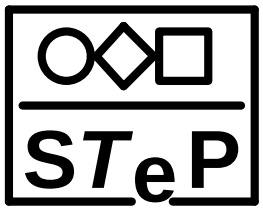

STeP

The Stanford Temporal Prover
Educational Release

Version 1.3

USER'S MANUAL

Nikolaj Bjørner, Anca Browne, Eddie Chang, Michael Colón,
Bernd Finkbeiner, Arjun Kapur, Zohar Manna,
Henny B. Sipma, and Tomás E. Uribe

October, 1997

Computer Science Department
Stanford University
Stanford, California 94305

Contents

1 Introduction **1**

1.1 Reading this Manual . . . 2
1.2 Starting STeP . . . 3
1.3 Command-line options . . . 4
1.4 Interacting with STeP . . . 4
1.5 Terms and Definitions . . . 7
1.6 Feedback . . . 11
1.7 Acknowledgements . . . 11

2 Systems and Specifications **13**

2.1 Types . . . 13
2.2 Expressions . . . 14
2.3 Specification Files . . . 18
2.4 Declarations . . . 19
2.4.1 Type Declarations . . . 19
2.4.2 Value and Macro Declarations . . . 21
2.4.3 Variable Declarations . . . 22
2.4.4 Rewrite and Simplification Rules ⋆ . . . 24
2.4.5 Channel Operations ⋆ . . . 25
2.5 SPL Programs . . . 26
2.6 Transition Systems . . . 30
2.6.1 Fair Transition Systems . . . 30
2.6.2 Modular Transition Systems . . . 32
2.6.3 Interface Declaration . . . 32
2.6.4 Transition System Syntax . . . 35
2.6.5 Clocked Transition Systems ⋆⋆ . . . 36
2.6.6 Hybrid Transition Systems ⋆⋆ . . . 37
2.6.7 Example . . . 39
2.7 When Parsing Fails . . . 39

3 The Top-Level Interface **43**

3.1 Menu Options . . . 45
3.2 Action Buttons . . . 57

4 Verification and Logical Rules **59**
4.1 The Proof Search 59
4.2 Verification Rules 60
4.3 WPC and Strengthening 63
4.4 Logical Rules 63
4.5 The Model Checker 64

5 The Verification Diagram Editor **67**
5.1 Definitions 67
5.2 Wait-for Diagrams 68
5.3 Invariance Diagrams 68
5.4 Chain Diagrams ⋆ 69
5.5 Compound Nodes 70
5.6 Verification Diagram Editor: Interface 72
5.7 Hierarchical Verification Diagrams ⋆ 75

6 Automatic Theorem-Proving **77**
6.1 The Automatic Simplifier 77
6.2 BDD Simplification 79
6.2.1 BDD split ⋆ 80
6.3 Tactics ⋆ 80
6.3.1 Composition Options 81
6.3.2 Top-level Prover Tactics 83
6.3.3 Interactive Prover Tactics 84
6.3.4 Interrupting Tactics 85
6.3.5 Commonly used Tactics 86

7 Automatic Generation of Invariants **87**
7.1 Local Invariants 87
7.2 Linear Invariants 88
7.3 Polyhedral Invariants ⋆ 88

8 The Interactive Prover ⋆ **89**
8.1 Proof Structure 89
8.2 User Interface 90
8.3 Interactive Prover Rules 93

9 Basic Tutorial **101**
9.1 Preliminaries 101
9.2 MUX-SEM: Mutual Exclusion 102
9.2.1 Using B-INV 103
9.2.2 Using G-INV 103
9.2.3 Using the Model Checker 107
9.2.4 Using Verification Diagrams 108
9.2.5 Using Linear Invariants 109
9.3 Two-process Bakery 109

9.3.1 Static analysis 110
9.3.2 Mutual exclusion 110
9.3.3 1-bounded overtaking 110
9.3.4 Response 112

Appendix

A Computational Model **115**
A.1 Fair Transition Systems 115
A.2 Computations 116
A.3 SPL semantics 116

B Linear-Time Temporal Logic **123**

C Differences with the Book **125**
C.1 SPL Programs 125
C.2 Control Locations in STeP 125
C.3 Type Declarations 131
C.4 Initialization 131
C.5 Parameterized Programs 132

D Distribution and Installation **133**
D.1 Example Programs 134

E STeP Settings **137**

Chapter 1

Introduction

The Stanford Temporal Prover, STeP, supports the computer-aided formal verification of concurrent and reactive systems based on temporal specifications. Reactive systems maintain an ongoing interaction with their environment, and their specifications are typically expressed as constraints on their behavior over time. STeP is not restricted to finite-state systems, but combines *model checking* with *deductive methods* to allow for the verification of a broad class of systems, including parameterized (N-component) circuit designs, parameterized (N-process) programs, and programs with infinite data domains.

Finite-state systems can be automatically verified using model checking alone. For large finite-state or for infinite-state systems, for which model checking is not possible or feasible, STeP also uses deductive methods: *Verification rules* [Manna and Pnueli, 1995b] reduce temporal properties to first-order verification conditions. *Verification diagrams* [Manna and Pnueli, 1994] are a visual language for guiding, organizing and displaying proofs. They can be used to construct proofs hierarchically, starting from a high-level sketch and proceeding incrementally through layers of greater detail.

STeP also includes techniques for automatic *invariant generation* [Bjørner *et al.*, 1997a]. Deductive verification almost always relies on finding, for a given program and specification, suitably strong invariants and intermediate assertions. The system generates bottom-up invariants automatically by analyzing the program text. Combining them with the property to be proved, sufficiently detailed invariants can often be obtained to carry through the verification process.

Finally, STeP provides a collection of simplification and decision procedures that automatically check the validity of a large class of first-order and temporal formulas. This degree of automated deduction can handle many of the verification conditions that arise in deductive verification.

Figure 1.1 shows an overview of STeP. The main inputs are a reactive system (which can be a description of hardware or software) and a property to be proven about the system, represented by a temporal logic formula. Verification can be performed by the model checker or by deductive means. Model checking is automatic; deductive verification is to a large extent automatic for simple *safety* properties, while *progress* properties require more user guidance, usually given in the form of a verification diagram. In all cases, the automatic prover is responsible for generating and proving the required verification conditions. An interactive Gentzen-style theorem prover is used to establish those verification conditions

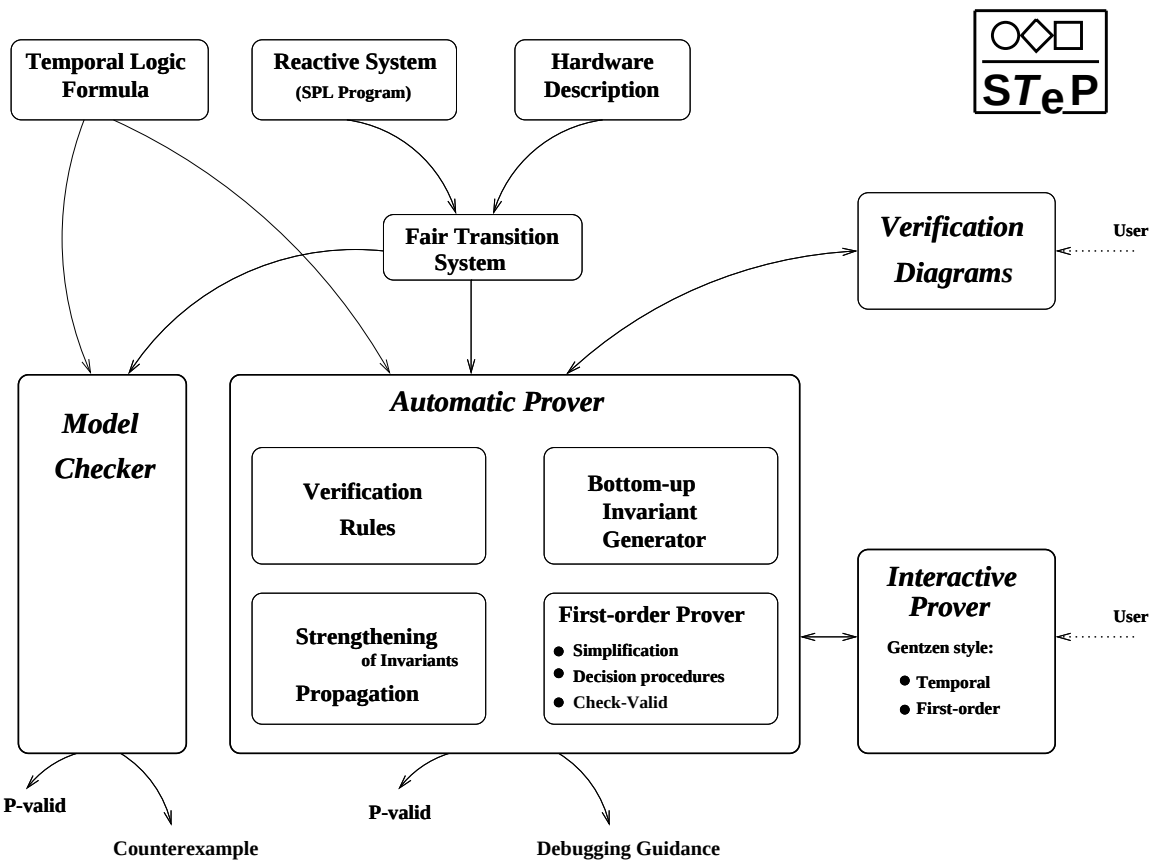

Figure 1.1: An overview of the STeP system.

that are not proved automatically.

STeP: The Educational Version

The current release of STeP is the *educational version*, and is a companion to the textbook [Manna and Pnueli, 1995b]. Like the book, it concentrates on safety properties, and is intended for use by students, instructors and researchers of formal methods. This is the updated manual for release version 1.3.

1.1 Reading this Manual

This manual is a complete reference guide for using STeP. For more background on STeP's framework for temporal specification and program verification, see [Manna and Pnueli, 1991] and [Manna and Pnueli, 1995b].

Following are some suggestions on reading this manual:

- **Getting started:** Section 1.2 describes how to get STeP up and running. Section 1.4 gives an overview of STeP and a flavor of how a typical session proceeds. Section 1.5

defines the main terms used throughout this manual.

- **The Tutorial:** This manual includes a basic tutorial on STeP, in Chapter 9. This tutorial is the recommended method for readers to familiarize themselves with the STeP interface and the different styles of verification in STeP.

  After the tutorial, you should be able to verify some of the other programs included in the `examples` directory as well as your own simple programs, using the rest of this manual as a reference.

Some sections are marked with a "$\star$" to indicate that they should be read only after the unmarked sections are understood. For instance, first-time users can postpone studying features such as tactics (Section 6.3) and user-defined rewrite and simplification rules (Section 2.4.4), until they feel comfortable with the rest of the system. Sections marked with "$\star\star$" are intended for advanced users and may be safely omitted by the beginning student.

Chapters 2-8 are the main reference for using the system. For those readers unfamiliar with [Manna and Pnueli, 1995b], Appendixes A and B summarize STeP's program and property specification languages. Appendix C presents the most important differences between STeP and [Manna and Pnueli, 1995b].

1.2 Starting STeP

The following are the steps for starting STeP:

1. Set your `STEP_DIR` environment variable to the full directory path where STeP is installed.[1] If this variable is not set, you will not be able to run STeP. (If you are installing the system, you can edit the `bin/STeP` file to correctly set this value for all users.)

2. Make sure that the machine where you plan to run STeP has permission to open windows on your display. To make sure this is the case, execute, in your X-windows console, the command "`xhost +`".

3. Add `STEP_DIR` to your *path*. In this way, when executing STeP you don't have to type `STEP_DIR/STeP`, but just `STeP`.

4. Execute STeP, by running `STeP`. If your `DISPLAY` variable and `xhost` permissions are set correctly, the STeP main window will appear on your screen.

A number of other environment variables can be optionally set according to each user's preferences; these are described in Appendix E.

Appendix D gives details on the installation of STeP. If you have problems installing or running STeP, please send mail to `step-bugs@cs.stanford.edu`.

[1] Environment variables are set with "`setenv` *variable value*" Unix commands, which you can add to your `.cshrc` file for convenience.

1.3 Command-line options

STeP can be invoked with a variety of command-line options to preload systems and specifications, or to use STeP as a stand-alone SPL translator or simplifier. The options are a sequence consisting of the following flags and arguments:

-display <display-address>:0 set the display

-prog <program> load program or transition system

-spec <specification(s)> load specification(s)

-spl2trans <program> compile spl program to transition system in std_out

-modelcheck <logfile> Apply the model checker to the loaded specifications.

-simplify <options> <specification(s)> simplify specification files where options are:

- **-a** Arithmetical simplification
- **-ac** Use AC utilities (not many available)
- **-b** Bdd simplification
- **-c** Conditional simplification rules
- **-e** Use basic equality (congruence closure)
- **-f** Fast simplify
- **-s** Standard simplification (default)
- **-v** First-order validity checking

-exit Do not process further arguments or open STeP

-help print usage information (this list)

1.4 Interacting with STeP

STeP has three main interface components: the *Top-level Prover*, from which verification sessions are managed and verification rules invoked; the *Interactive Prover*, used to prove the validity of first-order and temporal-logic formulas; and the *Verification Diagram Editor*, for the creation of Verification Diagrams. Figure 1.2 shows all STeP's interfaces with the program MUX-SEM loaded.

Verification Session—An Overview

A verification session begins by loading a program or a transition system that describes the system of interest and entering a temporal-logic formula that expresses one or more system properties to be proved.

The formula becomes the *top* or *root goal* of a *proof tree*. There are now several ways to proceed. The most common route is to apply a *verification rule* (Section 4.2), which reduces the formula to simpler first-order or temporal formulas, thus generating the first

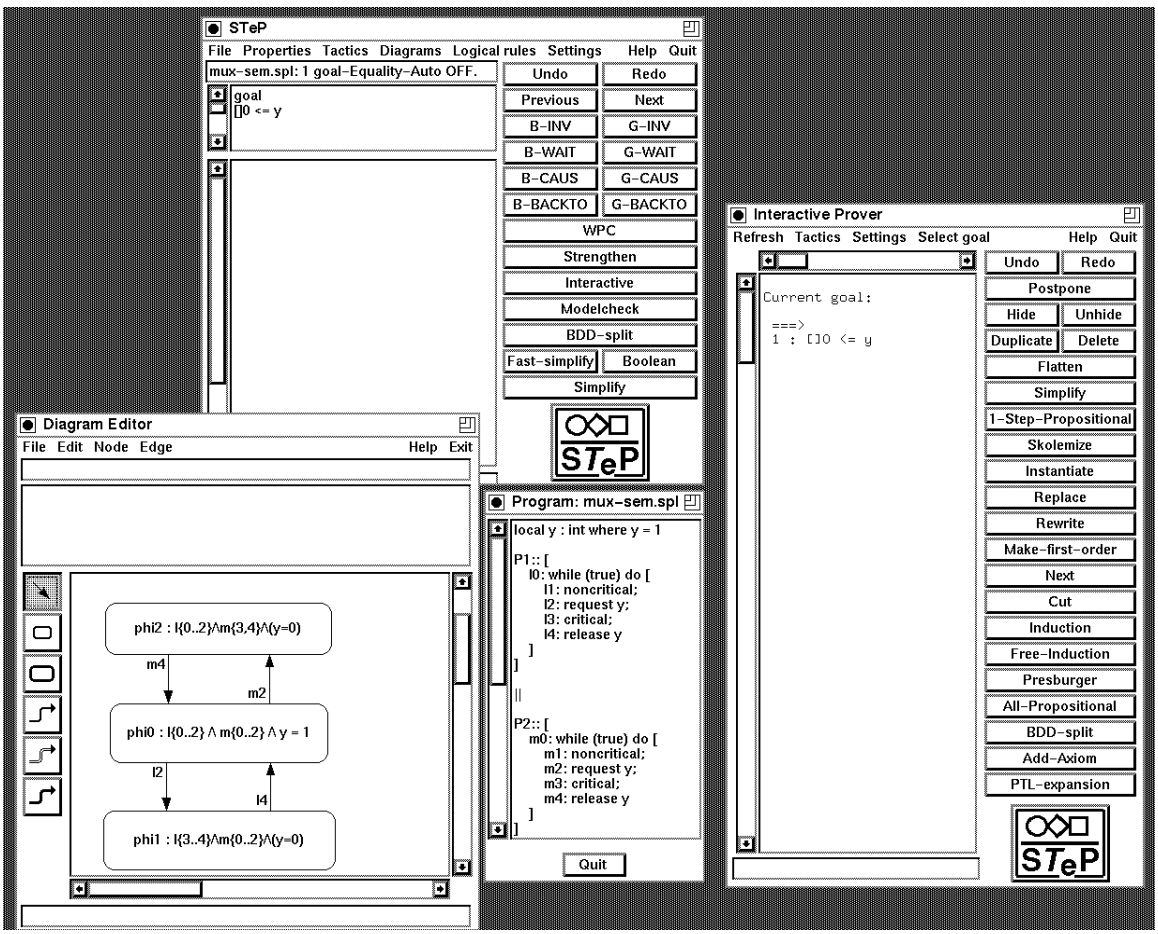

Figure 1.2: Overview of STeP's interfaces

level of *subgoals* of the proof tree. These subgoals are numbered accordingly: 1.1, 1.2, 1.3, etc. If all subgoals are established automatically by the Simplifier (i.e., the Simplifier can reduce them to *true*), all branches are *closed* and the proof is finished.

If the Simplifier is not able to prove one of the subgoals, we have one of the following possibilities:

1. The subgoal is not valid for the given program. In this case, the proof cannot succeed.

2. The property is valid for the program, but not generally valid. In this case you will have to (a) strengthen the original formula, or (b) first prove some auxiliary properties of the program, relative to which the problematic verification condition becomes valid. In the latter case you may start a new proof tree for the auxiliary property and return to the original property after the auxiliary proof is finished. This approach reflects the incremental proof style advocated in [Manna and Pnueli, 1995b].

   Another way to establish auxiliary properties is through the automatic generation of

invariants. STeP includes three such methods, which generate properties that often help in proving verification conditions (Chapter 7).

3. The property is generally valid, but the simplifier is unable to reduce it to *true*. In this case, the Interactive Prover (Chapter 8) can be used to establish the given property with user guidance. If this is successful, the given subgoal is closed. Partial results may also be returned to the Top-level Prover and incorporated in the current proof tree as a new subgoal.

Another way to verify a program property is to use a *verification diagram*. A verification diagram represents a set of verification conditions sufficient to establish a given property. Diagrams are created by the Verification Diagram Editor (Chapter 5) and subsequently transferred to the Top-level Prover.

