jEdit as far as possible. In particular, theory sources are passed directly from the editor to the prover, without indirection via physical files.

Despite the flexibility of URLs in jEdit, local files are particularly important and are accessible without protocol prefix. Here the path notation is that of the Java Virtual Machine on the underlying platform. On Windows the preferred form uses backslashes, but happens to accept forward slashes like Unix/POSIX. Further differences arise due to Windows drive letters and network shares.

The Java notation for files needs to be distinguished from the one of Isabelle, which uses POSIX notation with forward slashes on *all* platforms.³ Moreover, environment variables from the Isabelle process may be used freely, e.g. `$ISABELLE_HOME/etc/symbols` or `$POLYML_HOME/README`. There are special shortcuts: `~` for `$USER_HOME` and `~~` for `$ISABELLE_HOME`.

Since jEdit happens to support environment variables within file specifications as well, it is natural to use similar notation within the editor, e.g.

³Isabelle/ML on Windows uses Cygwin file-system access and Unix-style path notation.

in the file-browser. This does not work in full generality, though, due to the bias of jEdit towards platform-specific notation and of Isabelle towards POSIX. Moreover, the Isabelle settings environment is not yet active when starting Isabelle/jEdit via its standard application wrapper, in contrast to `isabelle jedit` run from the command line (§1.3).

Isabelle/jEdit imitates `$ISABELLE_HOME` and `$ISABELLE_HOME_USER` within the Java process environment, in order to allow easy access to these important places from the editor. The file browser of jEdit also includes *Favorites* for these two important locations.

Path specifications in prover input or output usually include formal markup that turns it into a hyperlink (see also §3.4). This allows to open the corresponding file in the text editor, independently of the path notation.

Formally checked paths in prover input are subject to completion (§3.5): partial specifications are resolved via directory content and possible completions are offered in a popup.

Prover IDE functionality

3.1 Document model

The document model is central to the PIDE architecture: the editor and the prover have a common notion of structured source text with markup, which is produced by formal processing. The editor is responsible for edits of document source, as produced by the user. The prover is responsible for reports of document markup, as produced by its processing in the background.

Isabelle/jEdit handles classic editor events of jEdit, in order to connect the physical world of the GUI (with its singleton state) to the mathematical world of multiple document versions (with timeless and stateless updates).

3.1.1 Editor buffers and document nodes

As a regular text editor, jEdit maintains a collection of *buffers* to store text files; each buffer may be associated with any number of visible *text areas*. Buffers are subject to an *edit mode* that is determined from the file name extension. The following modes are treated specifically in Isabelle/jEdit:

mode	file extension	content
isabelle	.thy	theory source
isabelle-ml	.ML	Isabelle/ML source
sml	.sml or .sig	Standard ML source

All jEdit buffers are automatically added to the PIDE document-model as *document nodes*. The overall document structure is defined by the theory nodes in two dimensions:

1. via **theory imports** that are specified in the *theory header* using concrete syntax of the **theory** command [4];
2. via **auxiliary files** that are loaded into a theory by special *load commands*, notably **ML_file** and **SML_file** [4].

In any case, source files are managed by the PIDE infrastructure: the physical file-system only plays a subordinate role. The relevant version of source text is passed directly from the editor to the prover, via internal communication channels.

3.1.2 Theories

The *Theories* panel (see also figure 3.1) provides an overview of the status of continuous checking of theory nodes within the document model. Unlike batch sessions of `isabelle build` [12], theory nodes are identified by full path names; this allows to work with multiple (disjoint) Isabelle sessions simultaneously within the same editor session.

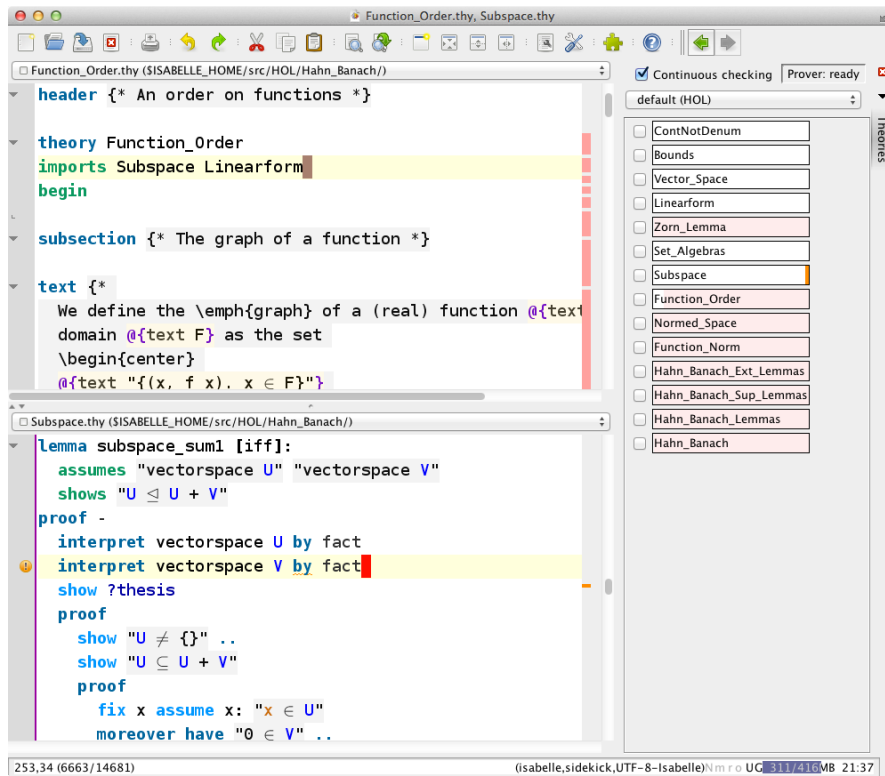


Figure 3.1: Theories panel with an overview of the document-model, and some jEdit text areas as editable view on some of the document nodes

Certain events to open or update editor buffers cause Isabelle/jEdit to resolve dependencies of theory imports. The system requests to load additional files into editor buffers, in order to be included in the document model for

further checking. It is also possible to let the system resolve dependencies automatically, according to the system option `jedit_auto_load`.

The visible *perspective* of Isabelle/jEdit is defined by the collective view on theory buffers via open text areas. The perspective is taken as a hint for document processing: the prover ensures that those parts of a theory where the user is looking are checked, while other parts that are presently not required are ignored. The perspective is changed by opening or closing text area windows, or scrolling within a window.

