

System-on-Chip Environment (SCE Version 2.2.0 Beta): Manual

Lukai Cai
Andreas Gerstlauer
Samar Abdi
Jerry Peng
Dongwan Shin
Haobo Yu
Rainer Dömer
Daniel D. Gajski

Technical Report CECS-TR-03-45
December 2003

System-on-Chip Environment (SCE Version 2.2.0 Beta): Manual

Lukai Cai
Andreas Gerstlauer
Samar Abdi
Jerry Peng
Dongwan Shin
Haobo Yu
Rainer Dömer
Daniel D. Gajski

Technical Report CECS-TR-03-45
December, 2003

Center for Embedded Computer Systems
University of California, Irvine
Irvine, CA 92697-3425, USA
+1 (949) 824-8919

<http://www.cecs.uci.edu>

Contents

1	Introduction	1
1.1	Architecture Explorer	1
1.2	Network Explorer	2
1.3	Communication Synthesizer	2
2	Overview	3
2.1	Modules	4
2.1.1	Main Window	6
2.1.2	Input Dialogs	6
2.1.3	Display Windows	7
2.2	Requirements	8
2.3	Interfaces	8
2.3.1	Internal Interfaces	8
2.3.2	External Interfaces	9
2.4	Performance	10
3	Windows/GUI	10
3.1	Menu Bar	11
3.1.1	File Menu	11
3.1.2	View Menu	12
3.1.3	Project Menu	12
3.1.4	Synthesis Menu	12
3.1.5	Windows	13
3.2	Project Window	14
3.2.1	Models Tab	15
3.2.2	Imports Tab	15
3.2.3	Sources Tab	16
3.3	Design Window	16
3.3.1	Hierarchy Tab	19
3.3.2	Behaviors Tab	19
3.3.3	Channels Tab	21
3.3.4	View Pane	22
3.4	Output Window	22
3.5	Workspace	23
3.6	Message Boxes	23
3.6.1	Error Dialogs	23
3.6.2	Information Dialogs	23

4	Functionality	24
4.1	Application	25
4.1.1	Preferences Editing	26
4.2	Project Handling	28
4.2.1	Project Creation	31
4.2.2	Project Opening	31
4.2.3	Project Saving	32
4.2.4	Project Settings Editing	34
4.2.5	Design Adding	34
4.2.6	Design Opening	35
4.2.7	Design Deletion	35
4.2.8	Design Renaming	36
4.2.9	Description Changing	36
4.2.10	Project Closing	36
4.3	File Handling	36
4.3.1	File Opening	37
4.3.2	File Saving	38
4.3.3	File Closing	40
4.3.4	File Import	40
4.3.5	Design Property Viewing	42
4.3.6	SCE Exiting	42
4.4	Design-Entity Handling	42
4.4.1	Entity Renaming	43
4.4.2	Entity Retyping	43
4.4.3	Entity Deletion	44
4.4.4	Hierarchy Displaying	44
4.5	Synthesis	45
4.5.1	Architecture Exploration	45
4.5.2	Network Exploration	53
4.5.3	Communication Synthesis	63
4.5.4	Decision Import	69
4.6	Window Management	72
	References	73
A	Manual Pages	74
A.1	scc - SpecC Compiler	75

B	Project File XML Schema	82
B.1	Element <code>sce</code>	82
B.1.1	Element <code>sce/project</code>	83
B.1.2	Element <code>sce/compiler</code>	83
B.1.3	Element <code>sce/models</code>	84
B.1.4	Element <code>sce/imports</code>	86
B.1.5	Element <code>sce/sources</code>	86
B.2	ComplexType <code>CompilerOption</code>	87
B.3	ComplexType <code>ModelItem</code>	88
B.3.1	Element <code>ModelItem/item</code>	90
B.3.2	Element <code>ModelItem/option</code>	91
B.4	ComplexType <code>SimpleItem</code>	91

List of Figures

1	SCE environment.	2
2	Overview of the software architecture of SCE.	3
3	Module diagram of SCE.	5
4	Main Window of SCE.	10
5	Project Window (Models tab).	14
6	Project Window (Imports tab).	16
7	Project Window (Sources tab)	17
8	Design Window (Hierarchy tab).	18
9	Design Window (Behaviors tab)	20
10	Design Window (Channels tab).	21
11	Output Window	22
12	Error dialog	24
13	Information dialog	24
14	Edit Preferences dialog (Compiler tab).	27
15	Edit Preferences dialog (Database tab).	29
16	Database Selection dialog.	29
17	Project Open dialog.	32
18	Project Save dialog.	33
19	Project Settings dialog.	34
20	File Open dialog.	37
21	File Save dialog.	39
22	File Importing dialog.	41
23	Design Property dialog.	41
24	PE Allocation dialog.	47
25	PE Selection dialog.	48
26	Architecture Refinement dialog.	52
27	Network Allocation dialog (Bus tab).	55
28	Bus Selection dialog.	56
29	Network Allocation dialog (CE tab).	57
30	CE Selection dialog.	58
31	Network Allocation dialog (Connectivity tab).	59
32	Network Refinement dialog.	62
33	Bus Parameter dialog.	66
34	Communication Refinement dialog.	68
35	Import Decisions dialog.	71

System-on-Chip Environment (SCE Version 2.2.0 Beta): Manual

L. Cai, A. Gerstlauer, S. Abdi, J. Peng, D. Shin, H. Yu, R. Dömer, D. Gajski

Center for Embedded Computer Systems

University of California, Irvine

Technical Report CECS-TR-03-45

December, 2003

1 Introduction

SCE represents a new technology that allows designers to capture system specification as a composition of C-functions. These are automatically refined into different models required at each step of the design process. Therefore designers can devote more effort to the creative part of designing and the tools can create models for validation and synthesis. The end result is that the designers do not need to learn new system level design languages (SystemC, SpecC, Superlog, etc.) or even the existing Hardware Description Languages (Verilog, VHDL).

Consequently, the designers have to enter only the golden specification of the design and make design decisions interactively in SCE. The models for simulation, synthesis and verification are generated automatically.

SCE contains three system-level synthesis tools: Architecture Explorer, Network Explorer, and Communication Synthesizer. This report is the manual of these three tools.

1.1 Architecture Explorer

Architecture Explorer (AE) is a tool that provides users a highly productive environment to enable rapid and extensive architecture exploration. The task of architecture exploration is to select the most cost-effective system architecture and mapping of specification to the architecture to satisfy all design constraints.

AE is the first tool in the SCE design flow. AE starts with the specification model captured by the designer. It then allows the users to specify the computation architecture by selecting system components (software and hardware processors, memories) from the database of processing elements (PEs). Following this, the designer can perform hardware-software co-design by partitioning the behaviors and variables to selected PEs and memo-

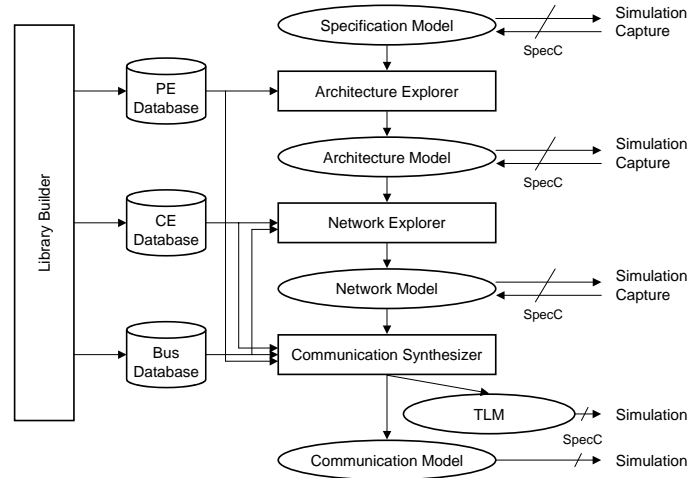


Figure 1: SCE environment.

ries in order to implement the computational part of the specification on the defined computation architecture. AE then generates and exports the new architecture model.

1.2 Network Explorer

Network Explorer (NE) is a tool that provides users a highly productive environment to enable rapid and extensive network exploration. The task of network exploration is to select the most cost-effective system communication topology to satisfy all design constraints.

NE is the second tool in the SCE design flow. NE starts with the architecture model produced by the Architecture Explorer. It then allows the users to specify the network topology by selecting communication components (buses and communication elements (CE)) from the databases of busses and CEs and by connecting components in the architecture (busses, PEs, memories, and CEs). NE then generates and exports the new network model.

1.3 Communication Synthesizer

Communication Synthesizer (CS) is a tool that provides users a highly productive environment to enable rapid and efficient communication synthesis. The task of communication synthesis is to select the detailed parameters of communication on each bus segment and to generate transaction-level or final bus-functional, pin- and timing-accurate communication

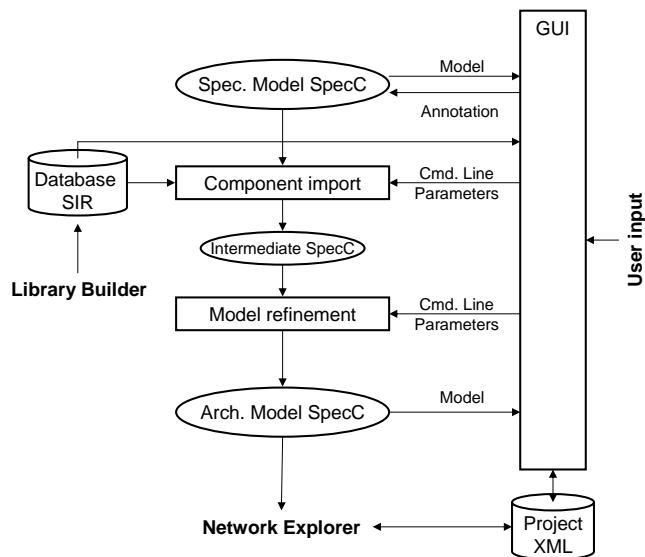


Figure 2: Overview of the software architecture of SCE.

models.

CS is the third tool in the SCE design flow. CS starts with the network model produced by the Network Explorer. It then allows the users to specify the communication details by setting the communication parameters for each bus segment in the system communication architecture. CS then generates and exports the new communication model or transaction-level model.

2 Overview

Figure 2 shows an overview of the software architecture of the Architecture Explorer, Network Explorer, or Communication Synthesizer. AE/NE/CS requires an input model of the intended design in the form of valid SpecC code [1] and conforming to the Specification Model Reference Manual [3]. The input model is the specification model for AE, the architecture model for NE, and the network model for CS. AE reads information about available components from a database of processing elements (PEs). NE reads information about available components from databases of buses and communication elements (CEs). CS reads information about available components from databases of processing elements (PEs), buses, and communication elements (CEs). The database is read in binary form as a file in SpecC Internal Representation (SIR) format as generated by the Library Builder tool that is part of SCE. Through SCE's graphical user interface (GUI), the user can browse the

specification/architecture/network model and PE/CE/bus database, and the user is provided with a GUI to enter design decisions to come up with different design decisions. For AE, the decisions are computation architectures and mappings of the specification onto those architectures. For NE, the decisions are communication topologies. For CS, the decisions are bus parameters.

For each candidate design decision as determined by the user, the AE/NE/CS automatically generates an output model that exactly reflects the design decisions made. The output model is the architecture model for AE, the network model for NE, and the communication/transaction-level model for CS. The output models are generated in the form of SpecC code following the Architecture Model Reference Manual, the Network Model Reference Manual, and the Communication Model Reference Manual. Information about design models and their relationships are tracked by the SCE and can be stored in project files that can be read and written by SCE in a custom XML format. Generated architecture models along with the corresponding project file(s) can then be passed to the following tools in the SCE design flow.

Internally, the SCE consists of separate components for component import, architecture/network/communication refinement, and graphical user interface (GUI). The main SCE application is the SCE GUI which in turn calls import and refinement components as needed. Component import and refinement are command line tools that are called and executed by the GUI where the GUI supplies the correct command line parameters, captures the output and handles (normal or abnormal) results. The GUI reads and displays design models, lets the user browse the database, and provides facilities to enter design decisions. Design decisions are stored by the GUI as annotations in the input design model. For generation of output design models, the GUI first passes the annotated design to the component import tool which, as requested by the GUI, imports the necessary component models out of the database into the design and writes an intermediate design model. The GUI then calls the refinement tool on the intermediate design model. The refinement tool in turn generates the final output architecture model which can be read, displayed and browsed through the GUI.

The run time of automatic model generation is in the order of minutes. Therefore with SCE, users can rapidly experiment with dozens, even hundreds of alternative design decisions to make the optimal design decisions.

2.1 Modules

The block diagram and flow chart of the SCE is displayed in Figure 3. The blocks represent SCE modules where each block is associated with one or more GUI elements. GUI elements and hence SCE blocks can be further classified into three types: the Main Window, input dialogs and display windows.

Besides above three types of windows, there are two types of arrows in Figure 3: solid

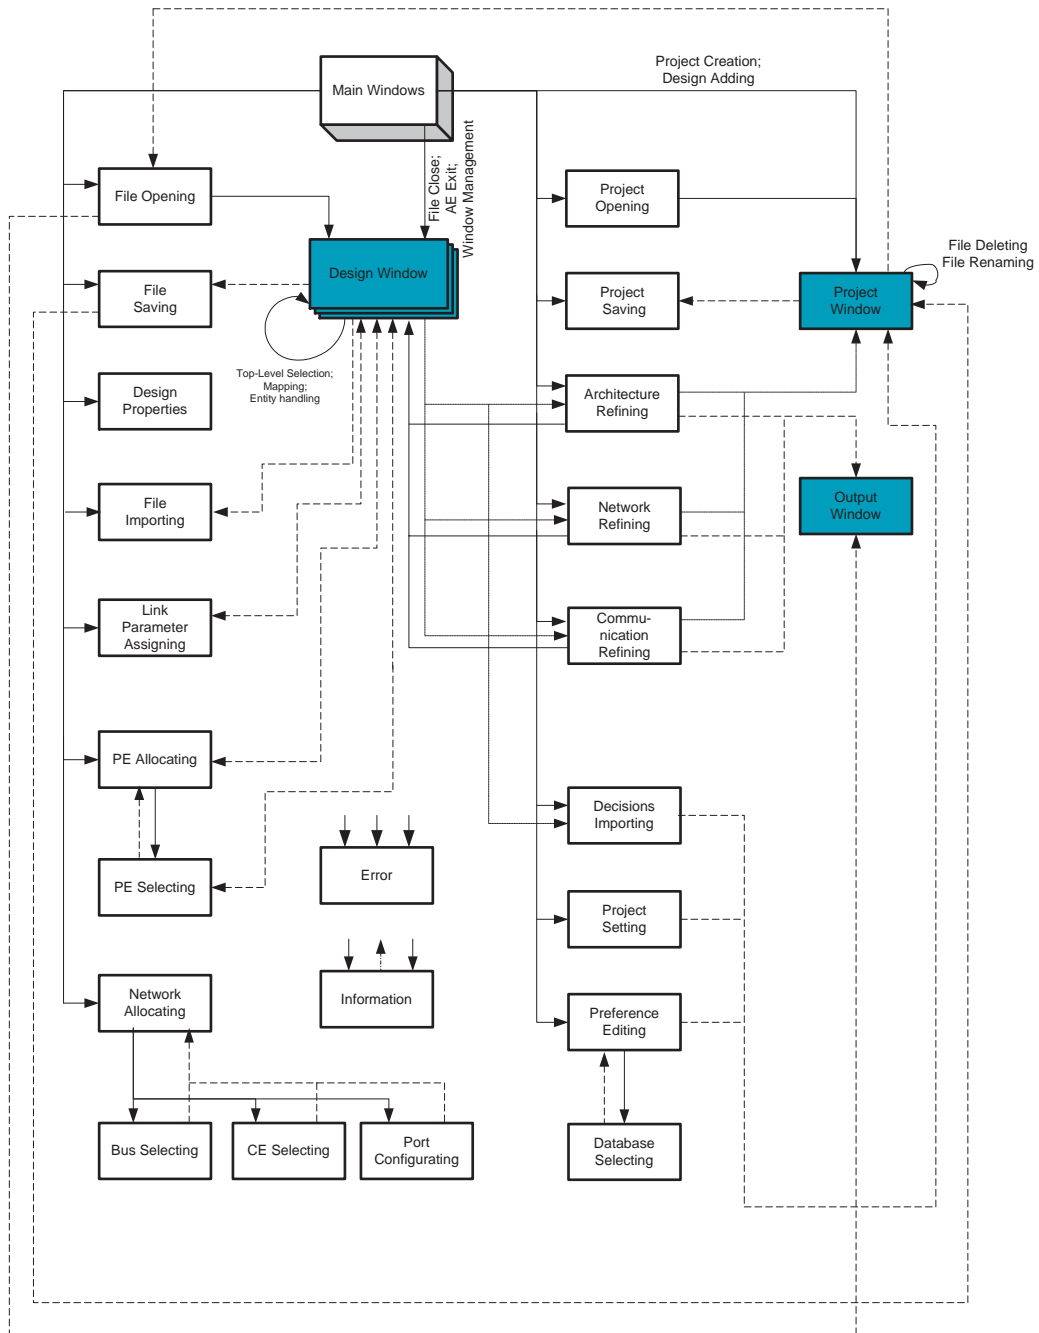


Figure 3: Module diagram of SCE.

arrows and dashed arrows. Solid arrows represent the pop-up relations. The window at the arrow head will be triggered by clicking a button or selecting a menu item in the window at the arrow tail. For example, by clicking *Open* button in *File* column in the menu of window Main Window, users will trigger *File Open* dialog. The text that appears on some arrows represents the arrow's task name. On the other hand, dashed arrows represent information updates. The data in the window at the arrow head will be updated by clicking a button in the window at the arrow tail. In order to simplify Figure 3, the arrows for error and information dialogs are not connected. They are popped up when the error or information occurs, which will be described in Section 4.

2.1.1 Main Window

The main window (shown as 3-D block in Figure 3), is the default, top-level window of the GUI and hence the SCE application. The main window contains a menu bar and a status bar through which the application can be controlled by the user and through which feedback is provided to the user about the state of the application.

2.1.2 Input Dialogs

Input dialogs (shown as regular blocks) allow users to input the design or project decisions by selecting, editing, or typing the details. They are all pop-up dialogs. Users directly/indirectly pop up all the input dialogs by clicking buttons or selecting menu items. Input dialogs include:

- (a) File Open dialog for selecting the design file which users want to open.
- (b) File Save dialog for saving current file as a new file with the specified name.
- (c) Design Property dialog for displaying the information related to the process of SCE, such as command history and processing status of SCE.
- (d) File Import dialog for selecting the design file which users want to import.
- (e) PE Allocation dialog for adding PEs to and removing PEs from the design.
- (f) PE Selection dialog for selecting PEs from the PE library (for input dialog PE Allocating).
- (g) Network Allocation dialog for defining the network topology of the design.
- (h) Bus Selection dialog for selecting busses from the bus database.
- (i) CE Selection dialog for selecting communication elements (CE) from the CE database.

- (j) Port Configuration dialog for editing ports of behaviors.
- (k) Bus Parameter dialog for assigning parameters for busses in the design.
- (l) Error dialogs for displaying of error messages of design tasks.
- (m) Information dialogs for providing informational feedback to the user.
- (n) Project Open dialog for reading a new project from disk.
- (o) Project Save dialog for saving current project as a new project on disk with the specified name.
- (p) Network Refinement dialog for selecting sub-tasks for network refinement.
- (q) Architecture Refinement dialog for selecting sub-tasks for architecture refinement.
- (r) Communication Refinement dialog for selecting sub-tasks for communication refinement.
- (s) Decision Import dialog for import the design decisions from other opened designs.
- (t) Project Settings dialog for setting project preferences.
- (u) Edit Preferences dialog for setting SCE preferences.
- (v) Database Selection dialog for selecting the preferred database.

2.1.3 Display Windows

Display windows (shown as shaded windows in Figure 3) graphically display the information of files, projects, and the process status of Architecture Explorer. Display windows are sub-windows under the main window. Display windows include:

- (a) Design Window displays the contents and the attributes of the designs saved in the opened files. The displayed information includes hierarchy of behaviors, execution sequence of behaviors, and variable/port details of the selected behavior.
- (b) Project Window displays the project information such as hierarchy of design models and included source files.
- (c) Output Window displays the captured output of the SCE command line tools called by the GUI.

2.2 Requirements

The SCE requires a system with the following runtime platform:

Host machine Intel-compatible x86 PC, 500MHz or higher, recommended minimum 128MB of RAM, 200MB of free hard disk space.

Operating system RedHat Enterprise Linux WS, Version 3

In order to compile the source code of SCE, the following additional software packages have to be installed (in binary form including necessary header files):

- SIR library, Version 2.2.x (UC Irvine)
- Qt library and toolkit, Version 3.3.x (Trolltech, Inc.)
- PyQt library, Version 3.11 or higher (Riverbank Computing, Ltd.)

Note that in order to redistribute the libraries together with the compiled SCE, commercial licenses of the above tools have to be obtained as necessary.