Finally, any property or subgoal may be established by the Model Checker, through the exploration of the state space of the system (Section 4.5). Although model checking is only guaranteed to succeed for finite-state programs whose state space is not too large, it is also generally useful to identify *erroneous* properties, since it provides a counterexample computation if the property is not valid for the given system.[2]

STeP can also be used as a theorem prover, independently of any particular system verification task. In this case you only need to load a specification, and only the Simplifier and the Interactive Prover are used.

Verification Session—A Quick Tour

To make things more concrete we will give an example of an actual STeP session. For a step-by-step description of the commands needed to run this example, see the basic STeP tutorial in Chapter 9.

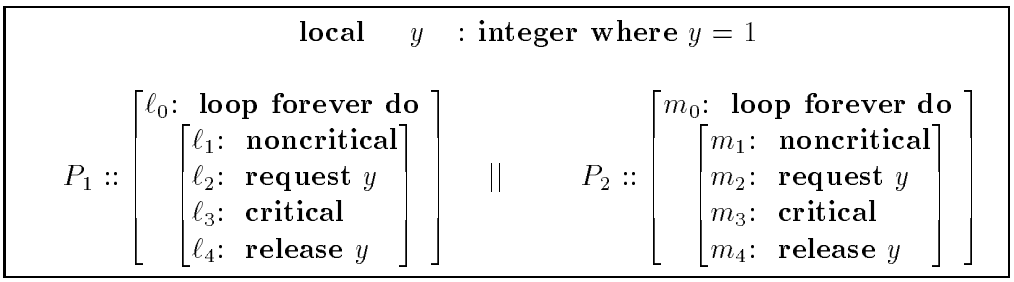

Figure 1.3: Program MUX-SEM (mutual exclusion by semaphores)

Consider program MUX-SEM, shown in Figure 1.3. It implements mutual exclusion using the semaphore y; **request** decrements the semaphore, when it is positive, and **release** increments the semaphore. We want to prove the following properties of this program:

$$\varphi_1 : \Box(y \geq 0)$$

[2]Restrictions may be placed on the system to make model checking more feasible—see Section 4.5.

$$\varphi_2 : \Box(\neg(\ell_3 \wedge m_3))$$

The first property, stating that y is always greater than or equal to 0, is proved automatically using rule B-INV (Section 4.2) and decision procedures for linear arithmetic. The property is then added to the list of background properties and may be used in subsequent proofs.

The second property expresses mutual exclusion and is not *inductive*, that is, it is not preserved by all transitions. In particular, the verification conditions for transitions ℓ_2 and m_2 are not valid, and indeed when we apply rule B-INV the Simplifier does not reduce them to *true*. We now have several options. Since this is a finite-state program with a small number of states, model checking has a good chance of succeeding; indeed, the Model Checker can quickly verify this property.

Another option is to generate some auxiliary invariants. One of the *linear invariants* automatically generated for this system is:

$$I : \Box(\ell_0 + \ell_1 + \ell_2 = m_3 + m_4 + y)$$

Unfortunately, property φ_2 is not inductive relative to this invariant, so the Simplifier is still unable to prove the verification conditions for ℓ_2 and m_2. Repeating the application of B-INV with all linear invariants active for simplification, you will notice that simplification is significantly slower, since the background properties are being used in simplification.

However, property φ_2 is directly implied by the invariant I and the property φ_1 we proved earlier. Therefore, we can apply rule MON-I (Section 4.4) directly to the top-level goal, and finish the proof after one simplification step.

Another option is to use rule G-INV (Section 4.2), supplementing the verification with a strengthened assertion, e.g.,

$$\ell_3 + \ell_4 + m_3 + m_4 + y = 1$$

This assertion implies mutual exclusion and is inductive (the locations have been *arithmetized*, e.g., the value of ℓ_3 is 1 if control resides at ℓ_3, 0 otherwise). All verification conditions are reduced to *true* by the Simplifier, thus completing the proof.

Yet another option is a proof by verification diagram. The lower left corner of Figure 1.2 shows a state-partition diagram for program MUX-SEM that implies mutual exclusion; the property $\Box(\neg(\ell_3 \wedge m_3))$ is proved automatically by invoking the *Verification Diagram Rule* on this diagram. This produces 35 different verification conditions, all of which are reduced to *true*.

1.5 Terms and Definitions

STeP Sub-systems

In STeP we identify the following sub-systems, each of which operates relatively independently of the rest:

Top-level Prover (Chapter 3) The Top-level Prover is the main interface to a STeP session, from which all other activities are initiated.

Interactive Prover (Chapter 8) The Interactive Prover is invoked from the Top-level Prover to prove a particular goal or subgoal. When the Interactive Prover is active the Top-level Prover is suspended. The Interactive Prover provides a Gentzen-style proof system, which is guided by the user. While the Interactive Prover is normally used to complete the proof of a given subgoal, partial results from the Interactive Prover can also be transferred back to the Top-level Prover.

Verification Diagram Editor (Chapter 5) The Verification Diagram Editor is also invoked from the Top-level Prover, but can be run concurrently with it. It is used to create and edit verification diagrams, which can be given to the Top-level Prover to use in a proof.

Model Checker (Section 4.5) The Model Checker can be invoked from the Top-level Prover on any given temporal subgoal. If the Model Checker establishes the P-validity of the formula, the corresponding goal is proved. Otherwise, a counterexample computation (i.e. one that violates the formula) is returned.

Simplifier (Section 6.1) The Simplifier is a collection of decision procedures and rewrite rules used to reduce and prove general verification conditions. The Simplifier may be invoked by the user, both from the Top-level Prover and the Interactive Prover, to simplify single goals or subgoals. In the Top-level Prover, the Simplifier uses all axioms, invariants, and previously proven properties that are currently active. The power of the Simplifier, i.e., the trade-off between power and speed, is controlled by a number of global settings.

STeP Activities

We identify a few (nested) levels of activities, each with its own environment:

STeP session At the top level is the STeP session, which lasts for the entire period STeP is running. Simplification settings are part of the STeP session environment. They can be changed at any time, but are not affected by loading new systems or moving on to new goals.

System verification session The next level down is the *system verification session*, which lasts from the time a program or system is loaded until a new one is loaded. Proven properties are accumulated as long as the session lasts. Thus both the system and all the properties proven about it, as well as the axioms loaded for it, are part of the system verification session environment. All properties proven are valid relative to the current system. A system verification session may be terminated explicitly by loading a new system or resetting STeP.

Theorem-proving session At the same level as the system verification session is the *theorem-proving session*. The difference is that a theorem-proving session environment does not include a system. Thus, properties are proved to be generally valid in this case, and only the Simplifier and the Interactive Prover can be used to prove formulas.

Goal session A system verification or theorem-proving session may contain multiple *goal sessions*. Each unfinished proof of a goal is represented by a tree with the goal as its root. At any given time, only one tree is being worked on. The user can move from one tree to the next in the middle of a proof, and later return to the unfinished proof, e.g., when some auxiliary properties have been established. A goal session ends when the goal is proved valid or is abandoned by the user.

Properties

We use different terms for properties depending on their status in a verification or theorem-proving session. We list the terms here in alphabetical order:

Axiom An *axiom* is provided by the user and is assumed to be valid throughout a system verification or theorem-proving session. An axiom is considered to be a background property.

Background property A *background property* is assumed or proven to be valid for the program in a system verification session, and assumed or proven to be generally valid in a theorem-proving session. The user controls their use in simplification by activating and de-activating them. Background properties may also be deleted altogether from a verification session.

If active, background properties are used by the Simplifier in the simplification process of goals and subgoals at the Top-level Prover, and can be added to the current sequent in the Interactive Prover.

Background properties include axioms entered by the user, properties already proved in the same verification session, and automatically generated invariants.

Note that the set of background properties is lost upon termination of a system-verification or theorem-proving session.

Goal A *goal* is a formula that we want to prove generally valid or valid over a given program, and becomes the root of a proof tree.

Invariant An *invariant* is a formula of the form $\Box p$, stating that p is true at every state of a program or system computation. Invariants can be independently proved, or automatically generated by one of three methods: local, linear, or polyhedral invariant generation. Proven invariants are added to the set of background properties, and will be used by the simplifier when simplifying new formulas. They can also be used as axioms in the interactive prover.

Subgoal A *subgoal* is a formula that results from the application of a verification rule to a goal or another subgoal. It is a node of the proof tree. An *open* subgoal is a leaf of the proof tree that has not been reduced to *true*.

Verification condition A *verification condition* is a formula that results from the application of a verification rule to a goal or a subgoal. Verification conditions are a subset of the subgoals.

System Description and Specification

STeP accepts a system description in two different formats, namely as an SPL program or as a fair transition system. Specifications are entered as linear-time temporal logic formulas.

SPL program A system description may be given in the form of an SPL program (*Simple Programming Language* program). The syntax of well-formed SPL programs is given in Section 2.5. SPL programs are parsed into fair transition systems. The semantics of SPL are given in Appendix A.

Fair transition system An alternative way to provide a system description is by means of a fair transition system. A *fair transition system* is a tuple with the following components:

- V, a set of system variables. These include all program variables.
- Θ, the initial condition.
- $\mathcal{T}$, a set of transitions.

Each transition is a relation between states, specifying how the system can move from one state to the next. Section 2.6 describes how transition systems are specified in STeP. Appendix A describes fair transition systems in more detail.

Auxiliary and system variables Auxiliary variables are declared in specifications. System variables get declared in system descriptions and participate in the transition relations of the described transition system.

Flexible and rigid variables The value of a *flexible* variable may change over time, while that of a *rigid* variable is assumed to be fixed throughout a computation. Auxiliary variables in STeP can be declared to be flexible or rigid. System variables declared as `in` are rigid, while those declared as `out` and `local` are flexible.

Note that the flexible or rigid status of variables has a significant effect on whether terms under the scope of temporal operators can be replaced by others during interactive and automatic theorem-proving.[3]

Future operators Future temporal operators are $\Box$ (Henceforth), $\Diamond$ (Eventually), $\mathcal{U}$ (Until), $\mathcal{W}$ (Waiting-for also called Unless) and $\bigcirc$ (Next).

In STeP's ASCII representation, these are `[]`, `<>`, `Until`, `Awaits`, and `()` (see Appendix B), respectively.

Past operators Past operators are $\boxminus$ (So-far), $\diamondminus$ (Once), $\mathcal{S}$ (Since), $\mathcal{B}$ (Back-to), and $\ominus$ (Previously).

In STeP's ASCII representation, these are `[-]`, `<->`, `Since`, `Backto`, and `(-)` (see Appendix B), respectively.

Past formula A *past formula* is a temporal-logic formula that does not contain any future operators, that is, it only contains state formulas and past operators.

[3] For example it is in general not possible to instantiate an existentially bound rigid variable by a term containing flexible variables.

1.6 Feedback

If you have problems running STeP or find bugs in the system, please send mail to

`step-bugs@cs.stanford.edu`.

If you have questions or feedback, send mail to

`step-comments@cs.stanford.edu`.

There is a mailing list for STeP users, which will be used to distribute updated information about the system. To be added or removed from the list, send mail to

`step-request@cs.stanford.edu`.

1.7 Acknowledgements

STeP has been designed and implemented by Nikolaj Bjørner, Anca Browne, Eddie Chang and Tomás Uribe. Henny Sipma, Bernd Finkbeiner, Mark Pichora, and Arjun Kapur have collaborated with the development and documentation of STeP.

The Verification Diagram editor was developed by Michael Colón. The interface tools were developed by Michael Colón and Hsiao-Lan Liao, and the SPL compiler was implemented by Jaejin Lee. Mark Stickel provided the AC-matching and unification facilities. Harish Devarajan implemented the decision procedure for full Presburger arithmetic. Assistance and feedback has been provided by Anuchit Anuchitanukul, Luca de Alfaro, Jeff Kamerer, Uri Lerner, and Yoshihiro Yamada. We would also like to thank Paul Neves and the students of CS256 at Stanford. The STeP logo was designed by Anuchit Anuchitanukul.

STeP has mainly been implemented in Standard ML of New Jersey, using its X-windows utilities eXene for the user interface. David McQueen, Lal Geore, Hantao Zhang, and John Reppy have been helpful assisting us in using SML/NJ and eXene. Xavier Leroy introduced us to ML.

The polyhedral invariant generation uses the polyhedra manipulation package developed at VERIMAG/CNRS (Grenoble) by Nicolas Halbwachs, Yann-Eric Proy and Herve Leverge. The model checker uses the PTL satisfiability tester by Hugh McGuire at Stanford.

The STeP project has been supported in part by the National Science Foundation under grants CCR-92-23226 and CCR-95-27927, the Defense Advanced Research Projects Agency under NASA grant NAG2-892, the United States Air Force Office of Scientific Research under grant F49620-93-1-0139, the ARO under grant DAAH04-95-1-0317, the ARO under MURI grant DAAH04-96-1-0341, and by Army contract DABT63-96-C-0096 (DARPA). In particular, implementation of verification support for hybrid systems was supported by the MURI grant.

The STeP project is supervised by Professor Zohar Manna.

Chapter 2

Systems and Specifications

A system can be input to STeP as an SPL program or directly as a transition system; SPL programs are parsed into transition systems. The computational model associated with transition systems and SPL programs is presented in Appendix A. The specification language is linear-time temporal logic. Its semantics is presented in Appendix B.

An SPL program or transition system is always loaded from a file. The properties to be proven may be entered directly or loaded from a file.

This chapter presents the STeP format for describing systems and their specifications. Since systems and specifications share the same syntax for variable and type declarations and logical expressions, we first describe the general syntax for types (Section 2.1) and expressions (Section 2.2). We then describe the overall style of specifications (Section 2.3) and the syntax for declarations (Section 2.4). Finally, we present the syntax for describing systems either as SPL programs (Section 2.5) or as transition systems (Section 2.6).

We use extended Backus-Naur form (EBNF) to describe the syntax of STeP's input. Terminal symbols are written in `typefont`, non-terminals in *italics*. Alternatives are separated by |, optional parts enclosed in [ ], zero or more occurrences of a construct are indicated by enclosing in { }, and { }+ is used for one or more occurrences.

Comments can be included in a file using the `%` character (indicating a comment until the end of the line), or using `(*` and `*)` for multi-line comments.

2.1 Types

The *basic types* in STeP are booleans, integers, rationals, range types, underspecified types, and datatypes. New compound types can be created using *type constructors*. The type constructors in STeP, listed below, define types for arrays, channels, and tuples of types in the usual way. A type can be enclosed in parenthesis to distinguish for instance a tuple of three elements, e.g., `int * rat * bool`, from a nested tuple of pairs, `(int * rat) * bool`.

type	::=	*id*	type identifier
	|	*basety*	base type
	|	`array` *type* `of` *type*	array type
	|	`channel of` *type*	synchronous channel
	|	*type* `set`	set type

	|	`channel [1..] of` *type*	unbounded asynchronous channel
	|	`channel [1 ..` *int-const* `] of` *type*	bounded asynchronous channel
	|	`[` *range* `]`	integer range type
	|	`real[` *range* `]` | `real(` *range* `]`	
	|	`real(` *range* `)` | `real[` *range* `)`	open and closed real ranges
	|	{ *id* : *type* | *expn* }	subtype
	|	*type* { `*` *type*}+	tuple type
	|	`{` *id* : *type* {`,` *id*: *type* } `}`	record type
	|	`(` *type* `)`	parenthesized type
	|	`[` *range* {`,` *range* } `]`	tuple of integer ranges
range	::=	*expn* .. *expn*	closed range
	|	.. *expn*	range from $-\infty$ to *expn*
	|	*expn* ..	range from *expn* to $+\infty$
basety	::=	`bool`	boolean type
	|	`int`	integer type
	|	`rat`	rational type
	|	`real`	real type

The type constructor $*$ has higher precedence than all other type constructors. Thus, the type

```
channel [1..]  of bool * int * int
```

is parsed as

```
channel [1..]  of (bool * int * int).
```

Note that `(bool * int * int)`, `((bool * int) * int)` and `(bool * (int * int))` are three different types.

New underspecified types, datatypes and enumeration types can be declared in specification files; see Section 2.4.1.

2.2 Expressions

Expressions are constructed from constants, identifiers, function application, array references, tuple expressions, binary infix expressions, prefix expressions and quantification. The following well-formedness constraints are imposed:

- The boolean type is treated as a subtype of the integers, where the boolean value `false` corresponds to the integer value 0, and `true` corresponds to 1. Therefore any expression of boolean type can be used in an integer context. This is referred to as *arithmetization.*[1] Similarly, the integers (type `int`) are treated as a subtype of rationals (type `rat`).

[1] Users should note that this arithmetization is very convenient, but at times confusing, see Section 2.7.

- The operators `/\` (conjunction), `\/` (disjunction), `-->` (implication), `<-->` (bi-implication), `!`, `~` (negation), and the temporal operators `Until`, `Awaits`, `Since`, `Backto`, `==>`, `<==>`, `()`, `[]`, `<>`, `(-)`, `[-]`, `<->` take boolean arguments, and the resulting expression is of type boolean.
- The operators `+,-,*` require integer or rational arguments and produce a value of the corresponding type. Division `/` requires arguments of type integer or rational, and results in a value of type rational. The operators `mod` and `div` require arguments of type integer and result in a value of type integer.
- Variables bound by quantifiers may not be array types.