The *Theories* panel provides some further options to influence the process of continuous checking: it may be switched off globally to restrict the prover to superficial processing of command syntax. It is also possible to indicate theory nodes as *required* for continuous checking: this means such nodes and all their imports are always processed independently of the visibility status (if continuous checking is enabled). Big theory libraries that are marked as required can have significant impact on performance, though.

Formal markup of checked theory content is turned into GUI rendering, based on a standard repertoire known from IDEs for programming languages: colors, icons, highlighting, squiggly underlines, tooltips, hyperlinks etc. For outer syntax of Isabelle/Isar there is some traditional syntax-highlighting via static keyword tables and tokenization within the editor. In contrast, the painting of inner syntax (term language etc.) uses semantic information that is reported dynamically from the logical context. Thus the prover can provide additional markup to help the user to understand the meaning of formal text, and to produce more text with some add-on tools (e.g. information messages with *sendback* markup by automated provers or disprovers in the background).

3.1.3 Auxiliary files

Special load commands like `ML_file` and `SML_file` [4] refer to auxiliary files within some theory. Conceptually, the file argument of the command extends the theory source by the content of the file, but its editor buffer may be loaded / changed / saved separately. The PIDE document model propagates changes of auxiliary file content to the corresponding load command in the theory, to update and process it accordingly: changes of auxiliary file content are treated as changes of the corresponding load command.

As a concession to the massive amount of ML files in Isabelle/HOL itself, the content of auxiliary files is only added to the PIDE document-model on demand, the first time when opened explicitly in the editor. There are further

tricks to manage markup of ML files, such that Isabelle/HOL may be edited conveniently in the Prover IDE on small machines with only 4–8 GB of main memory. Using `Pure` as logic session image, the exploration may start at the top `$ISABELLE_HOME/src/HOL/Main.thy` or the bottom `$ISABELLE_HOME/src/HOL/HOL.thy`, for example.

Initially, before an auxiliary file is opened in the editor, the prover reads its content from the physical file-system. After the file is opened for the first time in the editor, e.g. by following the hyperlink (§3.4) for the argument of its `ML_file` command, the content is taken from the jEdit buffer.

The change of responsibility from prover to editor counts as an update of the document content, so subsequent theory sources need to be re-checked. When the buffer is closed, the responsibility remains to the editor: the file may be opened again without causing another document update.

A file that is opened in the editor, but its theory with the load command is not, is presently inactive in the document model. A file that is loaded via multiple load commands is associated to an arbitrary one: this situation is morally unsupported and might lead to confusion.

Output that refers to an auxiliary file is combined with that of the corresponding load command, and shown whenever the file or the command are active (see also §3.2).

Warnings, errors, and other useful markup is attached directly to the positions in the auxiliary file buffer, in the manner of other well-known IDEs. By using the load command `SML_file` as explained in `$ISABELLE_HOME/src/Tools/SML/Examples.thy`, Isabelle/jEdit may be used as fully-featured IDE for Standard ML, independently of theory or proof development: the required theory merely serves as some kind of project file for a collection of SML source modules.

3.2 Output

Prover output consists of *markup* and *messages*. Both are directly attached to the corresponding positions in the original source text, and visualized in the text area, e.g. as text colours for free and bound variables, or as squiggly underlines for warnings, errors etc. (see also figure 3.2). In the latter case, the corresponding messages are shown by hovering with the mouse over the highlighted text — although in many situations the user should already get some clue by looking at the position of the text highlighting, without the text itself.

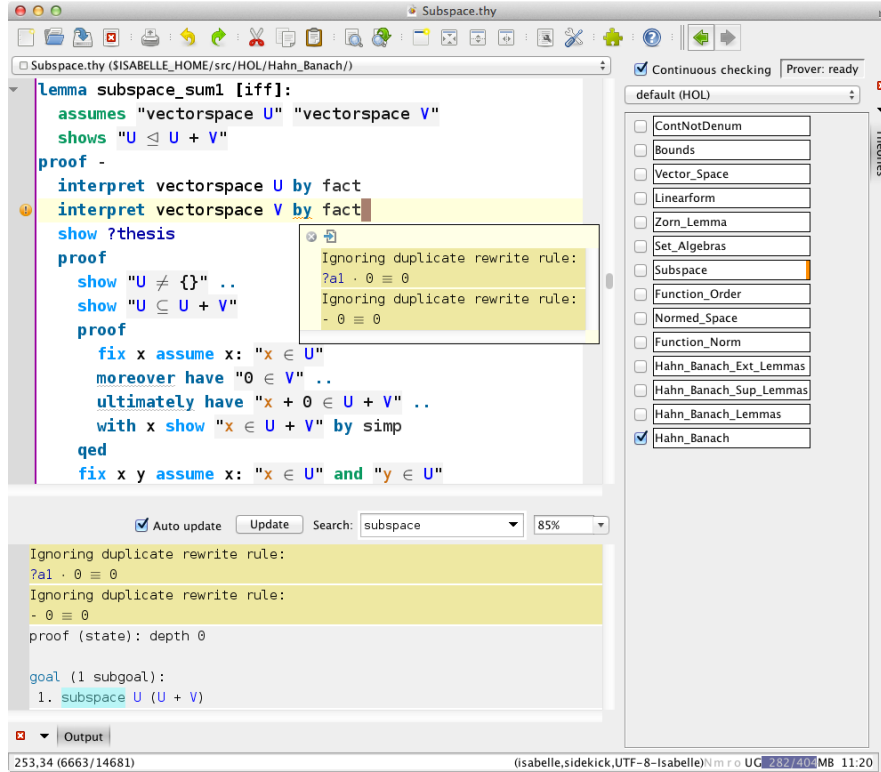


Figure 3.2: Multiple views on prover output: gutter area with icon, text area with popup, overview area, Theories panel, Output panel

The “gutter area” on the left-hand-side of the text area uses icons to provide a summary of the messages within the adjacent line of text. Message priorities are used to prefer errors over warnings, warnings over information messages, but plain output is ignored.

The “overview area” on the right-hand-side of the text area uses similar information to paint small rectangles for the overall status of the whole text buffer. The graphics is scaled to fit the logical buffer length into the given window height. Mouse clicks on the overview area position the cursor approximately to the corresponding line of text in the buffer. Repainting the overview in real-time causes problems with big theories, so it is restricted according to the system option `jedit_text_overview_limit` (in characters).

Another course-grained overview is provided by the *Theories* panel, but without direct correspondence to text positions. A double-click on one of the theory entries with their status overview opens the corresponding text buffer, without changing the cursor position.

In addition, the *Output* panel displays prover messages that correspond to a given command, within a separate window.