2.3 Interfaces

SCE interfaces can be separated into internal interfaces for information exchange between SCE components and external interfaces for information exchange with other tools both within SCE and outside of SCE.

2.3.1 Internal Interfaces

Components inside SCE exchange information through design models, command line parameters, logging output, and exit codes.

Design models Design models are exchanged between SCE components in the form of SpecC files stored on any file system supported by the underlying operating system (Linux). In addition to the SpecC code for the design models themselves, SCE components exchange information via annotations for design decisions and design meta-information.

Command line parameters When calling command line tools, the SCE GUI will pass information for controlling the tool in the form of command line parameters.

Logging output Command line tools will produce logging output during execution. This logging output is captured by the GUI and displayed to the user in the GUI's output window.

Exit codes Command line tools signal status information (success or error codes in the case of tool failure) to the GUI in the form of their exit codes. The GUI analyzes command line tool exit codes and translates them into necessary information or error messages.

2.3.2 External Interfaces

The SCE exchanges information with other tools and with the designer via design models, databases, project files and via SCE's graphical user interface:

Design models Design models are exchanged between SCE's tools in the form of SpecC source code stored as files in any file system supported by the underlying operating system (Linux). SpecC source files are stored as text files in DOS or Unix end-of-line format where SCE will be able to read both formats and to export files in DOS form. Generally, SpecC source code is stored in plain ASCII format. However, SCE will be able to transparently handle Kanji-encoded comments and strings.

SCE will generally be able to import any valid SpecC code that is parsable according to the syntax and grammar of SpecC 2.0 (based on standard ANSI C) as defined in the SpecC Language Reference Manual (LRM), Version 2.0 [1]. Note that inside SCE no object code or executables are ever created and therefore models imported into SCE can include foreign code that depends on libraries outside (i.e. does not have to be linkable on) the underlying host platform (Linux). In contrast, since SCE will preprocess the SpecC source files on the host platform and inline any included code in its exported models, pre-prepared, clean header files with all specific code to be included have to be supplied together with the SCE import SpecC models as necessary for external tools.

In addition, individual models imported into SCE will have to conform to the specific rules and guidelines defined in the specification documents for each type of model.

Databases Databases are generated by the library builder and stored as a collection of binary SpecC Internal Representation (SIR) files managed via a specific file system hierarchy on top of the general underlying file system. SCE components then read component models from these SIR database files.

Project files Project files are stored as XML files in any file system supported by the underlying operating system (Linux). Project XML files are text files using Unix text file format. Project files use a custom XML format that is common to all tools in the SCE environment, i.e. the project XML file format is shared among the SCE tools and any SCE tool is able to read, write and modify project files generated by or used as input to any other SCE tool. For the XML schema describing the SCE project file format see Appendix B.

User interface All user input is entered in SCE through a graphical user interface (GUI). The SCE GUI is built on top of the X11 windowing system and as such can be run on any local or remote X window server.

2.4 Performance

The SCE will guarantee that for a typical design with less than 10,000 lines of code, less than 10 PEs, and less than 100 behaviors, variables and channels, automatic generation of the refined model will take less than 5 minutes.

3 Windows/GUI

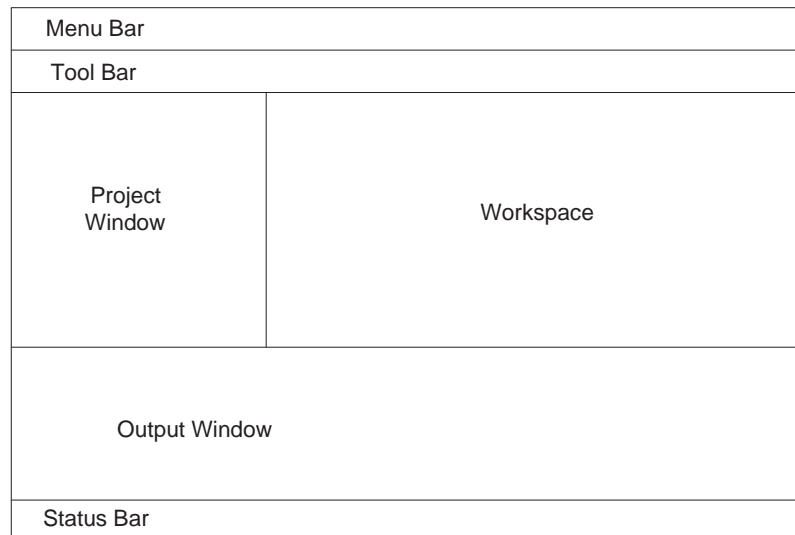


Figure 4: Main Window of SCE.

The primary GUI of SCE is the Main Window, which is displayed in Figure 4. The Main Window consists of six parts:

- (a) A Menu Bar that contains four columns of commands. Each column is a drop-down menu (see Section 3.1).
- (b) A Tool Bar that contains a list of short-cut icons. Each icon represents a command in the menu bar.
- (c) A Project Window (see Section 3.2).

- (d) A Workspace that contains a number of opened Design Windows (see Section 3.3).
- (e) An Output Window (see Section 3.4).
- (f) A Status Bar that displays the current status of SCE, such as “Loading...” or “Ready”.

In this section, we introduce organization-related and display-related details of Menu Bar, Project Window, Display Windows, and Output Window. Some windows contain drop-down menus or pop-up menus. The menus further contain design commands. The usage and functionality behind the commands will be described later in Section 4.

3.1 Menu Bar

The Menu Bar contains five main menu entries: *File*, *View*, *Project*, *Synthesis*, and *Windows*. Each main menu entry is a drop-down menu which contains a number of commands. In general, unless otherwise noted, selecting a main menu entry will apply the corresponding action to the currently active design, i.e. to the design window in the workspace that currently has the input focus. If there is no currently active design window, menu commands will silently fail (do nothing).

3.1.1 File Menu

The *File* menu contains eight commands:

Open Selecting *Open* will allow loading and opening of an existing design file (see Section 4.3.1).

Close Selecting *Close* will close the currently active design (see Section 4.3.3).

Save Selecting *Save* will save the currently active design file (see Section 4.3.2).

Save As Selecting *Save As* will save the currently active design as a new file (see Section 4.3.2).

Save All Selecting *Save All* will save all the opened files (see Section 4.3.2).

Import Selecting *Import* will import a design file into the currently active design (see Section 4.3.4).

Properties Selecting *Properties* will display the properties of the currently active design (see Section 4.3.5).

Exit Selecting *Exit* will exit from and quit SCE (see Section 4.3.6).

3.1.2 View Menu

The *View* menu contains three commands:

Show Variables Selecting *Show Variables* will toggle displaying of variables in the currently active design window (see Section 4.4.4).

Show Channels Selecting *Show Channels* will toggle displaying of channels in the currently active design window. (see Section 4.4.4)

Hierarchy Selecting *Hierarchy* will pop up a submenu that allows to choose between different modes for the amount of detail displayed in the design hierarchy in the currently active design window (see Section 4.4.4).

Preferences Selecting *Preferences* will allow viewing and modifying of application preferences (see Section 4.1.1).

3.1.3 Project Menu

The *Project* menu contains seven commands:

New Selecting *New* will create a new project and open it (see Section 4.2.1).

Open Selecting *Open* will open an existing project file (see Section 4.2.2).

Close Selecting *Close* will close the current project (see Section 4.2.10).

Save Selecting *Save* will save the current project file (see Section 4.2.3).

Save As Selecting *Save As* will save the current project as a new project file (see Section 4.2.3).

Settings Selecting *Settings* will allow viewing and modifying of project settings (see Section 4.2.4).

3.1.4 Synthesis Menu

The *Synthesis* menu contains three commands:

Allocate PE Selecting *Allocate PE* will allow allocation and selection of PEs/memories from the PE database (see Section 4.5.1).

Allocate Network Selecting *Allocate Network* will allow users to define the communication topology (see Section 4.5.2).

Assign Bus Parameters Selecting *Assign Bus Parameters* will allow users to select the parameters for each bus segment in the design (see Section 4.5.3).

Import Decisions Selecting *Import Design* will allow importing of design decisions from another opened design (see Section 4.5.4).

Architecture Refinement Selecting *Architecture Refinement* will perform architecture refinement to generate the refined architecture model (see Section 4.5.1).

Network Refinement Selecting *Network Refinement* will perform network refinement to generate the refined network model (see Section 4.5.2).

Communication Refinement Selecting *Communication Refinement* will perform communication refinement to generate the refined transaction-level or communication model (see Section 4.5.3).

3.1.5 Windows

The *Windows* menu contains six commands:

Close Selecting *Close* will close the currently active design window in the Workspace.

Close All Selecting *Close All* will close all design windows in the Workspace.

Next Selecting *Next* will switch focus to and raise the next design window in the Workspace.

Previous Selecting *Previous* will switch focus to and raise the previous design window in the Workspace.

Tile Selecting *Tile* will rearrange the design windows in the Workspace in a tiled style.

Cascade Selecting *Cascade* will rearrange the design windows in the Workspace in a cascaded style.

Project Manager Selecting/unselecting *Project Manager* will display or undisplay the Project Window.

Output Window Selecting/unselecting *Output Window* will display or undisplay the Output Window.

At the bottom of the *Windows* menu, the names of all opened design windows will be listed. Selecting the name of a design will switch focus to and raise (bring to the front) the corresponding design window.

Details about the usage and functionality of these commands are introduced in Section 4.6.

3.2 Project Window

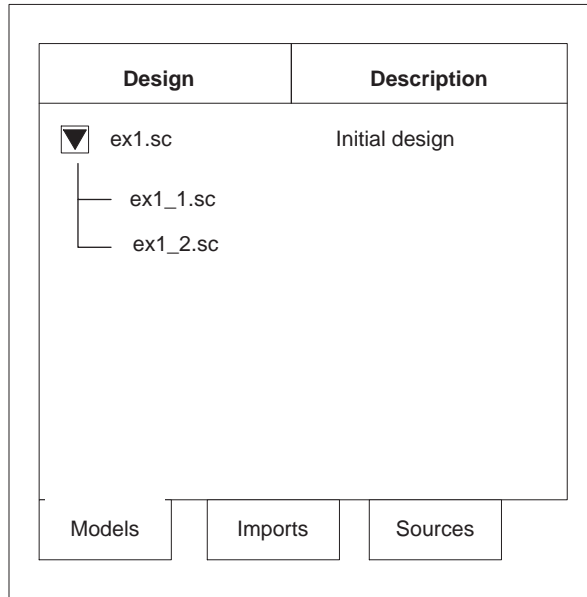


Figure 5: Project Window (Models tab).

The Project Window is a sub-window of the Main Window, displaying project information. In general, a project holds meta-information about a set of design files and their relationship, e.g. a parent-child relationship in case a design was generated from another design through refinement. Hence, each file represents a design of the project at one abstraction level. In addition, the project can hold project-specific settings for the design environment, such as the compiling and parsing environment (paths).

At any time, the SCE application can keep one project open and active in memory. The name of this currently active project is displayed in the title bar of the Main Window. The Project Window displays the hierarchical information of design files in the currently active project. If there is no active project, the Project Window is disabled.

The Project Window can be detached or docked. Users can drag the window (by its title bar or handle) to the desired place. If the Project Window is detached, it can be floating and displayed anywhere on the desktop. If the Project Window is docked, it has to be attached to any of the borders of the Main Window.

The Project Window contains three tabs: *Models*, *Imports*, and *Sources*. By clicking the tab at the bottom of the window, the corresponding tab will be activated and brought to the front.

3.2.1 Models Tab

Figure 5 shows the screenshot of the Models tab. The Models tab contains two columns. Column *Design* displays the name of each design file in the project. If one design is derived or generated from another, the derived design is displayed as the child design of the previous one. The derivation hierarchy is indicated through connecting lines of a tree structure.

Users can select a design by clicking the row in which the design is displayed. Double-clicking on a row will open the corresponding design. If the design is currently opened and loaded, double-clicking will activate (switch focus to and raise to the front) the corresponding design window in the Workspace (see Section 3.5).

The Column *Description* shows an optional description of the design. By clicking into the column, the user can edit the description text of the selected design directly in the corresponding table cell.

Right-clicking on a row will open a context menu pop-up for the selected design. The context menu contains four commands:

Open Selecting *Open* will open the selected file (see Section 4.3.1). If the selected design is already opened, the corresponding design window in the workspace will be activated.

Delete Selecting *Delete* will remove the selected file from the project and optionally delete the file on disk (see Section 4.2.7).

Rename Selecting *Rename* will rename the selected file (see Section 4.2.8).

Change Description *Change Description* will trigger editing of the description of the selected file. (see Section 4.2.9)

Note that renaming and changing of description actions can also be triggered by clicking into the corresponding column of the selected file.

3.2.2 Imports Tab

Figure 6 shows the screenshot of Imports tab. The Imports tab displays a lists of imported design names. The list contains the union of all the sub-designs that have been imported directly or indirectly into any of the design files that are part of the project.

Users can select an imported design by clicking the row in which the imported design is displayed. Double-clicking on an imported design opens the corresponding design in the workspace.

Right-clicking on an imported design opens a context-menu pop-up for the selected design. The context menu contains one command:

Open Selecting *Open* will open the import design file (see Section 4.3.1).

Note that the opening action is equivalent to double-clicking on the imported design.

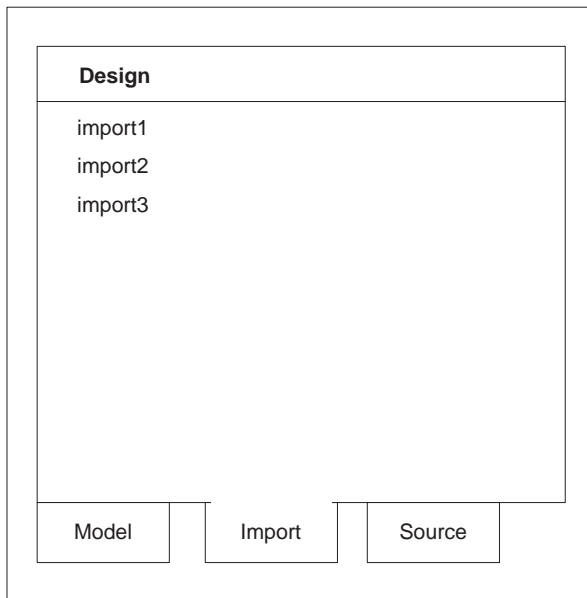


Figure 6: Project Window (Imports tab).

3.2.3 Sources Tab

Figure 7 shows the screenshot of Sources tab. The Sources tab displays a lists of names of source files. The list contains the union of all SpecC source files that are sources for the design files that are part of the project.

The Sources tab does not allow any action on the files and is for informative purposes only.

3.3 Design Window

The Design Window displays the content and the attributes of an opened design, and it allows browsing and navigation of the design hierarchy. The screenshot of the Design Window is displayed in Figure 8.

The Design Window consists of two parts: the side bar and the view pane. The side bar displays the basic information for navigation of the structure of the design. It further consists of three tabs: *Hierarchy*, *Behaviors*, and *Channels*. The view pane displays the detailed information for the entity that has been selected in the side bar.

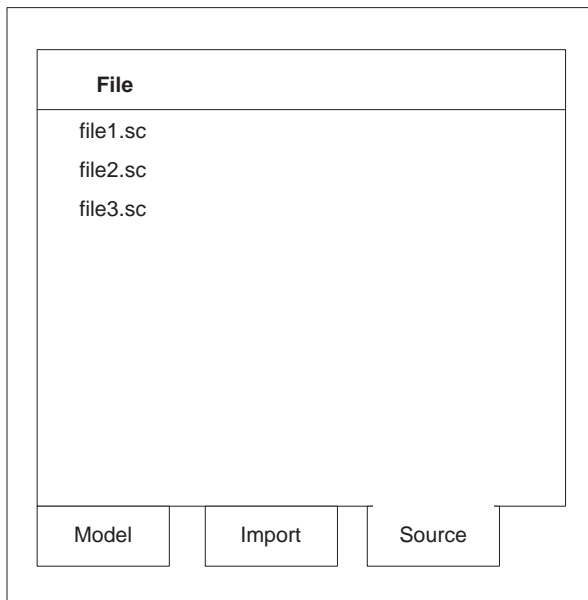


Figure 7: Project Window (Sources tab)

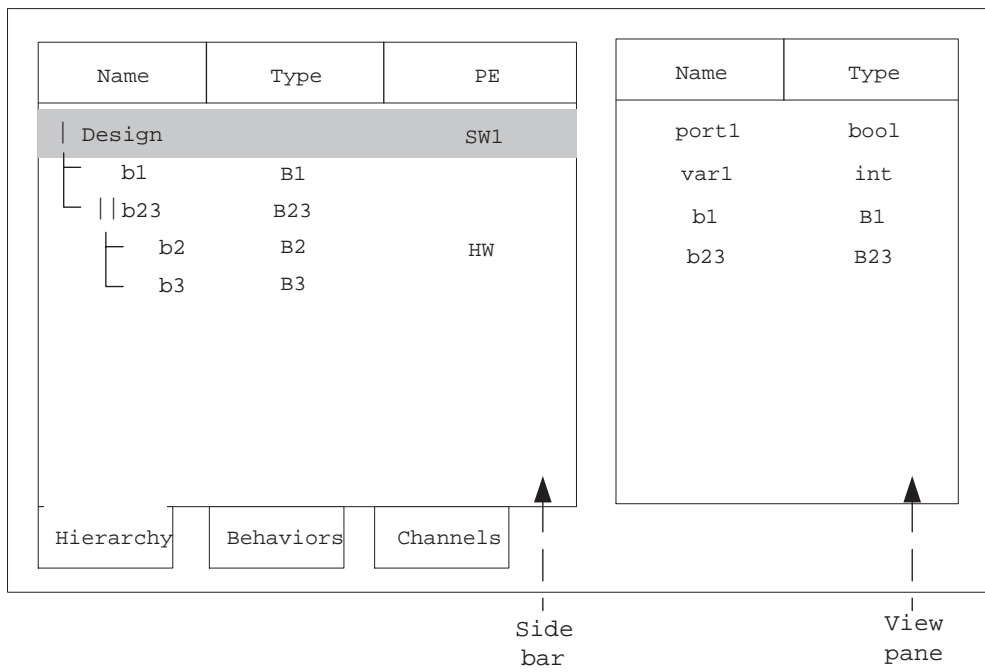


Figure 8: Design Window (Hierarchy tab).

3.3.1 Hierarchy Tab

The Hierarchy tab in the side bar are illustrated in Figure 8. In the column *Name*, the design hierarchy of behaviors, channels (optional) and variables (optional) is displayed in a tree form. The hierarchy tree indicates the entity type (variable, channel or sequential, parallel, FSM, or leaf behavior) visually through icons in the *Name* column. Entities are sorted according to their calling order in the design, i.e. according to their execution order in case of sequential or FSM behavior compositions. Displaying of channels and variables in the Design Window Sidebar Hierarchy tab can be toggled via the View menu and the toolbar (see Section 4.4.4).

At the root of the hierarchy tree in the Hierarchy tab, behavior (and optionally channel) types/classes are listed. At lower levels of the hierarchy, sub-behavior (and optionally channel and/or variable) instances inside the respective parent behavior (or channel) are listed. For any but the roots of the hierarchy tree, the column *Name* shows the name and the column *Type* shows the type of the respective design instances. For classes at the root of the tree, the *Name* column shows the name of the respective class (i.e. the type) and the *Type* column is empty. For example (Figure 8), behavior class *Design* contains two child behavior instances *b1* (of type *B1*) and *b23* (of type *B23*), which are executed sequentially identified by symbol |.

If PE allocation information is available, the Hierarchy tab contains a column *PE* which shows the PE mapping information. If no PE allocation information is available, the *PE* column is not shown. If the PE mapping is empty, this implies that the behavior in this row is mapped to the same PE to which its parent is mapped.

Right-clicking on a design entity row in the Hierarchy tab opens a context menu pop-up for the selected entity. The context menu contains four commands:

Rename Selecting *Rename* will rename the selected entity (see Section 4.4.1).

Delete Selecting *Delete* will delete the selected entity (see Section 4.4.3).

Change Type Selecting *Change Type* will allow plug-and-play to change the type of the selected entity. *Change Type* is not shown for the roots of the hierarchy tree.

Set as Top Level Selecting *Set as Top Level* will set the selected behavior as the top level design behavior (see Section 4.5.1).

3.3.2 Behaviors Tab

The Behaviors tab in the sidebar (Figure 9) lists all behavior types/classes in the design. The name of each behavior type/class is shown in a column *Name*. Behavior types are sorted by name and the sort order can be toggled by clicking into the *Name* column header.

Name	Type
B1	
B2	
B23	
B3	
Design	

Name	Type
port1	bool
var1	int
b1	B1
b23	B23

Hierarchy

Behaviors

Channels

Figure 9: Design Window (Behaviors tab)

Right-clicking on a behavior row in the Behaviors tab opens a context-menu pop-up for the selected behavior class. The context menu of the *Behaviors* tab contains three commands:

Rename Selecting *Rename* will rename the selected entity (see Section 4.4.1).