The full grammar for expressions that can be used in specifications is as follows:

expn	::=	*bool-const*	boolean value, true or false
	|	*int-const*	integer value
	|	*id*	identifier
	|	`<< >>`	the empty channel
	|	`<<` *expns* `>>`	enumerated channel
	|	`{` *id* `:` *type* `|` *expn* `}`	set comprehension
	|	*id* `{` *index* { `,` *index* } `}`	disjunction of locations id can be an array identifier
	|	*id*`(`*expns*`)`	function application
	|	*expr*`[`*expns*`]`	array reference
	|	`Deriv(` *expr* `)`	first derivative of expr
	|	`(`*expn*`)`	parenthesized expression
	|	`(`*expn*`,`*expns*`)`	a tuple of expressions
	|	`{` *id* `=` *expn* {`;` *id* `=` *expn* } *[;]* `}`	a record expression
	|	`#` *(int-const | id) expn*	tuple/record projection
	|	`@` *(int-const | id)* `(`*expn*`,` *expn*`)`	tuple/record update
	|	*expn infix expn*	infix operator
	|	*prefix expn*	unary prefix operator
	|	`if` *expn* `then` *expn* `else` *expn*	if then else
	|	`let` *id* `=` *expn* { `;` *id* `=` *expn*} *[;]* `in` *expn*	let binding
	|	*bind bv . expn*	binding
	|	*expn′*	primed expression
expns	::=	*expn* | *expn* `,` *expns*	
bv	::=	*ids*	globally or locally declared
	|	*ids* `:` *type [kind] [,* *bv]*	bound variables
index	::=	*int-const* | *int-const* `..` *int-const*	indexing locations
infix	::=	`/\` | `\/` | `-->` | `<-->`	binary boolean operators
	|	`Until` | `Awaits` | `Since` | `Backto`	binary temporal operators

	|	`==>` | `<==>`	entailment and congruence
	|	`+` | `*` | `-` | `/` | `mod` | `div`	arithmetical operators
	|	`=` | `!=` | `<` | `>` | `>=` | `<=`	predicate operators
prefix	::=	`!` | `~`	negation
	|	`-`	minus
	|	`()` | `[]` | `<>`	future operators
	|	`(-)` | `[-]` | `<->`	past operators
	|	`append` | `head` | `tail` | `length`	channel operations
	|	`update`	array update
	|	`Deriv`	
bind	::=	`Forall` | `Exists` | `Exists!`	logical binding
	|	`Sum` | `Prod`	summation and product
	|	`Integrate`	integration
	|	`Lambda`	lambda binding
kind	::=	`Rigid` | `Flexible`	flexibility of variable
bool-const	::=	`true` | `false`	
int-const	::=	*digit* { *digit* }	
id	::=	*alpha* { *alpha* | *digit* | `_` }	
alpha	::=	*upper or lower case letter*	
digit	::=	*number between 0 and 9*	

location disjunction: As an abbreviated notation for the disjunction of consecutive locations, the specification may contain expressions of the form

```
l{1..5,7,9..11}[i]
```

which is shorthand for

```
    l1[i] \/ l2[i] \/ l3[i]  \/ l4[i] \/ l5[i]
 \/ l7[i] \/ l9[i] \/ l10[i] \/ l11[i]
```

channel operations are described separately in Section 2.4.5.

array update: Arrays are multidimensional tables. Individual table entries can be updated by the built-in function `update`.

Example: The expression entered below:

```
in    N, M : int where N > 0 /\ M > 0
local A    : array [1..N, 1 ..M] of int
  update(A,5,1,1)
```

is identical to the array `A`, except for position `[1,1]`, where it has been updated to the value `5`.

The simplifier uses the following simplification rule automatically (so it really does not have to be declared):

```
variable a,x : [1..N]
variable b,y : [1..M]
variable e   : int
SIMPLIFY : (update(A,e,a,b))[x,y] --->
                    if a = x /\ b = y then e else A[x,y]
```

See section 2.4.4 for how to declare your own simplification rules.

tuple projection: The i'th element of a tuple t can be accessed as:

$$\#\ i\ t$$

Tuple indices are counted from 1, i.e., the first element in a tuple has index 1.

record projection: The elements of a record expression of t type $\{id_1\!:\ type_1, \ldots, id_n\!:\ type_n\}$ can be accessed as:

$$\#\ id_i\ t$$

binding: The binding operators `Forall`, `Exists`, and `Exists!` correspond to $\forall$, $\exists$ and $\exists!$ (unique existence) respectively. They bind the variables in their scope, so for instance in

```
Forall x : int .  p(x)
```

the variable `x` is bound by the quantifier `Forall`. The binding operators `Sum` and `Prod` correspond to the arithmetic operators Σ and Π and bind variables similarly.

The binding operator `Integrate` corresponds to the better known mathematical symbol $\int$. Its argument must be a subset of reals.

Lambda abstraction is partially supported via the binding operator `Lambda`, which corresponds λ. Whenever e has type σ in `Lambda` x `:` $\tau.e$, then the type of the resulting expression is `array` τ `of` σ. The simplifier supports lambda abstraction by always reducing (`Lambda` $x : \tau.e)[e']$ to $e[e'/x]$.

As the syntax indicates, the type field of the bound variables is optional. If no type is given, where the variable is bound, the variable must have been declared before using a variable declaration. If the variable has not been previously declared the parser still tries to infer the type of the variable using polymorphic type inference.

priming: The primed version of a flexible variable refers to its value at the next state. The parser accepts primed expressions, and automatically distributes the prime inwards to all its flexible variables. Double priming is not supported and results in a parse-error.

The precedence of the various operators is given in Table 2.1.

Symbols	*Associativity*
`of`	*left*
`array`	*left*
bind bv .	*nonassoc*
`then`	*nonassoc*
`else`	*nonassoc*
`<-->, <==>`	*right*
`-->, ==>`	*right*
`\/`	*right*
`/\`	*right*
`Until Awaits Since Backto`	*right*
`() [] <> (-) <-> [-]`	*nonassoc*
`= < > >= <= !=`	*left*
`~ !`	*nonassoc*
`mod div`	*left*
binary `-, +`	*left*
`/ *`	*left*
unary `-`	*right*
id `head tail append length assign`	*left*

Table 2.1: Operator precedences from lowest to highest

2.3 Specification Files

A specification file usually contains three parts: declarations, axioms and properties.

1. The *declarations* define new types, values, macros, rewrite rules, and auxiliary variables.

2. The *axioms* define properties that may be assumed by STeP in a verification session. Each axiom becomes a background property, and may be activated or deactivated from the Top-level Prover.

3. The *properties* are those formulas you want to prove true for a program or system, or are state-valid in general. These become top-level *goals* in the Top-level Prover.

The format of the specification file is described by the following grammar:

specification	::=	*[declarations] expn*	
	|	`SPEC` {*declaration* | *axiom* | *property*}	
declarations	::=	{ *declaration [;]* } *declaration ;*	
axiom	::=	`AXIOM` *desc : [module-env* `|=]` *expn*	named axiom
property	::=	`PROPERTY` *desc : [module-env* `|=]` *expn*	named property specification
module-env	::=	*id*	A module instance
	|	*module-env* `||` *module-env*	module composition

The nonterminal *desc* is an informal description of the corresponding axiom or property, and can be any sequence of characters except colon (:).

Note that STeP does not guarantee the consistency of the given set of axioms, so the specification of axioms should be done with caution.

A specification file may contain a single formula, which by default becomes the property to be proven, or may contain multiple axioms and properties. In addition, the specification file may contain declarations of macros, values, and auxiliary variables. An example of a file with multiple axioms and properties is shown in Figure 2.1.

```
SPEC (* Greatest common divisor spec file *)

value COMMUTATIVE gcd :  int*int --> int

AXIOM gcd1:  []Forall m,n:int .  (m != n --> gcd(m,n) = gcd(m-n,n))

AXIOM gcd2:  []Forall m:int .  (m > 0 --> gcd(m,m) = m)
AXIOM gcd3:  []Forall m:int .  (m < 0 --> gcd(m,m) = -m)

PROPERTY aux1:  [](y1 > 0)
PROPERTY aux2:  [](y2 > 0)
PROPERTY aux3:  [](gcd(y1,y2) = gcd(a,b))

PROPERTY partial correctness:  18 ==> g = gcd(a,b)
```

Figure 2.1: Specification file with multiple axioms and properties

2.4 Declarations

Declarations can be in one of the following forms, which we describe below:

declaration	::=	*type-decl* | *datatype-decl*	declaration of types
	|	*value-decl* | *macro-decl*	declaration of constructs
	|	*aux-var-decl* | *system-var-decl*	declaration of variables
	|	*simplify* | *cond-simplify* | *rewrite* | *order*	rewrite rules

2.4.1 Type Declarations

Type declarations are used to define abbreviations for compound types, new underspecified types, and new *datatypes*.

Type abbreviations are declared as `type` *id* `=` *type*, where *type* is a *type expression* (see Section 2.1). Omitting `=` and the type expression declares *id* as a new, underspecified type:

type-decl	::=	`type` *id* `=` *type*	type declaration
	|	`type` *id* `=` `{` *ids* `}`	enumeration type
	|	`type` *id*	underspecified type identifier

Datatype Declarations ⋆

Datatypes are declared using `==` instead of a single `=`, with the following syntax:

datatype-decl	::=	`type` *id* `==` *constructors* { `and` *id* `==` *constructors* }	datatype declaration
constructors	::=	*constructor* { | *constructor* }	
constructor	::=	*id* [`::` *destructors*]	constructor *id* and fields
destructors	::=	*id* `:` *type* { `,` *id* `:` *type* }	datatype destructors

The enumeration type declaration

`type` *id* `=` $\{id_1, \ldots, id_n\}$

is shorthand for the datatype declaration

`type` *id* `==` $id_1 \mid \ldots \mid id_n$.

Example: A *tree* datatype can be defined recursively with a *forest* datatype as follows:

```
type ident
type tree   ==  Node ::  id: ident, children: forest
and forest  ==  Nil
            |   Cons ::  first: tree, next: forest
```

The type `ident` is defined as a new underspecified type. A forest is a set of trees represented as a list—either `Nil` or the `Cons` of a tree and a forest. A tree itself is a `Node` containing an id and a set of children, represented as a forest.

The above datatype definition *implicitly* declares the following function symbols:

```
value Node: ident * forest --> tree
value is_Node: tree --> bool
value id: tree -> ident
value children: tree -> forest
value Nil: forest
value is_Nill: forest --> bool
value Cons: tree * forest --> forest
value is_Cons: forest --> bool
value first: forest --> tree
value next: forest --> forest
```

The *constructors* (`Node`, `Nil`, `Cons`) are used to create new elements of the corresponding type, while the *destructors* (`id`, `children`, `first`, `next`) are used to access the

different components of the datatypes. The *testers* (`is_Node`, `is_Nil`, `is_Cons`) are used to identify the top-most constructor.

For the destructors, the following simplification rules, explained in Section 2.4.4, are automatically inferred, and do not have to be declared:

```
variable i:  ident
variable ch,n:  forest
variable t:  tree
SIMPLIFY : id(Node(i,ch))        --->  i
SIMPLIFY : children(Node(i,ch))  --->  ch
SIMPLIFY : first(Cons(t,n))      --->  t
SIMPLIFY : next(Cons(t,n))       --->  n
```

These rules are used automatically by the simplification procedures. Note that no simplification rules are given for `first(Nil)` and `next(Nil)` although these expressions are well typed. Their interpretation is left *underspecified*, but they are not erroneous.

For the constructors, axioms asserting injectivity are also asserted automatically and used by the simplifier, so they don't have to be declared:

```
Forall t:tree .  Forall f:forest .  Nil != Cons(t,f)
Forall i,j:ident .  Forall c1,c2:forest .
     Node(i,c1) = Node(j,c2) --> i = j /\ c1 = c2
Forall t1,t2:tree .  Forall sib1,sib2:forest .
     Cons(t1,sib1) = Cons(t2,sib2) --> t1 = t2 /\ sib1 = sib2
```

For the testers, axioms making sure that the testers identify their datatype constructor are asserted automatically and used by the simplifier, so they do not have to be declared either:

```
Forall t:  tree, f:  forest .  is_Nil(Nil) /\ !is_Nil(Cons(t,f))
Forall t:  tree, f:  forest .  !is_Cons(Nil) /\ is_Cons(Cons(t,f))
Forall t:  tree .  is_Node(t)
```

The complete axiomatization of datatypes also includes *well-founded induction schemas* (see [Manna and Waldinger, 1993].) These can be used when proving properties by induction in the Interactive Prover.

2.4.2 Value and Macro Declarations

Program and specification files may contain value and macro declarations. Their syntax is given by the following grammar:

value-decl	::=	`value` *[ac] ids* : *[type* `-->]` *type*	value declaration
macro-decl	::=	`macro` *id* : *[type* `-->]` *type*	macro declaration
		`where` *id* `[(` *ids* `)]` = *expn*	

	|	`macro` *id [(* *typed-ids* *)] [:* *type] =* *expn*	
ac	::=	`AC` | `COMMUTATIVE` | `ASSOCIATIVE`	
ids	::=	*id* {, *id* }	identifier list
typed-ids	::=	*id [:* *type]* {,*id [:* *type]*}	

Uninterpreted function symbols are declared as values. Binary function symbols can have additional attributes used by the Simplifier: the keyword `AC` states that the function is both associative and commutative; `COMMUTATIVE` and `ASSOCIATIVE` state that the function is only commutative or associative, respectively. For example, in Figure 2.1 the `gcd` function was defined to be commutative. Thus the simplifier assumes that `Forall a,b: int . gcd(a,b) = gcd(b,a)` besides the other axioms.

Macros, which may take optional arguments, are declared with the keyword `macro`. A `where` clause defines the expression corresponding to the macro. Each macro takes a fixed number of arguments and is expanded by the parser through purely syntactic substitution. A macro definition can use other previously defined macros, which will be expanded in turn.

Example: Figure 2.2 shows an example of value and macro declarations as they may appear in a program or a specification file. The value declarations define the integer constant c and function symbol f, which takes a pair of arguments of boolean type and returns an integer. The macro declarations indicate that the symbol m should be expanded to the expression $c < c + 1$, and the symbol g with integer arguments (i, j) should be expanded into $i + i + j + 5$. Implicit typing, as well as partial typing is possible as illustrated with macros `h` and `k`.

```
value c:int
value f:bool * bool --> int
macro m:bool where m = c < c+1
macro g:int * int --> int where g(i,j) = i+i+j+5
macro h(a,b) = g(g(a,b),b)
macro k(x:int) = x + 1
```

Figure 2.2: Value and macro definitions

Example: Figure 2.3 shows a specification file to prove 1-bounded overtaking for program `mux-pet1` (Figure 3.4). The proof by rule G-WAIT requires the input of four auxiliary assertions. Having defined these assertions in the specification file allows the user to enter them by name, `phi0`, `phi1`, `phi2` or `phi3`, in the dialog window, reducing the chance of typing errors and saving time if the proof needs to be repeated multiple times.

2.4.3 Variable Declarations

aux-var-decl	::=	`variable` *ids* : *type [kind]*	auxiliary variable declaration

```
% 1-bounded overtaking for mux-pet1
%
% for proof by rule G-WAIT use the following intermediate assertions:
%
% phi0:  l4
% phi1:  l3 /\ ((m0\/m1\/m2\/m5) \/ (m3 /\ !(s=1)))
% phi2:  l3 /\ m4
% phi3:  l3 /\ m3 /\ s=1
SPEC
macro phi0 = l4
macro phi1 = l3 /\ (m{0..2,5} \/ (m3 /\ !(s=1)))
macro phi2 = l3 /\ m4
macro phi3 = l3 /\ m3 /\ s=1

PROPERTY 1-bounded overtaking:
    l3 ==> ((!m4) Awaits ((m4) Awaits ((!m4) Awaits l4)))
```

Figure 2.3: Specification file for proving `mux-pet1` 1-bounded overtaking

kind	::=	`Rigid` | `Flexible`	
system-var-decl	::=	*mode ids* : *type* [`where` *expn*]	system variable declaration
	|	[`external`] `clock` *ids* [`where` *expn*]	system clock
mode	::=	[`external`] `in` | [`external`] `out` | `local`	
ids	::=	*id* {, *id* }	list of identifiers

Auxiliary variables are declared with the keyword `variable`. Auxiliary variables can only be declared in the specification file, or in a directly entered goal, and are not part of the system being verified. Each auxiliary variable can be optionally declared to be `Rigid` (its value does not change over time) or `Flexible` (its value may change from one point in time to another). By default, auxiliary variables are rigid. Auxiliary variables are often used as part of macros and rewrite rules.

For convenience, a specification may also include declarations for system variables; these are further described in Section 2.5. Here we just mention that clock variables can be modified by the module where they are declared. Variables of type `in`, `out` and `clock` are always visible. Their declaration can be preceded by the optional `external` keyword, making them shared variables that can be modified by other modules. Variables declared as `local` and non-external `out` are private, and cannot be changed by other modules. Non-external `in` variables are treated as constants.

Example: The declarations

```
variable x,y,z : int
```

```
variable a,b,c : bool Flexible
variable i     : [1..N]
```

declare the rigid integer-type variables x, y and z, the flexible boolean variables a, b and c, and the rigid variable i of type range [1..N].

Previously declared auxiliary variables can be used in quantification, so the type of the quantified variable type does not have to be declared within the quantifier (see Section 2.2).

2.4.4 Rewrite and Simplification Rules ⋆

Specifications can also include *simplification and rewrite rules* to efficiently build in equational axioms, as well as *ordering relations* to control automatic rewriting.

User-defined rewrite rules can be specified in one of two ways: rewrite rules declared with the `REWRITE` keyword are applied *under the user's control*, step-by-step, in the Interactive Prover. *Simplification rules*, declared with the `SIMPLIFY` keyword, are applied automatically and exhaustively whenever the Simplifier is invoked from either the Top-level Prover or the Interactive Prover. Thus, the set of user-defined simplification rules is expected to *terminate*, while the set of rewrite rules, under user control, does not have to be terminating. Conditional simplification rules can be declared by adding a condition to the rule declared by `SIMPLIFY`. They get get invoked when Conditional simplification rules are enabled via the Simplifier Settings described in Figure 3.11.

An ordering relation, given by the `ORDER` keyword, specifies a linear order over uninterpreted function symbols, as a list. This order is used when converting equalities into rewrite rules during simplification; roughly speaking, if $t_1 = t_2$ or $t_2 = t_1$ is an equality encountered in contextual simplification, t_1 will be rewritten to t_2 if t_1 is built from a function symbol that appears earlier in the order.[2]

The syntax for simplification and rewrite rules is given by the following grammar:

rewrite	::=	`REWRITE` *desc* : *expn* `--->` *expn*	rewrite rule
simplify	::=	`SIMPLIFY` *desc* : *expn* `--->` *expn*	simplification rule
cond-simplify	::=	`SIMPLIFY` *desc* : *expn* `--->` *expn* `if` *expn*	conditional simplification
order	::=	`ORDER` *ids*	

The nonterminal *desc* can be any sequence of characters except colon (:), and is used as an informal description of the simplification or rewrite rule. The auxiliary variables free in the expression to the left of `--->` should be a superset of the free variables of the right hand side expression. In connection with rewrite and simplification rules, system variables are treated as constants (constructors).

Example:

Figure 2.4 shows an example of two rewrite rules. The rule `Peano 1`, rewrites $6 + 0$ to 6, and the rule `ad hoc` rewrites, $f(4 + 2)$ to `if y then 4+2 else -(4+2)`.

[2]The Simplifier uses the *recursive path ordering on terms* to orient equalities, see [Dershowitz and Jouannaud, 1990].

```
variable x:int
local y:bool
value f:int --> int
REWRITE Peano 1:  x + 0 ---> x
REWRITE ad hoc:  f(x) ---> if y then x else -x
```

Figure 2.4: Rewrite rules in a specification file

As with axiomatizations, the user is responsible for the consistency of the set of rewrite and simplification rules. The reader interested in learning more about rewrite rules and their applications should consult, for example, [Dershowitz and Jouannaud, 1990]. An example of simplification rules appears in Section 2.4.1.