The cursor position in the presently active text area determines the prover command whose cumulative message output is appended and shown in that window (in canonical order according to the internal execution of the command). There are also control elements to modify the update policy of the output wrt. continued editor movements. This is particularly useful with several independent instances of the *Output* panel, which the Dockable Window Manager of jEdit can handle conveniently.

Former users of the old TTY interaction model (e.g. Proof General) might find a separate window for prover messages familiar, but it is important to understand that the main Prover IDE feedback happens elsewhere. It is possible to do meaningful proof editing within the primary text area and its markup, while using secondary output windows only rarely.

The main purpose of the output window is to “debug” unclear situations by inspecting internal state of the prover.¹ Consequently, some special messages for *tracing* or *proof state* only appear here, and are not attached to the original source.

In any case, prover messages also contain markup that may be explored recursively via tooltips or hyperlinks (see §3.4), or clicked directly to initiate certain actions (see §3.6 and §3.7).

3.3 Query

The *Query* panel provides various GUI forms to request extra information from the prover. In old times the user would have issued some diagnostic command like **find_theorems** and inspected its output, but this is now integrated into the Prover IDE.

A *Query* window provides some input fields and buttons for a particular query command, with output in a dedicated text area. There are various query modes: *Find Theorems*, *Find Constants*, *Print Context*, e.g. see figure 3.3. As usual in jEdit, multiple *Query* windows may be active at the same time: any number of floating instances, but at most one docked instance (which is used by default).

The following GUI elements are common to all query modes:

¹In that sense, unstructured tactic scripts depend on continuous debugging with internal state inspection.

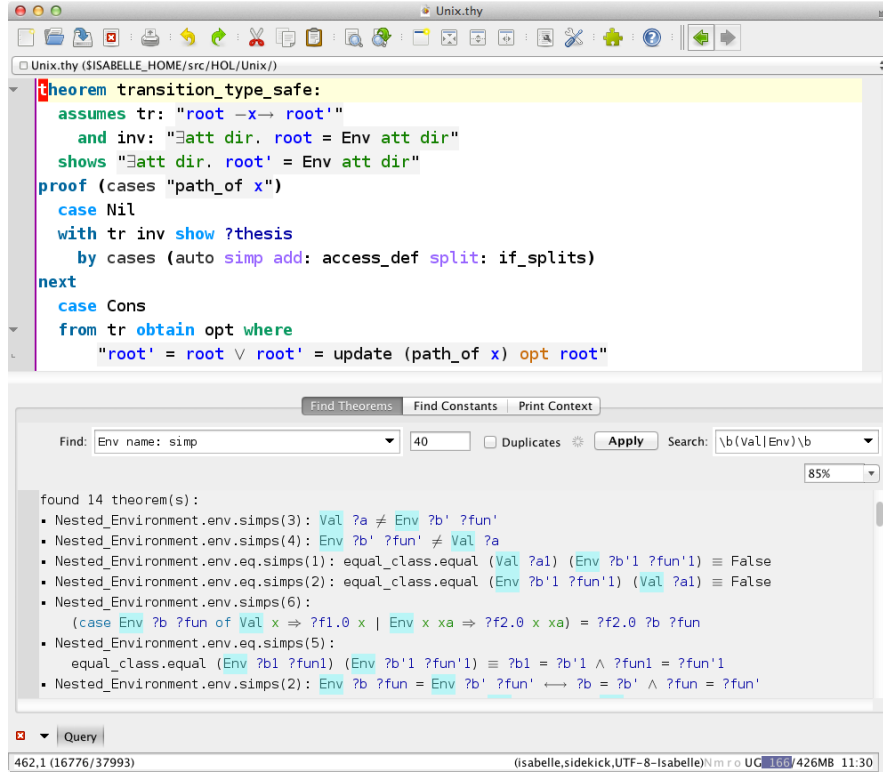


Figure 3.3: An instance of the Query panel

- The spinning wheel provides feedback about the status of a pending query wrt. the evaluation of its context and its own operation.
- The *Apply* button attaches a fresh query invocation to the current context of the command where the cursor is pointing in the text.
- The *Search* field allows to highlight query output according to some regular expression, in the notation that is commonly used on the Java platform.² This may serve as an additional visual filter of the result.
- The *Zoom* box controls the font size of the output area.

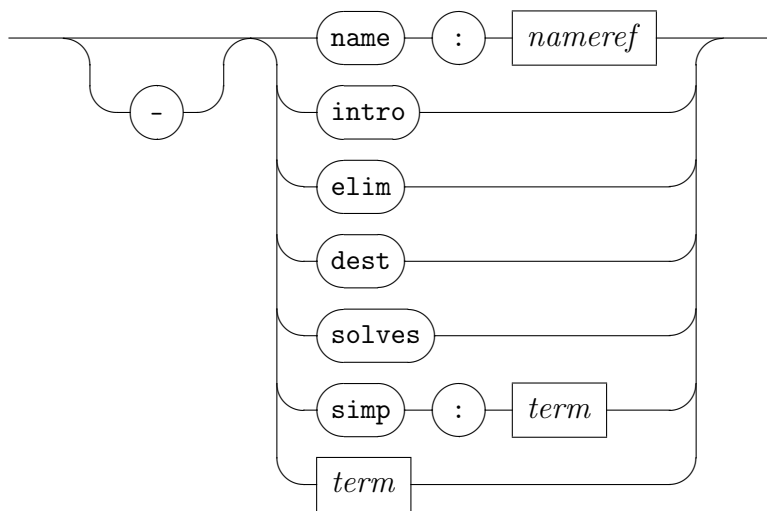
All query operations are asynchronous: there is no need to wait for the evaluation of the document for the query context, nor for the query operation itself. Query output may be detached as independent *Info* window, using a menu operation of the dockable window manager. The printed result usually

²<http://docs.oracle.com/javase/7/docs/api/java/util/regex/Pattern.html>

provides sufficient clues about the original query, with some hyperlink to its context (via markup of its head line).

3.3.1 Find theorems


The *Query* panel in *Find Theorems* mode retrieves facts from the theory or proof context matching all of given criteria in the *Find* text field. A single criterium has the following syntax:



See also the Isar command `find_theorems` in [4].

3.3.2 Find constants

The *Query* panel in *Find Constants* mode prints all constants whose type meets all of the given criteria in the *Find* text field. A single criterium has the following syntax:



See also the Isar command `find_consts` in [4].

3.3.3 Print context

The *Query* panel in *Print Context* mode prints information from the theory or proof context, or proof state. See also the Isar commands **print_context**, **print_cases**, **print_term_bindings**, **print_theorems**, **print_state** described in [4].