Delete Selecting *Delete* will delete the selected entity (see Section 4.4.3).

Set as Top Level Selecting *Set as Top Level* will set the selected behavior as the top level design behavior (see Section 4.5.1).

3.3.3 Channels Tab

Name	Type
Channel1	
Channel2	

Name	Type
var1	int
var2	bool

Hierarchy Behaviors Channels

Figure 10: Design Window (Channels tab).

In the Channels tab (Figure 10), all the channel types/classes in the design are listed. The name of each channel type/class is shown in a column *Name*. Channel types are sorted by name and the sort order can be toggled by clicking into the *Name* column header.

Right-clicking on a channel row in the Channels tab opens a context-menu pop-up for the selected channel class. The context menu of the *Channels* tab contains two commands:

Rename Selecting *Rename* will rename the selected entity (see Section 4.4.1).

Delete Selecting *Delete* will delete the selected entity (see Section 4.4.3).

3.3.4 View Pane

When clicking on a design entity row in any of the sidebar tabs, the corresponding row is selected and details of the selected entity are displayed in the view pane. The details shown are the contents of the respective entity and they include contained sub-entities including ports, methods, and variable, channel and child behavior instances. Both the name and the type of sub-entities are displayed in the *Name* and *Type* columns of the view pane, respectively. Elements of the view pane list are sorted by class (port, variable, behavior, channel). Within each class and among classes, the sort order can be set by clicking on the *Name* or *Type* column headers.

3.4 Output Window

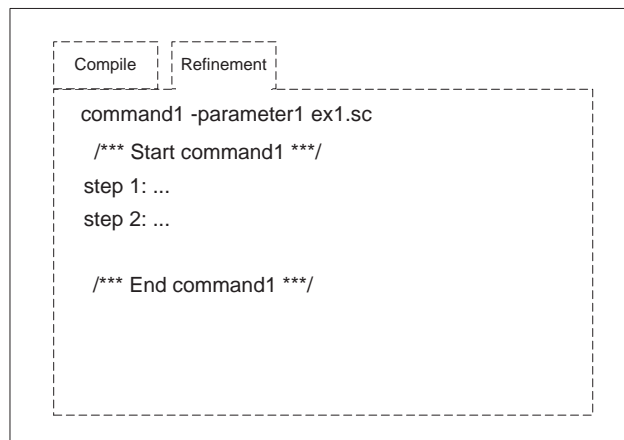


Figure 11: Output Window

The Output Window displays the information related to the process of SCE, such as logged status, diagnostic and error output of background commands. The screenshot of Output Window is displayed in Figure 11. The Output Window contains two tabs: *Compile* and *Refinement*. The *Compile* tab displays the log messages generated during preprocessing and parsing of SpecC code when opening, loading and importing design files. The *Refinement* tab displays the log messages generated by the command line tools spawned by the main application GUI during design refinement. The Output Window is for informational purposes only and doesn't contain any button, box or context menu that users can click or edit.

The Output Window can be detached or docked. Users can drag the window (by its title bar or handle) to the desired place. If the Output Window is detached, it can be floating and displayed anywhere on the desktop. If the Output Window is docked, it has to be attached to any of the borders of the Main Window.

3.5 Workspace

In general, multiple designs can be open and loaded in the SCE application. The design windows for all currently opened and loaded designs are shown in the Workspace. Within the Workspace, design windows can be minimized, maximized, resized and closed freely via their title bar, title bar icons and handles on their window frames. Closing a design window closes the corresponding design (file).

At any time, there is exactly one active design window in the Workspace. The active window is the one that has the input focus and it is visualized by highlighting its title bar. Unless otherwise noted, all menu, toolbar or other commands apply to the currently active design window. Clicking into a design window activates the corresponding window and raises it to the front of the Workspace. A newly opened design windows automatically becomes the active window.

3.6 Message Boxes

As a result of certain actions, the SCE application will pop up message box dialogs for feedback to or input from the user about handling of special situations. Message boxes are used to provide informative messages and to ask simple questions. In general, there are two types of message boxes: error dialogs and information dialogs.

3.6.1 Error Dialogs

If the application encounters an abnormal error situation in which user notification about the failure of the initiated action is required, an Error dialog will be popped up (Figure 12). The Error dialog displays an error message at the top-half of the Error dialog. At the bottom-half, an Error dialog contains one button: *Ok*. Clicking *Ok* will close the Error dialog and original dialog (if any) that prompted the message. After clicking, the original action that prompted the message is aborted and cancelled.

3.6.2 Information Dialogs

If the application encounters an abnormal situation in which user notification is required and the user is given several choices on how to continue, an Information dialog will be popped up (Figure 13). An information message and associated question is displayed at the top-half of the dialog. The bottom-half of the dialog contains three buttons: *Yes*, *No*,




Figure 12: Error dialog

and *Cancel*. Clicking *Yes* will accept the recommendation and do the corresponding action. Clicking *No* will not accept the recommendation and will not do the corresponding action but will continue the original action that prompted the message in the first place. Finally, clicking *Cancel* will not do the recommended action and will also cancel the original action that prompted the message. Clicking one of above three buttons will close the Information dialog and original dialog (if any) that prompted the message.




Figure 13: Information dialog

4 Functionality

The functionality of SCE can be classified to six categories: application, project handling, file handling, design-entity handling, architecture exploration, and window management.

In this section, sub-windows or sub-menus described in Section 3 are referred using the following format: *Win* : *Sub*. *Win* refers to display windows:

- *Main* represents the Main Window.
- *Project* represents the Project Window.
- *Design* represents the Design Window.

Sub refers to drop-down menus or sub windows (tabs):

- For the Main Window, *Sub* is either *File*, *View*, *Project*, *Synthesis*, or *Windows* (drop-down menus introduced in Section 3.1).
- For the Project Window, *Sub* is either *Models*, *Imports*, or *Sources* (tabs introduced in Section 3.2).
- For the Design Window, *Sub* is either *Hierarchy*, *Behaviors*, or *Channels* (sidebar tabs introduced in Section 3.3).

For example, *Project :: Models* refers to the Models tab in the Project Window.

Main menu or context menu commands described in Section 3 are referred to using the following format: *Win :: Sub* \Rightarrow *Command* where *Command* refers to a command. For example, *Main::File* \Rightarrow *Open* refers to the *Open* command in the *File* menu of the Main Window menu bar. On the other hand, *Project::Models* \Rightarrow *Open* refers to the *Open* command in the context menu of the Project Window Models tab.

4.1 Application

The main application of SCE supports a set of persistent application settings. Application settings are persistently stored across different invocations of the tool. In fact, application settings are shared among all tools in the SCE environment, i.e. they are persistent across invocation of different tools at different times.

Application settings are stored in both system-wide and user-specific locations. System-wide application settings affect all users of SCE applications on the system. They are stored in a file on disk in a location that is configurable during compile time of SCE. User-specific application settings, on the other hand, are stored in a file in the user's Linux home directory. The application first reads the system-wide and then the user-specific settings, i.e. user-specific settings can override (if given) system-wide settings and if no user-specific settings are given, application settings default to the system-wide settings. If no system-wide settings are available, compiled-in defaults are used.

Application settings in general provide the standard settings (paths, etc.) to use by default for the different parts of SCE applications. Note that application settings can be overwritten or extended by project-specific settings (see Section 4.2). Application settings include:

Compiler settings A set of options for preprocessing and parsing SpecC source files. When opening/loading or importing a design file, the SpecC compiler ('scc', see Appendix A.1) is used internally to compile the SpecC source file into SCE's internal SpecC Internal Representation (SIR) [2] format. Via the compiler settings, the options for preprocessing and parsing passed to the SpecC compiler are specified. Specifically, compiler settings contain the following:

Standard include path An ordered list of directories in which to search for include files during preprocessing.

Standard import path An ordered list of directories in which to search for imported files during parsing.

Macro defines An ordered list of preprocessor macro definitions.

Macro undefines An ordered list of preprocessor macro undefines.

Compiler options Additional compiler switches passed literally to the SpecC compiler. Possible compiler switches are switches for setting warning and verbosity levels.

Database paths Location of the database SIR files for the PE database, the CE database, and the bus database.

All paths in the application settings are relative to the current working directory when starting the application, i.e. relative paths in the settings are converted into absolute paths by prepending the working directory during startup of the application.

In terms of application settings, SCE supports functions to view and edit application settings/preferences.

4.1.1 Preferences Editing

Operation Preference editing allows viewing and setting of the application settings of SCE. Users start editing the preferences of SCE by selecting *Main :: Project ⇒ Preference*. This will pop-up the Edit Preferences dialog, which allow users to browse and specify the compiler and database settings. The Edit Preferences dialog is illustrated in Figure 14. A dashed rectangular represents a tab. A solid rectangular represents an edit box. A rounded rectangular represents a button.

There are two tabs in the Edit Preferences dialog: Compiler and Database. By clicking the tab at top-left corner of the window, users can select either of them for viewing and editing.

- (a) The Compiler tab allows viewing and editing of compiler settings. The screenshot for Compiler tab is shown in Figure 14. The Compiler tab contains line edit boxes for all compiler settings. The text in the *Include Path* and *Import Path* lines defines the directory lists (separated by colons) for the standard include and standard import paths, respectively. The text in the *Defines* and *Undefines* lines define the list of macro defines and undefines (separated by semicolons), respectively. Finally, the text in the *Options* line defines the compiler options/switches.

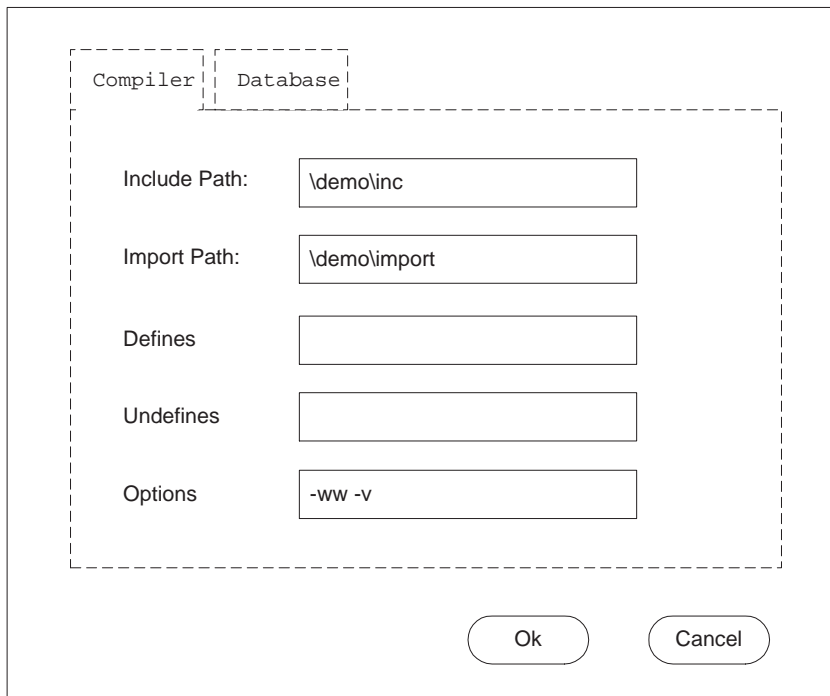


Figure 14: Edit Preferences dialog (Compiler tab).

- (b) The Database tab allows for viewing and selecting of database file paths. The screenshot for Database tab is shown in Figure 15. Users can type in the file names and paths of PE, CE, and Bus databases in *PE Database*, *CE Database*, and *Bus Database* line edit boxes. Besides typing in the databases file names, users can also select the names by using *Select* buttons next to the edit boxes. Clicking *Select* button will pop up a Database Selection dialog displayed in Figure 16

Database Selection dialog allows users to choose and select existing database files on disk to use for each of the three databases. In the Database Selection dialog, the solid rectangular represents the edit box. The dashed rectangular represents the display box. The round rectangulars represents buttons.

In the Database Selection dialog, users should first specify the database directory in *Look-in* box. The content of the directory will be automatically displayed in the display box in the center. The database type in the *File Type* box defaults to SIR files for databases but can be chosen by the user. All the database files with the specified type will be displayed in the display box. Users further type in the database name in *Database Name* box. Finally, by clicking *Open* button, the database with the specified name will be selected. If users click *Cancel* button, then the action of database selection will be cancelled. Either clicking *Open* or *Cancel* button will close the Database Selection dialog.

Buttons *Ok* and *Cancel* appear at the bottom of the Edit Preference dialog. If users click the *Ok* button, all the edited preferences are saved. If users click the *Cancel* button, all the edited preferences are discarded. Either clicking *Ok* or *Cancel* button will close Preference dialog.

4.2 Project Handling

Project handling deals with project issues. Project handling functionality is common and shared across all SCE tools. It allows for tracking of design meta-data over the whole lifetime of a design. A project acts as a unified container that holds all information related to a certain design at various levels of abstraction, i.e. it contains all the information that describes the organization of design files that are part of the project. Furthermore, a project contains project-specific settings that can override or extend application-specific settings (see Section 4.1) for compiler paths, options, etc. Specifically, a project contains the following information:

Design models A tree of design files and their relationship. If a design has been generated from another design through refinement, it is a child of the source design in the tree. For each model, the tree stores the design name, the location of the design's files on disk, the abstraction level, and the command used to generate the model.

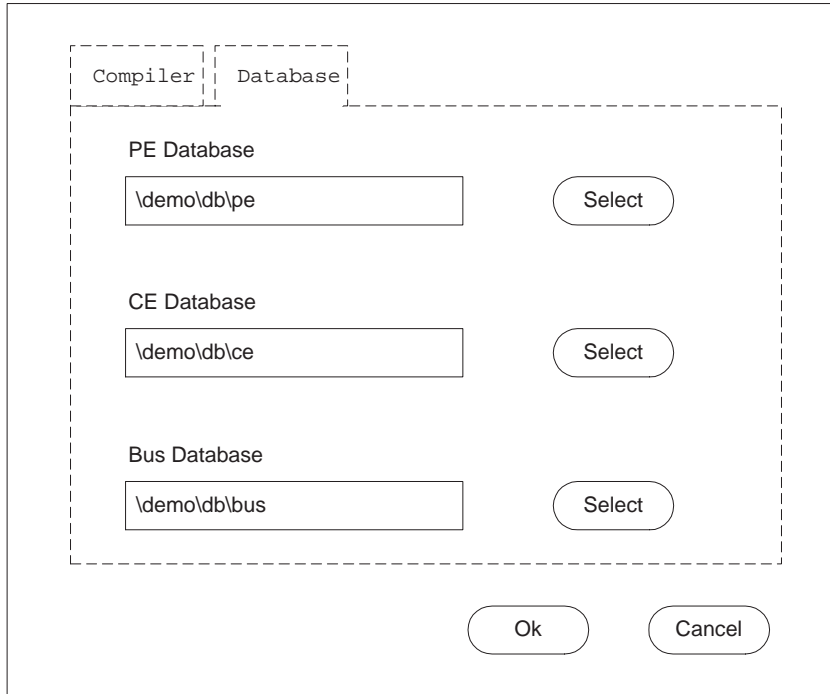


Figure 15: Edit Preferences dialog (Database tab).

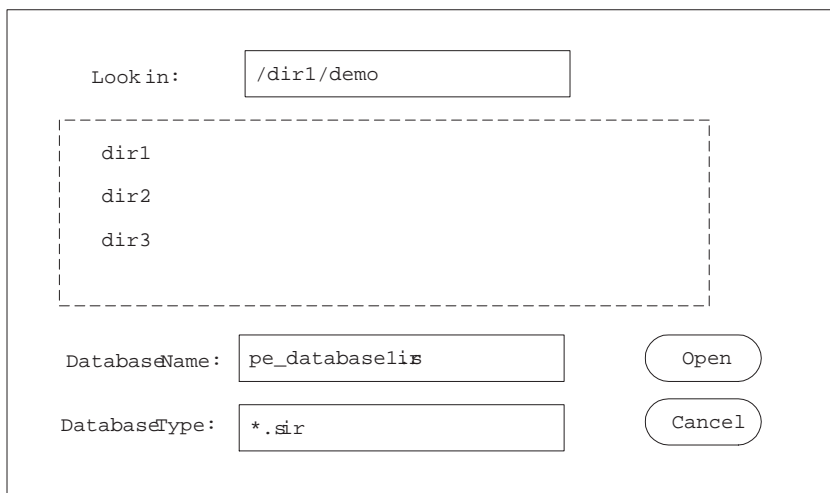


Figure 16: Database Selection dialog.

Imports A list of imported design files. The list of imports contains the union of all (sub-)designs imported by any of the models that are part of the project.

Sources A list of source files. The list of sources contains the union of all SpecC source files from which the models that are part of the project have been compiled. For each source file, the location (path) of the file on disk is stored in the project.

Compiler settings A set of project-specific options for preprocessing and parsing SpecC source files. Compiler settings contain include paths, import paths, compiler options, and macro defines and undefines. Project-specific compiler settings generally overwrite or extend the corresponding application-specific settings. In the case of paths, project paths are prepended to the standard paths defined in the application settings (i.e. they are prepended to the directory search list). In all other cases, options or macro defines/undefined are appended to the compiler command line after the standard options and macros defined in the application settings.

All paths in the project settings are defined to be relative to the location of the project file, i.e. relative paths in a project file are converted into absolute paths by appending the project file's directory during loading/opening of a project file. During saving/writing of project files, absolute paths are in turn converted back to relative paths if they point to a location below the target project file directory.

Projects are stored as text files on disk in a custom XML format. See Appendix B for the XML schema describing the project file format. The project file format is the same for all tools in the SCE environment, i.e. a project file can be read, modified and written by any SCE tool.

Projects can be read from and saved as project files at any time in the SCE application. At any time, however, at maximum only one project can be open and loaded. While a certain project is open and loaded, its settings apply to all actions performed during that time. In addition, certain actions will automatically update and add data in the currently opened and loaded project.

In order to deal with management of projects, SCE supports a set of project handling functions. Specifically, project handling consists of the following functions:

- (a) Project Creation to create a new projects (see Section 4.2.1).
- (b) Project Opening to open and load existing projects from project files on disk (see Section 4.2.2).
- (c) Project Saving to save the current project into a project file (see Section 4.2.3).
- (d) Project Settings to edit the settings of the opened project (see Section 4.2.4).
- (e) Design Adding to adds new design files into the opened project (see Section 4.2.5).

- (f) Design Opening to open a design model or import that is part of the project (see Section 4.2.6).
- (g) Design Deletion to delete a design model file from the project and optionally the disk (see Section 4.2.7).
- (h) Design Renaming to rename a design model and design file in the project and on disk (see Section 4.2.8).
- (i) Description Changing to change the description of a design model in the opened project (see Section 4.2.9).
- (j) Project Closing to close the current project (see Section 4.2.10).

4.2.1 Project Creation

Operation Users can create a new project by selecting *Main :: Project ⇒ New*.

Error/Information Messages Assuming before project creation, users have opened another project in SCE, the currently opened project has been modified and the opened project is not saved yet. When users select *Main :: Project ⇒ New*, an Information dialog will be popped up querying the user whether he wants to save the current project first before creating a new one. If the user accepts the recommendation, a Project Saving action (see Section 4.2.3) is performed first.

4.2.2 Project Opening

Operation Users open an existing project by selecting *Main :: Project ⇒ Open*. The selection will pop-up the Project Open dialog in which the user can choose and select an existing project file on disk to open and load. The screenshot of Project Open dialog is shown in Figure 17. The solid rectangular represents the edit box. The dashed rectangular represents the display box. The round rectangulars represent buttons.

Users should first specify the project directory in *Look-in* box. The content of the directory will be automatically displayed in the display box in the center. The file type defaults to project files (with a '.sce' suffix) but users can specify any file type in the *Project Type* box. All the project files with the specified type will be displayed in the display box. Users further select the project name in the *Project Name* box. Finally, by clicking *Open* button, the project with the specified name will be opened. If users click the *Cancel* button, the action of project opening will be cancelled. Either clicking *Open* or *Cancel* button will close the Project Open dialog.

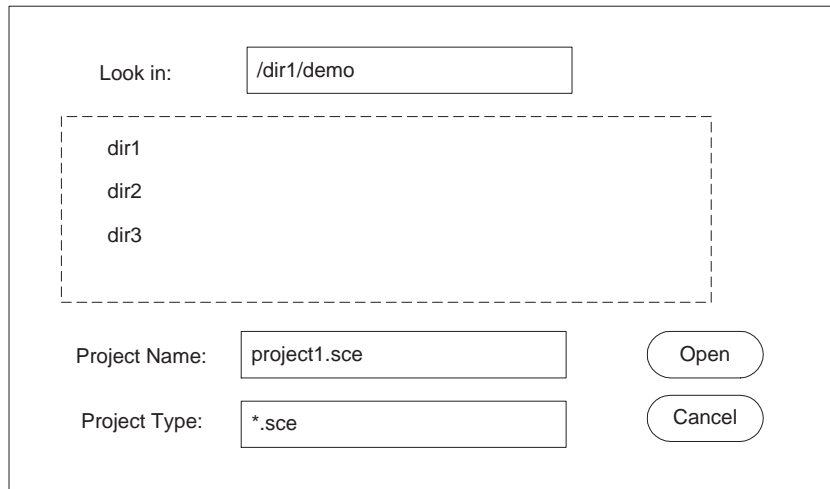


Figure 17: Project Open dialog.