2.4.5 Channel Operations ⋆

Asynchronous channels are modeled as queues, i.e., new elements are appended to the tail of the queue, and elements are extracted from the head.

The basic operations on channels, `head`, `tail`, `append`, and `length` are best described in the specification file that users should provide when they want to establish properties of channels. That is, STeP does not include the partial axiomatization given below; instead, specifications should provide separate axiomatizations for each type of asynchronous channel used.

```
%% Sample channel specification
SPEC
type t
type chan = channel [1..] of t
local ch : chan
value pick : chan * int --> t
variable x : t

AXIOM : [] Forall i : int .
            pick(append(ch,x),i) =
                 if i-1 = length(ch) then x else pick(ch,i)
AXIOM : [] Forall i : int . pick(tail(ch),i) = pick(ch,i+1)
SIMPLIFY : head(ch)             ---> pick(ch,1)
SIMPLIFY : length(append(ch,x)) ---> 1 + length(ch)
SIMPLIFY : length(tail(ch))     --->
                        if length(ch) = 0 then 0 else (length(ch)) - 1
```

Hence the operation `head` extracts the first element from the buffer, `tail` returns the buffer excluding its first element, `append` inserts a new element to the end of the buffer, and `length` measures the number of elements residing in the buffer. The auxiliary function `pick`, axiomatized in the specification, provides us with a formal definition of the `head`,

`tail`, and `append` operations. We notice that taking `head` of an empty buffer only simplifies to some expression involving `pick` and its value is therefore underspecified.

2.5 SPL Programs

The syntax of SPL programs follows that of traditional imperative languages such as Pascal. In addition to the basic constructs found in these languages, SPL supports nondeterminism by means of the *selection* statement `or` and parallel composition by means of the *cooperation* statement `||`. Parallel processes can interact through shared variables such as semaphores, as well as by synchronous and asynchronous channels. Execution of parallel processes is assumed to proceed by interleaving (see Appendix A for a complete account of SPL semantics).

Program variables

As in most imperative languages, program variables must be declared prior to their use, in the appropriate scope. SPL provides an additional *mode* associated with the declaration of program variables. A variable declared as `in` is an input to the program; it may not be modified by any of the program statements, and is assumed to have a fixed value throughout the program execution. Variables declared as `local` or `out` may be changed by the program.

Variables may be initialized or constrained with an optional `where` clause. The given condition is assumed to hold only at the start of the program, not each time the program enters the block where the variable is declared.

Parameterized programs

SPL allows you to write *parameterized* (N-process) programs, in which an arbitrary number of identical processes are declared. Both parameterized composition and parameterized selection are supported.

Example:

Figure 2.5 shows a deadlock-free solution to the N-process *Dining Philosophers* problem. The program contains `N` instances of the infinite loop `l0-l8`, each with its unique value of `i`.

Grouped statements

A *grouped statement* is a set of statements enclosed by `<<` and `>>`. A grouped statement is translated into a single transition, that is, one atomically executable unit. Grouped statements may only contain basic statements, and may only be combined into composite statements by selection and concatenation. A grouped statement may contain at most one `send` or `receive` statement on the same channel.

Example: To request resources `r` and `s` simultaneously one can enter

```
in N : int where N >= 2
local f :  array [1..N] of int where Forall i:[1..N] . (f[i] = 1)
local r :  int where r = N - 1

|| P[i:[1..N]] ::
        l0:  loop forever do [
              l1 :  noncritical;
              l2 :  request r;
              l3 :  request f[i];
              l4 :  request f[(i mod N) + 1];
              l5 :  critical;
              l6 :  release f[i];
              l7 :  release f[(i mod N) + 1];
              l8 :  release r
        ]
```

Figure 2.5: Program DINE (dining philosophers)

```
<< request r; request s >>
```

which can also be encoded as

```
guard r > 0 /\ s > 0 do (r,s) := (r-1,s-1)
```

Program statements

The following grammar provides the syntax of SPL programs:

program	::=	{ *decl [;]* } *composite_stmt [; label]*	
composite_stmt	::=	*composite_stmt* ; *composite_stmt*	concatenation
	|	*composite_stmt* `||` *composite_stmt*	cooperation statement
	|	*composite_stmt* `or` *composite_stmt*	selection statement
	|	*[label] stmt*	
stmt	::=	*basic_group_stmt*	
	|	`request` *variable*	request statement
	|	`release` *variable*	release statement
	|	`noncritical`	noncritical statement
	|	`critical`	critical statement
	|	`choose` *variable*	choose
	|	`produce` *variable*	produce statement
	|	`consume` *variable*	consume statement

	|	`guard` *p_expn* `do` *assignment*	guarded assignment
	|	`if` *p_expn* `then` *composite_stmt*	1-way conditional
	|	`if` *p_expn* `then` *composite_stmt* `else` *composite_stmt*	2-way conditional
	|	`when` *p_expn* `do` *composite_stmt*	when statement
	|	`while` *p_expn* `do` *composite_stmt*	while statement
	|	`repeat` *composite_stmt* `until` *p_expn*	repeat statement
	|	`loop forever do` *composite_stmt*	infinite loop
	|	*[id : :]* `[` *program* `]`	block statement
	|	`<<` *comp_group_stmt* `>>`	grouped statement
	|	`or` *binding composite_stmt*	parameterized selection
	|	`||` *binding composite_stmt*	parameterized cooperation
comp_group_stmt	::=	*basic_group_stmt*	
	|	`if` *p_expn* `then` *comp_group_stmt*	1-way conditional
	|	`if` *p_expn* `then` *comp_group_stmt* `else` *comp_group_stmt*	2-way conditional
	|	`when` *p_expn* `do` *comp_group_stmt*	when statement
	|	*comp_group_stmt* `or` *comp_group_stmt*	selection statement
	|	*comp_group_stmt* `;` *comp_group_stmt*	concatenation
basic_group_stmt	::=	`skip`	skip statement
	|	*assignment*	
	|	`await` *p_expn*	await statement
	|	*variable* `<==` *p_expn*	send statement
	|	*variable* `==>` *variable*	receive statement
assignment	::=	*left-val* `:=` *p_expn*	
left-val	::=	*p_expn*	assignable value
variable	::=	*id*	variable
	|	*id* `[` *p_expns* `]`	array variable
binding	::=	*id* `[` *params* `]` `: :`	named binding
	|	*params* `.`	unnamed binding
params	::=	*param* {`,` *param*}	list of parameters
param	::=	*id* `:` `[` *range* `]`	parameter
label	::=	*id* `:`	label
ids	::=	*id* { `,` *id* }	list of identifiers
decl	::=	*type-decl* | *datatype-decl*	SPL declaration
	|	*value-decl* | *macro-decl*	

	|	*system-var-decl*	
system-var-decl	::=	*mode ids* : *type* [`where` *expn*]	basic variable declaration
mode	::=	`in` | `out` | `local`	mode of variable

Assignable expressions are built from system variables declared as `local` or `out`, and array references to `local` or `out` arrays. Composite assignable expressions are obtained by collecting simpler assignable expressions into tuples.

Example: An example of a composite assignment swaps the values of `x` and `y` as follows:

```
(x,y) := (y,x)
```

Expressions

The expressions allowed in SPL programs are a subset of the expressions allowed in specifications. In particular, the following are *not* allowed in SPL programs:

- primed variables
- `Prod` and `Sum`
- temporal operators
- location expressions

Furthermore, when binding variables by quantifiers inside an SPL program, their types must always be declared locally, and the optional *kind* is not allowed (they are always `Rigid`).

The syntax of SPL expressions is then as follows:

p_expn	::=	*bool-const*	boolean value, true or false
	|	*int-const*	integer value
	|	*id*	identifier
	|	`<< >>`	the empty channel
	|	`<<` *p_expns* `>>`	enumerated channel
	|	*id*(*p_expns*)	function application
	|	*expr*`[`*p_expns*`]`	array reference
	|	(*p_expn*)	parenthesized expression
	|	(*p_expn*,*p_expns*)	a tuple of expressions
	|	`{` *id* = *expn* {; *id* = *expn* } [;] `}`	a record expression
	|	`#` (*int-const* | *id*) *p_expn*	tuple/record projection
	|	`@` (*int-const* | *id*) (*p_expn*, *p_expn*)	tuple/record update
	|	*p_expn infix p_expn*	infix operator
	|	*prefix p_expn*	unary prefix operator
	|	`if` *p_expn* `then` *p_expn* `else` *p_expn*	if then else
	|	*bind ids : type* . *p_expn*	binding of the variables *ids*

p_expns	::=	*p_expn* | *p_expn* , *p_expns*	
infix	::=	`/\` | `\/` | `-->` | `<-->`	binary boolean operators
	|	`+` | `*` | `-` | `/` | `mod` | `div`	arithmetical operators
	|	`=` | `!=` | `<` | `>` | `>=` | `<=` | `<>`	predicate operators
prefix	::=	`!` | `~` | `-`	negation and minus
	|	`append` | `head` | `tail` | `length`	channel operations
	|	`update`	array update
bind	::=	`Forall` | `Exists` | `Exists!`	binding operators
bool-const	::=	`true` | `false`	
int-const	::=	*digit* { *digit* }	
id	::=	*alpha* { *alpha* | *digit* | `_` }	

2.6 Transition Systems

Transition systems are the basic system representation in STeP. SPL programs are parsed into fair transition systems; an alternative to SPL is to describe transition systems directly. Appendix A includes the basic definitions of transition systems. For a more leisurely introduction, see [Manna and Pnueli, 1991].

2.6.1 Fair Transition Systems

A transition system specification contains:

- Declarations of types and free variables.
- An optional *initial condition*, which is assumed to hold at the initial state of every computation.
- A set of transitions.

Transitions are expressed in terms of, and are usually equivalent to (see below), *transition relations*. A transition relation describes how the system can change from one state to the next. Strictly speaking, a transition relation can be expressed by a formula over the set of primed and unprimed system variables; however, for added convenience, transition relations in STeP can be specified by a combination of five different fields. Each field is optional, and STeP inserts default values for absent ones. The five fields are the following:

- The `modvar` field: contains a set of variables that are changed when the transition is taken. To indicate that the transition relation allows a variable to change value nondeterministically, one can include it in the `modvar` field and not mention it elsewhere.

By default the only variables that are changed by the transition relation are the ones primed in the `modrel` field and the ones to the left-hand side of assignments in the `assign` field described below. The `modvar` field does not need to contain those variables.

- The `define` field: contains definition of auxiliary variables used within the transition relation. It enables the user to introduce shorthands for expressions appearing in the `enable`, `modrel` field, as well as on the right hand sides of the assignments in the `assign` field. The expressions being abbreviated in the `define` field may not contain primed variables. The assignments made in the `define` field are processed sequentially. Thus, `define x := f(x), x := g(x)`, can be abbreviated as `define x := g(f(x))`, but not as `define x := f(g(x))`.

- The `enable` field: the *enabling condition*, a formula that may not contain any primed variables. The transition can only be taken if the enabling condition is satisfied. By default, the enabling condition is *true*.

- The `modrel` field: contains the *modifying relation*, an arbitrary relation over primed and unprimed variables that is assumed to hold when the transition is taken. Its default value is *true*.

- The `assign` field: lists the assignments for those variables whose next-state values can be described as a function of the current-state (unprimed) variables. The assignments listed in the `assign` field are made in parallel. Thus, `assign x := y, y := x` swaps the values of `x` and `y` rather than only assigning `y` to `x`.

The user can enter these fields in any order desired, though the parser will first process all the `define` fields before processing anything else.

An assignment `x := t` (in the assign field) can be equivalently expressed as a conjunct `x' = t` in the modifying relation, but assignment relations need not also appear in the `modrel` field. In fact, it is preferable to only include them in the `assign` field. For example, rather than having `x'=x+1` as a conjunct in the `modrel` field, you should include `x:=x+1` in the `assign` field.

In general, variables that appear on the left-hand side of an assignment statement should not appear primed in the modifying relation (there is the risk of unwillingly specifying an unsatisfiable transition relation). Still, there is the possibility that the user will include a conjunct of the form `x'=x+1` in the `modrel` field, or will mention a primed variable that is also assigned to. The parser will automatically discover these cases and normalize the transition relation, while printing a warning to the console.

A *transition* consists of one or more *transition relations*. This partition allows a more convenient representation of certain transition relations, e.g., those for the `while` and `if-then-else` SPL constructs.[3] A transition is considered enabled if at least one of its

[3]The verification rules generate one verification condition for each transition relation, which may be simpler to manipulate than the verification condition for the entire transition. See Appendix A for the semantics of `while` and conditional statements. See [Manna and Pnueli, 1995b] for an example where transition relations are used to simplify verification conditions.

constituent transition relations is enabled. It is considered taken when one of its transition relations is taken. A transition can be taken only if it is enabled.

Transitions that only differ with respect to some parameters can be collected in one *parameterized* transition. The optional parameters are listed after the transition's name and declare one transition for each value assignment to the parameters.

A transition contains an optional *fairness requirement*, which can be `Just`, `Compassionate` or `NoFairness`. The default fairness requirement is `NoFairness`.

Example: The transition system in Figure 2.6 implements Euclid's algorithm for computing the greatest common divisor of two positive integers x and y.

```
Transition System % Euclid's algorithm
in a,b:  int where a > 0 /\ b > 0
local x,y:  int
out gcd:  int
Initially x = a /\ y = b
Transition t1 Just:
    enable x > y
    assign x := x - y
Transition t2 Just:
    enable x < y
    assign y := y - x
Transition t3 Just:
    enable x = y
    assign gcd := x
```

Figure 2.6: Euclid's algorithm for greatest common divisor

The system illustrates that no control locations need to be associated with the transitions when it is entered directly as a transition system. This contrasts with the transitions generated from SPL programs. An SPL program corresponding to the above transition system is shown in Figure 2.7. It generates the transition system shown in Figure 2.8.

2.6.2 Modular Transition Systems

A fair transition system can be composed of several modules. Syntactically we describe a system as a list of modules, where the last module defines the meaning of the system. Each module can be instantiated within modules that are defined *after* it in the list.

2.6.3 Interface Declaration

- **Input Parameters.** The value of these variables is observed by the module, but not modified.

```
in a,b:  int where a > 0, b > 0
local x,y:  int where x = a, y = b
out gcd:  int
l0:  [loop forever do
      [
        t1:  guard x > y do x := x - y
      or
        t2:  guard x < y do y := y - x
      or
        t3:  guard x = y do gcd := x
      ]
    ]
```

Figure 2.7: SPL program for Euclid's algorithm

```
Transition System
in a,b:  int
local x,y:  int
out gcd:  int
control pi0 :  [0..0]
Initially a > 0 /\ b > 0 /\ x = a /\ y = b /\ pi0 = 0
Transition t1 Just:
    enable pi0 = 0 /\ x > y
    assign x := x - y
Transition t2 Just:
    enable pi0 = 0 /\ x < y
    assign y := y - x
Transition t3 Just:
    enable pi0 = 0 /\ x = y
    assign gcd := x
```

Figure 2.8: Transition system obtained from the Euclid SPL program (Figure 2.7)

transition-system ::= `Module System` { *module* }$^+$
module ::= `Module` *id* : *interface-decl* [`local` *decl*] [*initial*] *body* `EndModule`
body ::= [*transition* | `include` *composition*]*

interface-decl ::= { *parameter-decl* }
parameter-decl ::= [`external`] `in` *params*
[`external`] `out` *params* |
`export` *exported-transitions* |
fmodule
params ::= *id* : *type* { , *id* : *type* }
exported-transitions ::= *id* { , *id* }
fmodule ::= `formal module` *id* : *interface-decl* `EndModule`

  - *External input parameters.* The value of these variables is not changed, by the module or its environment. These inputs maybe thought of as static.
  - *Internal input parameters.* The module observes the value of these variables: They may be modified by the environment, but not by the module. These are the default type of input parameters. These inputs may be thought of as dynamic.

- **Output Parameters.** The value of these variables can be modified by the module.
  - *Internal output parameters.* The value of these variables is only controlled by the module. The values are observable to the environment, but cannot be changed by the environment. These are the default kind of output parameters and may be thought of as owned.
  - *External output parameters.* The value of these variables can be modified by both the module and the environment. These outputs maybe thought of as shared variables.

- **Exported Transitions.** A list of transitions that are made accessible to the instantiating module.

- **Formal Modules.** A list of formal modules, for each an identifier and a formal interface declaration is given.

We define a set of operators on modules. Most commonly used is instantiation of a module within another one. In the simplest case instantiations are just listed in the body of the module; they are then, by default, composed by parallel composition without synchronization. More sophisticated compositions are made possible by the following syntax:

1. **Parallel Composition.**

   The result of the parallel composition is a module that contains the transitions of both modules, where exported transitions of the two modules that have the same argument label are synchronized.

```
composition      ::=  ( composition )
                   |  composition || composition
                   |  composition \ ( Transition | in | out ) ( id { , id } )
                   |  composition [ substlist ]
                   |  module-instance
module-instance  ::=  id : ( id [ where { inst-decl }+ ] )
inst-decl        ::=  in substexprlist
                   |  out substlist
                   |  transition substlist
                   |  module id = id
substlist        ::+  id = assignable-expr { , id = assignable-expr }
renaminglist     ::=  id = id { , id = id }
substexprlist    ::=  id = expr { , id = expr }
```

2. **Module Instantiation.** Parameters of a module can be instantiated with actual parameters as follows:

   - external (static) input parameters are instantiated by expressions over constants and static input parameters.
   - internal (dynamic) input parameters are instantiated by arbitrary expressions.
   - internal (owned) output parameters are instantiated by expressions over private variables and owned output parameters.
   - external (shared) output parameters can be instantiated by expressions over private variables and output parameters.
   - formal modules are instantiated by compositions of concrete and formal modules.

3. **Hiding.**

   - Hiding of *variables*. Parameters are transformed into local variables.
   - Hiding of *transitions*. Transitions are removed from the list of exported transitions.

4. **Relational Renaming of Transitions.** New labels are associated with transitions; it is possible to duplicate transitions by assigning more than one new label.

2.6.4 Transition System Syntax

The syntax of transition systems is given by the following grammar:

```
transition-system ::=  Transition System
                           { decl } [initial] { transition+ }
                     |  Clocked Transition System
                           { decl } [initial] [progress] { transition }+
                     |  Hybrid Transition System
                           { decl } [initial] [progress]
```

```
                    { transition }+ { activity }+
decl            ::= type-decl | datatype-decl
                |   value-decl | macro-decl
                |   system-var-decl

system-var-decl ::= mode ids : type [where expn]
                |   clock ids [where expn]
                |   [external] continuous ids [where expn]

mode            ::= in | out | local

initial         ::= Initially expn
progress        ::= Progress expn
transition      ::= Transition name [ [ params ] ] [fairness] : relation { ; relation }
relation        ::= { modvar  ids |
                      define  { id := expn }+ |
                      enable  expn |
                      modrel  expn |
                      assign  assignment { , assignment } }

activity        ::= Activity name : activity-relation
activity-relation ::= { modvar  ids |
                        enable  expn |
                        modrel  expn |
                        assign  assignment { , assignment } }

name            ::= id
params          ::= id : type { , id : type }
fairness        ::= Justice | Compassionate | NoFairness
assignment      ::= left-val := expn
left-val        ::= expn
```

Assignable expressions are built from system variables declared as `local`, `out` or `clock`, as well as array references or tuple projections of `local` or `out` arrays/tuples.