3.4 Tooltips and hyperlinks

Formally processed text (prover input or output) contains rich markup information that can be explored further by using the **CONTROL** modifier key on Linux and Windows, or **COMMAND** on Mac OS X. Hovering with the mouse while the modifier is pressed reveals a *tooltip* (grey box over the text with a yellow popup) and/or a *hyperlink* (black rectangle over the text with change of mouse pointer); see also figure 3.4.

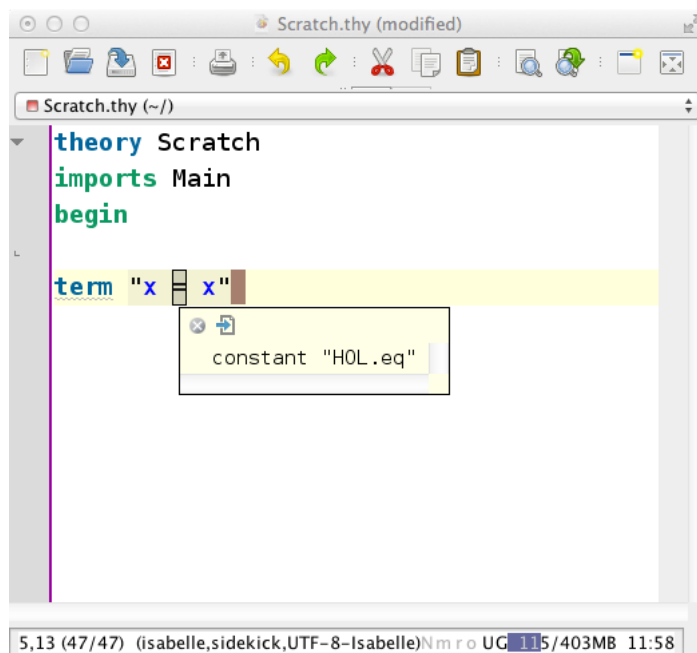


Figure 3.4: Tooltip and hyperlink for some formal entity

Tooltip popups use the same rendering mechanisms as the main text area, and further tooltips and/or hyperlinks may be exposed recursively by the same mechanism; see figure 3.5.

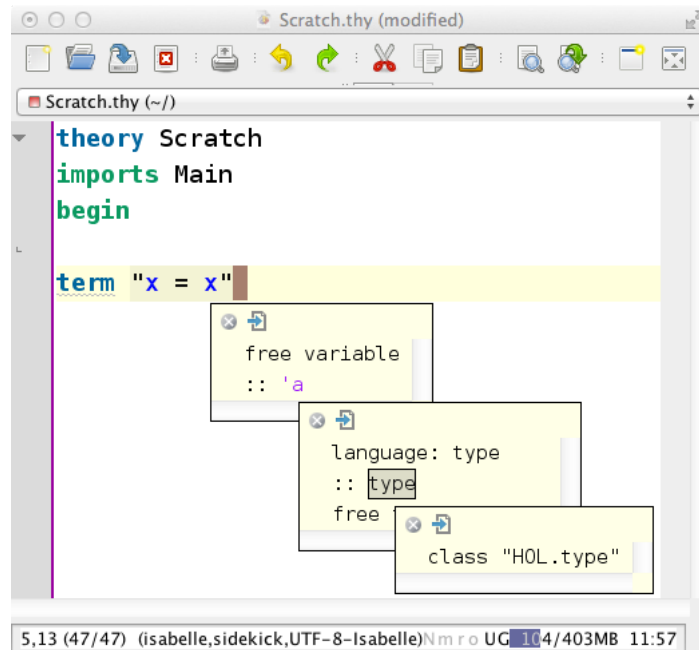


Figure 3.5: Nested tooltips over formal entities

The tooltip popup window provides some controls to *close* or *detach* the window, turning it into a separate *Info* window managed by jEdit. The **ESCAPE** key closes *all* popups, which is particularly relevant when nested tooltips are stacking up.

A black rectangle in the text indicates a hyperlink that may be followed by a mouse click (while the **CONTROL** or **COMMAND** modifier key is still pressed). Such jumps to other text locations are recorded by the *Navigator* plugin, which is bundled with Isabelle/jEdit and enabled by default, including navigation arrows in the main jEdit toolbar.

Also note that the link target may be a file that is itself not subject to formal document processing of the editor session and thus prevents further exploration: the chain of hyperlinks may end in some source file of the underlying logic image, or within the Isabelle/ML bootstrap sources of Isabelle/Pure.

3.5 Completion

Smart completion of partial input is the IDE functionality *par excellence*. Isabelle/jEdit combines several sources of information to achieve that. De-

spite its complexity, it should be possible to get some idea how completion works by experimentation, based on the overview of completion varieties in §3.5.1. The remaining subsections explain concepts around completion more systematically.

Explicit completion is triggered by the action `isabelle.complete`, which is bound to the keyboard shortcut `C+b`, and thus overrides the jEdit default for `complete-word`.

Implicit completion hooks into the regular keyboard input stream of the editor, with some event filtering and optional delays.

Completion options may be configured in *Plugin Options / Isabelle / General / Completion*. These are explained in further detail below, whenever relevant. There is also a summary of options in §3.5.6.

The asynchronous nature of PIDE interaction means that information from the prover is delayed — at least by a full round-trip of the document update protocol. The default options already take this into account, with a sufficiently long completion delay to speculate on the availability of all relevant information from the editor and the prover, before completing text immediately or producing a popup. Although there is an inherent danger of non-deterministic behaviour due to such real-time parameters, the general completion policy aims at determined results as far as possible.

3.5.1 Varieties of completion

Built-in templates

Isabelle is ultimately a framework of nested sub-languages of different kinds and purposes. The completion mechanism supports this by the following built-in templates:

- ‘ (single ASCII back-quote) supports *quotations* via text cartouches. There are three selections, which are always presented in the same order and do not depend on any context information. The default choice produces a template “`⟨□⟩`”, where the box indicates the cursor position after insertion; the other choices help to repair the block structure of unbalanced text cartouches.
- @{ is completed to the template “`@{□}`”, where the box indicates the cursor position after insertion. Here it is convenient to use the wildcard “`__`” or a more specific name prefix to let semantic completion of name-space entries propose antiquotation names.

With some practice, input of quoted sub-languages and antiquotations of embedded languages should work fluently. Note that national keyboard layouts might cause problems with back-quote as dead key: if possible, dead keys should be disabled.

Syntax keywords

Syntax completion tables are determined statically from the keywords of the “outer syntax” of the underlying edit mode: for theory files this is the syntax of Isar commands.

Keywords are usually plain words, which means the completion mechanism only inserts them directly into the text for explicit completion (§3.5.3), but produces a popup (§3.5.4) otherwise.