Error/Information Messages If the specified project doesn't exist before clicking *Open* button, then clicking *Open* button has no effect.

In case of errors reading the project file from disk (file errors, wrong file format), an error dialog with a corresponding error message is popped up. Upon confirming the error, the Project Opening action is cancelled.

Assuming before project opening, users have opened another project in SCE, the opened project is modified and the opened project is not saved yet. When users open a different project, the Information dialog will be popped up to recommend users to save the previous project first and, if the recommendation is accepted, a Project Saving action will be performed. This is the same as the case in task Figure 4.2.1.

4.2.3 Project Saving

Operation Users can save the current project by one of the following two methods:

- (a) Selecting *Main :: Project ⇒ Save*. The project will be saved using the current project name.

If the saved project is unnamed (a new project created by task Project Creating), then selecting *Main :: Project ⇒ Save* will do the same action as selecting *Main :: Project ⇒ Save As* (see below).

- (b) Users can save the current project under any (new) name by selecting *Main :: Project ⇒ Save As*. The selection will pop-up the Project Save dialog in which user

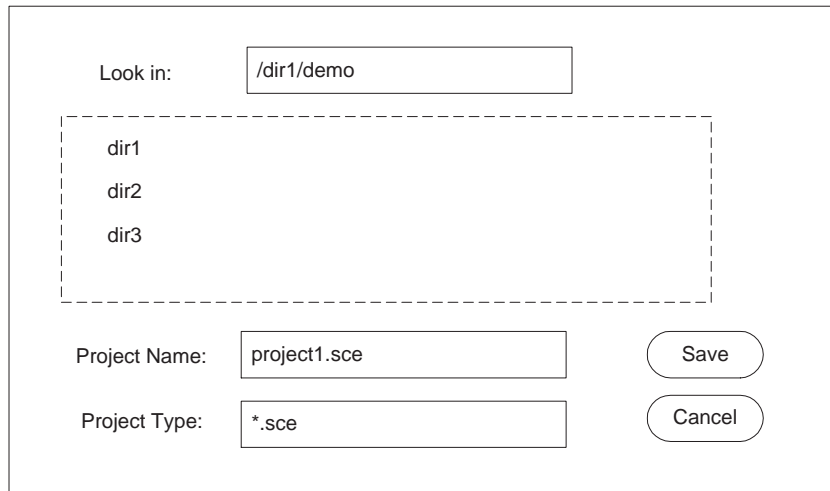


Figure 18: Project Save dialog.

can choose the directory and file name to save the project under. The screenshot of the Project Save dialog is displayed in Figure 18. The solid rectangular represents the edit box. The dashed rectangular represents the display box. The round rectangulars represent the buttons.

In the Project Save dialog, users should first specify the project directory in *Look-in* box. The content of the directory will be automatically displayed in the display box in the center. The file type defaults to project files (‘.sce’ suffix) but users can specify any file type in the *Project Type* box. All the project files with the specified type will be displayed in the display box. Users then select the project name in *Project Name* box. Finally, by clicking the *Save* button, the current project will be saved in a project file with the specified name. If users click *Cancel* button, then the action of project saving will be cancelled. Either clicking *Save* or *Cancel* button will close the Project Save dialog.

Error/Information Messages When selecting *Main :: Project ⇒ Save As* and specifying the file name of an existing file on disk, an Information dialog will pop up asking the user whether he wants to overwrite the existing file. If the user declines this, the Project Saving action will be cancelled.

When selecting *Main :: Project ⇒ Save* or *Main :: Project ⇒ Save As*, errors may occur (file errors, e.g. if no space is available on the hard disk). In this case, an Error dialog as shown in Figure 12 will be popped up, corresponding error messages will be displayed, and the Project Saving action will be cancelled.

4.2.4 Project Settings Editing

Operation Project setting allows users to edit project settings. Unlike application preferences editing in Section 4.1.1, project setting apply only to the current project. Users start project settings editing by selecting *Main :: Project ⇒ Setting*. The selection will pop up the Project Settings dialog, which is displayed in Figure 19. In the Project Settings dialog, user can view and edit the compiler settings stored in the project. The dialog contains line edit boxes for all compiler settings. The text in the *Include Path* and *Import Path* lines defines the directory lists (separated by colons) for the project-specific include and import paths, respectively. The text in the *Defines* and *Undefines* lines define the list of macro defines and undefines (separated by semicolons), respectively. Finally, the text in the *Options* line defines the project's compiler options/switches.

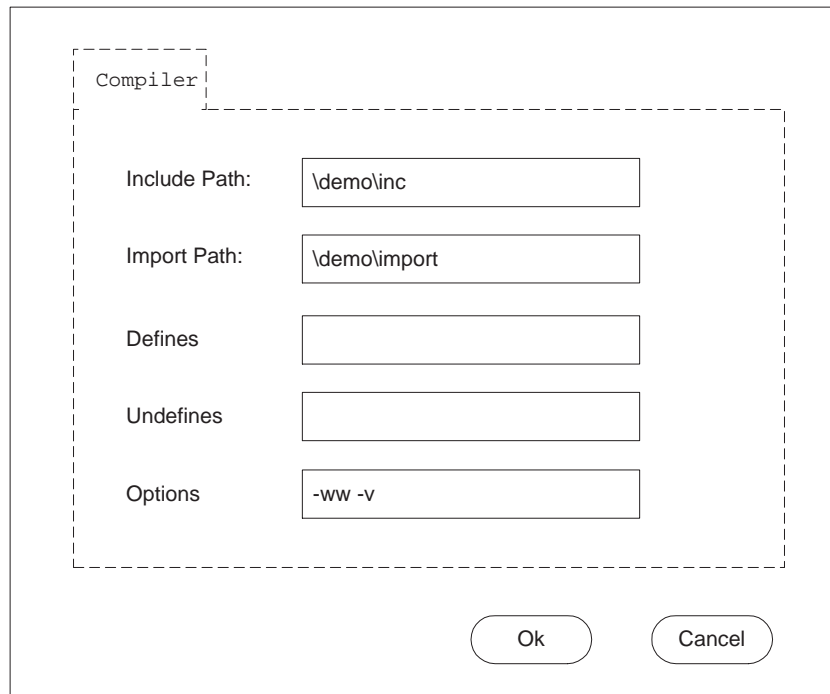


Figure 19: Project Settings dialog.

4.2.5 Design Adding

Operation Users can add any currently opened design file to the project. In order to do that, users select *Main :: Project ⇒ Add Design*. After clicking, the design corresponding

to the currently active design window will be added to the project. The design will be added to the project as a new root in the forest of design model trees. In addition, any imported designs and source files of the new design will be merged into the list of imports and sources that are part of the project.

Error/Information Messages If users try to add a file which is already in the current project, an error dialog with a corresponding error message will be popped up (as illustrated in Figure 12).

4.2.6 Design Opening

Operation Double-clicking on a design model in the Project Window models or import tabs (*Project :: Models* or *Project :: Imports*), a File Opening action on the given design file will be performed (see Section 4.3.1 for details, including error/information messages), i.e. a corresponding design window will be opened in the workspace. If the selected design is already opened in the workspace, its design window will be raised to the top and made active. The Design Opening action can also be triggered via corresponding context menu entries in the model and import tabs (selecting *Project :: Models* \Rightarrow *Open* or *Project :: Imports* \Rightarrow *Open*).

4.2.7 Design Deletion

Operation Users can delete files from the current project and optionally from disk. Selecting the corresponding entry in the context menu of a design model in the Project Window models tab (*Project :: Models* \Rightarrow *Delete*) will delete the selected design from the project and optionally from disk. In case of a model with children in the model tree, the user will also be given the option to recursively delete all the model's children.

Error/Information Messages After selecting *Project :: Models* \Rightarrow *Delete*, an Information dialog will be popped up to query to user whether he wants to also delete the corresponding model's files on disk. If the selected file has children, then another Information dialog querying the user whether he wants to also recursively delete all children and children's children of the selected model. If the selects recursive deletion, an Information dialog similar to the initial Information dialog to query about deletion of corresponding files on disk will pop up for each child model.

If file deletion on disk is selected, an error dialog may pop up in case of disk/file errors. Upon confirmation of the error, the Project Deletion action will be aborted.

4.2.8 Design Renaming

Operation By selecting *Project :: Models* ⇒ *Rename* or by clicking into the *Design* column of the Project Window models tab, users can rename the file name displayed in the *Design* column. Renaming is performed in place inside the column cell itself by opening a corresponding text edit box. Renaming can be aborted by pressing the *Esc* key. Pressing *Enter* accepts the newly entered name and renames the design both in the project and on disk. If the design is loaded and opened, the corresponding design window in the Workspace will also be automatically renamed.

Error/Information Messages If the new design name entered by the user is the name of a design already existing in the project, an Error dialog with a corresponding error message will pop up and, after confirmation, the Design Renaming action will be aborted.

If renaming the file on disk results in an error (file error), a corresponding Error dialog will be popped up and the Design Renaming will be aborted.

4.2.9 Description Changing

Operation Users can change the file description displayed in the *Description* column of *Project :: Models* tab by selecting the *Project :: Models* ⇒ *Change Description* context menu entry or by clicking into the corresponding *Description* column in the row representing the selected file. After clicking, the corresponding cell in the *Description* column of *Project :: Models* is editable in place. Editing can be aborted with *Esc* and is accepted by pressing *Enter*.

4.2.10 Project Closing

Operation Users can close the current project by selecting *Main :: Project* ⇒ *Close*.

Error/Information Messages If the current project is modified and not yet saved, selecting *Close* will pop up an Information dialog which recommends to save the current project first. If the user accepts the recommendation, a Project Saving action (Section 4.2.3) is performed before closing the project.

4.3 File Handling

File handling deals with issues relating to manipulation of design files within SCE. File handling includes opening, saving, and closing of actual design model files on disk. File handling is closely related to Design Windows (Section 3.3) and Design Window Management (Section 4.6). In general, there is a one-to-one association between design models,

design files on disk and design windows in the Workspace. Each Design Window represents a view onto one loaded design file which in turn stores the data of one design model, and vice versa. For example, both File Closing (Section 4.3.3) and Window Closing (Section 4.6) will close the design file and the design window and unload the design from SCE's memory. Specifically, File Handling consists of the following tasks:

- (a) File Opening to open and load existing design files from disk (see Section 4.3.1).
- (b) File Saving to save a design into a design file on disk (see Section 4.3.2).
- (c) File Closing to close a design file (see Section 4.3.3).
- (d) File Import to import an existing design file from disk into a currently opened design (see Section 4.3.4).
- (e) Design Property to display the current design file's properties (see Section 4.3.5).
- (f) SCE Exiting to exit the SCE application (see Section 4.3.6).

4.3.1 File Opening

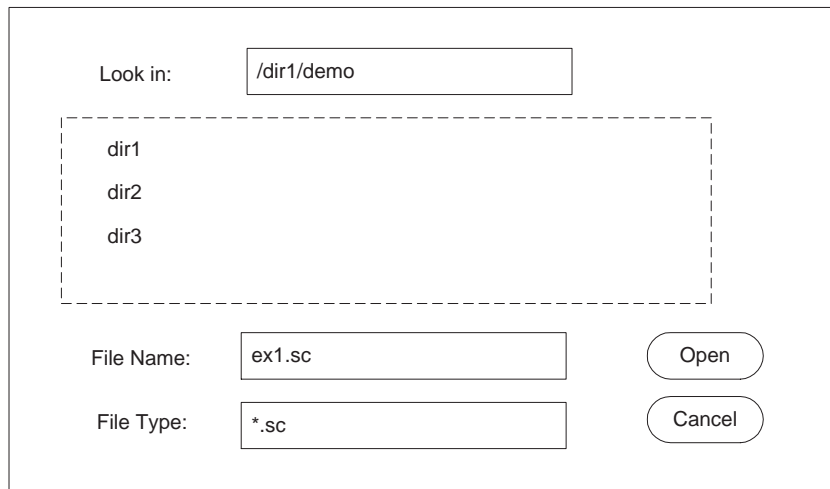


Figure 20: File Open dialog.

Operation Users can open an existing design file on disk in different ways:

- (a) Selecting *Main :: File ⇒ Open* will pop up the File Open dialog in which the user can choose and select an existing file on disk to open and load. The File Open dialog is illustrated in Figure 20. The solid rectangular represents the edit box. The dashed rectangular represents the display box. The round rectangulars represent the buttons.

Users should first specify the directory of the file in *Look-in* box. The content of the directory will be automatically displayed in the display box in the center. The file type defaults to SpecC source files (‘.sc’ suffix) but user can specify any file type in the *File Type* box. All the files with the specified type will be displayed in the display box. Users then further select the file name in the *File Name* box. Finally, by clicking the *Open* button, the file with the specified name will be open. If users click the *Cancel* button, the action of file opening will be cancelled. Either clicking *Open* or *Cancel* button will close the File Open dialog.

- (b) Double-clicking on a design in the Project Window models or imports tabs (*Project :: Models* or *Project :: Imports*) or selecting the corresponding context-menu entries (*Project :: Models ⇒ Open* or *Project :: Imports ⇒ Open*) will open and load the design file for the selected design model.

Opening and loading a design file in either way will result in a corresponding new Design Window popping up in the Workspace. The new Design Window will automatically be made the active one and raised to the front of the Workspace.

If the selected design file is already opened and loaded, it will not be loaded again from disk and File Opening will only result in activating and raising the corresponding Design Window in the Workspace.

Error/Information Messages In the case of selecting from disk, if the specified file doesn’t exist before clicking the *Open* button, then clicking *Open* has no effect.

In case of errors reading the design file from disk (file errors, wrong file format), an error dialog with a corresponding error message is popped up. Upon confirming the error, the File Opening action is cancelled.

4.3.2 File Saving

Operation Users can save opened and loaded design files (Design Windows in the Workspace) by one of the following three methods:

- (a) Selecting *Main :: File ⇒ Save* will save the file of the currently active Design Window using its current name.
- (b) Users can save the file of the currently active Design Window under any (new) name by selecting *Main :: File ⇒ Save As*. The selection will pop up the File Save dialog

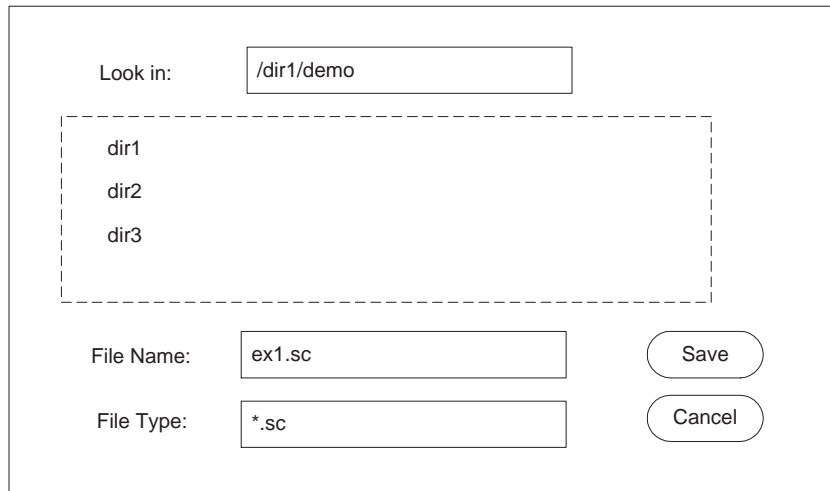


Figure 21: File Save dialog.

in which user can choose the directory and file name to save the design under. The screenshot of the File Save dialog is shown in Figure 21. The solid rectangular represents the edit box. The dashed rectangular represents the display box. The round rectangulars represent the buttons.

In the File Save dialog, users should first specify the directory of the file in *Look-in* box. The content of directory will be automatically displayed in the display box in the center. The file type defaults to SpecC source files (‘.sc’ suffix) but users can specify any file type in the *File Type* box. All the files with the specified type will be displayed in the display box. Users then further select the file name in *File Name* box. Finally, by clicking *Save* button, the current opened file will be saved as the file with the specified name. If users click *Cancel* button, then the action of File Saving will be cancelled. Either clicking *Save* or *Cancel* button will close the File Save dialog.

- (c) Selecting *Main :: File ⇒ Save All* will save the files of all currently opened Design Windows in the Workspace using their current names.

Error/Recommendation Messages. When selecting *Main :: File ⇒ Save As* and specifying the file name of an existing file on disk, an Information dialog will pop up asking the user whether he wants to overwrite the existing file. If the user declines this, the File Saving action will be cancelled.

When selecting *Main :: File ⇒ Save*, *Main :: File ⇒ Save As*, or *Main :: File ⇒ Save All*, errors may occur (file errors, e.g. if no space is available on the disk). In this case, an

Error dialog as shown in Figure 12 will be popped up, corresponding error messages will be displayed, and the File Saving action will be cancelled.

4.3.3 File Closing

Operation Users can close the file and window of the currently active Design Window in the Workspace by selecting *Main :: File ⇒ Close*. Closing a file will unload the design from memory and will close the corresponding Design Window in the Workspace.

Error/Information Messages If the current design is modified and not yet saved, selecting *Close* will pop up an Information dialog which recommends to save the current design first. If the user accepts the recommendation, a File Saving action (Section 4.3.2) is performed before closing the file.

4.3.4 File Import

Operation Users can import an existing design file on disk into a currently opened and loaded design. Importing a design will merge the design's contents into the currently opened design (equivalent to a SpecC `import` statement). All the design entities in the imported files can then be used in the current design. For example, in order to do plug-and-play of behaviors, users can replace a behavior in a design with an imported compatible behavior (see Section 4.4.2 for the Changing Type action).

In order to import a file users should first active the target Design Window in the Workspace and then select *Main :: File ⇒ Import*. The selection will pop up File Importing dialog in which users can select and choose an existing design file on disk to open and import into the currently active Design Window. The File Importing dialog is shown in Figure 22. The solid rectangular represents the edit box. The dashed rectangular represents the display box. The round rectangulars represent the buttons.

In the File Importing dialog, users should first specify the directory of the file in *Look-in* box. The content of the directory will be automatically displayed in the display box in the center. The file type defaults to SpecC source files (‘.sc’ suffix) but users can specify any file type in the *File Type* box. All the files with the specified type will be displayed in the display box. Users then further select the file name in the *File Name* box. Finally, by clicking the *Open* button, the file with the specified name will be imported. If users click *Cancel* button, then the action of File Importing will be cancelled. Either clicking *Open* or *Cancel* button will close the File Import dialog.

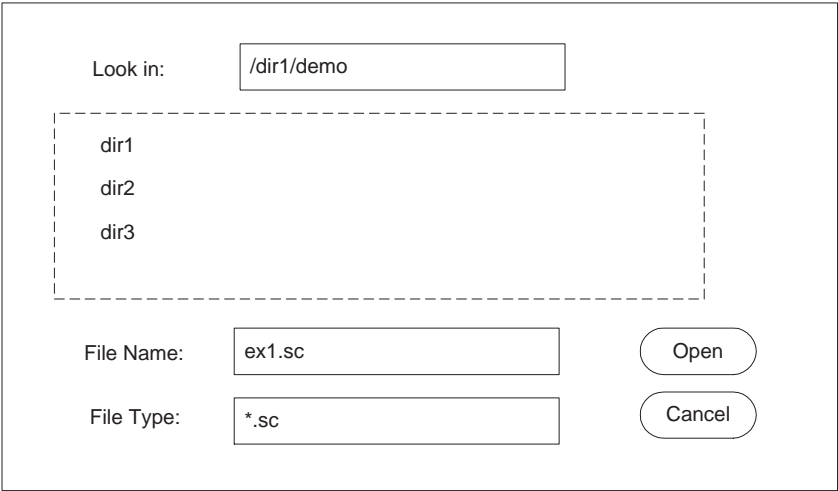


Figure 22: File Importing dialog.



Figure 23: Design Property dialog.

4.3.5 Design Property Viewing

Operation Users can display and view a design's properties by selecting *Main :: Project ⇒ Properties*. The selection will pop up a Design Properties dialog which is illustrated in Figure 23. The dashed rectangular represents the display box. The round rectangular represents the button.

In the Design Property dialog, the box *Design Name* displays the design name (*project1* in Figure 23), the Box *File Name* displays the name of the design file (e.g. */demo/ex1.sc*), and the box *Change Log* displays the history of SCE refinement commands applied to the design. The left part of the *Change Log* specifies which commands were executed. The right part of *Change Log* specifies the exact command lines with parameters. Clicking button *Ok* will close Design Properties dialog .

4.3.6 SCE Exiting

Operation Selecting *Main :: File ⇒ Exit* will exit the SCE application and close the SCE GUI.

The current project, if any, will be automatically saved by triggering a Project Saving action (see Section 4.2.3). Note that this can result in a Project Save dialog popping up in case the current project is unnamed and modified.