Example: A legal assignment statement is given in Figure 2.9, where the second component of `X` is updated to `a[3] + 1` and simultaneously `a[3]` is updated to `2`.

2.6.5 Clocked Transition Systems ⋆⋆

Clocked transition systems [Kesten *et al.*, 1996] can be used to model discrete or continuous real-time systems, and are partially supported in the educational release of STeP. Clocked transition systems differ from fair transition systems in the following ways:

- There are designated clock variables `T`, the master clock, (global time) and `tick` (time progress measure) which may not be changed by the standard transitions. Only `T` may

```
Transition System
local X : int * int * bool
local a :  array [1..5] of int
Transition t:
    assign (#2 X,a[3]) := (a[3] + 1,2)
```

Figure 2.9: Legal assignments

be used in the standard transitions. Initially `T` is set to 0. These variables are declared automatically when a clocked transition system is specified.

- Fairness conditions on transitions are irrelevant, and therefore only the condition `NoFairness` is allowed.
- The user can specify a special *progress condition*. It restricts the amount that time is allowed to progress between the standard transitions.

The progress condition is written as:

`Progress` *expn*

where *expn* is assumed to be a first-order formula $F(clock_1, \ldots, clock_n, T)$, with the auxiliary clocks $clock_1, \ldots, clock_n$ and the global clock `T` as free variables. The progress condition corresponds to a transition of the form:

```
Transition tick:
    modvar      tick
    enable      tick > 0
           /\   Forall t :  rat .
                     (0 <= t /\ t <= tick -->  F(clock_1+t,...,clock_n+t,T+t))
    assign      (clock_1,...,clock_n,T) := (clock_1+tick,...,clock_n+tick,T+tick)
```

which is generated automatically by the parser. It is helpful to think about this transition as one that allows time to progress from `T` to `T + tick`, whenever the progress condition F is satisfied for all intermediary values between `T` and `T + tick`.

Clock variables are always of type `rat`.

The verification of the classic railroad crossing real-time system using STeP is described in [Bjørner *et al.*, 1997b].

2.6.6 Hybrid Transition Systems ⋆⋆

Hybrid transition systems [Manna and Pnueli, 1995a] can be used to model hybrid systems, i.e. systems that contain continuous components: the values of some state variables may change continuously over time. Hybrid transition systems are an extension of *clocked transition systems* (Section 2.6.5). The evolution of continuous variables is described by *activities*.

To support the specification of derivatives, the expression syntax has been extended with the construct `Deriv(`*expn*`)`, which denotes the first derivative of an expression.

Hybrid transition systems have the following components.

- A set of variables, which is partitioned into three kinds: *discrete variables*, declared as in Fair Transition Systems, *clocks*, declared as in Clocked Transition Systems, and *piece-wise continuous variables*, declared with the keyword `continuous`. Like clocked transition systems, hybrid transition systems have a master clock, designated by `T` of type clock.
- An (optional) initial condition. The conjunct `T = 0` is automatically added to the initial condition.
- An (optional) progress condition, written as `Progress` *expn*, where *expn* is a first-order formula $\Pi(T, clock_1, \ldots, clock_n, v_1, \ldots, v_m)$ with $T, clock_1, \ldots, clock_n, v_1, \ldots, v_m$ appearing as free variables, where $v_1, \ldots, v_n$ are continuous variables or their derivatives.
- A set of transitions, specified in the same way as for fair transition systems and clocked transition systems. Transitions must not modify the master clock.
- A set of activities. An activity is expressed by an activity relation. It describes how the continuous variables evolve over time if certain conditions hold. An activity relation contains the same four fields as a transition relation.
  - `modvar` same as for transition systems.
  - `enable` this field contains the *enabling condition*. It cannot contain any primed variables, clocks or continuous variables. The activity applies if the enabling condition is satisfied.
  - `assign` this field contains a number of assignments. The left-hand side of the assignment should be a derivative of a continuous variable. The right-hand side of the assignment may refer to the (free) variable `t`, as the independent time variable. Time is considered to start at 0 at the start of an activity. Differential inclusions can be expressed by

    `Deriv(x) := { y: Real | p <= y` $\wedge$ `y <= q}`

  - `modrel` this field is currently not support for activities.

The hybrid transition system is compiled into a transition system. For each activity α of the form

$$\begin{array}{lll} \texttt{enable} & : & En \\ \texttt{assign} & : & Deriv(d_i) := \{x \mid lb_i \leq x \wedge x \leq ub_i\}, \quad \text{for } i = 1 \ldots m \\ & : & Deriv(c_j) := f_j, \quad \text{for } j = 1 \ldots n \end{array}$$

a transition named α is created with the following fields, assuming the system has K clocks, $cl_1, \ldots, cl_K$, and progress condition $\Pi(cl_k, T, c_j, d_i)$ (for $i = 1 \ldots m,\ j = 1 \ldots n,\ k = 1 \ldots K$):

$$
\begin{array}{lcl}
\texttt{modvar} & : & cl_1, \ldots, cl_k, c_1, \ldots, c_n, d_1, \ldots, d_m \\
\texttt{enable} & : & En \wedge\ tick > 0 \wedge \\
 & & \forall t, e_i.(0 \leq t \wedge t \leq tick \wedge e_i \geq d_i + \int_0^t lb_i dt \wedge e_i \leq d_i + \int_0^t ub_i dt \rightarrow \\
 & & \Pi(cl_k + t, T + t, c_j + \int_0^t f_j dt, e_i) \\
\texttt{assign} & : & cl_k := cl_k + tick \\
 & & T := T + tick \\
 & & c_j := c_j + \int_0^{tick} f_j dt \\
\texttt{modrel} & : & d_i' \leq d_k + \int_0^{tick} ub_i dt \wedge d_i' \geq d_k + \int_0^{tick} lb_i dt
\end{array}
$$

2.6.7 Example

Figure 2.10 shows an example of a hybrid transition system. It is a transcription of the water level monitor described in [Alur *et al.*, 1993]. This example is included in the example directory.

2.7 When Parsing Fails

It is common for programs and specifications to fail to parse on the first try. The error messages generated by the parser typically indicate the line number and character range where the error was detected, separated by a ":", and provide some extra information about the offending part.

Unfortunately, the error messages are not always as illuminating as we would like. Here are some hints for use when things don't work as expected.

- The different kinds of arrows can be confusing:
  - `<->` is the temporal operator $\diamondsuit$ ("*once*");
  - `<-->` is double implication
  - `<=` is "less than or equals"
  - `-->` is normal implication, but
  - `-->` is also used in function types
  - `==>` is the temporal "entails" relation (`p==>q` stands for `[](p-->q)`), but
  - `==>` is also used for channel operations in SPL
  - `<==>` is a built-in macro, expanded to `[](p<-->q)`
  - `--->` is used to specify rewrite rules
- A period (".") is expected after quantifiers.
- Variables declared in the specification file and quantified variables are rigid by default. To specify a flexible variable in quantification, just write

  `Forall x : int Flexible` . *expn*.

Even after parsing a file, watch out for the following:

```
Hybrid Transition System Waterlevel Controller

type valveStates = {on0,on1,off0,off1}

local s:valveStates where s = on0
clock x where x = 0
continuous y where y = 1

Progress
    (s = on0 --> y <= 10) /\
    (s = on1 --> x <= 2) /\
    (s = off0 --> y >= 5) /\
    (s = off1 --> x <= 2)

Transition on01:
    enable s = on0 /\y = 10
    assign s := on1, x := 0

Transition onoff:
    enable s = on1 /\x = 2
    assign s := off0

Transition off01:
    enable s = off0 /\y = 5
    assign s := off1, x := 0

Transition offon:
    enable s = off1 /\x = 2
    assign s := on0

Activity Aon:
    enable s = on0 \/ s = on1
    assign Deriv(y) := 1

Activity Aoff:
    enable s = off0 \/ s = off1
    assign Deriv(y) := -2
```

Figure 2.10: The Waterlevel Controller hybrid transition system

- Arithmetization:

  `1 <= i <= N` is parsed as `(1 <= i) <= N`, hence as "the truth value of `1 <= i` is less than or equal to `N`", and not as `1 <= i /\ i <= N`.

  `x=y=0` is parsed as `(x=y) = 0`, hence as "the truth value of x=y is false", and not as the expected `x=0 /\ y=0`.

- Rigid vs. flexible variables (see above).

- Scope of quantifiers and declarations:

  Binding operators have maximal scope. Hence

```
Forall x : int . p(x) --> q
```

  is parsed as

```
Forall x : int . (p(x) --> q)
```

  and not as

```
(Forall x : int . p(x)) --> q
```

  although `x` is not free in q.

- When in doubt, add parentheses!

Chapter 3

The Top-Level Interface

Starting STeP brings up the Top-level Prover window (Figure 3.1), a point-and-click X-windows interface. The Top-level Prover window has the following main components:

1. At the top of the main STeP window is the *main menu*, used to load systems and specifications, save proof searches, view systems, select properties to be used in proofs, online help, and exit STeP.

2. Under the main menu is the *status line*. It displays the name of the current system file loaded, if any; the number of unfinished proof trees, the number of open goals in the current proof tree, the current simplifier settings, and whether to automatically apply Simplify, Check-Valid or none of these.

3. Below the status line is the *current goal window*, where the current unproven goal or subgoal is displayed. This window has a scroll bar at the left. Verification rules are applied to the formula that appears in this window.

4. Below the current goal window is the *output window*, which displays the results of applying rules and generating invariants. It will also state when a proof is complete, and generally leave a record of the system's activities. This window can be reset with the Clear text window item under the Settings pull-down menu.

5. Under the output window, at the bottom of the top-level prover, is STeP's *message/help line*. In this window, STeP indicates what action is being done, which dialog window is waiting for input, or if an action has terminated successfully or not. It will also display a brief description of any feature that is being selected. For example, if Undo is selected, this window will display "`Backtrack to parent`".

6. At the right side of the top-level window is the *action region*, a double column of buttons that invoke verification rules and decision procedures and apply them to the the current goal.

   Besides the (temporal) verification rules, buttons in this column invoke the automatic Simplifier or start up the Interactive Prover. The Previous and Next buttons cycle through the open goals in the current search.

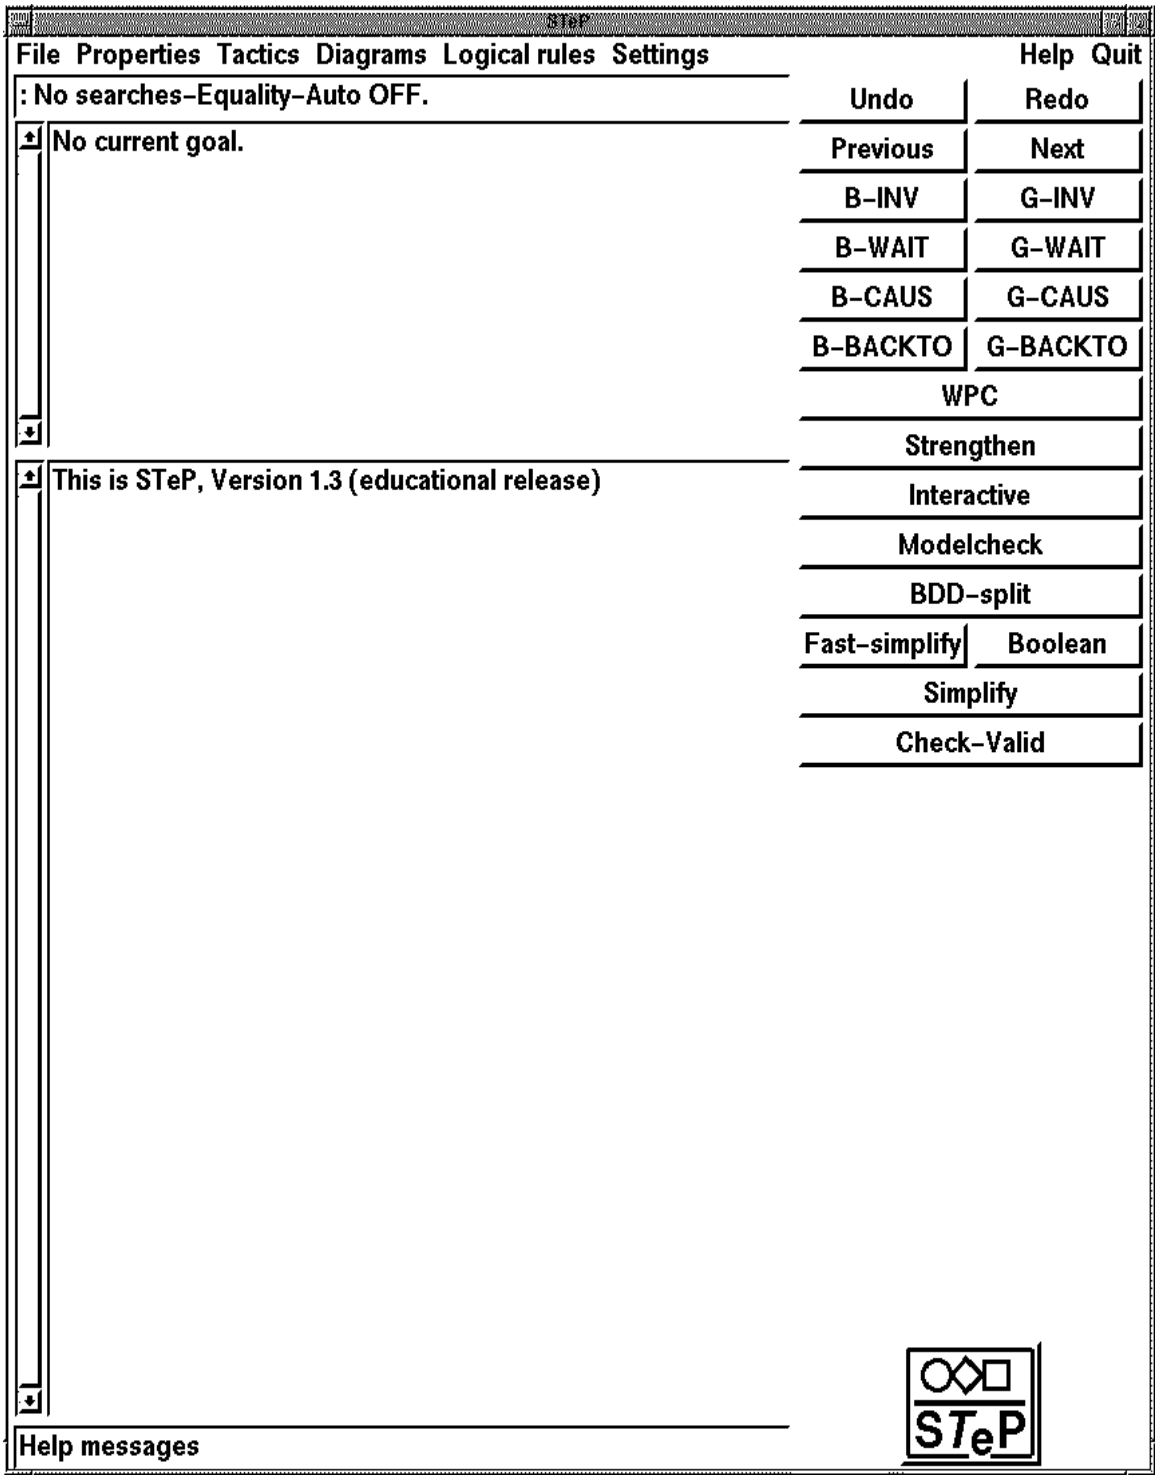

Figure 3.1: Top-level Prover window

7. In the lower right-hand corner of the window is a button with the STeP logo. This button functions as the *interrupt button*, and can be used to halt most of the automatic operations performed by STeP.

3.1 Menu Options

The File Menu

The File pull-down menu (Figure 3.2) is used to load system description and specification files, and to save finished and unfinished proofs.

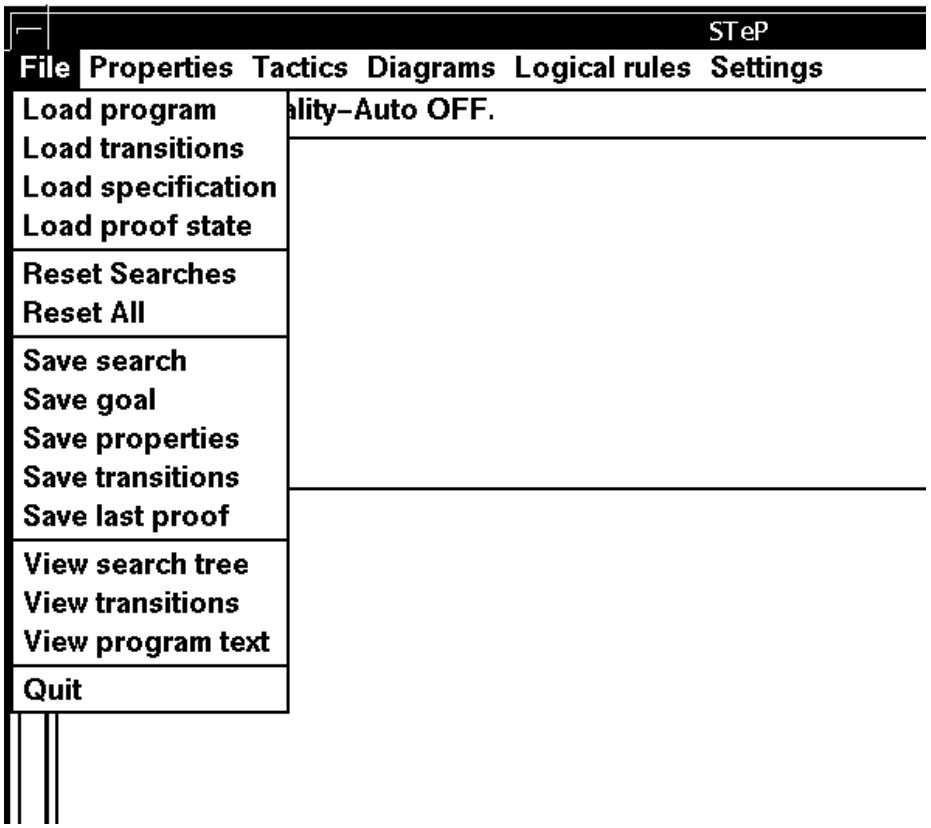

Figure 3.2: Top-Level Prover File Menu

Several options in this menu activate the *File Browser*, a utility that allows you to move around in the directory hierarchy and select files.