At the point where outer syntax keywords are defined, it is possible to specify an alternative replacement string to be inserted instead of the keyword itself. An empty string means to suppress the keyword altogether, which is occasionally useful to avoid confusion, e.g. the rare keyword **simproc_setup** vs. the frequent name-space entry *simp*.

Isabelle symbols

The completion tables for Isabelle symbols (§2.3) are determined statically from `$ISABELLE_HOME/etc/symbols` and `$ISABELLE_HOME_USER/etc/symbols` for each symbol specification as follows:

completion entry	example
literal symbol	<code>\<forall></code>
symbol name with backslash	<code>\forall</code>
symbol abbreviation	<code>ALL</code> or <code>!</code>

When inserted into the text, the above examples all produce the same Unicode rendering \forall of the underlying symbol `\<forall>`.

A symbol abbreviation that is a plain word, like `ALL`, is treated like a syntax keyword. Non-word abbreviations like `-->` are inserted more aggressively, except for single-character abbreviations like `!` above.

Symbol completion depends on the semantic language context (§3.5.2), to enable or disable that aspect for a particular sub-language of Isabelle. For example, symbol completion is suppressed within document source to avoid confusion with \LaTeX macros that use similar notation.

Name-space entries

This is genuine semantic completion, using information from the prover, so it requires some delay. A *failed name-space lookup* produces an error message that is annotated with a list of alternative names that are legal. The list of results is truncated according to the system option `completion_limit`. The completion mechanism takes this into account when collecting information on the prover side.

Already recognized names are *not* completed further, but completion may be extended by appending a suffix of underscores. This provokes a failed lookup, and another completion attempt while ignoring the underscores. For example, in a name space where `foo` and `foobar` are known, the input `foo` remains unchanged, but `foo_` may be completed to `foo` or `foobar`.

The special identifier “`__`” serves as a wild-card for arbitrary completion: it exposes the name-space content to the completion mechanism (truncated according to `completion_limit`). This is occasionally useful to explore an unknown name-space, e.g. in some template.

File-system paths

Depending on prover markup about file-system path specifications in the source text, e.g. for the argument of a load command (§3.1.3), the completion mechanism explores the directory content and offers the result as completion popup. Relative path specifications are understood wrt. the *master directory* of the document node (§3.1.1) of the enclosing editor buffer; this requires a proper theory, not an auxiliary file.

A suffix of slashes may be used to continue the exploration of an already recognized directory name.

Spell-checking

The spell-checker combines semantic markup from the prover (regions of plain words) with static dictionaries (word lists) that are known to the editor.

Unknown words are underlined in the text, using `spell_checker_color` (blue by default). This is not an error, but a hint to the user that some action may be taken. The jEdit context menu provides various actions, as far as applicable:

```
isabelle.complete-word  
isabelle.exclude-word  
isabelle.exclude-word-permanently  
isabelle.include-word  
isabelle.include-word-permanently
```

Instead of the specific `isabelle.complete-word`, it is also possible to use the generic `isabelle.complete` with its default keyboard shortcut `C+b`.

Dictionary lookup uses some educated guesses about lower-case, upper-case, and capitalized words. This is oriented on common use in English, where this aspect is not decisive for proper spelling, in contrast to German, for example.

3.5.2 Semantic completion context

Completion depends on a semantic context that is provided by the prover, although with some delay, because at least a full PIDE protocol round-trip is required. Until that information becomes available in the PIDE document-model, the default context is given by the outer syntax of the editor mode (see also §3.1.1).

The semantic *language context* provides information about nested sub-languages of Isabelle: keywords are only completed for outer syntax, symbols or antiquotations for languages that support them. E.g. there is no symbol completion for ML source, but within ML strings, comments, antiquotations.

The prover may produce *no completion* markup in exceptional situations, to tell that some language keywords should be excluded from further completion attempts. For example, `:` within accepted Isar syntax loses its meaning as abbreviation for symbol `∈`.

The completion context is *ignored* for built-in templates and symbols in their explicit form “`\<foobar>`”; see also §3.5.1. This allows to complete within broken input that escapes its normal semantic context, e.g. antiquotations or string literals in ML, which do not allow arbitrary backslash sequences.

3.5.3 Input events

Completion is triggered by certain events produced by the user, with optional delay after keyboard input according to `jedit_completion_delay`.

Explicit completion works via action `isabelle.complete` with keyboard shortcut `C+b`. This overrides the shortcut for `complete-word` in jEdit,

but it is possible to restore the original jEdit keyboard mapping of `complete-word` via *Global Options / Shortcuts* and invent a different one for `isabelle.complete`.

Explicit spell-checker completion works via `isabelle.complete-word`, which is exposed in the jEdit context menu, if the mouse points to a word that the spell-checker can complete.

Implicit completion works via regular keyboard input of the editor. It depends on further side-conditions:

1. The system option `jedit_completion` needs to be enabled (default).
2. Completion of syntax keywords requires at least 3 relevant characters in the text.
3. The system option `jedit_completion_delay` determines an additional delay (0.5 by default), before opening a completion popup. The delay gives the prover a chance to provide semantic completion information, notably the context (§3.5.2).
4. The system option `jedit_completion_immediate` (enabled by default) controls whether replacement text should be inserted immediately without popup, regardless of `jedit_completion_delay`. This aggressive mode of completion is restricted to Isabelle symbols and their abbreviations (§2.3).
5. Completion of symbol abbreviations with only one relevant character in the text always enforces an explicit popup, regardless of `jedit_completion_immediate`.

3.5.4 Completion popup

A *completion popup* is a minimally invasive GUI component over the text area that offers a selection of completion items to be inserted into the text, e.g. by mouse clicks. Items are sorted dynamically, according to the frequency of selection, with persistent history. The popup may interpret special keys `ENTER`, `TAB`, `ESCAPE`, `UP`, `DOWN`, `PAGE_UP`, `PAGE_DOWN`, but all other key events are passed to the underlying text area. This allows to ignore unwanted completions most of the time and continue typing quickly. Thus the popup serves as a mechanism of confirmation of proposed items, but the default is to continue without completion.

The meaning of special keys is as follows:

key	action
ENTER	select completion (if <code>jedit_completion_select_enter</code>)
TAB	select completion (if <code>jedit_completion_select_tab</code>)
ESCAPE	dismiss popup
UP	move up one item
DOWN	move down one item
PAGE_UP	move up one page of items
PAGE_DOWN	move down one page of items

Movement within the popup is only active for multiple items. Otherwise the corresponding key event retains its standard meaning within the underlying text area.