Error/Recommendation Messages. If there are any open Design Windows that are modified and not yet saved, an Information dialog will pop up for each window, querying the user whether he wants to save the corresponding design. The user will be able to cancel the whole Exit action via the corresponding dialog button. If the user accepts the recommendation to save the file, a File Saving action will be triggered (see Section 4.3.2). Note that the File Saving action can trigger additional Error dialogs which in turn can abort the whole Exit operation in case of file errors during saving.

Automatic Project Saving can result in errors and corresponding Error dialogs popping up. In turn, the whole Exit operation can be aborted in those cases.

4.4 Design-Entity Handling

Design-Entity Handling deals with manipulation of design entities such as behaviors, channels, variables, and channels. In general, Design entities are manipulated directly in the currently active Design Window in the Workspace. Design-Entity Handling is part of the refinement process and allows the user to perform some typical refinement tasks in a manual fashion. Basically, Design-Entity Handling tasks are tasks that do not change the semantics of the design. Rather, they are tasks to apply, for example, cosmetic or syntactic changes like behavior renaming to the design. Specifically, Design-Entity Handling consists of the following tasks:

- (a) Entity Renaming to rename the selected entity (see Section 4.4.1).
- (b) Entity Retyping to change the type of the selected entity (see Section 4.4.2).
- (c) Entity Deletion to delete the selected entity from the design (see Section 4.4.3).
- (d) Hierarchy Displaying to toggle between different modes for displaying of the design hierarchy (see Section 4.4.4).

4.4.1 Entity Renaming

Operation The names of design entities including behaviors, behavior instances, channels, channel instances, variables, and ports are displayed in *Name* column of the Design Window tabs. Users can change these names of entities. In order to do it, users first select the corresponding row in *Design :: Hierarchy*, *Design :: Behaviors* or *Design :: Channels* tabs. Selecting *Design :: Hierarchy* \Rightarrow *Rename*, *Design :: Behaviors* \Rightarrow *Rename*, or *Design :: Channels* \Rightarrow *Rename* context-menu commands will then allow users to edit the names directly in the cell by opening a corresponding text edit box in place. Renaming can be aborted by pressing the *Esc* key. Pressing *Enter* accepts the newly entered name and renames the entity in the design.

Error/Information Messages If renaming the entity in the design leads to an error (e.g. if an entity with the same name already exists), an Error dialog with a corresponding error message will pop up and subsequently the operation will be aborted.

4.4.2 Entity Retyping

Operation Users can change the type of design entity instances including variable, channel and behavior instances that are displayed in the *Type* column of Design Window hierarchy tab. In order to do it, users first need to select the corresponding row in the *Design :: Hierarchy* tab. Selecting the *Design :: Hierarchy* \Rightarrow *Change Type* context-menu command will then allow users to chose a type for the instance entity from a list of compatible types directly in the cell by opening a drop-down combo box with entries for all compatible types in place. Compatibility in SpecC is defined by a match in the list of port types and the list of implemented interfaces, e.g. two behaviors are compatible if they have the same list of ports in terms of order and type of ports. Users can choose any of the offered types from the drop-down list. Selecting a new type will then change the type of the instance entity in the design.

Error/Information Messages If retyping of the entity in the design leads to an error, an Error dialog with a corresponding error message will pop up and subsequently the operation will be aborted.

4.4.3 Entity Deletion

Operation Users can delete unused design entities from the design. In order to do it, users first need to select the corresponding row in *Design :: Hierarchy*, *Design :: Behaviors*, or *Design :: Channels* tabs. Selecting *Design :: Hierarchy* \Rightarrow *Delete*, *Design :: Behaviors* \Rightarrow *Delete* or *Design :: Channels* \Rightarrow *Delete* context-menu commands will then delete the selected entity from the design.

Error/Information Messages If deleting the entity in the design leads to an error (e.g. if the entity is used by another entity, i.e. if some part of the design depends on that entity), an Error dialog with a corresponding error message will pop up and subsequently the operation will be aborted.

4.4.4 Hierarchy Displaying

Operation Users can toggle between different modes for displaying the design hierarchy in the Design Window hierarchy tab. Selecting *Main :: View* \Rightarrow *Hierarchy* will open a sub-menu pop-up containing for entries corresponding to the four different display modes:

Show All Displays a forest of instantiation hierarchy trees where all uninstantiated behaviors are roots and the complete instantiation tree for all of them is shown.

Show Testbench Displays the complete instantiation hierarchy tree for the *Main* behavior class root only.

Show Design Displays the instantiation hierarchy tree only for the root class that is currently the top-level design behavior. This basically shows the current system being designed.

Show Architecture Displays the current top-level behavior as the only root class and only displays first-order (direct) children of (instances in) that behavior. This basically shows only the top, system level of the current system being designed.

Display modes are mutually exclusive, i.e. selecting one display mode will turn off the previous mode and switch to the new mode instead.

In addition to switching hierarchy display modes, users can toggle displaying of both variables and channels in the hierarchy tree. Selecting *Main :: View :: Show Variables* or *Main :: View :: Show Channels* toggles between display of variables and channels in

the hierarchy tab of the current Design Window. If channels are displayed, both channel instances and uninstantiated channels are shown in the forest of instantiation hierarchy trees. If variables are displayed, both global variables and variable instances inside behaviors (and optionally channels) are shown in the forest of instantiation hierarchy trees.

4.5 Synthesis

4.5.1 Architecture Exploration

Architecture Exploration is the main task of the Architecture Explorer tool. Architecture Exploration deals with the process of implementing a specification on a computation architecture consisting of PEs and memories in order to generate a respective architecture model for the design. Architecture exploration therefore helps designers to allocate PEs/memories, map design entities to the allocated PEs/memories and to generate the Architecture Model. Specifically, Architecture Exploration consists of four tasks:

- (a) Top-Level Selection to select the behavior representing the top level of the system to be designed (see Section 4.5.1).
- (b) PE Allocation to allocate and select PEs/memories from the PE database in order to assemble the system architecture (see Section 4.5.1).
- (c) Mapping to map the design's computation entities to the selected PEs/memories (see Section 4.5.1).
- (d) Architecture Refinement to automatically generate an Architecture Model from the given Specification Model based on the decision made during PE Allocation and Mapping (see Section 4.5.1).

In order to perform Architecture Exploration, not all the tasks described above need to be done. However, some tasks must be executed and must be executed in a certain order. These mandatory tasks and their execution sequence are:

1. Project Creation or Project Opening.
2. Preferences Editing and Project Settings Editing.
3. File Opening.
4. Design Adding.
5. Top-level Selection.
6. PE Allocation.

7. Mapping.
8. Architecture Refinement.
9. Project Saving and/or AE Exiting.

Note that steps 6, 7, and 8 can be performed repeatedly in a loop in order to generate multiple candidate architecture design models in one AE session.

Top-Level Selection Users can select the top-level behavior which represents the system design to be implemented. All the children of the top-level behavior are considered to be part of the system design. On the other hand, All the parents/siblings of the top-level behavior are considered to belong to the testbench outside of the actual design.

Operation Users first select the desired behavior by clicking its row in the *Design :: Hierarchy* or *Design :: Behaviors* tabs. Then, users set the selected behavior as the top-level behavior by selecting the *Design :: Hierarchy ⇒ Set As Top – Level* or *Design :: Behaviors ⇒ Set As Top – Level* commands in the context-menu for the respective behavior.

After selection, the instantiation hierarchy tree with the current top-level behavior at its root will be visually emphasized/highlighted in the Design Window’s hierarchy tab (through use of an italic font for entity texts).

Error/Information Messages All allocation information is stored in the design as annotations at the top-level design behavior (see Section 4.5.1). When switching the top-level from one behavior to another, if an allocation exists at the current behavior, an Information dialog is popped up querying the user whether he wants to copy the existing allocation over to the new top-level behavior. If the user accepts the recommendation and if then an allocation already exists at the new top-level behavior, another Information dialog is popped up querying the user whether he wants to overwrite the allocation information at the new top-level behavior. If the user declines the recommendation, the previous allocation at the new top-level behavior is kept.

PE Allocation Users can select PEs/memories out of the PE database in order to allocate and assemble the system architecture. PE allocation information is stored in the design itself as an allocation table that is annotated at the top-level design behavior (see Section 4.5.1). As a consequence, different allocation tables at different top-level behaviors can exist in the same design, reflecting the fact that incremental design will require changes in the allocation as design progresses from one part of the system to another.

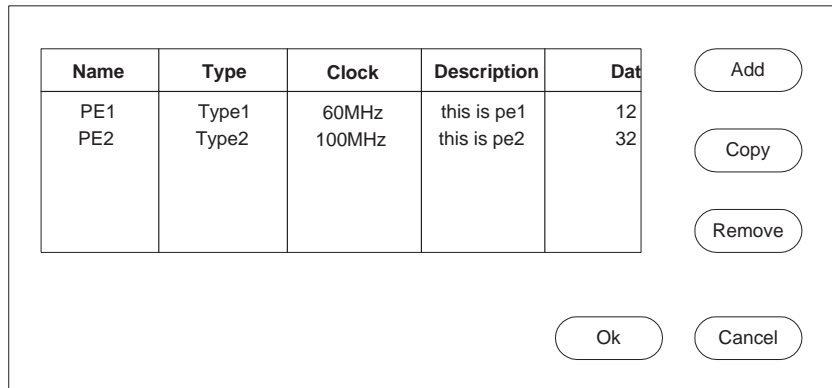


Figure 24: PE Allocation dialog.

Operation In order to do PE Allocation, users first select *Main :: Synthesis* ⇒ *Allocate PEs*. As a result, the current allocation is read from the design and a PE Allocation dialog is popped up. In case of errors reading the allocation from the design (e.g. wrong allocation table format), a new, empty allocation table is used in the dialog.

The screenshot of the PE Allocation dialog is shown in Figure 24. The round rectangles represent the buttons. In the PE Allocation dialog, the table shows the list of currently allocated PEs. The header of the table indicates the meaning of each column, such as PE's name and type. Each row in the table represents an allocated PE/memory that is part of the current system target architecture. For each PE, its name, its type, its attributes and an editable description are shown in the respective columns of the table. The list of allocated PEs can be sorted by any column and in ascending or descending order each by clicking into the corresponding column header. By default, the list is sorted by ascending names.

In the PE Allocation dialog, users can perform the following actions:

PE Adding In order to add a PE into the design, users click the button *Add* to pop up the PE Selection dialog which opens and loads the PE database and allows users to select an additional PE out of the PE database. The screenshot of the PE Selection dialog is shown in Figure 25.

At the left of the PE Selection dialog is a PE category table. Each row represents one category of PEs in the database. For example, row *Processors* contains all the general-purpose processors in the database.

By clicking and selecting one row in the table at left, users will be shown all the PEs in the selected category in the table at the right. Each row of the table at the right represents one type of PE in the database under the selected category. The name of the PE type (*Component* column) and other attributes of the PE type are displayed in

separate columns. Users can select the desired PE type by clicking the corresponding row.

There are two buttons at the bottom of the PE Selection dialog: *Ok* and *Cancel*. By clicking the *Ok* button, an additional PE of the selected PE type and with an automatically determined name is added into the design's allocated architecture. In addition, PEs can be allocated by double-clicking into the desired PE type row in the PE Selection dialog (equivalent to selecting the row and pressing *Ok*). Clicking the *Cancel* button aborts and cancels PE selection without changes to the PE allocation. Either clicking *Ok* or *Cancel* button will close the PE Selection dialog and return to the PE Allocation dialog.

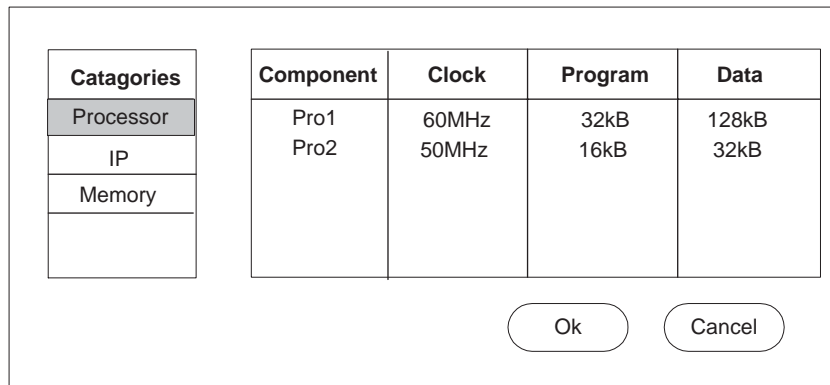


Figure 25: PE Selection dialog.

PE Copying In order to duplicate an existing PE in the design's PE allocation, users can select a PE by clicking the corresponding row in the allocation table and click the button *Copy*. Clicking *Copy* will add a new PE instance with an automatically determined name and with the same type, attributes and description as the currently selected PE to the design's allocation.

PE Deletion . In order to remove a PE from the design's allocation, users can select the target PE to be removed in the allocation table and click the *Remove* button. Clicking *Remove* will remove the selected PE from the list of allocated PEs.

PE Editing Users can edit the PE name and PE description in the PE Allocation dialog by clicking into the respective *Name* or *Description* column of the corresponding PE. Clicking into any of these cells will allow editing of the respective text in the cell by opening a text edit box in place. Pressing the *Esc* key during editing aborts the edit operation. Pressing *Enter* accepts the entered text and changes the PE name or description in the allocation accordingly.

Allocation Closing There are two buttons at the bottom of the PE Allocation dialog: *Ok* and *Cancel*. By clicking the *Ok* button, the allocation displayed in the PE Allocation dialog's allocation table is saved into the design. Clicking *Cancel* will abort and cancel PE Allocation. When cancelling, all modifications made to the allocation table will be lost and no data will be saved in the design. Either clicking *Ok* or *Cancel* button will close PE Allocation dialog.

Error/Information Messages There are several possible errors during PE Allocation:

- (a) Before PE Allocation, selecting *Main :: Synthesis ⇒ Allocate PEs* if no top-level behavior is selected in the design will pop up an Error dialog to that effect and will abort the PE Allocation operation.
- (b) During PE Editing, if users try to give PEs a name which is already used as the name of another PE in the design, an Error dialog will be popped up, a corresponding error message will be shown, and the editing operation will be aborted and cancelled.
- (c) During PE Adding, when clicking the *Add* button the PE database stored on disk will be opened and loaded in order to read the list of available PEs from the database. In case of errors during database opening (e.g. file errors or wrong file format), an Error dialog will be popped up and the PE Adding operation will be aborted.

Furthermore, when adding a PE to the allocation via the PE Selection dialog, the selected PE type is read from the database. In case of database read errors (file errors, database format errors) during this operation, an Error dialog will be popped up and the PE Adding operation aborted.
- (d) During PE Allocation, clicking the *Ok* button in the PE Allocation dialog will write the allocation table back to the design. In case of errors, an Error dialog is popped up and PE Allocation is aborted completely.

Mapping In order to implement the computation in the specification model on the allocated computation architecture consisting of PEs and memories, users have to be able to map the behaviors, variables and channels in the specification onto the allocated PEs. Hence, Mapping consists of separate Behavior Mapping, Variable Mapping and Channel Mapping tasks:

Behavior Mapping Behavior Mapping allows for mapping of behavior types/classes in the design onto allocated PEs, i.e. behavior mapping information is stored as annotations at the behavior classes in the design. In order to be able to map a behavior onto a PE, the PE out of the database must allow execution of arbitrary code on it. Users can explicitly map every behavior type on a PE. Explicit mapping will map all instances

of that behavior onto the selected PE. If instances should be mapped to different PEs, appropriate copies of the behavior have to be made outside of AE first. If a behavior is not mapped to any PE, all of its instances will be implicitly (and recursively) mapped onto the same PE as the parent behavior class in which they are instantiated in. If different instances are implicitly mapped to different PEs, appropriate copies of the behavior in each PE will be automatically generated during refinement. Note that user must map all the behaviors under the top-level behavior to PEs either implicitly (by mapping the top-level behavior itself) or explicitly.

Variable Mapping Variable Mapping allows for mapping of variable instances in the design into local memories of allocated regular PEs or into allocated global, shared memory PEs. Variable mapping information is stored as annotations attached to variable definitions inside a behavior class. As a result, multiple incarnations of the same variable in different instances of the parent behavior will all share the same mapping information. If a variable is not explicitly mapped by the user, during refinement a local copy of the variable will be created in each PE accessing the variable. Refinement will also automatically insert necessary code (additional behaviors inside PEs and channels between PEs) for synchronization and message passing to keep copies updated and synchronized such that shared semantics are preserved. Implicit mapping is not supported for variables that are shared among concurrent behaviors mapped to different PEs. If a variable is explicitly mapped into local memory of a regular PE or into a shared, global memory PE, all of its incarnations will be moved there and other PEs will access the variable through a memory interface. Explicit mapping of variables is only supported for target PEs out of the database that support external accesses via a memory interface.

Channel Mapping Channel Mapping allows for mapping of complex channel instances (of *c_queue* or *c_semaphore* type) in the specification onto PEs in the allocated target architecture. Channel mapping information is stored as annotations attached to channel instantiations inside a behavior class. As a result, multiple incarnations of the same channel (instance) in different instances of the parent behavior will all share the same mapping information. If a complex channel (instance) is explicitly mapped to a PE, implementations for all of its incarnations will be generated in that PE during refinement. If a complex channel (instance) is not explicitly mapped, implementations for each of its incarnations will be generated in the PE the parent behavior of that incarnation is mapped to. In both cases, other PEs accessing the channel incarnation will communicate with the target PE the channel incarnation is implemented in through additional automatically inserted simple channels. In order to be able to map a channel onto a PE, the PE out of the database must allow execution of arbitrary code on it.

Operation Users can map behaviors, variables and channels in the design to allocated regular or memory PEs via the additional *PE* column in the Design Window hierarchy tab. Note that the *PE* column is only shown if PE allocation information is available (see Section 4.5.1): By default, the *PE* column shows the current mapping information for each entity in the design. In case of errors reading the mapping information from the design (e.g. wrong annotation format), an empty, implicit (i.e. lack of explicit) mapping will be assumed.

In order to explicitly map an entity, users should click into the *PE* column of the respective entity in the *Design :: Hierarchy* tab. If the desired entities are not shown in the hierarchy tab, users should first enable display of variables or channels by selecting *Main : View ⇒ Show Variables* or *Main : View ⇒ Show Channel*.

Clicking into the *PE* column of the Design Window hierarchy tab will open a drop-down combo box in place with entries for all possible target PEs in order allow users to select a target PE to map the entity to directly in the cell. In the combo box, users will be able to choose from all possible PEs that the specific entity can be mapped to (see above for enforced restrictions). In addition, the combo box contains an empty entry to chose in order to remove any existing explicit mapping and switch to implicit mapping for that entity.

Selecting an entry from the combo box will write the corresponding mapping into the design. If there are multiple incarnations of an explicitly mapped entity (behavior class, variable definition or channel instance), the hierarchy tab display will be updated after changing the mapping to reflect the new mapping for all entity's incarnations in the *PE* column.

After mapping, the behaviors mapped to different PEs will be shown with different colors.

Architecture Refinement Architecture Refinement executes the implementation decisions made in the other Architecture Exploration tasks by refining the current Specification Model into an automatically generated Architecture Model based on and reflecting the decision made during PE Allocation and Mapping.

Operation In order to do refinement, users can select the *Main :: Synthesis ⇒ Architecture Refinement* menu entry to pop up the Architecture Refinement dialog. The screenshot of the Architecture Refinement dialog is shown in Figure 26. The small square with check mark represent the check boxes. The round rectangulars represent buttons.

In the Architecture Refinement dialog, users can select whether individual sub-tasks of the architecture refinement process will be performed or not. By checking or unchecking the check boxes tasks are turned on and off and partially or completely refined models can be generated. By default, all tasks are turned on. The three sub-tasks of architecture refinement are:

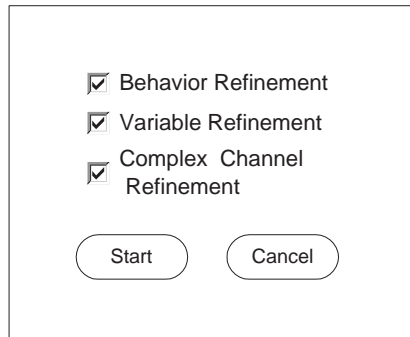


Figure 26: Architecture Refinement dialog.

- (a) *Behavior Refinement* which introduces PE behaviors from the database, groups the original behaviors under the new PEs, and inserts synchronization and message passing to preserve execution semantics.
- (b) *Variable Refinement* which (re-)distributes variables into PEs, generates necessary PE memory interfaces, and updates accesses to shared variables inside leaf behaviors.
- (c) *Channel Refinement* which generates implementations of complex channels inside PEs, creates adapters for accesses to complex channel implementations, and inserts necessary simple channels in between.

If *Behavior Refinement* is turned off, both *Variable Refinement* and *Channel Refinement* are turned off and can not be turned on. In this case, no output model is generated and only input validation is performed. In other cases, exactly one output model is generated in which optionally variables and channels are not refined but left untouched.