Activating the File Browser brings up a window with several panels as shown in Figure 3.3. The name of the window reflects the function that called it (in the figure, **Load program**). The window under `File name` shows the current file selected, and you can enter the filename here directly. The two buttons below allow you to display all files (*), or only those with the default extension for the current function (*.spl , in this case, since

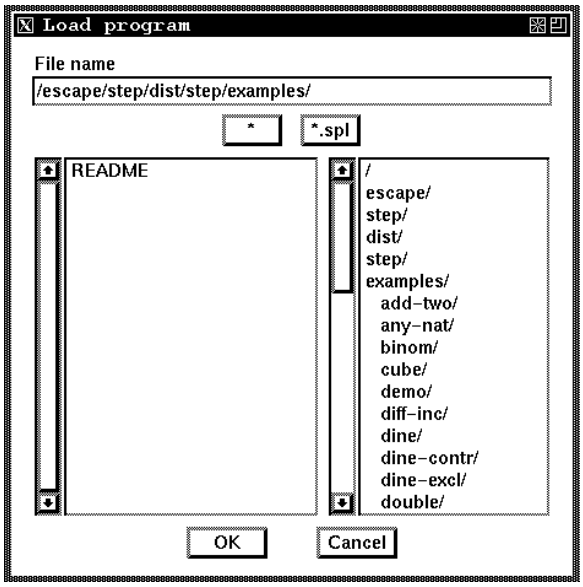

Figure 3.3: File Browser

the default extension for programs is `.spl`). The left panel below these buttons shows the files in the selected directory, and the right panel shows the directory hierarchy. To move to another directory, click on the directory name in the right panel. To select a file, click on the filename in the left panel. To load or save a file, click OK after selecting the file or entering its name.

If a file does not parse correctly, a **Read Error** window will appear with an indication of the first parsing error encountered. The file can be corrected and read again by clicking Try same file again, or the operation canceled with Cancel.

Load program Loads an SPL program from a file. The default extension for a program file is `.spl`. The program will be parsed into a fair transition system. The program text is displayed in a separate window (setting the `STEP_SHOW_LOADED` environment variable to OFF disables this—see Appendix E). The program text window may be deleted at any time by clicking Quit, and redisplayed with the View program text option below.

Figure 3.4 shows the program window for program MUX-PET1. The program window also displays the correspondence between the program locations and the value of the internal control location variable. Control locations are explained in detail in Appendix C.2.

If Load program is selected while another verification session is still in progress, a confirmation window will first appear. The entire verification session environment is reset when a new program is loaded, i.e., all background properties and unproven goals are deleted.

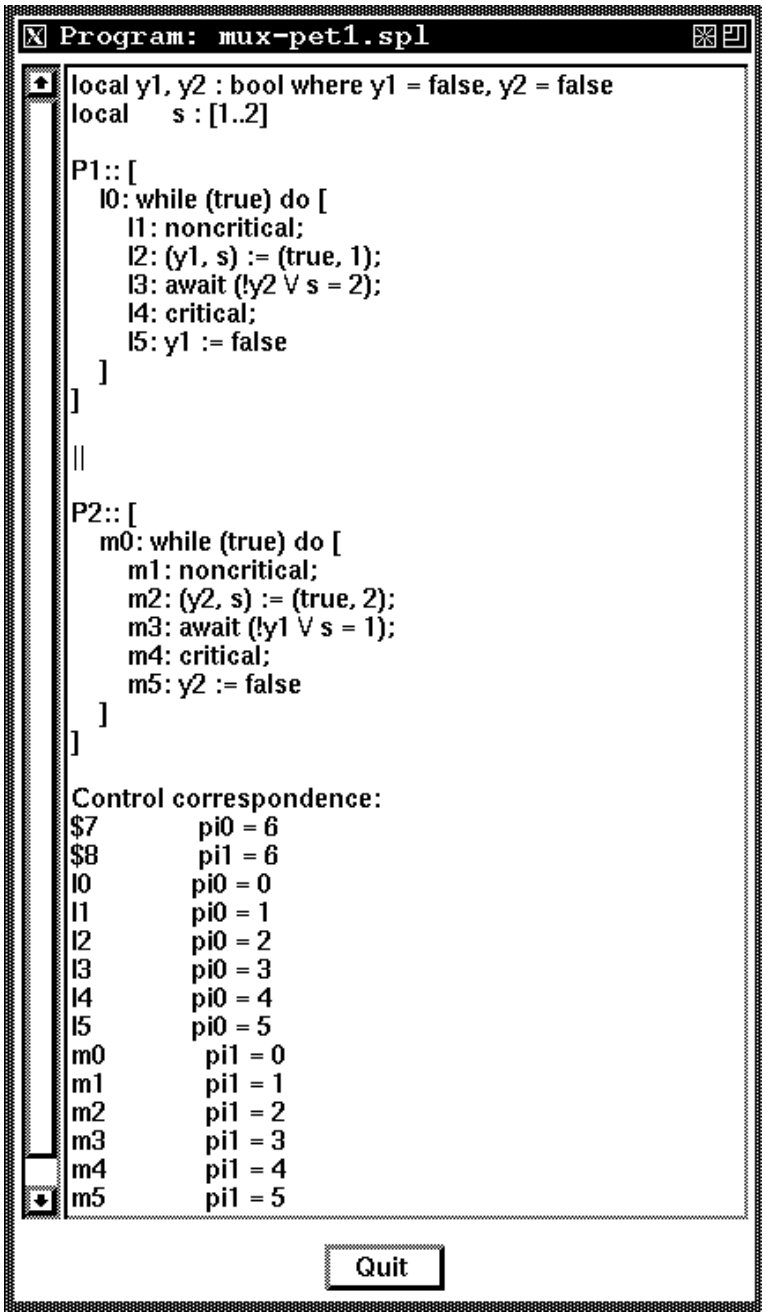

Figure 3.4: Program Window for program MUX-PET1

Load transitions The Load transitions option works like Load program, but expects a transition system file. The default extension for transition system files is `.fts`. The format for transition systems was given in Section 2.6.

As with programs, the transition system is displayed in a separate window (unless the `STEP_SHOW_LOADED` environment variable is `OFF`). Loading a new transition system terminates the current verification session.

Load specification Loads a specification (a list of axioms and properties to be proved) from a file. The default extension of a specification file is `.spec`. The input format of specification files was given in Section 2.3. Each property becomes the root of a separate proof tree, and the first one appears in the current goal window. If a specification file is loaded during another goal session, the current proof search is suspended and the first property in the specification file will become the new current goal, as the root of a new proof search. The old proof search, however, remains in the system, and may be returned to later (e.g. with the select search option).

Load proof state Loads a previously saved proof-state and resumes one proof, from where it has been saved. This option allows the user to store a proof search in a file, exit STeP and return to the proof when convenient later. One should choose the `.state` suffix when saving a search (see the menu option Save search) to the format suitable for reloading the state. When selecting Load proof state STeP assumes that the same program has been loaded, as on the time of saving.

Reset Searches Selecting Reset Searches deletes all pending proof trees and goals, as well as all background properties. However, the current program or transition system is retained. Thus, it resets a verification session.

Reset All Terminates a system-verification or theorem-proving session: it removes all current proof trees, background properties, and the current system description, if any.

Save search Saves the current proof tree to a file, in ASCII format, for off-line inspection (with the `.search` suffix), as a rerunnable tactic (with the `.tactic` suffix), and for instantaneous reloading (with the `.state` suffix). Note that the current proof tree is necessarily unfinished, i.e., it still has open subgoals. Once all subgoals are closed the proof tree is removed from the list, and the next proof tree, if any, becomes the current one. To save the last complete proof, use Save last proof, described below.

Save goal Saves the current goal to a file. The file will contain declarations for all free variables.

Save properties Saves all background properties to a file. The background properties include all properties proven and all invariants generated during the current program verification session, as well as all axioms asserted by the user.

Save transitions Saves the current transition system to a file; this is useful to see how SPL programs are converted to transition systems.

Save last proof Saves the finished proof of the most recently established goal.

View search tree Allows the user to navigate in the so-far developed proof search via a graphical interface of the proof tree. The current goal is drawn with a white node, other pending goals, and goals that have been backtracked from are grey, whereas intermediary nodes are black. By selecting nodes in the graphical interface one can change the current goal. The graphical display gets continually updated when navigating using the conventional verification rules.

View transitions Displays a separate window containing the current transition system. It can also be selected when the system description is loaded as an SPL program. The transition window can be deleted at any time by clicking on the Quit button of this window.

View program text Displays a separate window with the text of the current SPL program, similar to View transitions above.

Quit Terminates the STeP session, after confirmation from the user.

The Properties Menu

The Properties pull-down menu (Figure 3.5) is used to manage the background properties, move between different proof trees, generate invariants, and enter new goals and axioms directly.

Activate/Deactivate This option displays a window with all the background properties, i.e., all axioms entered, invariants generated, and properties proven during the current system verification session. (See for example Figure 3.6, which displays the axioms, some invariants and a proven property `aux1` of the program GCD.) By default all properties are active, which means that they are used by the Simplifier. This window allows you to deactivate properties, or delete them altogether. A property is deactivated by clicking on the switch, and removed by clicking on the skull icon. The buttons Select All and Deselect All select or deselect all the listed properties.

No window will appear when there are no background properties.

Select goal This menu option allows you to move to another open goal in the current proof tree. When selected, it displays a window listing all open goals in the current proof tree (see, for example, Figure 3.7). Each goal has a select button. By clicking on Select the corresponding goal is made the current one. The Cancel button returns to the current goal.

Select axiom To assert the result of an unfinished proof-search as an axiom in the current search, this menu option will display the set of unfinished searches which can be selected one by one. STeP keeps internal track of dependencies such that the user cannot make the result of a search depend on itself.

Next search Moves to the next proof tree with open subgoals. The previous proof tree is kept in its current state, and may be returned to later. This option is useful if you reach a dead-end in a proof and realize you need to prove another property first.

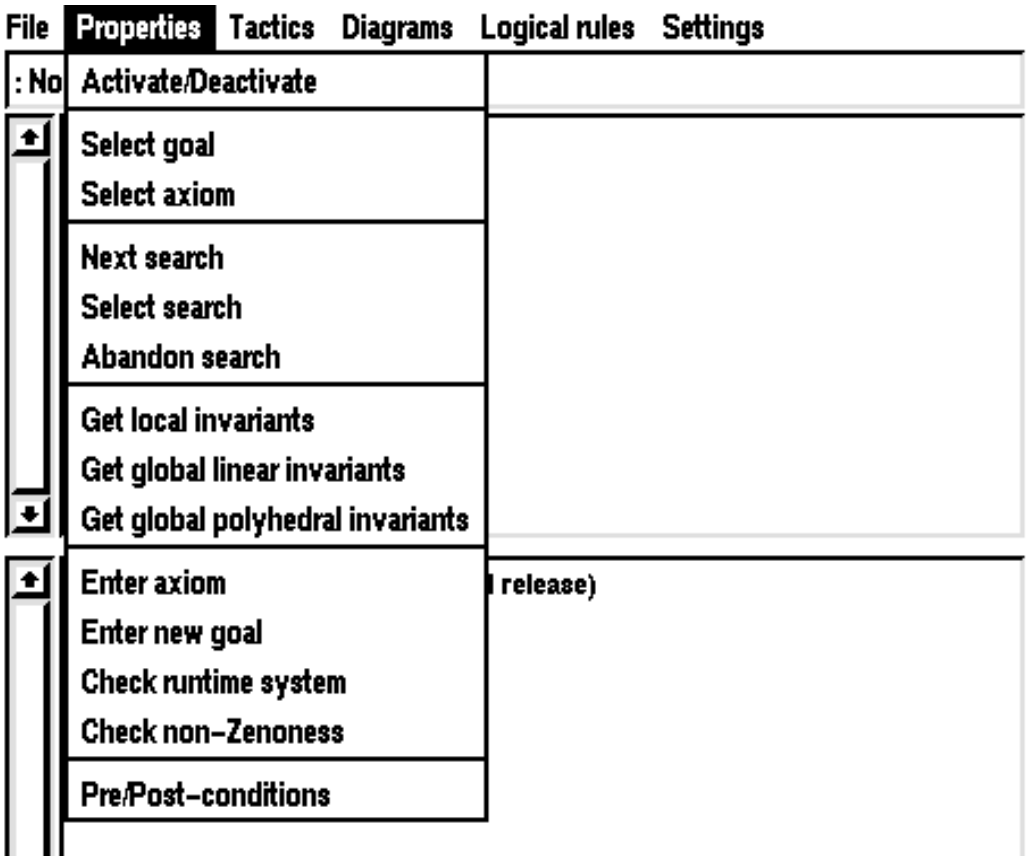

Figure 3.5: Properties menu

Select search Generalizes Next search above to move to any other proof tree. A window listing the root nodes of all proof trees is displayed. See Figure 3.8 for an example. By clicking on the Select button, the corresponding proof tree is made the current proof tree.

No window will appear if there is nothing to select.

Abandon search Deletes the current proof tree and makes the next proof tree in the list, if any, the current proof tree. A confirmation window appears before this action is taken.

The following options generate program invariants. When the invariant generation is done, the invariants are displayed in the output window and added to the list of background properties. By default they are active, and thus used in the simplification of goals. They may be deactivated with the Activate/Deactivate menu option (above). When the invariant generation is finished a message will appear in the message window.

Newly generated invariants are simplified using the basic simplification and decision procedures, independently of previous background properties. Invariant generation and