3.5.5 Insertion

Completion may first propose replacements to be selected (via a popup), or replace text immediately in certain situations and depending on certain options like `jedit_completion_immediate`. In any case, insertion works uniformly, by imitating normal jEdit text insertion, depending on the state of the *text selection*. Isabelle/jEdit tries to accommodate the most common forms of advanced selections in jEdit, but not all combinations make sense. At least the following important cases are well-defined:

No selection. The original is removed and the replacement inserted, depending on the caret position.

Rectangular selection of zero width. This special case is treated by jEdit as “tall caret” and insertion of completion imitates its normal behaviour: separate copies of the replacement are inserted for each line of the selection.

Other rectangular selection or multiple selections. Here the original is removed and the replacement is inserted for each line (or segment) of the selection.

Support for multiple selections is particularly useful for *HyperSearch*: clicking on one of the items in the *HyperSearch Results* window makes jEdit select all its occurrences in the corresponding line of text. Then explicit completion can be invoked via **C+b**, e.g. to replace occurrences of `-->` by `→`.

Insertion works by removing and inserting pieces of text from the buffer. This counts as one atomic operation on the jEdit history. Thus unintended completions may be reverted by the regular `undo` action of jEdit. According to normal jEdit policies, the recovered text after `undo` is selected: `ESCAPE` is required to reset the selection and to continue typing more text.

3.5.6 Options

This is a summary of Isabelle/Scala system options that are relevant for completion. They may be configured in *Plugin Options / Isabelle / General* as usual.

- `completion_limit` specifies the maximum number of name-space entries exposed in semantic completion by the prover.
- `jedit_completion` guards implicit completion via regular jEdit key events (§3.5.3): it allows to disable implicit completion altogether.
- `jedit_completion_select_enter` and `jedit_completion_select_tab` enable keys to select a completion item from the popup (§3.5.4). Note that a regular mouse click on the list of items is always possible.
- `jedit_completion_context` specifies whether the language context provided by the prover should be used at all. Disabling that option makes completion less “semantic”. Note that incomplete or severely broken input may cause some disagreement of the prover and the user about the intended language context.
- `jedit_completion_delay` and `jedit_completion_immediate` determine the handling of keyboard events for implicit completion (§3.5.3). A `jedit_completion_delay > 0` postpones the processing of key events, until after the user has stopped typing for the given time span, but `jedit_completion_immediate = true` means that abbreviations of Isabelle symbols are handled nonetheless.
- `jedit_completion_path_ignore` specifies “glob” patterns to ignore in file-system path completion (separated by colons), e.g. backup files ending with tilde.
- `spell_checker` is a global guard for all spell-checker operations: it allows to disable that mechanism altogether.

- `spell_checker_dictionary` determines the current dictionary, taken from the colon-separated list in the settings variable `JORTHO_DICTIONARIES`. There are jEdit actions to specify local updates to a dictionary, by including or excluding words. The result of permanent dictionary updates is stored in the directory `$ISABELLE_HOME_USER/dictionaries`, in a separate file for each dictionary.
- `spell_checker_elements` specifies a comma-separated list of markup elements that delimit words in the source that is subject to spell-checking, including various forms of comments.

3.6 Automatically tried tools

Continuous document processing works asynchronously in the background. Visible document source that has been evaluated may get augmented by additional results of *asynchronous print functions*. The canonical example is proof state output, which is always enabled. More heavy-weight print functions may be applied, in order to prove or disprove parts of the formal text by other means.

Isabelle/HOL provides various automatically tried tools that operate on outermost goal statements (e.g. **lemma**, **theorem**), independently of the state of the current proof attempt. They work implicitly without any arguments. Results are output as *information messages*, which are indicated in the text area by blue squiggles and a blue information sign in the gutter (see figure 3.6). The message content may be shown as for other output (see also §3.2). Some tools produce output with *sendback* markup, which means that clicking on certain parts of the output inserts that text into the source in the proper place.

The following Isabelle system options control the behavior of automatically tried tools (see also the jEdit dialog window *Plugin Options / Isabelle / General / Automatically tried tools*):

- `auto_methods` controls automatic use of a combination of standard proof methods (*auto*, *simp*, *blast*, etc.). This corresponds to the Isar command **try0** [4].

The tool is disabled by default, since unparameterized invocation of standard proof methods often consumes substantial CPU resources without leading to success.

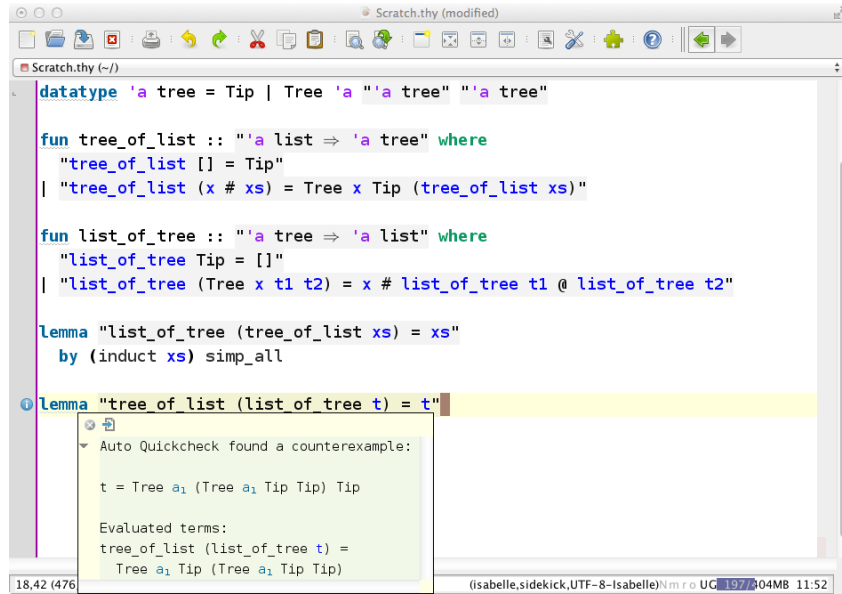


Figure 3.6: Result of automatically tried tools

- `auto_nitpick` controls a slightly reduced version of **nitpick**, which tests for counterexamples using first-order relational logic. See also the Nitpick manual [2].

This tool is disabled by default, due to the extra overhead of invoking an external Java process for each attempt to disprove a subgoal.

- `auto_quickcheck` controls automatic use of **quickcheck**, which tests for counterexamples using a series of assignments for free variables of a subgoal.

This tool is *enabled* by default. It requires little overhead, but is a bit weaker than **nitpick**.

- `auto_sledgehammer` controls a significantly reduced version of **sledgehammer**, which attempts to prove a subgoal using external automatic provers. See also the Sledgehammer manual [1].