User can then start the architecture refinement process by clicking the *Start* button. If users click the *Cancel* button, the Architecture Refinement operation will be aborted and cancelled. Either clicking *Start* or *Cancel* buttons will close the Architecture Refinement dialog.

After clicking the *Start* button, the architecture refinement command line components of the Architecture Explorer will be executed in the background. Any diagnostic, status and informative output of the architecture refinement tools will be shown in the Refinement tab of the Output Window (*Output :: Refine*).

When the architecture refinement process is finished, the newly generated architecture model is automatically opened and loaded, and a corresponding new Design Window is created in the Workspace. The new Design Window is automatically activated and raised to the front. In addition, the new architecture model is automatically added to the current

project (see Design Adding, Section 4.2.5) as a child of the specification model it was generated from.

While the architecture refinement background tools are running, the majority of the main AE GUI is disabled. However, users can abort/kill execution of the background tools by selecting *Main :: Synthesis ⇒ Stop*. After clicking, the current architecture refinement background task is aborted.

Error/Information Messages If the architecture refinement background tools abort with an error (e.g. unmapped behaviors in the design) or are killed via the *Main :: Synthesis ⇒ Stop* menu entry, an Error dialog with a corresponding error message will pop up. Specifically, the architecture refinement background tools check for and can produce the following classes of errors:

- No top level behavior.
- No or invalid PE allocation.
- No or invalid PE models in the database.
- Unsupported (inter-PE) channels in the specification.
- Unsupported (partitioned) behavior types in the specification.
- Unsupported (global or partitioned) shared variable and/or port types in the specification.
- No or invalid behavior or complex channel mapping.
- Invalid variable mapping or no mapping for variables shared between concurrent behaviors on different PEs.

Upon confirming the error, the remainder of the Architecture Refinement Operation will be cancelled.

If the design the new model was generated from is not in the project, the new model is not added to the project and an Error dialog to that effect will be popped up.

4.5.2 Network Exploration

Network Exploration is the main task of the Network Explorer tool. Network Exploration deals with the process of implementing end-to-end communication channels between PEs in the architecture model over point-to-point channels in a network of PEs, CEs and busses in order to generate the respective network model. Network Exploration therefore helps designers to allocate and define a communication topology and to generate a Network Model. Specifically, Network Exploration consists of two tasks:

- (a) Network Allocation to allocate, select and define the communication network topology (see Section 4.5.2).
- (b) Network Refinement to automatically generate a Network Model from the given Architecture Model based on the decisions made during Network Allocation (see Section 4.5.2).

In order to perform the network exploration, not all the tasks described above need to be done. However, some tasks must be executed and must be in a certain order. These mandatory tasks and their execution sequence are:

1. Project Creation or Project Opening.
2. Preferences Editing and Project Settings Editing.
3. File Opening.
4. Design Adding.
5. Network Allocation.
6. Network Refinement.
7. Project Saving and/or NE Exiting.

Note that steps 5 and 6 can be performed repeatedly in a loop in order to generate multiple candidate communication design models in one NE session.

Network Allocation In order to define the system network topology, users can allocate and connect (together with already allocated PEs) busses and CEs (bridges, transducers) out of the bus and CE databases, respectively. Network allocation and connection information is stored in the design itself as allocation and connection tables that are annotated at the top-level design behavior (see Section 4.5.1). As a consequence, different allocation/connection tables at different top-level behaviors can exist in the same design, reflecting the fact that incremental design will require changes in the topology as design progresses from one part of the system to another.

Operation In order to do Network Allocation, users first select *Main :: Synthesis* \Rightarrow *Allocate Network*. As a result, the current allocation and connectivity is read from the design and a Network Allocation dialog is popped up. In case of errors reading the allocation from the design (e.g. wrong allocation table format), new, empty allocation and connection tables are used in the dialog.

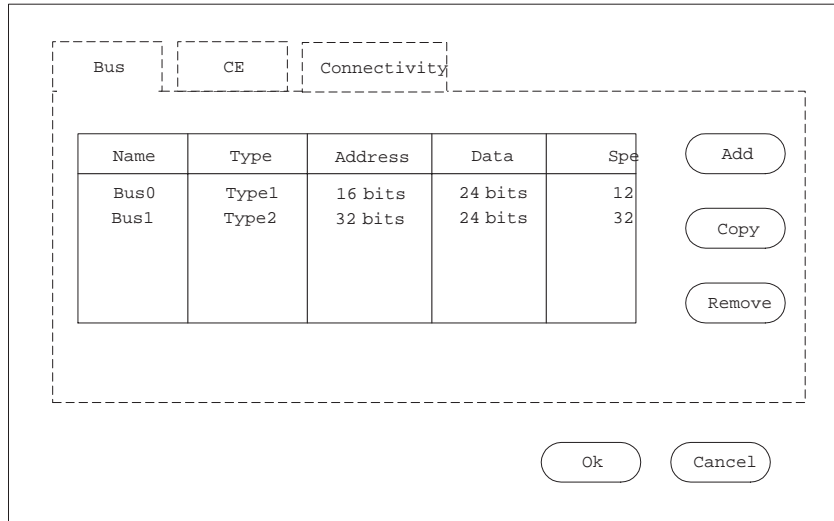


Figure 27: Network Allocation dialog (Bus tab).

The screenshot of the Network Allocation dialog is shown in Figure 27 where the round rectangulars represent buttons. The Network Allocation dialog has three tabs, *Bus*, *CE* and *Connectivity*, in which the user can allocate busses, allocate CEs, or define the PE/CE to bus connectivity, respectively. Clicking the tabs at the top of the dialog will raise and shown the corresponding tab in front.

Note that as a result of previous architecture exploration, busses might be pre-allocated and pre-connected in the Architecture Model generated by the Architecture Explorer. The Architecture Explorer will automatically pre-allocate and pre-connect necessary and mandatory busses for PEs that come with pre-defined busses (i.e. CPUs). This information is then passed from the Architecture Explorer to the Network Explorer through the Architecture Model. As a result, such busses and their connectivity to PEs might show up when first opening the Network Allocation dialog for a new Architecture model.

In the *Bus* tab, a table with the list of currently allocated busses will be shown (Figure 27). The header of the table indicates the meaning of each column, such as bus name and type. Each row in the table represents an allocated bus that is part of the current system target architecture. For each bus, its name, its type, its attributes and an editable description are shown in the respective columns of the table. The list of allocated busses can be sorted by any column and in ascending or descending order each by clicking into the corresponding column header. By default, the list is sorted by ascending names. In the *Bus* tab of the Network Allocation dialog, users can perform the following actions:

Bus Adding In order to add a bus to the design, users click the button *Add* to pop up the

Bus Selection dialog which opens and loads the bus database and allows users to select an additional bus out of the bus database. The screenshot of the Bus Selection dialog is shown in Figure 28.

At the left of the Bus Selection dialog is a bus category table. Each row represents one category of busses in the database. For example, row *Standard* contains all the standard busses in the database.

By clicking and selecting one row in the table at left, users will be shown all the busses in the selected category in the table at the right. Each row of the table at the right represents one type of bus in the database under the selected category. The name of the bus type (*Bus* column) and other attributes of the bus type are displayed in separate columns. Users can select the desired bus type by clicking the corresponding row.

There are two buttons at the bottom of the Bus Selection dialog: *Ok* and *Cancel*. By clicking the *Ok* button, an additional bus of the selected bus type and with an automatically determined name is added into the design's allocated network architecture. In addition, busses can be allocated by double-clicking into the desired bus type row in the Bus Selection dialog (equivalent to selecting the row and pressing *Ok*). Clicking the *Cancel* button aborts and cancels Bus selection without changes to the bus allocation. Either clicking *Ok* or *Cancel* button will close the Bus Selection dialog and return to the Network Allocation dialog.

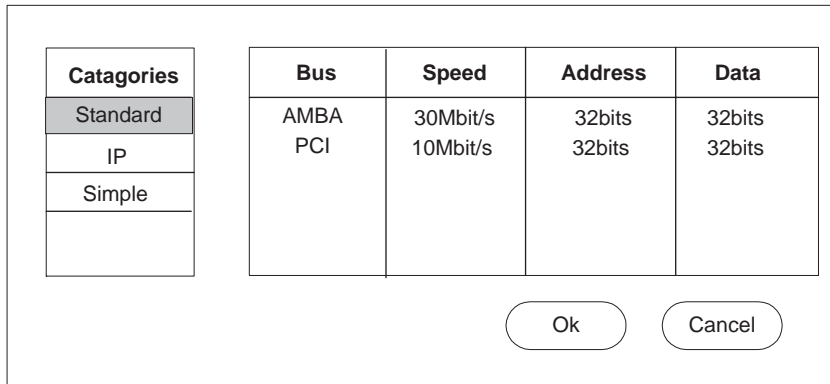


Figure 28: Bus Selection dialog.

Bus Copying In order to duplicate an existing bus in the design's bus allocation, users can select a bus by clicking the corresponding row in the allocation table and click the button *Copy*. Clicking *Copy* will add a new bus instance with an automatically determined name and with the same type, attributes and description as the currently

selected bus to the design's allocation.

Bus Deletion . In order to remove a bus from the design's allocation, users can select the target bus to be removed in the allocation table and click the *Remove* button. Clicking *Remove* will remove the selected bus from the list of allocated busses. Busses that currently have PEs or CEs connected to them will not be available for deletion (the *Remove* button will be inactive and grayed out for them).

Bus Editing Users can edit the bus name and bus description in the bus allocation tab by clicking into the respective *Name* or *Description* column of the corresponding bus. Clicking into any of these cells will allow editing of the respective text in the cell by opening a text edit box in place. Pressing the *Esc* key during editing aborts the edit operation. Pressing *Enter* accepts the entered text and changes the bus name or description in the allocation accordingly.

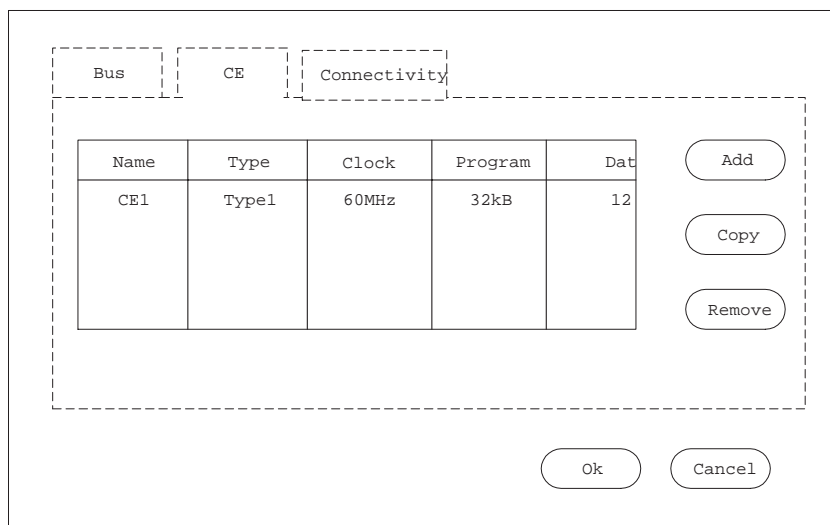


Figure 29: Network Allocation dialog (CE tab).

In the *CE* tab, a table with the list of currently allocated CEs will be shown (Figure 29). The header of the table indicates the meaning of each column, such as CE name and type. Each row in the table represents an allocated CE that is part of the current system target architecture. For each CE, its name, its type, its attributes and an editable description are shown in the respective columns of the table. The list of allocated CEs can be sorted by any column and in ascending or descending order each by clicking into the corresponding column header. By default, the list is sorted by ascending names. In the *CE* tab of the Network Allocation dialog, users can perform the following actions:

CE Adding In order to add a CE to the design, users click the button *Add* to pop up the CE Selection dialog which opens and loads the CE database and allows users to select an additional CE out of the CE database. The screenshot of the CE Selection dialog is shown in Figure 30.

At the left of the CE Selection dialog is a CE category table. Each row represents one category of CEs in the database. For example, row *Transducer* contains all the transducers in the database.

By clicking and selecting one row in the table at left, users will be shown all the CEs in the selected category in the table at the right. Each row of the table at the right represents one type of CE in the database under the selected category. The name of the CE type (*CE* column) and other attributes of the CE type are displayed in separate columns. Users can select the desired CE type by clicking the corresponding row.

There are two buttons at the bottom of the CE Selection dialog: *Ok* and *Cancel*. By clicking the *Ok* button, an additional CE of the selected CE type and with an automatically determined name is added into the design's allocated network architecture. In addition, CEs can be allocated by double-clicking into the desired CE type row in the CE Selection dialog (equivalent to selecting the row and pressing *Ok*). Clicking the *Cancel* button aborts and cancels CE selection without changes to the CE allocation. Either clicking *Ok* or *Cancel* button will close the CE Selection dialog and return to the Network Allocation dialog.

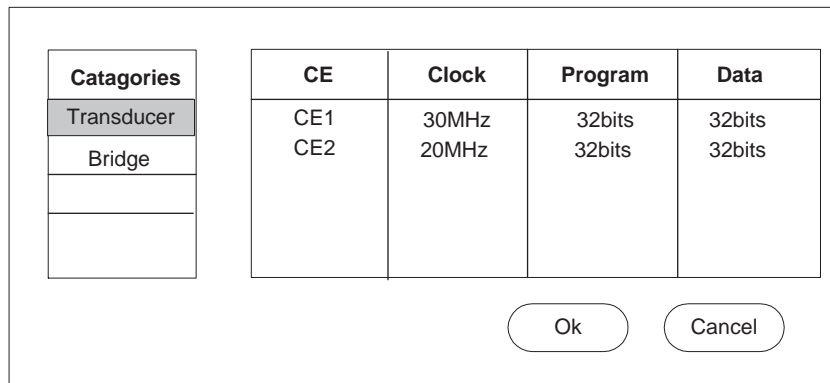


Figure 30: CE Selection dialog.

CE Copying In order to duplicate an existing CE in the design's CE allocation, users can select a CE by clicking the corresponding row in the allocation table and click the button *Copy*. Clicking *Copy* will add a new CE instance with an automatically determined name and with the same type, attributes and description as the currently

selected CE to the design's allocation.

CE Deletion . In order to remove a CE from the design's allocation, users can select the target CE to be removed in the allocation table and click the *Remove* button. Clicking *Remove* will remove the selected CE from the list of allocated busses.

CE Editing Users can edit the CE name and CE description in the CE allocation tab by clicking into the respective *Name* or *Description* column of the corresponding CE. Clicking into any of these cells will allow editing of the respective text in the cell by opening a text edit box in place. Pressing the *Esc* key during editing aborts the edit operation. Pressing *Enter* accepts the entered text and changes the CE name or description in the allocation accordingly.

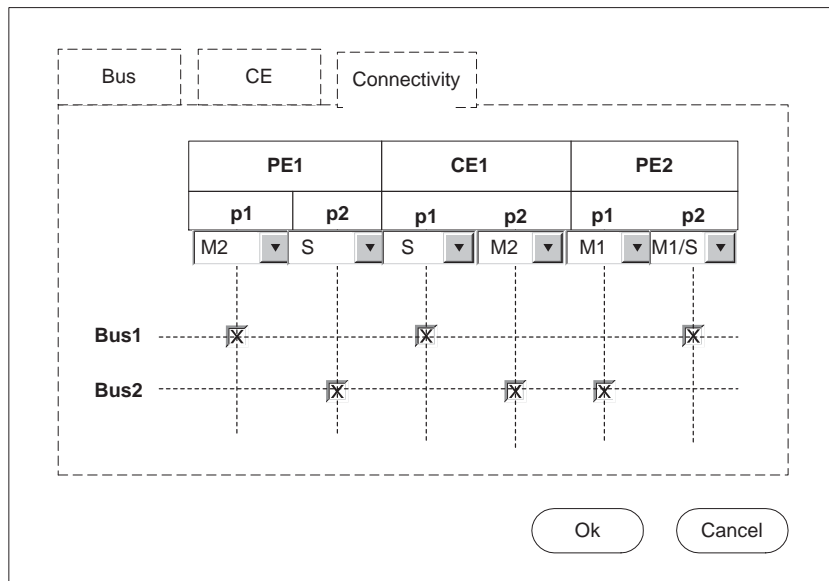


Figure 31: Network Allocation dialog (Connectivity tab).

In the *Connectivity* tab, a matrix that displays connectivity of PEs and CEs to busses is shown (Figure 31). The top rows/headers of the connectivity table show all the currently allocated PEs and CEs together with the list of logical/interface ports for each PE/CE in the design. For each PE/CE and port, its name is shown in the respective header row. The list of PEs/CEs is automatically updated whenever a CE is added or deleted from the CE allocation in the *CE* tab. The rows of the connectivity matrix in the *Connectivity* tab list the currently allocated busses in the design by their names. The list of busses is automatically updated whenever a bus is added or deleted from the bus allocation in the *Bus* tab. For

each bus, the matrix visually shows the connectivity of ports to this bus by marking the corresponding columns. In the *Connectivity* tab of the Network Allocation dialog, users can perform the following actions:

Port Adding/Deletion The list of ports/interfaces for each PE and CE is determined by the list of available ports defined by the PE/CE database model. For synthesizable PEs/CEs with undefined or variable port lists in the database, the user can add or remove ports from the PE's or CE's port list. Right-clicking onto a PE's or CE's port list or port will pop up a context menu with the following entries:

Add port Selecting *Add port* will add an additional port with an automatically determined name to the corresponding PE or CE. This menu entry is not available (grayed out) for PEs or CEs with fixed port lists.

Remove port Selecting *Remove port* will remove the selected port from the corresponding PE or CE. This menu entry is not available (grayed out) for fixed/mandatory ports.

Clicking onto the name of a non-fixed port will open a text edit box in place in order to allow the user to edit the port name directly in the cell. Pressing *Esc* will abort editing while *Enter* accepts the new name.

Port Connecting In the table, the users can then connect ports of a PEs/CEs to busses by clicking into the corresponding table cell at the intersection of the port and bus. If the selected port is already connected to a different bus, the connection will be moved (i.e. the previous connection be removed and the port re-connected to the new bus). Ports with pre-defined, fixed protocols can only be connected to matching busses and clicking into any other bus row will have no effect. Pre-connected ports with mandatory, pre-defined and fixed bus connections can not be connected to any other bus, i.e. clicking into any row will have no effect.

For each port connection, the user can further specify the interface type of the port. Generally, ports can be connected to busses either as bus master, bus slave or as combined bus master/slave. If the bus from the database supports arbitration and allows for multiple masters on the bus, any master connection must also specify the arbiter port to use where the bus database model defines the set of possible master ports supported by the bus. Clicking into a valid (see above) connection cell of the connectivity matrix will open a drop-down combo box in that cell in place with entries for all possible connection types for that port/bus combination. In general, users will be able to choose in the combo box from the following connection types: slave, name of any available (not connected to any other PE/CE port) bus master port, or any combination of slave with any available master port. Selecting an entry from the combo box will change the connection type and update the text displayed in the cell accordingly.

There are two buttons at the bottom of Network Allocation dialog: *Ok* and *Cancel*. By clicking the *Ok* button, the network allocation and connectivity is saved back into the design. If users click the *Cancel* button, Network Allocation will be cancelled and all modifications to the allocation and connectivity are discarded. Either clicking *Ok* or *Cancel* button will close the Network Allocation dialog.

Error/Information Messages There are several possible errors during Network allocation in general: if no top-level behavior is selected in the design when selecting *Main :: Synthesis* \Rightarrow *Allocate Network* an Error dialog to that effect will be popped up and the Network Allocation operation will be aborted. Furthermore, clicking the *Ok* button in the Network Allocation dialog will write the allocation and connection tables back to the design. In case of errors, an Error dialog is popped up and Network Allocation is aborted.

There are several errors that can happen specifically during bus or CE allocation:

- (a) During Bus or CE Editing, if users try to give busses or CEs a name which is already used as the name of another bus or CE in the design, an Error dialog will be popped up, a corresponding error message will be shown, and the editing operation will be aborted and cancelled.
- (b) During Bus or CE Adding, when clicking the *Add* button the bus or CE database stored on disk will be opened and loaded in order to read the list of available busses/CEs from the database. In case of errors during database opening (e.g. file errors or wrong file format), an Error dialog will be popped up and the Bus or CE Adding operation will be aborted.

Furthermore, when adding a bus or CE to the allocation via the Bus or CE Selection dialog, the selected bus/CE type is read from the database. In case of database read errors (file errors, database format errors) during this operation, an Error dialog will be popped up and the Bus/CE Adding operation aborted.

Finally, several errors can happen during connectivity definition:

- (a) During port name editing, if users try to give ports a name which is already used as the name of another port of the same PE/CE, a corresponding Error dialog will be shown and the renaming operation cancelled.
- (b) If, upon closing of the Network Allocation dialog by clicking the *Ok* button, there are any PE/CE ports that are not connected, a corresponding Error dialog will be shown and the user will be returned to the Network Allocation dialog.

Network Refinement Network Refinement executes the implementation decisions made in the other Network Exploration tasks by refining the current Architecture Model into an automatically generated Network Model based on and reflecting the decision made during Network Allocation.