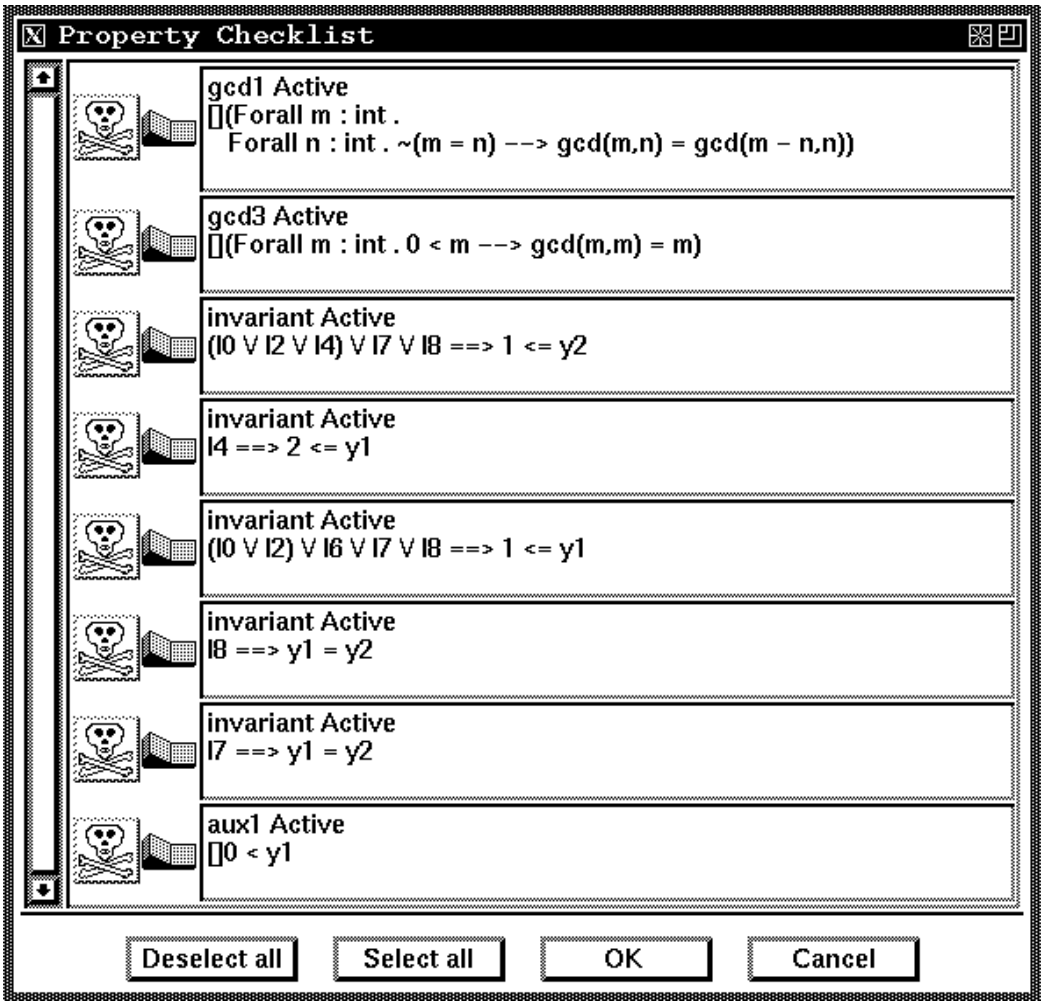

Figure 3.6: Activate/Deactivate Window

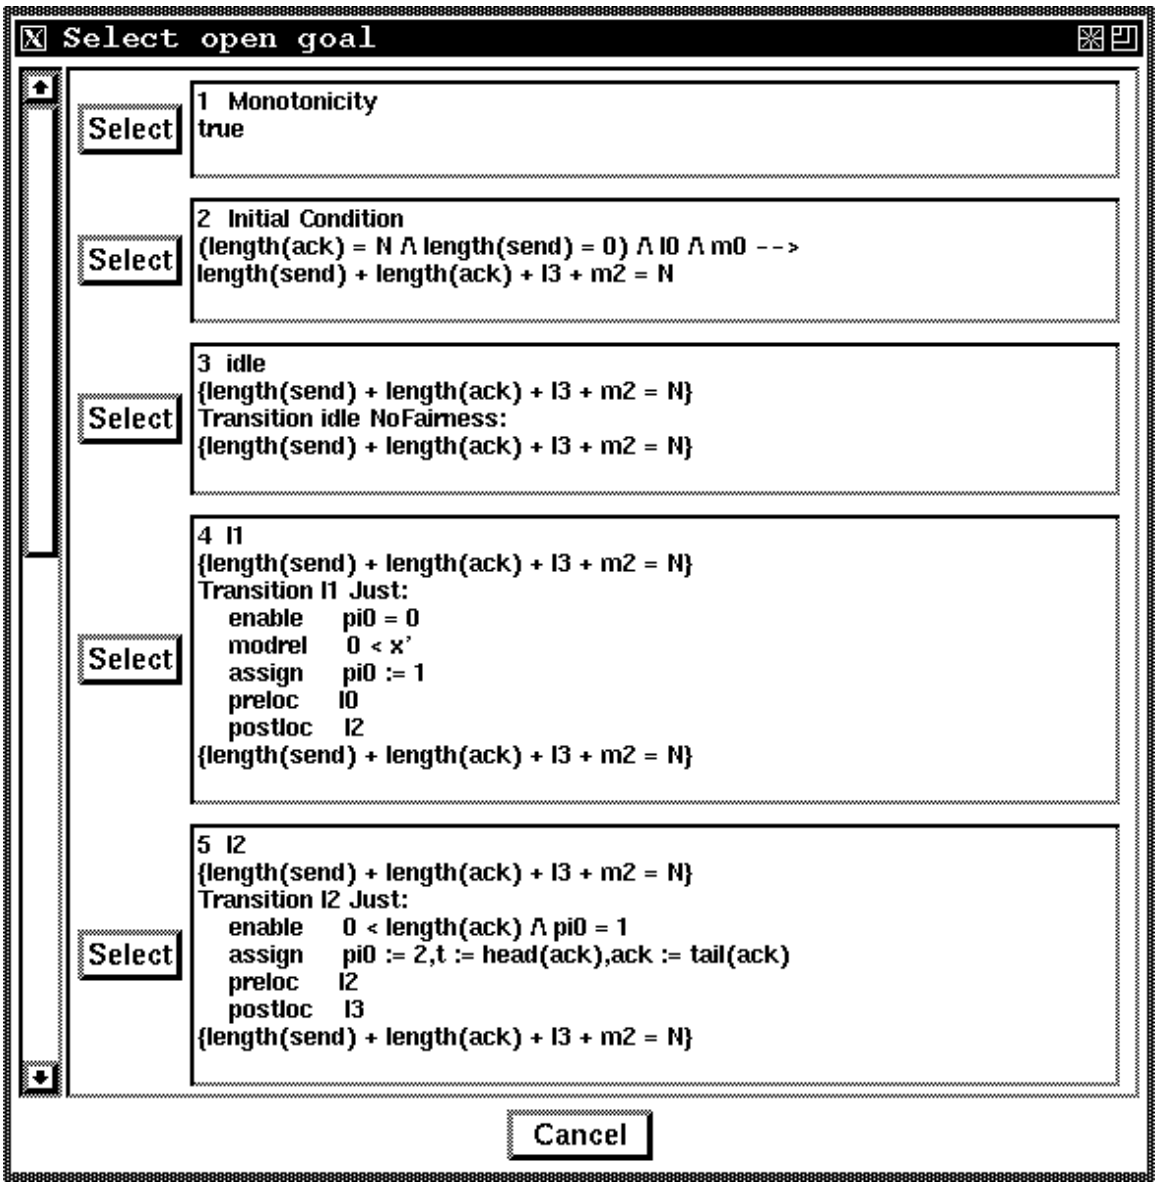

Figure 3.7: Goal selection window

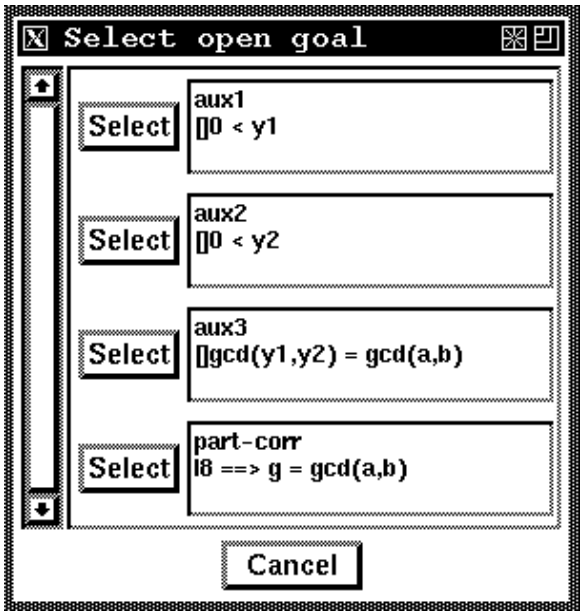

Figure 3.8: Proof tree selection window

simplification can be halted by clicking on the interrupt button (the STeP logo).

Get local invariants Activates the local invariant generator; requires an SPL program.

Get global linear invariants Activates the linear invariant generator; requires an SPL program or fair transition system.

Get global polyhedral invariants Activates the polyhedral invariant generator, after prompting the user for a number of parameters that control the generation method. See Chapter 7 for details. Polyhedral invariants are generated for SPL programs and fair transition systems.

Enter axiom Used to enter a property that the system may assume to be true during the current program verification session, i.e., it is added to the list of background properties. By default, the new axiom is active, and will thus be used by the Simplifier.

Note: STeP does not guarantee the consistency of the set of background axioms, nor does it require you to prove any of them, so this feature should be used with caution.

Enter new goal Used to enter a new goal to be proven. If the formula parses correctly it will become the root of a new proof tree, which is made the current proof tree while the previous one is put on hold.

Check runtime system Generates a set of goals that, if proved valid, ensure that the transition system will not generate any runtime errors such as division by 0 or array references out of range.

Pre/Post-conditions This option allows you to calculate *weakest liberal precondition* for as well as the *strongest postcondition* an arbitrary formula and transition in the current system (see Section 4.3).

The dialog window shown in Figure 3.9 (upper window) will appear; after entering the transition name and formula, clicking on the OK button will display the weakest liberal precondition or strongest postcondition in a separate window (shown in Figure 3.9's lower window). The formula may be saved to a file with the Save button.

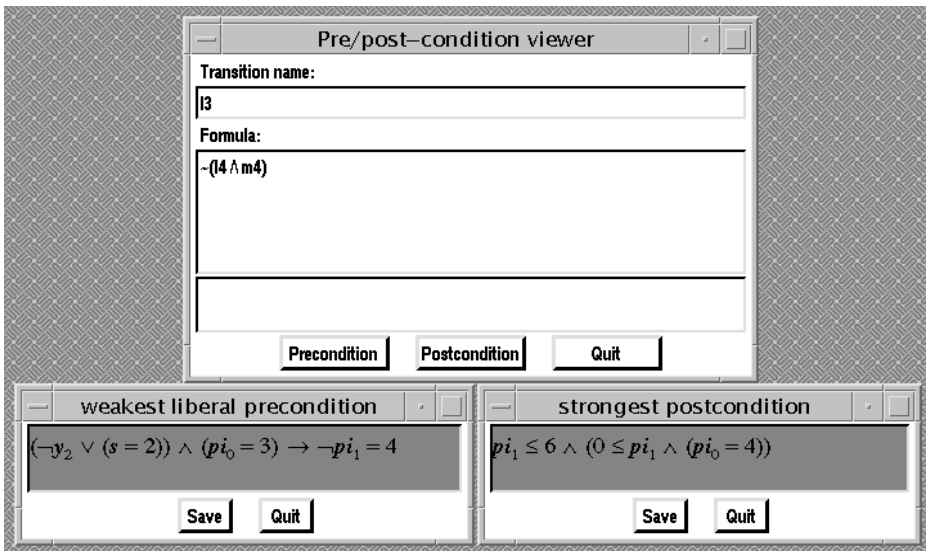

Figure 3.9: Weakest Pre- and Strongest Post-condition: entry and output window

This function has no effect on the proof tree or on the background properties; it serves a purely informative role. Weakest liberal precondition is also provided as an inference rule, useful for strengthening invariants, and is described in Section 4.3.

When a transition τ has relation $\rho_\tau(\overline{x},\overline{x}')$ then the weakest liberal precondition of $\varphi(\overline{x})$ is the equivalent of $\forall\overline{x}'.(\rho_\tau(\overline{x},\overline{x}') \rightarrow \varphi(\overline{x}'))$. Likewise the strongest postcondition is the equivalent of: $\exists\overline{x}_0.(\rho_\tau(\overline{x}_0,\overline{x}) \wedge \varphi(\overline{x}_0))$. Before displaying the pre and post-conditions the viewer will perform trivial simplifications, which may eliminate some of the quantifiers indicated here.

The Tactics Menu

The Tactics pull-down menu is used to run tactics loaded from a file or entered directly. Tactics can be run in two modes: *batchmode* and *interactive*. Batchmode tactics are applied without any user interaction; they may not always terminate, but can be interrupted by clicking on the STeP icon (the *interrupt button*).

For interactive tactics, a window appears with the first atomic step of the tactic. The Apply button applies the displayed step, and the next one will appear. In this way, you

can apply tactic steps until the proof is finished, or at any point interrupt the tactic by clicking on the Quit button.

Tactics and their syntax are presented in Section 6.3.

Load batch tactic Reads a tactic from a file using the File Browser and applies it in batch mode.

Load interactive tactic Reads a tactic from a file and runs it in interactive mode.

Enter batch tactic Prompts the user to enter a tactic, which is then parsed and run in batch mode. Error messages from parsing the tactic are displayed in the bottom of the input window.

Enter interactive tactic Prompts the user to enter a tactic, which is then parsed and run in interactive mode.

The Diagrams Menu

The Diagrams pull-down menu is used to start up the Verification Diagram Editor and use verification diagrams in the current verification session. The Verification Diagram Editor may be used in parallel with the Top-level Prover, and is described in more detail in Chapter 5.

Edit diagram Starts the Verification Diagram Editor in a separate window.

Verification Diagram rule Uses the current verification diagram as a verification rule for the current goal or subgoal. It generates subgoals corresponding to all the verification conditions implicit in the diagram, if the diagram is applicable. Thus, the verification diagram can be seen as a specialized verification rule.

A verification diagram is chosen to be the current one by the Current-to-verify menu option in the Verification Diagram Editor.

The Logical rules Menu

The Logical rules pull-down menu (Figure 3.10) provides inference rules that simplify temporal properties. Thus, they can be used to derive new properties from previously proven ones. Unlike verification rules, these rules are independent of the system being verified. Section 4.4 presents a full description of these rules.

The Settings Menu

The Settings pull-down menu (Figure 3.11) controls some of the global STeP environment variables, such as automatic simplification of subgoals and the strength of the Simplifier.

Automatic Simplification: Simplify/Check Valid/OFF Sets and resets automatic simplification; see Section 4.1.

Simplification Flags Brings up a window used to set the default strength of the Simplifier. See Section 6.1 for more details.

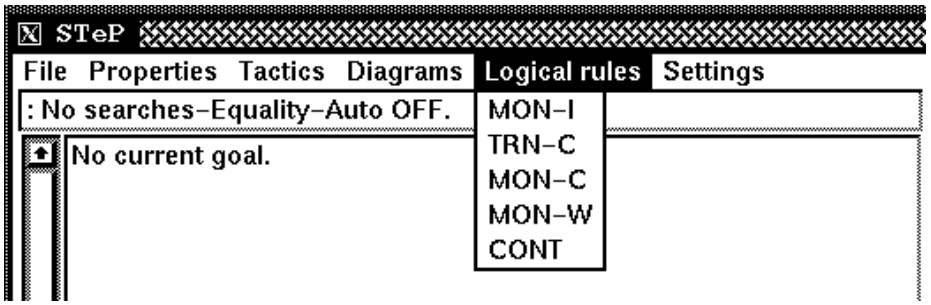

Figure 3.10: Logical Rules menu

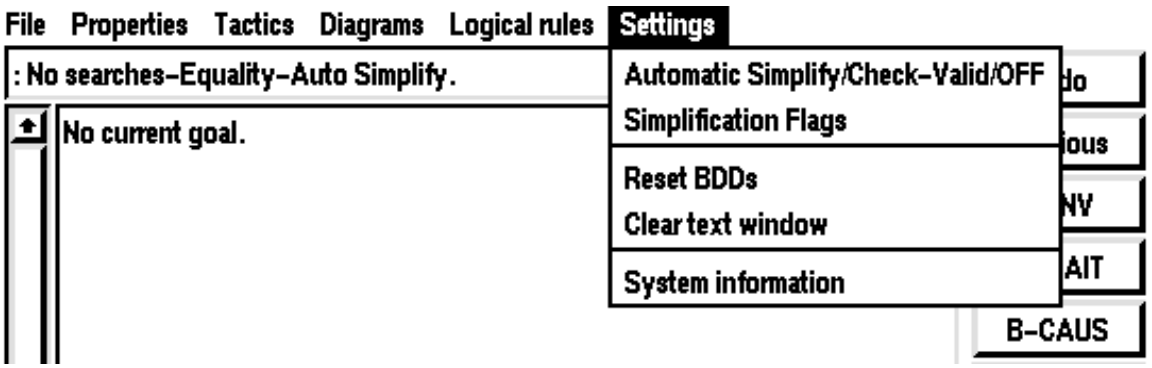

Figure 3.11: Settings menu

Reset BDDs Resets the internal BDD package; if the package is used often, this option can be used sporadically to save memory, at some (usually small) short-term efficiency loss (see Section 6.2). This option should also be used if the number of BDD nodes reaches the maximum given by the `STEP_BDD_NODES` environment variable (see Section E).

Clear text window Clears the output window, but has no other effect on the system.

System information Prints out miscellaneous statistics (in the output window) for the current verification session, mostly related to the BDD package (see Section 6.2).

The Help menu

On-line help is available via your favorite WWW browser when clicking on the Help menu near the right corner of the top-level interface. The hyper-text help pages give an alternative and often quicker introduction to STeP's facilities.

3.2 Action Buttons

The right-hand column of the Top-level Prover contains rules to manipulate and construct the main proof search. The first four buttons manipulate the structure of the proof tree:

Undo This button makes the parent of the current subgoal into the current goal. The previous goal is remembered, to enable a Redo, until another verification rule is applied to the new goal.

Redo Redoes a previous Undo, that is, restores the part of the proof tree eliminated by Undo.

Previous Postpones the current goal, and turns the previous subgoal in the subgoal queue into the current subgoal.

Next Postpones the current subgoal and moves to the next subgoal in the subgoal queue.

The next set of buttons invoke verification rules as described in [Manna and Pnueli, 1995b]. These rules are not applicable if there is no current program or transition system.

Verification rules:

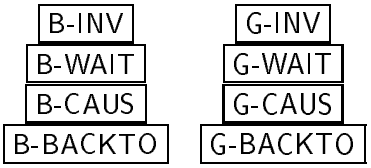

See Section 4.2 for a description of the main verification rules in their basic and general versions. Section 4.2 also describes the WPC and Strengthen proof rules.

The remaining buttons in the action region invoke the Interactive Prover and the Simplifier as proof rules, applied to the current goal:

Interactive Invokes the Interactive Prover on the current goal. The Interactive Prover window will appear with the current goal as the single consequent in the sequent. Once the Interactive Prover is active, the Top-level Prover is suspended, to be reactivated when the Interactive Prover session is finished.

Upon exiting the Interactive Prover you have the option to return the result to the Top-level Prover. If the goal was proved, then the current goal is closed, and the next goal in the proof tree, if any, becomes the current goal. If the goal was not proved, a formula equivalent to the remaining subgoals is returned to the Top-level Prover as a new subgoal.

The Interactive Prover is described in Chapter 8.

Modelcheck Invokes the Model Checker on the current goal. Before the Model Checker starts, you can enter additional assumptions to limit the search space. If the goal is proved without any additional assumptions, the current goal is closed, and the next goal in the proof tree, if any, becomes the current goal. If the goal is not proved, it remains unchanged.

If the goal is proved under additional assumptions, it is not added to the set of background properties.

The Model Checker results are reported in a logfile. See Section 4.5 for a more detailed description of the modelchecker.

BDD-split, Fast-simplify, Boolean, Check-Valid, Simplify These buttons invoke various versions of the automatic Simplifier; while BDD-split can generate a set of subgoals, the others can only reduce the formula to a simpler one. See Section 6.1 for details.

STeP **logo** The STeP logo serves as the interrupt button, and can be used to exit from each of the following time-consuming operations:

- Simplification of verification conditions
- Generation and simplification of invariants
- Model checking

Chapter 4

Verification and Logical Rules

The main function of the Top-level Prover is to apply *verification rules* to temporal properties, reducing them to simpler verification conditions. These verification conditions become subgoals in the top-level proof tree. The verification rules depend on the system being verified. In contrast, a set of *logical rules* reduces temporal properties independent of the system being verified.

This chapter describes the general structure of the proof search and the rules used to construct it in the Top-level Prover.

4.1 The Proof Search

The structure of the proof search is similar in the Top-level Prover and the Interactive Prover: the proof is a tree, with the original goal as its root node. Nodes are labeled with subgoals, called *verification conditions* in the Top-level Prover and *sequents* in the Interactive Prover, and the application of different rules generates new subgoals (or children nodes) for the current node.

Since more than one proof rule can be applied at each node, the process is a *search*. The proof is finished when there are no subgoals left; the number of subgoals is reduced whenever a simplification step reduces a subgoal to *true*, which is the only formula that requires no proof.

In the Top-level Prover, the rules applied are temporal verification rules, logical rules, or simplification steps. In addition, rules that can strengthen a verification condition are available, such as weakest precondition. Finally, goals in the Top-level Prover can also be closed by sessions of the Interactive Prover.

The Background Properties:

Background properties include axioms, automatically generated invariants, and previously established properties. When a top-level proof is completed, the property is added to the list of background properties.

In the Top-level Prover, simplification is carried out relative to the properties and axioms that are active at the time (except for the Fast-simplify option, discussed in Section 6.1).

In the Interactive Prover, these background properties have to be explicitly added by the user to the current sequent before the simplification rule is called.

Navigating the Proof Tree:

The options Postpone (in the Interactive Prover), Previous, Next (in the Top-level Prover), and Select goal (in both provers) can be used to jump from one open node to another. The Undo and Redo buttons alternate between a given node and its children.

In the Top-level Prover, any number of proof searches can be conducted simultaneously, sharing the same set of background properties. The Next search button can be used to jump from one proof search to the next.

Automatic Simplification

A verification rule can generate a large number of verification conditions, depending on the size of the underlying transition system. After having applied a verification rule a message in the message window will display how many subgoals were generated. If automatic simplification is OFFthe first subgoal will be displayed in the specification window. All the generated subgoals can be viewed by selecting the Select goal option of the Properties pull-down menu.

If automatic simplification is ON, the Simplifier will be applied to each subgoal. Only those subgoals that fail to simplify to *true* will remain. Simplifications steps which fail to simplify to *true* will be undone to enable subsequent application of rules such as Strengthen and WPC; these are only applicable to unsimplified verification conditions, since they require knowing the particular transition that generates the verification condition.

When automatic simplification is OFF, a verification rule only generates subgoals, but does not simplify them. These then have to be manually simplified by successive clicks on Simplify, or with the Interactive Prover. If you unintentionally invoke a verification rule with automatic simplification OFF, you can click on Undo, turn automatic simplification ON, and apply the rule again.