This tool is disabled by default, due to the relatively heavy nature of Sledgehammer.

- `auto_solve_direct` controls automatic use of **solve_direct**, which checks whether the current subgoals can be solved directly by an existing theorem. This also helps to detect duplicate lemmas.

This tool is *enabled* by default.

Invocation of automatically tried tools is subject to some global policies of parallel execution, which may be configured as follows:

- `auto_time_limit` (default 2.0) determines the timeout (in seconds) for each tool execution.
- `auto_time_start` (default 1.0) determines the start delay (in seconds) for automatically tried tools, after the main command evaluation is finished.

Each tool is submitted independently to the pool of parallel execution tasks in Isabelle/ML, using hardwired priorities according to its relative “heaviness”. The main stages of evaluation and printing of proof states take precedence, but an already running tool is not canceled and may thus reduce reactivity of proof document processing.

Users should experiment how the available CPU resources (number of cores) are best invested to get additional feedback from prover in the background, by using a selection of weaker or stronger tools.

3.7 Sledgehammer

The *Sledgehammer* panel (figure 3.7) provides a view on some independent execution of the Isar command **sledgehammer**, with process indicator (spinning wheel) and GUI elements for important Sledgehammer arguments and options. Any number of Sledgehammer panels may be active, according to the standard policies of Dockable Window Management in jEdit. Closing such windows also cancels the corresponding prover tasks.

The *Apply* button attaches a fresh invocation of **sledgehammer** to the command where the cursor is pointing in the text — this should be some pending proof problem. Further buttons like *Cancel* and *Locate* help to manage the running process.

Results appear incrementally in the output window of the panel. Proposed proof snippets are marked-up as *sendback*, which means a single mouse click inserts the text into a suitable place of the original source. Some manual editing may be required nonetheless, say to remove earlier proof attempts.

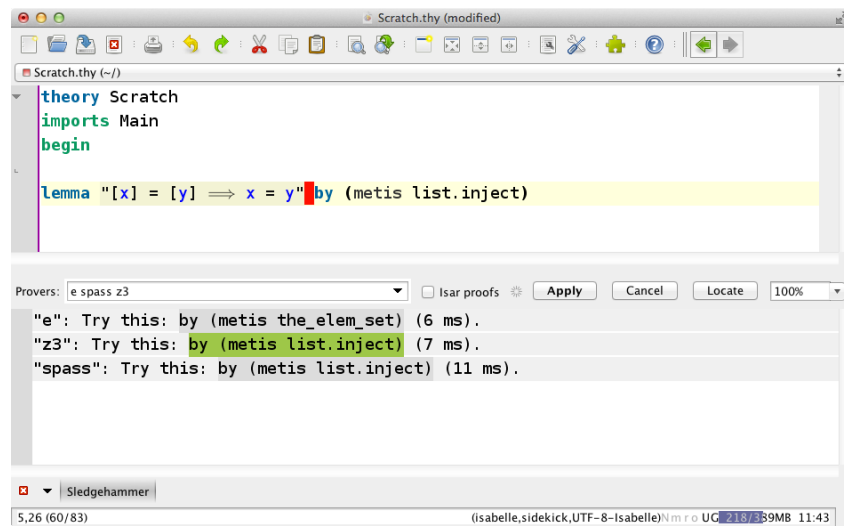


Figure 3.7: An instance of the Sledgehammer panel

Miscellaneous tools

4.1 Timing

Managed evaluation of commands within PIDE documents includes timing information, which consists of elapsed (wall-clock) time, CPU time, and GC (garbage collection) time. Note that in a multithreaded system it is difficult to measure execution time precisely: elapsed time is closer to the real requirements of runtime resources than CPU or GC time, which are both subject to influences from the parallel environment that are outside the scope of the current command transaction.

The *Timing* panel provides an overview of cumulative command timings for each document node. Commands with elapsed time below the given threshold are ignored in the grand total. Nodes are sorted according to their overall timing. For the document node that corresponds to the current buffer, individual command timings are shown as well. A double-click on a theory node or command moves the editor focus to that particular source position.

It is also possible to reveal individual timing information via some tooltip for the corresponding command keyword, using the technique of mouse hovering with `CONTROL/COMMAND` modifier key as explained in §3.4. Actual display of timing depends on the global option `jedit_timing_threshold`, which can be configured in *Plugin Options / Isabelle / General*.

The *Monitor* panel visualizes various data collections about recent activity of the Isabelle/ML task farm and the underlying ML runtime system. The display is continuously updated according to `editor_chart_delay`. Note that the painting of the chart takes considerable runtime itself — on the Java Virtual Machine that runs Isabelle/Scala, not Isabelle/ML. Internally, the Isabelle/Scala module `isabelle.ML_Statistics` provides further access to statistics of Isabelle/ML.

4.2 Low-level output

Prover output is normally shown directly in the main text area or secondary *Output* panels, as explained in §3.2.

Beyond this, it is occasionally useful to inspect low-level output channels via some of the following additional panels:

- *Protocol* shows internal messages between the Isabelle/Scala and Isabelle/ML side of the PIDE editing protocol. Recording of messages starts with the first activation of the corresponding dockable window; earlier messages are lost.

Actual display of protocol messages causes considerable slowdown, so it is important to undock all *Protocol* panels for production work.

- *Raw Output* shows chunks of text from the `stdout` and `stderr` channels of the prover process. Recording of output starts with the first activation of the corresponding dockable window; earlier output is lost.

The implicit stateful nature of physical I/O channels makes it difficult to relate raw output to the actual command from where it was originating. Parallel execution may add to the confusion. Peeking at physical process I/O is only the last resort to diagnose problems with tools that are not PIDE compliant.

Under normal circumstances, prover output always works via managed message channels (corresponding to `writeln`, `warning`, `Output.error_message` in Isabelle/ML), which are displayed by regular means within the document model (§3.2).

- *Syslog* shows system messages that might be relevant to diagnose problems with the startup or shutdown phase of the prover process; this also includes raw output on `stderr`.

A limited amount of syslog messages are buffered, independently of the docking state of the *Syslog* panel. This allows to diagnose serious problems with Isabelle/PIDE process management, outside of the actual protocol layer.

Under normal situations, such low-level system output can be ignored.