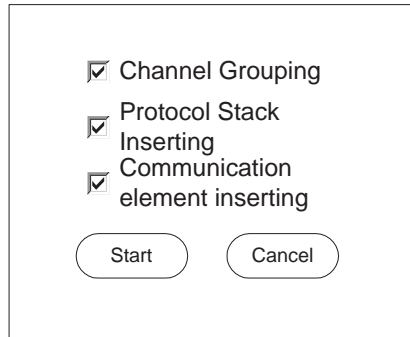


Figure 32: Network Refinement dialog.

Operation In order to do refinement, users can select the *Main :: Synthesis* \Rightarrow *Network Refinement* menu entry to pop up the Network Refinement dialog. The screenshot of the Network Refinement dialog is shown in Figure 32. The small square with check mark represent the check boxes. The round rectangulars represent buttons.

In the Network Refinement dialog, users can select whether individual sub-tasks of the network refinement process will be performed or not. By checking or unchecking the check boxes tasks are turned on and off and partially or completely refined models can be generated. By default, all tasks are turned on. The three sub-tasks of network refinement are:

- (a) *Protocol insertion* which generates implementations for the upper network layers (presentation and session layers) of the protocol stack into the system PEs.
- (b) *Channel merging* which automatically groups and merges sequentially accesses channels over one shared channel.
- (c) *CE insertion* which inserts communication elements such as transducers and bridges between PEs and generates necessary additional network protocol implementations inside PEs and CEs.

The *Channel merging* and *CE insertion* sub-tasks are automatically turned off and can not be turned on if *Protocol insertion* is disabled. If all tasks are turned off, no output model is

generated and only input validation is performed. In other cases, exactly one output model at varying levels of refinement is generated.

User can then start the network refinement process by clicking the *Start* button. If users click the *Cancel* button, the Network Refinement operation will be aborted and cancelled. Either clicking *Start* or *Cancel* buttons will close the Network Refinement dialog.

After clicking the *Start* button, the network refinement command line components of the Network Explorer will be executed in the background. Any diagnostic, status and informative output of the network refinement tools will be shown in the Refinement tab of the Output Window (*Output :: Refine*).

When the network refinement process is finished, the newly generated network model is automatically opened and loaded, and a corresponding new Design Window is created in the Workspace. The new Design Window is automatically activated and raised to the front. In addition, the new network model is automatically added to the current project (see Design Adding, Section 4.2.5) as a child of the architecture model it was generated from.

While the network refinement background tools are running, the majority of the main NE GUI is disabled. However, users can abort/kill execution of the background tools by selecting *Main :: Synthesis ⇒ Stop*. After clicking, the current network refinement background task is aborted.

Error/Information Messages If the network refinement background tools abort with an error or are killed via the *Main :: Synthesis ⇒ Stop* menu entry, an Error dialog with a corresponding error message will pop up. Specifically, the network refinement background tools check for and can produce the following classes of errors:

- No top level behavior.
- No or invalid network allocation and topology information.
- No or invalid CE and bus models in the database.
- Unsupported channels or channel connectivity structure in the system architecture at the top level.

Upon confirming the error, the remainder of the Network Refinement Operation will be cancelled.

If the design the new model was generated from is not in the project, the new model is not added to the project and an Error dialog to that effect will be popped up.

4.5.3 Communication Synthesis

Communication Synthesis is the main task of the Communication Synthesizer tool. Communication synthesis deals with the process of implementing point-to-point communication

channels between PEs and CEs in a network model over actual busses and bus protocols in order to generate a respective communication model for the design. Communication synthesis therefore helps designers to determine the communication details on each bus and to generate a Transaction-Level or Communication Model. Specifically, Communication Synthesis consists of two tasks:

- (a) Bus Parameter Assignment to assign the parameters to the communication performed over each bus in the system (see Section 4.5.3).
- (b) Communication Refinement to automatically generate a Transaction-Level or Communication Model from the given Network Model based on the decisions made during Bus Parameter Assignment (see Section 4.5.3).

In order to perform Communication Synthesis, not all the tasks described above need to be done. However, some tasks must be executed and must be executed in a certain order. These mandatory tasks and their execution sequence are:

1. Project Creation or Project Opening.
2. Preferences Editing and Project Settings Editing.
3. File Opening.
4. Design Adding.
5. Bus Parameter Assignment.
6. Communication Refinement.
7. Project Saving and/or CS Exiting.

Note that steps 5 and 6 can be performed repeatedly in a loop in order to generate multiple candidate communication design models in one CS session.

Bus Parameter Assignment The network model defines the sets of logical point-to-point communication channels between PEs and CEs to be implemented over each bus segment in the system. In order to implement data transfers and synchronization associated with these logical point-to-point communication channels over the actual busses, users need to assign both distinct physical bus addresses and distinct bus interrupt lines to each channel as needed. In general, multiple slaves and multiple masters can be connected to each bus. Interrupt lines are needed for channels with synchronization from slave to master. On the other hand, synchronization from master to slave and any kind of data transfer between master and slave requires assignment of a corresponding bus address. In order to distinguish

channels going over the same bus, each channel has to be assigned a separate addresses and/or interrupt.

The network model can contain double-handshake channels, single-handshake channels and PE memory interfaces mapped to each bus. Since double-handshake channels require data transfers and two-way synchronization, they need to be assigned one address and one interrupt line each. Single-handshake channels require one-way synchronization with no data transfer only. Depending on their direction, they therefore need to be assigned either one address (from master to slave) or one interrupt line (from slave to master). Memory interfaces require a contiguous block of data transfers each with a size determined by the network model. During communication synthesis, a base address needs to be assigned to each memory interface which in turn will determine the range of addresses occupied by the memory interface on the respective bus. Corresponding address and interrupt assignments are stored as annotations attached to memory interfaces and channel instances inside the top-level behavior of the design.

Busses are taken out of the bus database during Communication Synthesis. Each bus in the database defines the range of addresses and the set of interrupt lines it can support. During Bus Parameter Assignment, users can choose parameters for each channel freely out of the set of supported addresses and interrupt lines. However, database models for each PE and CE connected to the bus can define ranges of addresses on the bus reserved for PE- or CE-internal use. During Bus Parameter Assignment, the user is restricted from selecting any of those reserved addresses. For example, CPUs can define restricted addresses on their busses reserved for CPU-internal accesses to its local memory or interrupt controller. Also, bus slave interfaces of bridges will always reserve the complete range of master side addresses mapped into the slave side.

Operation In order to assign addresses and interrupts, users can pop up the Bus Parameter dialog by selecting *Main :: Synthesis ⇒ Assign Bus Parameters*. As a result, the current interrupt and address assignments are read from the design and shown in the dialog. In case of errors reading the parameters from the design (e.g. wrong annotation table format), assignments in error are discarded silently.

The screenshot of Bus Parameter dialog is shown in Figure 33. The dialog contains a number of tabs, one for each bus in the system where the name of a tab is determined by the name of its bus. Clicking on a tab at the top of the dialog will raise and show the corresponding tab in the front.

Each tab contains a table with four columns: *Channel*, *Start Address*, *End Address*, and *Interrupt*. The table lists all channels, memory interfaces and PEs/CEs on the corresponding bus. The *Channel* column shows the name of each channel, interface or PE/CE. The columns *Start Address* and *End Address* show the range of addresses assigned to each channel, interface or PE/CE in hexadecimal format. Note that for a channel, start address

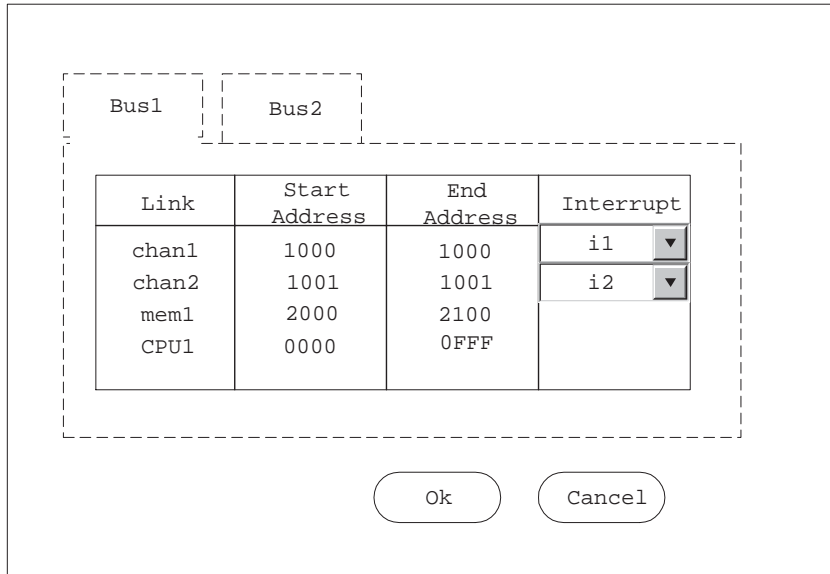


Figure 33: Bus Parameter dialog.

and end address are always the same. Channel and PE/CE address fields can be empty if no address is required or reserved. Finally, the *Interrupt* column shows the name of the bus interrupt line assigned to a channel. The interrupt field can be empty if not interrupt line is required. It is always empty for memory interfaces and PEs. The list of channels, memory interfaces and PEs/CEs can be sorted by any column and in ascending or descending order each by clicking into the corresponding column header. By default, the list is sorted by ascending start addresses.

For Bus Parameter Assignment, users can perform the following actions in the dialog:

Address Assignment Users can assign start (base) addresses to channels and memory interfaces in the list. In order to assign addresses, users click into the *Start Address* column for the respective channel or interface. The dialog will prevent editing and ignore column clicks for PEs/CEs and for channels that do not require an address. Otherwise, clicking will open a text edit box in place to allow entering of the desired address in hexadecimal format directly in the cell. The user can abort editing by pressing the *Esc* key. Pressing *Enter* accepts the entered text and assigns the new address to the channel/interface. Entering a new address will automatically update the *End Address* column for the corresponding channel/interface (for channels, the end address equals to the start address; for memories, the end address equals to start address + memory size - 1).

Interrupt Assignment Users can assign interrupts to certain channels in the list by clicking into the *Interrupt* column of the respective channel. Clicking into the *Interrupt* column for memory interfaces, PEs/CEs or channels that don't require an interrupt will be ignored. Otherwise, clicking will open a drop-down combo box in place of the cell with entries for all possible interrupt lines supported by the bus. In the combo box users can choose and select any of the bus interrupt lines to assign to the channel.

If the user assigns an address or interrupt that is already used by another channel, interface or PE/CE to any of the channels or interfaces, the conflicting entries will be highlighted (using red text color) in the corresponding cells of the Bus Parameter dialog.

There are two buttons at the bottom of the Bus Parameter dialog: *Ok* and *Cancel*. By clicking the *Ok* button, the assigned bus parameters are saved into the design. If users click the *Cancel* button, the Bus Parameter Assignment task will be cancelled. Either clicking *Ok* or *Cancel* buttons will close the Bus Parameter dialog.

Error/Information Messages When entering a new address, the Bus Parameter dialog performs validation of the entered text to ensure that it is in correct hexadecimal format and within range of supported bus addresses. In case of errors, a corresponding Error dialog is popped up and address editing eventually resumed.

When clicking the *Ok* button in the dialog, address and interrupt assignments are validated for conflicts. If there are conflicting/overlapping assignments, an Information dialog is popped up querying the user whether he is sure to commit the changes. If the user declines, he will be returned to the Bus Parameter dialog. If the user accepts (or if there are no conflicts in the first place), address and interrupt assignments annotations are written back into the design. If an error writing the annotations occurs, a corresponding Error dialog is popped up and the Bus Parameter Assignment aborted.

Communication Refinement Communication Refinement executes the implementation decisions made in the other Communication Synthesis tasks by refining the current Network Model into an automatically generated Transaction-Level or Communication Model based on and reflecting the decision made during Bus Parameter Assignment.

Operation In order to do refinement, users can select the *Main :: Synthesis ⇒ Communication Refinement* menu entry to pop up the Communication Refinement dialog. The screenshot of the Communication Refinement dialog is shown in Figure 34. The small square with check mark represent the check boxes. The round rectangulars represent buttons.

In the Communication Refinement dialog, users can select whether individual sub-tasks of the architecture refinement process will be performed or not. By checking or unchecking the check boxes tasks are turned on and off and partially or completely refined models can

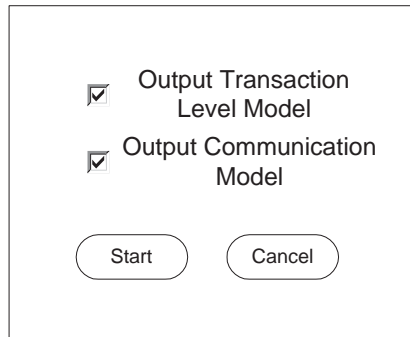


Figure 34: Communication Refinement dialog.

be generated. By default, all tasks are turned on, i.e. both transaction-level and communication models are generated for the design. The two sub-tasks of communication refinement are:

- (a) *Generate Transaction Level Model* which generates a transaction-level model at the output.
- (b) *Generate Communication Model* which generates a pin-accurate, bus-functional communication model at the output.

If both sub-tasks are turned on, two output models will be generated. On the other hand, if not sub-task is turned on, no output will be generated and only input validation will be performed. In all other cases, one output model of selected format will be generated.

User can then start the communication refinement process by clicking the *Start* button. If users click the *Cancel* button, the Communication Refinement operation will be aborted and cancelled. Either clicking *Start* or *Cancel* buttons will close the Communication Refinement dialog.

After clicking the *Start* button, the communication refinement command line components of the Communication Synthesizer will be executed in the background. Any diagnostic, status and informative output of the communication refinement tools will be shown in the Refinement tab of the Output Window (*Output :: Refine*).

When the communication refinement process is finished, the newly generated transaction-level or communication models are automatically opened and loaded, and corresponding new Design Windows are created in the Workspace. The new Design Windows are automatically activated and raised to the front. In addition, the new models are automatically added to the current project (see Design Adding, Section 4.2.5) as a child of the specification model they were generated from.

While the architecture refinement background tools are running, the majority of the main CS GUI is disabled. However, users can abort/kill execution of the background tools

by selecting *Main :: Synthesis ⇒ Stop*. After clicking, the current communication refinement background task is aborted.

Error/Information Messages If the communication refinement background tools abort with an error or are killed via the *Main :: Synthesis ⇒ Stop* menu entry, an Error dialog with a corresponding error message will pop up. Specifically, the communication refinement background tools check for and can produce the following classes of errors:

- No top level behavior.
- No or invalid network allocation and topology information.
- No or invalid bus models in the database.
- No (if required) or invalid bus-functional PE/CE models in the database.
- Unsupported channels, PEs/CEs or connectivity structure in the system architecture at the top level.
- No or invalid address and interrupt assignments.

Upon confirming the error, the remainder of the Communication Refinement Operation will be cancelled.

If the design the new models were generated from is not in the project, the new models are not added to the project and an Error dialog to that effect will be popped up.

4.5.4 Decision Import

In order to take over implementation decisions from a previously done design into a new design, SCE allows to import design decisions from one design into another. With this functionality, previously made design decisions can easily be transferred to a new design as a starting point for exploration.

Operation Users can import design decisions from any currently opened and loaded design into the currently active design by selecting the Main Menu command *Main :: Synthesis ⇒ Import Decisions*. As a result, the Import Decisions dialog will be popped up as shown in Figure 35. The solid rectangular with triangle symbol represents the drop-down menu. The square with check mark represents a select box. The round rectangulars represents buttons.

First, in order to select a design from which decisions are imported, users select the name of the source design in the *Source design* drop-down box. The combo box contains

all the names of currently opened Design Windows except for the currently active one. Users select one design from them as the source design.

Secondly, users can select decisions which will be imported to the current design. Users do it by checking the selection items as follows:

PE Allocation Copies the allocated PEs/memories from the source design to the currently active design. For all behaviors in the source design, PE allocation annotations are copied to the behavior (if any) with the same name in the target design.

Behavior Mapping Copies behavior mapping information from the source design to the currently active design. For all behaviors in the source design, PE mapping annotations are copied to the behavior (if any) with the same name in the target design.

Variable Mapping Copies variable mapping information from the source design to the currently active design. For all variable definitions in all behaviors in the source design, PE mapping annotations are copied to the variable (if any) with the same name in the behavior (if any) with the same name in the target design.

Channel Mapping Copies channel mapping information from the source design to the currently active design. For all channel instances in all behaviors in the source design, PE mapping annotations are copied to the channel instance (if any) with the same name in the behavior (if any) with the same name in the target design.

Bus Allocation Copies the allocated busses from the source design to the currently active design. For all behaviors in the source design, bus allocation annotations are copied to the behavior (if any) with the same name in the target design.

CE Allocation Copies the allocated busses from the source design to the currently active design. For all behaviors in the source design, CE allocation annotations are copied to the behavior (if any) with the same name in the target design.

Connectivity Copies the network topology connectivity from the source design to the currently active design. For all behaviors in the source design, network connectivity annotations are copied to the behavior (if any) with the same name in the target design.

Address assignment Copies the address assignments from the source design to the currently active design. For all channel instances in all behaviors in the source design, address assignment annotations are copied to the channel instance (if any) with the same name in the behavior (if any) with the same name in the target design.

Interrupt assignment Copies the interrupt assignments from the source design to the currently active design. For all channel instances in all behaviors in the source design,

interrupt assignment annotations are copied to the channel instance (if any) with the same name in the behavior (if any) with the same name in the target design.

The Import Decisions dialog includes an additional check item to select whether any existing annotations in the target design should be overwritten or kept.

There are two buttons on the Import Decisions dialog: *Import* and *Cancel*. By clicking the *Import* button, the design decisions of the selected items are imported to the selected design from the source design. If users clicks the *Cancel* button, the action of Decision Importing will be aborted and cancelled. Either clicking *Import* or *Cancel* button will close the Import Decisions dialog.

Source design:

Architecture Synthesis

- PE Allocation
- Behavior Mapping
- Variable Mapping
- Complex Channel Mapping

Network Synthesis

- Bus Allocation
- CE Allocation
- Topology Selection

Communication Synthesis

- Address Assignment
- Interrupt Assignment

Import Cancel

Figure 35: Import Decisions dialog.

Error/Information Messages Decision Importing requires reading and writing of annotations in the designs. In case of errors (e.g. wrong annotation format) upon pressing the *Import* button, an Error dialog will be popped up and the Decision Importing operation will be aborted.

4.6 Window Management

Window Management deals with the management of design windows in the Workspace. Window Management allows for closing, resizing, and arranging of multiple simultaneously opened design windows within the Workspace. Specifically, the tasks for Window Management are:

Window Closing Users can close the currently active design window in the Workspace by selecting *Main :: Window ⇒ Close*. In addition, any of the design windows can be closed by clicking on a respective icon in the window's title bar.

Users can close all the currently opened design windows in the Workspace by selecting *Main :: Window ⇒ Close All*.

In all cases, closing a design window triggers a File Closing action for the corresponding design file (see Section 4.3.3).

Window Arranging Users can automatically arrange design windows in the Workspace in a variety of manners. Selecting *Main :: Window ⇒ Tile* will rearrange the design windows in the Workspace in a tiled fashion. Selecting *Main :: Window ⇒ Cascade* will rearrange the design windows in the Workspace in a cascaded manner. Apart from that, windows can be freely resized and moved within the Workspace by dragging their title bar or borders. In addition, users can maximize and minimize design windows by clicking on a respective icon on the window's title bar.

Window Switching Selecting *Main :: Window ⇒ Next* or *Main :: Window ⇒ Previous* will switch the focus to and activate the next/previous design window in the list of opened windows. Using these actions, users can cycle through the list of windows. Design windows are ordered in the window list according to the order in which they were opened. In addition, users can activate and raise any of the opened design windows by clicking into the window.

Finally, the bottom of the *Main :: Window* menu contains entries for all currently opened design windows. Selecting any of these menu entries will activate and raise the corresponding design window.

Window Toggling Selecting *Main :: Window ⇒ Project Manager* or *Main :: Window ⇒ Output Window* will toggle (turn on and off) displaying of Project and Output Windows, respectively.

References

- [1] R. Dömer, A. Gerstlauer and D. D. Gajski. *SpecC Language Reference Manual, Version 2.0*, SpecC Technology Open Consortium (STOC), Japan, December 2002.
- [2] R. Dömer. *The SpecC Internal Representation (V2.0.3)*, Technical Report 03-21, University of California, Irvine, January 1999.
- [3] J. Peng, A. Gerstlauer, K. Ramineni, R. Dömer and D. D. Gajski. *System-On-Chip Specification Style Guide*, Technical Report CECS-TR-03-21, Center for Embedded Computer Systems, University of California, Irvine, 2003.

A Manual Pages

This appendix contains the documentation in the form of manual pages for external, third-party tools used by SCE.

A.1 **scc** - SpecC Compiler

NAME

scc – SpecC Compiler

SYNOPSIS

```
scc -h  
scc design [ command ] [ options ]
```

DESCRIPTION

scc is the compiler for the SpecC language. The main purpose of **scc** is to compile a SpecC source program into an executable program for simulation. Furthermore, **scc** serves as a general tool to translate SpecC code from various input to various output formats which include SpecC source text, SpecC binary files in SpecC Internal Representation format, and other compiler intermediate files.

Using the first command syntax as shown in the synopsis above, a brief usage information and the compiler version are printed to standard output and the program exits. Using the second command syntax, the specified *design* is compiled. By default, **scc** reads a SpecC source file, performs preprocessing and builds the SpecC Internal Representation (SIR). Then, C++ code is generated, compiled and linked into an executable file to be used for simulation. However, the subtasks performed by **scc** are controlled by the given *command* so that, for example, only partial compilation is performed with the specified *design*.

On successful completion, the exit value 0 is returned. In case of errors during processing, an error code with a brief diagnostic message is written to standard error and the program execution is aborted with the exit value 10.

For preprocessing and C++ compilation, **scc** relies on the availability of an external C++ compiler which is used automatically in the background. By default, the GNU compiler **gcc/g++** is used.

ARGUMENTS

design specifies the name of the design; by default, this name is used as base name for the input file and all output files;

COMMAND

The *command* has the format - *suffix1* 2 *suffix2*, where *suffix1* and *suffix2* specify the format of the main input and output file, respectively. This command also implies the compilation steps being performed. By default, the command `-sc2out` is used which specifies reading a SpecC source file (e.g. `design.sc`) and generating an executable file (e.g. `a.out`) for simulation. All necessary intermediate files (e.g. `design.cc`, `design.o`) are generated automatically.

Legal command suffixes are:

- sc* SpecC source file (default: *design.sc*)
- si* preprocessed SpecC source file (default: *design.si*)
- sir* binary SIR file in SpecC Internal Representation format (default: *design.sir*)
- cc* C++ simulation source file (default: *design.cc*)
- h* C++ simulation header file (default: *design.h*)
- cch* both, C++ simulation source file and C++ header file (default: *design.cc* and *design.h*)
- o* linker object file (default: *design.o*)
- out* executable file for simulation (default: *design*); however, with the `-ip` option, a shared library will be produced (default: *libdesign.so*)

OPTIONS

- `-v` | `-vv` | `-vvv` increase the verbosity level so that all tasks performed are logged to standard error (default: be silent); at level 1, informative messages for each task performed are displayed; at level 2, additionally input and output file names are listed; at level 3, very detailed information about each executed task is printed;
- `-w` | `-ww` | `-www` increase the warning level so that warning messages are enabled (default: warnings are disabled); four levels are supported ranging from only important warnings (level 1) to pedantic warnings (level 4); for most cases, warning level 2 is recommended (`-ww`);

- `-g` enable debugging of the generated simulation code (default: no debugging code); this option disables optimization;
- `-O` enable optimization of the generated simulation code (default: no optimization); this option disables debugging;
- `-ip` enable intellectual property (IP) mode; when generating a SIR binary or SpecC text file, only declarations of symbols marked public will be included (the public interface of an IP is created); when generating C++ code, non-public symbols will be output so that they will be invisible outside the file scope; when compiling or linking, the compiler and linker are instructed to create a shared library instead of an executable file (creation of an IP simulation library);
- `-n` suppress creation of new log information when generating the output SIR file (default: update log information); see also section ANNOTATIONS below;
- `-sl` suppress source line information (preprocessor directives) when generating SpecC or C++ source code (default: include source line directives);
- `-sn` suppress all annotations when generating SpecC source code (default: include annotations);
- `-st tabulator stepping` set the tabulator stepping for SpecC/C++ code generation; this setting is used for code indentation; a value of 0 will disable the indentation of the generated code (default: 4);
- `-sT system tabulator stepping` set the system tabulator stepping (`\t`) for SpecC/C++ code generation; if set, tab characters will be used for indentation; if a value of 0 is specified, only spaces will be used for indentation (default: 8);
- `-sw line wrapping` set the column for line wrapping; in code generation, any line longer than this value is subject to line wrapping; if a value of 0 is specified, no line wrapping will be performed (default: 70);
- `-i input file` specify the name of the input file explicitly (default: `design.suffix1`); the name `'-'` can be used to specify reading from standard input;

- o output file* specify the name of the final output file explicitly (default: *design.suffix2*); the name '-' can be used to specify writing to standard output;
- D* do not define any standard macros; by default, the macro *__SPECC__* is defined automatically (it is set to 1); furthermore, implementation dependent macros may be defined; this option suppresses the definition of all these macros;
- Dmacrodef* define the preprocessor macro *macrodef* to be passed to the preprocessor;
- U* do not undefine any macros; by default, few macros are undefined automatically (in order to allow C/C++ standard header files to be used); this option is implementation dependent;
- Uundef* undefine the preprocessor macro *undef* which will be passed to the preprocessor as being undefined; the macro *undef* will be undefined after the definition of all command-line macros; this allows to selectively suppress macros from being defined in the preprocessing stage;
- I* clear the standard include path; by default, the standard include path consists of the directory *\$/SPECC/inc*; this option suppresses the default include path;
- Idir* append *dir* to the include path (extend the list of directories to be searched for including source files); include directories are searched in the order of their specification; unless suppressed by option *-I*, the standard include path is automatically appended to this list; by default, only the standard include directories are searched;
- L* clear the standard library path; by default, the standard library path consists of the directory *\$/SPECC/lib*; this option suppresses the default library path;
- Ldir* append *dir* to the library path (extend the list of directories to be searched for linker libraries); the library path is searched in the specified order; unless suppressed by option *-L*, the standard library path is automatically appended to this list; by default, only the standard library path is searched;

- l* when linking, do not use any standard libraries; the default libraries are displayed when calling the compiler with the *-h* option; the *-l* option suppresses linking against these standard libraries;
- llib* pass *lib* as a library to the linker so that the executable is linked against *lib*; libraries are linked in the specified order; unless suppressed by option *-l*, the standard libraries are automatically appended to this list; by default, only standard libraries are used;
- P* reset the import path; clear the list of directories to be searched for importing files; by default, the current directory is searched first, followed by the standard import directory `$SPECC/import`; this option suppresses this standard import path;
- Pdir* append *dir* to the import path, extending the list of directories to be searched for importing files; import directories are searched in the order of their specification; unless suppressed by option *-P*, the standard search path is automatically appended to this list; by default, only the standard import path is searched;
- xpp preprocessor_call* redefine the command to be used for calling the C preprocessor (default: `"g++ -E -x c %p %i -o %o"`); the preprocessor call must contain three markers `%p`, `%i` and `%o`, which indicate the options and file names used in the call; in the specified string, the `%p` marker will be replaced with the list of specified preprocessor options; the `%i` and `%o` markers will be replaced with the actual input and output filenames, respectively;
- xcc compiler_call* redefine the command to be used for calling the C/C++ compiler (default: `"g++ -c %c %i -o %o"`); the compiler call must contain three markers `%c`, `%i` and `%o`, which indicate the options and file names used in the call; in the specified string, the `%c` marker will be replaced with the list of specified compiler options; the `%i` and `%o` markers will be replaced with the actual input and output filenames, respectively;
- xld linker_call* redefine the command to be used for calling the linker (default: `"g++ %i -o %o %l"`); the linker call must contain three markers `%l`, `%i` and `%o`, which indicate the options and file names used in the call; in the specified string, the `%l` marker will be replaced with the list of specified linker options; the `%i` and `%o` markers

will be replaced with the actual input and output filenames, respectively;

-xp preprocessor_option pass an option directly to the C/C++ preprocessor (default: none);

-xc compiler_option pass an option directly to the C/C++ compiler (default: none);

-xl linker_option pass an option directly to the linker (default: none);

ENVIRONMENT

SPECC is used to determine the installation directory of the SpecC environment where SpecC standard include files (directory `$SPECC/inc`), SpecC standard import files (directory `$SPECC/import`), and SpecC system libraries (directory `$SPECC/lib`) are located.

SPECC_LICENSE_FILE determines the license file (path and file name) to be used by the SpecC environment; if undefined, the environment variable *SPECC* is used as the path to the license file called "license.sce"; if neither *SPECC_LICENSE_FILE* nor *SPECC* exist, the file "license.sce" is searched in the current directory;

ANNOTATIONS

The following SpecC annotations are recognized by the compiler:

_SCE_LOG contains the log information of the SIR file; this global annotation is created and maintained automatically by the SpecC compiler and the SpecC tool set and can be used to determine the origin and the operations performed on the design model; *_SCE_LOG* is a composite annotation consisting of a list of log entries, ordered by time of creation; each log entry consists of a time stamp, command line, source file, version info, and an optional comment;

_SCC_RESERVED_SIZE for external behaviors and channels (IP components), this indicates the size reserved in the C++ class for internal use; the annotation type is unsigned int; if found at class definitions, this annotation is checked automatically for reasonable values; for IP declarations, the annotation can be created automatically with the *-ip* option;

_SCC_PUBLIC for global symbols, this annotation indicates whether the symbol is public and will be visible in a shared library; the annotation type is bool; this annotation only is recognized with the `-ip` option;

VERSION

The SpecC compiler `scc` is version 2.2.0.

AUTHOR

Rainer Doemer <doemer@ics.uci.edu>

COPYRIGHT

(c) 1997-2004 CECS, University of California, Irvine

SEE ALSO

`gcc(1)`, `g++(1)`, `sir_delete(1)`, `sir_depend(1)`, `sir_import(1)`, `sir_isolate(1)`,
`sir_list(1)`, `sir_note(1)`, `sir_rename(1)`, `sir_strip(1)`, `sir_tree(1)`, `sir_wrap(1)`

BUGS, LIMITATIONS

Variables of enumerator type cannot be initialized at the time of their declaration. The SpecC compiler issues a (false) error message in this case. As a simple work-around, however, enumerator variables can be initialized by use of standard assignment statements at the beginning of their lifetimes.

B Project File XML Schema

Elements	Complex types
sce	CompilerOption ModelItem SimpleItem

B.1 Element sce


diagram	
children	project compiler models imports sources
source	<pre> <xs:element name="sce"> <xs:annotation> <xs:documentation>Root</xs:documentation> </xs:annotation> <xs:complexType> <xs:sequence> <xs:element name="project" type="xs:string"/> <xs:element name="compiler"> <xs:complexType> <xs:sequence> <xs:element name="option" type="CompilerOption" minOccurs="0" maxOccurs="unbounded"/> </xs:sequence> </xs:complexType> </xs:element> </xs:sequence> </xs:complexType> </xs:element> </pre>


```

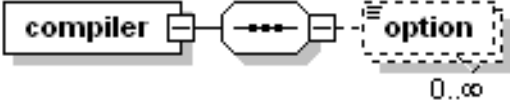
<xs:element name="models">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="item" type="ModelItem"
        minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
<xs:element name="imports">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="item" type="SimpleItem"
        minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
<xs:element name="sources">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="item" type="SimpleItem"
        minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:sequence>
</xs:complexType>
</xs:element>

```

B.1.1 Element sce/project

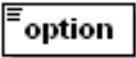
diagram	
type	xs:string
source	<xs:element name="project" type="xs:string"/>

B.1.2 Element sce/compiler

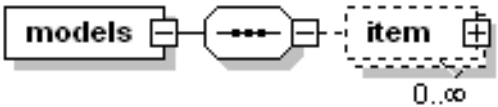
diagram	
children	option

source	<pre> <xs:element name="compiler"> <xs:complexType> <xs:sequence> <xs:element name="option" type="CompilerOption" minOccurs="0" maxOccurs="unbounded"/> </xs:sequence> </xs:complexType> </xs:element> </pre>
--------	---

Element sce/compiler/option

diagram													
type	CompilerOption												
attributes	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Use</th> <th>Default</th> <th>Fixed</th> <th>Annotation</th> </tr> </thead> <tbody> <tr> <td>name</td> <td></td> <td>required</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Name	Type	Use	Default	Fixed	Annotation	name		required			
Name	Type	Use	Default	Fixed	Annotation								
name		required											
source	<pre> <xs:element name="option" type="CompilerOption" minOccurs="0" maxOccurs="unbounded"/> </pre>												

B.1.3 Element sce/models

diagram	
children	item
source	<pre> <xs:element name="models"> <xs:complexType> <xs:sequence> <xs:element name="item" type="ModelItem" minOccurs="0" maxOccurs="unbounded"/> </xs:sequence> </xs:complexType> </xs:element> </pre>

Element sce/models/item

diagram						
type	ModelItem					
children	item option					
attributes	Name	Type	Use	Default	Fixed	Annotation
	name level	xs:string	required optional			<i>Abstraction level of the design model.</i>
	design	xs:string	optional			<i>Design name. Defaults to base file name.</i>
	input	xs:string	optional			<i>File name of the intermediate design model at the input of refinement after import.</i>
	type	xs:string	optional	<empty>		<i>Automatic refinement command line (comma-separated elements) used to generate model. Empty if manually generated.</i>
source	<pre><xs:element name="item" type="ModelItem" minOccurs="0" maxOccurs="unbounded"/></pre>					

B.1.4 Element `sce/imports`

diagram	
children	item
source	<pre><xs:element name="imports"> <xs:complexType> <xs:sequence> <xs:element name="item" type="SimpleItem" minOccurs="0" maxOccurs="unbounded"/> </xs:sequence> </xs:complexType> </xs:element></pre>

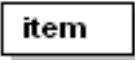
Element `sce/imports/item`

diagram													
type	SimpleItem												
attributes	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Use</th> <th>Default</th> <th>Fixed</th> <th>Annotation</th> </tr> </thead> <tbody> <tr> <td>name</td> <td>xs:string</td> <td>required</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Name	Type	Use	Default	Fixed	Annotation	name	xs:string	required			
Name	Type	Use	Default	Fixed	Annotation								
name	xs:string	required											
source	<pre><xs:element name="item" type="SimpleItem" minOccurs="0" maxOccurs="unbounded"/></pre>												

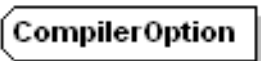
B.1.5 Element `sce/sources`

diagram	
children	item
source	<pre><xs:element name="sources"> <xs:complexType> <xs:sequence> <xs:element name="item" type="SimpleItem" minOccurs="0" maxOccurs="unbounded"/> </xs:sequence> </xs:complexType> </xs:element></pre>

Element `sce/sources/item`

diagram													
type	SimpleItem												
attributes	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Use</th> <th>Default</th> <th>Fixed</th> <th>Annotation</th> </tr> </thead> <tbody> <tr> <td>name</td> <td>xs:string</td> <td>required</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Name	Type	Use	Default	Fixed	Annotation	name	xs:string	required			
Name	Type	Use	Default	Fixed	Annotation								
name	xs:string	required											
source	<code><xs:element name="item" type="SimpleItem" minOccurs="0" maxOccurs="unbounded"/></code>												

B.2 ComplexType CompilerOption

diagram													
type	extension of xs:string												
properties	base xs:string												
used by	element <code>sce/compiler/options</code>												
attributes	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Use</th> <th>Default</th> <th>Fixed</th> <th>Annotation</th> </tr> </thead> <tbody> <tr> <td>name</td> <td></td> <td>required</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Name	Type	Use	Default	Fixed	Annotation	name		required			
Name	Type	Use	Default	Fixed	Annotation								
name		required											
source	<pre> <xs:complexType name="CompilerOption"> <xs:simpleContent> <xs:extension base="xs:string"> <xs:attribute name="name" use="required"> <xs:simpleType> <xs:restriction base="xs:string"> <xs:enumeration value="incpath"/> <xs:enumeration value="importpath"/> <xs:enumeration value="defines"/> <xs:enumeration value="undefines"/> <xs:enumeration value="options"/> </xs:restriction> </xs:simpleType> </xs:attribute> </xs:extension> </xs:simpleContent> </xs:complexType> </pre>												

B.3 ComplexType ModellItem

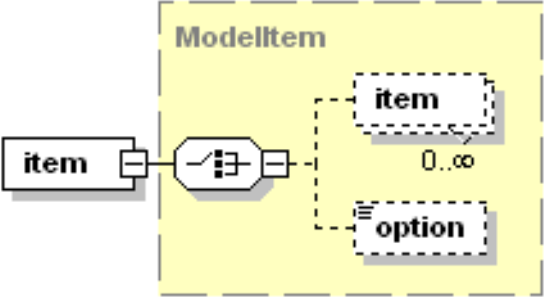
diagram						
type	extension of SimpleItem					
properties	base SimpleItem					
children	item option					
used by	elements sce/models/item ModellItem/item					
attributes	Name	Type	Use	Default	Fixed	Annotation
	name	xs:string	required			
	level		optional			<i>Abstraction level of the design model.</i>
	design	xs:string	optional			<i>Design name. Defaults to base file name.</i>
	input	xs:string	optional			<i>File name of the intermediate design model at the input of refinement after import.</i>
	type	xs:string	optional	<empty>		<i>Automatic refinement command line (comma-separated elements) used to generate model. Empty if manually generated.</i>
source	<xs:complexType name="ModellItem">					

```

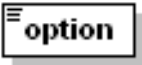
<xs:complexContent>
  <xs:extension base="SimpleItem">
    <xs:choice>
      <xs:element name="item" type="ModelItem" minOccurs="0"
        maxOccurs="unbounded"/>
      <xs:element name="option" minOccurs="0">
        <xs:complexType>
          <xs:simpleContent>
            <xs:extension base="xs:string">
              <xs:attribute name="name" type="xs:string"
                use="required" fixed="description"/>
            </xs:extension>
          </xs:simpleContent>
        </xs:complexType>
      </xs:element>
    </xs:choice>
    <xs:attribute name="level" use="optional">
      <xs:annotation>
        <xs:documentation>Abstraction level of the design model
        </xs:documentation>
      </xs:annotation>
      <xs:simpleType>
        <xs:restriction base="xs:string">
          <xs:pattern value="0x[a-z0-9A-Z]+L?"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:attribute>
    <xs:attribute name="design" type="xs:string" use="optional">
      <xs:annotation>
        <xs:documentation>Design name.
          Defaults to base file name.</xs:documentation>
      </xs:annotation>
    </xs:attribute>
    <xs:attribute name="input" type="xs:string" use="optional">
      <xs:annotation>
        <xs:documentation>File name of the intermediate design model
          at the input of refinement after import.
        </xs:documentation>
      </xs:annotation>
    </xs:attribute>
    <xs:attribute name="type" type="xs:string" use="optional" default="">
      <xs:annotation>
        <xs:documentation>Automatic refinement command line (comma-
          separated elements) used to generate model.
          Empty if manually generated</xs:documentation>
      </xs:annotation>
    </xs:attribute>
  </xs:extension>
</xs:complexContent>
</xs:complexType>

```


B.3.1 Element ModelItem/item

diagram						
type	ModelItem					
children	item option					
attributes	Name	Type	Use	Default	Fixed	Annotation
	name level	xs:string	required optional			<i>Abstraction level of the design model.</i>
	design	xs:string	optional			<i>Design name. Defaults to base file name.</i>
	input	xs:string	optional			<i>File name of the intermediate design model at the input of refinement after import.</i>
	type	xs:string	optional	<empty>		<i>Automatic refinement command line (comma-separated elements) used to generate model. Empty if manually generated.</i>
source	<pre><xs:element name="item" type="ModelItem" minOccurs="0" maxOccurs="unbounded"/></pre>					

B.3.2 Element Modelltem/option

diagram													
type	extension of xs:string												
attributes	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Use</th> <th>Default</th> <th>Fixed</th> <th>Annotation</th> </tr> </thead> <tbody> <tr> <td>name</td> <td>xs:string</td> <td>required</td> <td></td> <td></td> <td>description</td> </tr> </tbody> </table>	Name	Type	Use	Default	Fixed	Annotation	name	xs:string	required			description
Name	Type	Use	Default	Fixed	Annotation								
name	xs:string	required			description								
source	<pre><xs:element name="option" minOccurs="0"> <xs:complexType> <xs:simpleContent> <xs:extension base="xs:string"> <xs:attribute name="name" type="xs:string" use="required" fixed="description"/> </xs:extension> </xs:simpleContent> </xs:complexType> </xs:element></pre>												

B.4 ComplexType SimpleItem

diagram													
used by	<table border="1"> <thead> <tr> <th>elements</th> <th>sce/imports/item</th> <th>sce/sources/item</th> </tr> </thead> <tbody> <tr> <td>complexType</td> <td></td> <td>Modelltem</td> </tr> </tbody> </table>	elements	sce/imports/item	sce/sources/item	complexType		Modelltem						
elements	sce/imports/item	sce/sources/item											
complexType		Modelltem											
attributes	<table border="1"> <thead> <tr> <th>Name</th> <th>Type</th> <th>Use</th> <th>Default</th> <th>Fixed</th> <th>Annotation</th> </tr> </thead> <tbody> <tr> <td>name</td> <td>xs:string</td> <td>required</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Name	Type	Use	Default	Fixed	Annotation	name	xs:string	required			
Name	Type	Use	Default	Fixed	Annotation								
name	xs:string	required											
source	<pre><xs:complexType name="SimpleItem"> <xs:attribute name="name" type="xs:string" use="required"/> </xs:complexType></pre>												