Known problems and workarounds

- **Problem:** Odd behavior of some diagnostic commands with global side-effects, like writing a physical file.
Workaround: Copy/paste complete command text from elsewhere, or disable continuous checking temporarily.
- **Problem:** No direct support to remove document nodes from the collection of theories.
Workaround: Clear the buffer content of unused files and close *without* saving changes.
- **Problem:** Keyboard shortcuts C+PLUS and C+MINUS for adjusting the editor font size depend on platform details and national keyboards.
Workaround: Rebind keys via *Global Options / Shortcuts*.
- **Problem:** The Mac OS X key sequence COMMAND+COMMA for application *Preferences* is in conflict with the jEdit default keyboard shortcut for *Incremental Search Bar* (action `quick-search`).
Workaround: Rebind key via *Global Options / Shortcuts* according to national keyboard, e.g. COMMAND+SLASH on English ones.
- **Problem:** Mac OS X system fonts sometimes lead to character drop-outs in the main text area.
Workaround: Use the default IsabelleText font. (Do not install that font on the system.)
- **Problem:** Some Linux/X11 input methods such as IBus tend to disrupt key event handling of Java/AWT/Swing.
Workaround: Do not use X11 input methods. Note that environment variable XMODIFIERS is reset by default within Isabelle settings.

- **Problem:** Some Linux/X11 window managers that are not “re-parenting” cause problems with additional windows opened by Java. This affects either historic or neo-minimalistic window managers like *awesome* or *xmonad*.

Workaround: Use a regular re-parenting X11 window manager.

- **Problem:** Recent forks of Linux/X11 window managers and desktop environments (variants of Gnome) disrupt the handling of menu popups and mouse positions of Java/AWT/Swing.

Workaround: Use mainstream versions of Linux desktops.

- **Problem:** Full-screen mode via jEdit action `toggle-full-screen` (default keyboard shortcut F11) works on Windows, but not on Mac OS X or various Linux/X11 window managers.

Workaround: Use native full-screen control of the window manager (notably on Mac OS X).

Bibliography

- [1] J. C. Blanchette. *Hammering Away: A User's Guide to Sledgehammer for Isabelle/HOL*. <http://isabelle.in.tum.de/doc/sledgehammer.pdf>.
- [2] J. C. Blanchette. *Picking Nits: A User's Guide to Nitpick for Isabelle/HOL*. <http://isabelle.in.tum.de/doc/nitpick.pdf>.
- [3] M. Wenzel. *The Isabelle/Isar Implementation*. <http://isabelle.in.tum.de/doc/implementation.pdf>.
- [4] M. Wenzel. *The Isabelle/Isar Reference Manual*. <http://isabelle.in.tum.de/doc/isar-ref.pdf>.
- [5] M. Wenzel. Parallel proof checking in Isabelle/Isar. In G. Dos Reis and L. Théry, editors, *ACM SIGSAM Workshop on Programming Languages for Mechanized Mathematics Systems (PLMMS 2009)*. ACM Digital Library, 2009.
- [6] M. Wenzel. Asynchronous proof processing with Isabelle/Scala and Isabelle/jEdit. In C. S. Coen and D. Aspinall, editors, *User Interfaces for Theorem Provers (UITP 2010), FLOC 2010 Satellite Workshop*, ENTCS. Elsevier, July 2010.
- [7] M. Wenzel. Isabelle as document-oriented proof assistant. In J. H. Davenport, W. M. Farmer, F. Rabe, and J. Urban, editors, *Conference on Intelligent Computer Mathematics / Mathematical Knowledge Management (CICM/MKM 2011)*, volume 6824 of *LNAI*. Springer, 2011.
- [8] M. Wenzel. Isabelle/jEdit — a Prover IDE within the PIDE framework. In J. Jeuring et al., editors, *Conference on Intelligent Computer Mathematics (CICM 2012)*, volume 7362 of *LNAI*. Springer, 2012.
- [9] M. Wenzel. READ-EVAL-PRINT in parallel and asynchronous proof-checking. In *User Interfaces for Theorem Provers (UITP 2012)*, EPTCS, 2013.
- [10] M. Wenzel. Shared-memory multiprocessing for interactive theorem proving. In S. Blazy, C. Paulin-Mohring, and D. Pichardie, editors, *Interactive Theorem Proving - 4th International Conference, ITP 2013, Rennes, France, July 22-26, 2013. Proceedings*, volume 7998 of *Lecture Notes in Computer Science*. Springer, 2013.

- [11] M. Wenzel. Asynchronous user interaction and tool integration in Isabelle/PIDE. In G. Klein and R. Gamboa, editors, *Interactive Theorem Proving — 5th International Conference, ITP 2014, Vienna, Austria*, volume 8558 of *Lecture Notes in Computer Science*. Springer, 2014.
- [12] M. Wenzel and S. Berghofer. *The Isabelle System Manual*. <http://isabelle.in.tum.de/doc/system.pdf>.

Index

- auto_methods (system option), 32
- auto_nitpick (system option), 33
- auto_quickcheck (system option), 33
- auto_sledgehammer (system option), 33
- auto_solve_direct (system option), 33
- auto_time_limit (system option), 34
- auto_time_start (system option), 34
- complete-word (action), 25, 28
- completion_limit (system option), 27, **31**
- editor_chart_delay (system option), 36
- find_consts (command), 22
- find_theorems (command), 22
- isabelle.complete (action), 25, 28
- isabelle.complete-word (action), 28, 29
- isabelle.control-bold (action), 11
- isabelle.control-reset (action), 11
- isabelle.control-sub (action), 11
- isabelle.control-sup (action), 11
- isabelle.exclude-word (action), 28
- isabelle.exclude-word-permanently (action), 28
- isabelle.include-word (action), 28
- isabelle.include-word-permanently (action), 28
- isabelle.options (action), **5**
- jedit (tool), **5**
- jedit_completion (system option), 29, **31**
- jedit_completion_context (system option), **31**
- jedit_completion_delay (system option), 29, **31**
- jedit_completion_immediate (system option), 29, **31**
- jedit_completion_path_ignore (system option), **31**
- jedit_completion_select_enter (system option), **31**
- jedit_completion_select_tab (system option), **31**
- jedit_print_mode (system option), 6
- jedit_timing_threshold (system option), 36
- JORTHO_DICTIONARIES (setting), **32**
- ML_file (command), 15, 17
- nitpick (command), 33
- print_cases (command), 23
- print_context (command), 23
- print_state (command), 23
- print_term_bindings (command), 23
- print_theorems (command), 23
- quick-search (action), 38
- quickcheck (command), 33
- sledgehammer (command), 33, 34
- SML_file (command), 15, 17
- solve_direct (command), 33

- spell_checker (system option), **31**
- spell_checker_color (system option),
27
- spell_checker_dictionary (system
option), **32**
- spell_checker_elements (system op-
tion), **32**

- theory (command), 15
- toggle-full-screen (action), 39
- try0 (command), 32