4.2 Verification Rules

This section describes the verification rules available in STeP. In the following description, Θ refers to the initial condition of the current program or transition system, and $\mathcal{T}$ refers to its set of transitions (see Appendix A). For a formula p, p_0 is the *initial version* of p, that is, the first-order formula equivalent to p at the initial state of the computation.

As usual, the notation $\{p\}\tau\{q\}$, stands for the verification condition $(p \wedge \rho_\tau) \Rightarrow q'$. The *inverse verification condition*, $\{p\}\tau^{-1}\{q\}$, is equivalent to $\{\neg q\}\tau\{\neg p\}$, giving an alternate presentation of verification conditions. Below, we try to use the most convenient one in each case. The notation $\{p\}\mathcal{T}\{q\}$ stands for the set of verification conditions $\{p\}\tau\{q\}$ for each transition τ in the system; similarly, $\{p\}\mathcal{T}^{-1}\{q\}$ is the set $\{p\}\tau^{-1}\{q\}$ for each τ.

Verification conditions of the form $\alpha \Rightarrow \beta$ are replaced by $\alpha \rightarrow \beta$ when α and β are *assertions*, that is, state-formulas with no temporal operators. In this case, the P-validity

of $\alpha \rightarrow \beta$ implies that of $\alpha \Rightarrow \beta$. If this P-validity cannot be established, the invariant $\alpha \Rightarrow \beta$ may be proved separately, for an incremental proof.

- B-INV Applies the *basic invariance rule* to the current subgoal, which has to be of the form $\Box p$, for a past formula p. This rule generates the following subgoals:

$$\begin{array}{ll} p \Rightarrow p & \text{Monotonicity} \\ \Theta \rightarrow p_0 & \text{Initial condition} \\ \{p\}\tau\{p\}, & \text{for each } \tau \in \mathcal{T} \end{array}$$

  The monotonicity subgoal is generated because this rule is based on rule G-INV, below; this goal always trivially simplifies to *true*.

- G-INV Applies the *general invariance rule* to the current subgoal. The current goal has to be of the form $\Box q$, with q a past formula. After clicking, a dialog window will appear that allows you to enter a strengthened version p of q. The default is q. After entering the auxiliary assertion p, the following subgoals are generated:

$$\begin{array}{ll} p \Rightarrow q & \text{Monotonicity} \\ \Theta \rightarrow p_0 & \text{Initial condition} \\ \{p\}\tau\{p\} & \text{for each } \tau \in \mathcal{T} \end{array}$$

- B-WAIT Applies the basic *nested wait-for* rule. This rule is applicable only if the current goal is a nested wait-for formula, that is, it is of the form

$$p \Rightarrow q_n \,\mathcal{W}\, q_{n-1} \,\mathcal{W}\, q_{n-2} \ldots q_1 \,\mathcal{W}\, q_0,$$

  where $p, q_n, \ldots q_0$ are past formulas. Informally, the nested wait-for formula asserts that whenever p becomes true, when there is a, possibly empty, interval where q_n holds, followed by a q_{n-1} interval, up to q_0 becomes true, unless one of the previous q's remains true infinitely after. After clicking on this button the following subgoals are generated:

$$\begin{array}{ll} p \Rightarrow \bigvee_{i=0}^{n} q_i & \text{Premise N1} \\ q_i \Rightarrow q_i & \text{Monotonicity } i, \text{ for each } i = 0 \ldots n \\ \{q_i\}\mathcal{T}\{\bigvee_{j \le i} q_j\} & \text{for each } i = 1 \ldots n \end{array}$$

  The monotonicity conditions are generated because the implementation of B-WAIT is based on that of G-WAIT, below. They always trivially simplify to *true*.

- G-WAIT This button invokes the general *nested wait-for* rule. This rule is applicable only if the current goal is a nested wait-for formula, that is, it is of the form

$$p \Rightarrow q_n \,\mathcal{W}\, q_{n-1} \,\mathcal{W}\, q_{n-2} \ldots q_1 \,\mathcal{W}\, q_0,$$

  where $p, q_n, \ldots q_0$ are past formulas. A dialog window will appear which allows you to enter the auxiliary assertions $\varphi_n, \ldots, \varphi_0$. Their defaults are $q_n, \ldots, q_0$, respectively. After entering these assertions the following subgoals are generated:

$$\begin{array}{ll} p \Rightarrow \bigvee_{i=0}^{n} \varphi_i & \text{Premise N1} \\ \varphi_i \Rightarrow q_i & \text{Monotonicity } i, \text{ for each } i = 0 \ldots n \\ \{\varphi_i\}\mathcal{T}\{\bigvee_{j \le i} \varphi_j\} & \text{for each } i = 1 \ldots n \end{array}$$

- B-CAUS The basic causality rule is a specialization of the basic invariance rule. It is applicable to goals of the form

$$p \Rightarrow \diamondminus q,$$

where p and q are past formulas. It usually generates simpler verification conditions than its counterpart B-INV, which is also applicable to goals of this form, by taking into account the special form of the goal.

This rule generates the following subgoals:

$$\begin{array}{ll} p \Rightarrow p \ \vee \ q & \text{Monotonicity} \\ \Theta \rightarrow \neg(p)_0 \ \vee \ (q)_0 & \text{Initial condition} \\ \{p\}\mathcal{T}^{-1}\{p \vee q\} & \end{array}$$

- G-CAUS The general causality rule specializes the general invariance rule. It is applicable to goals of the form

$$p \Rightarrow \diamondminus q,$$

where p and q are past formulas. It usually generates simpler verification conditions than its counterpart G-INV, which is also applicable to goals of this form, by taking into account the special form of the formula to be proved.

After clicking on this button a dialog window will appear, which allows you to enter an intermediate formula φ (default q). After φ is entered, the following subgoals are generated:

$$\begin{array}{ll} p \Rightarrow \varphi \ \vee \ q & \text{Monotonicity} \\ \Theta \rightarrow \neg(\varphi)_0 \ \vee \ (q)_0 & \text{Initial condition} \\ \{\varphi\}\mathcal{T}^{-1}\{\varphi \vee q\} & \end{array}$$

- B-BACKTO The *basic back-to rule* establishes properties of the form

$$p \Rightarrow q_n \, \mathcal{B} \ \ldots \ \mathcal{B} \, q_0$$

for past formulas $p, q_n, \ldots, q_0$. It generates the following subgoals:

$$\begin{array}{ll} p \Rightarrow \bigvee_{i=0}^{n} q_i & \text{Premise N1} \\ q_i \Rightarrow q_i & \text{Monotonicity for each } i \text{ in } 1 \ldots n \\ \{q_i\}\mathcal{T}^{-1}\{\bigvee_{j \leq i} q_j\} & \text{for each } i \text{ in } 1 \ldots n \end{array}$$

- G-BACKTO The *general back-to* rule generalizes the above by allowing strengthening of the q_i's and weakening of p. To prove

$$p \Rightarrow q_n \, \mathcal{B} \ \ldots \ \mathcal{B} \, q_0,$$

given formulas $\varphi_0, \ldots, \varphi_n$, the following subgoals are generated:

$$\begin{array}{ll} p \Rightarrow \bigvee_{i=0}^{n} \varphi_i & \text{Premise N1} \\ \varphi_i \Rightarrow q_i & \text{Monotonicity for each } i \text{ in } 1 \ldots n \\ \{\varphi_i\}\mathcal{T}^{-1}\{\bigvee_{j \leq i} \varphi_j\} & \text{for each } i \text{ in } 1 \ldots n \end{array}$$

4.3 WPC and Strengthening

A basic and often successful heuristic for constructing inductive invariance properties is based on *weakest preconditions*.

- WPC The weakest precondition verification rule can be used to prove a verification condition that does not itself simplify to *true*. The rule is applicable to subgoals of the form

  $$\{\varphi\}\tau\{\psi\}$$

  and when applied it generates the subgoal

  $$\text{Weakest Precondition} \qquad \Box(\rho_\tau \rightarrow \psi')$$

  which implies the original verification condition.

  The symbol ρ_τ refers to the transition relation, and is the logical form of the transition τ. The formula ψ' is the primed version of ψ, in which each system variable occurring free in ψ has been replaced by its primed version. Appendix A describes transitions and transition relations in more detail.

- Strengthen The strengthening rule can be used to strengthen a property that is not inductive. The rule is applicable to subgoals of the form

  $$\{\varphi\}\tau\{\psi\}$$

  and when applied it generates the subgoal

  $$\Box(\varphi \wedge (\rho_\tau \rightarrow \psi')).$$

  while deleting the original verification condition and its siblings from the proof tree.

4.4 Logical Rules

The Logical rules pull-down menu provides inference rules that allow the derivation of new properties from previously proven properties.

If the current formula does not match the kind of formula that the given rule applies to, a message will appear in the bottom help window, and the given function will have no effect.

MON-I The *monotonicity rule for invariances* can be used to prove a property of the form $\Box p$, where p is a state formula. If the current goal is $\Box p$, the following subgoal is generated:

$$\bigwedge_i \varphi_i \;\rightarrow\; p$$

where φ_i ranges over all the currently active in the background properties. MON-I also applies to goals of the form p, where p is a state formula, and $\{p\}\tau\{q\}$, where p and q are state formulas.

TRN-C The *transitivity of causality rule* proves a causality property from two other causality properties. If the current goal has the form $p \Rightarrow \diamondsuit r$, this option displays a dialog window for entering an auxiliary assertion. After a new assertion q is entered (default is r), the following two subgoals are generated:

$$\begin{array}{ll} p \Rightarrow \diamondsuit q & \text{Left Transitivity} \\ q \Rightarrow \diamondsuit r & \text{Right Transitivity} \end{array}$$

MON-C The *monotonicity rule for causality* can prove a causality property from a stronger causality property. If the current goal has the form $p \Rightarrow \diamondsuit u$, for past formulas p and u, this option generates a dialog window for entering two auxiliary assertions. After entering the two assertions q and r (defaults are p and u) the following three subgoals are generated:

$$\begin{array}{rl} p \Rightarrow q & \text{Monotonicity 1} \\ q \Rightarrow \diamondsuit r & \text{Causality} \\ r \Rightarrow u & \text{Monotonicity 2} \end{array}$$

MON-W The *monotonicity rule for wait-for formulas* can be used to prove a wait-for formula from a stronger wait-for formula. If the current goal has the form $p \Rightarrow q_n \mathcal{W} \ldots \mathcal{W} q_0$, a dialog window will appear for entering $n+2$ auxiliary assertions. The default assertions are $p, q_n, \ldots, q_0$. After entering these assertions, $p_{new}, q_{n,new}, \ldots q_{0,new}$ and clicking the OK button, the following $n+3$ subgoals are generated:

$$\begin{array}{rl} p \Rightarrow p_{new} & \text{Monotonicity} \\ q_{n,new} \Rightarrow q_n & \text{Monotonicity } n \\ \vdots & \\ q_{0,new} \Rightarrow q_0 & \text{Monotonicity 0} \end{array}$$

and

$$p_{new} \Rightarrow q_{n,new} \mathcal{W} \ldots \mathcal{W} q_{0,new}$$

CONT The *contradiction rule* can be used to prove a general safety property by converting it into a causality formula.

If the current goal is of the form $\Box \neg p$, for past formula p, the following subgoal will be generated:

$$p \Rightarrow \diamondsuit \texttt{false} \qquad \text{Contradiction}$$

4.5 The Model Checker

Model checking determines if a given temporal formula holds for a given system by a systematic exploration of the state-space of the system. STeP includes a model checker for transition systems and formulas of linear-time temporal logic, which can be used to close subgoals or produce counterexamples to particular properties.

The STeP model checker checks the current goal under an optional additional first-order assumption. If the goal is found to be false, a counterexample is reported.

When the Modelcheck proof rule is invoked, a dialog window appears requesting the input of an auxiliary first-order assumption, by default *true*. You may enter restrictions on parameters or system variables with infinite domains, in order to make the program finite-state. For example, for the program MPX-SEM you may enter `M=2`, indicating that the number of processes you want to consider is two.

The goal is closed only if the Model Checker proves it with assumption *true*.

The Model Checker Preprocessor

The Model Checker first applies a preprocessor to the given system and the goal. If possible, a finite number of problem instances with constant-size arrays and constant-bounded parameters equivalent to the original goal is produced. If this is not possible, the Model Checker terminates, control returns to the Top-level Prover and a message is displayed. Some of the tasks accomplished by the preprocessor are:

- simplification of data types;
- elimination of unnecessary variables;
- instantiation of parameters and other variables;
- introduction of additional transitions corresponding to possible run-time errors;
- general simplification of the transition system and of the property, using the Simplifier.

Note that the additional assumption can be used to restrict the range of a variable on which array or parameter bounds depend.

The Core Model Checker

For each problem instance, an external core model checker is invoked. It searches the state-space and stops when

1. a run-time error is found,
2. a counterexample computation is found, or
3. the entire state space has been explored.

The main algorithm used by the core model checker is adapted from [Hojati *et al.*, 1993].

Restrictions

Not all programs and properties accepted by STeP can be model checked. The datatypes handled by the core model checker are: boolean, integer, and channels, and arrays over booleans and integers. The preprocessor can reduce some other datatypes to these. The model checker can only handle basic logic, arithmetic, and channel operations applied to these variables. Formulas given to the Model Checker cannot have temporal operators in the scope of binding operators.

Log file

The model checking results are reported in a *log file*, selected by the user before model checking begins.

After calling the core model checker, the log file will contain: the list of relevant system variables, the preprocessed transition system, some statistical information, and, if the formula proved to be invalid, a counterexample computation.

Environment Variables

The search space may be infinite if there are system variables that range over an infinite domain, and in this case the Model Checker may or may not terminate. Even if the system is finite-state, the state-space may prove too large for the time and memory available. The following settings can be used to limit the Model Checker resources:

- `STEP_MC_SPACE`: Maximum amount of memory, in Megabytes.
  Default: unlimited.
- `STEP_MC_TIME`: Maximum amount of user time, in minutes.
  Default: unlimited.
- `STEP_MC_CPU`: Maximum amount of CPU time, in minutes.
  Default: unlimited.

The Model Checker can be interrupted with STeP's *interrupt button* (the STeP logo). NOTE: This may leave a UNIX process running in the background; see the UNIX `kill` man pages for details on how to kill runaway processes.

Chapter 5

The Verification Diagram Editor

Verification diagrams are a visual representation of a proof. A verification diagram describes sets of states given by assertions and possible transitions between them. A well-formed verification diagram represents a set of verification conditions that are sufficient to establish a given temporal formula. Verification diagrams are described in detail in [Manna and Pnueli, 1994] and [Manna and Pnueli, 1995b].

STeP currently supports three types of verification diagrams:

- *Invariance diagrams*, to represent proofs of invariance formulas: $\Box p$.
- *Wait-for diagrams*, to represent proofs of nested wait-for formulas:

$$p \Rightarrow q_n \mathcal{W} q_{n-1} \ldots \mathcal{W} q_0.$$

- *Chain diagrams*, to represent proofs of response formulas $p \Rightarrow \Diamond q$ that do not require induction.

Here, p, q, $q_0, \ldots, q_n$ are state assertions.

Hierarchical verification diagrams (Section 5.7) support distributing a large verification diagram over several files.

5.1 Definitions

A verification diagram is a directed labeled graph constructed as follows:

- *Nodes* in the graph are labeled by assertions.
- *Edges* in the graph represent transitions between nodes. Each edge connects one node to another and is labeled by the name of a transition in the program. We refer to an edge labeled by τ as a τ-*edge*.
- One of the nodes may be designated as a *terminal node*. This node is distinguished by having a boldface boundary. No edges may depart from a terminal node.

Verification diagrams represent sets of verification conditions as follows: For every non-terminal node (labeled by) φ and transition τ, let $\varphi_1, \ldots, \varphi_k$ be the nodes to which φ is connected by τ-edges. The *verification condition associated with φ and τ* is given by:

$$\{\varphi\}\ \tau\ \{\varphi \vee \varphi_1 \vee \ldots \vee \varphi_k\}.$$

Note that if $k = 0$, i.e., no τ-edges depart from φ, the verification condition is:

$$\{\varphi\}\tau\{\varphi\}.$$

No verification conditions are associated with terminal nodes. A diagram is called *valid over program P (P-valid)* if all the verification conditions associated with nodes of the diagram are P-state valid. Thus, a valid diagram provides a succinct representation of the verification conditions of the form $\{\varphi\}\tau\{\psi\}$, for all combinations of assertions φ and ψ and transitions τ that appear in the diagram.

5.2 Wait-for Diagrams

A verification diagram is called a WAIT-FOR diagram if there is only one terminal node and it is labeled by φ_0 and the diagram is weakly acyclic, i.e., whenever node φ_i is connected by an edge to node φ_j, $i \geq j$. P-valid WAIT-FOR diagrams can be used to establish the P-validity of a nested wait-for formula.

A P-valid WAIT-FOR diagram establishes that the formula

$$\left(\bigvee_{j=0}^{m} \varphi_j\right) \Rightarrow \varphi_m \mathcal{W} \varphi_{m-1} \ldots \varphi_1 \mathcal{W} \varphi_0$$

is P-valid. To establish the P-validity of

$$p \Rightarrow q_m \mathcal{W} q_{m-1} \ldots \mathcal{W} q_0$$

the following additional verification conditions (referred to as *side verification conditions*) have to be established:

$$p \rightarrow \bigvee_{j=0}^{m} \varphi_j \qquad \text{(W1)}$$

$$\varphi_i \rightarrow q_i \quad \text{for } i = 0, 1, \ldots, m. \quad \text{(W2)}$$

5.3 Invariance Diagrams

A verification diagram is called an INVARIANCE diagram if it is exit-free, i.e., it has no terminal node. Unlike WAIT-FOR diagrams, INVARIANCE diagrams may contain cycles.

A P-valid INVARIANCE diagram with nodes $\varphi_1, \ldots, \varphi_m$ establishes that the formula

$$(\bigvee_{j=1}^{m} \varphi_j) \Rightarrow \Box(\bigvee_{j=1}^{m} \varphi_j)$$

is P-valid. To establish the validity of $\Box q$ the following two side verification conditions have to be established:

$$\Theta \rightarrow \bigvee_{j=1}^{m} \varphi_j \quad \text{(I1)}$$

$$\left(\bigvee_{j=1}^{m} \varphi_j\right) \rightarrow q \quad \text{(I2)}$$

5.4 Chain Diagrams ⋆

Verification diagrams can also represent fairness requirements. The main extension over the previous types of diagrams is the introduction of two additional types of edges, *double* and *solid*. A double edge corresponds to a helpful transition that is just, and a solid edge corresponds to a helpful transition that is compassionate. Thus, edges connecting nodes in a diagram can be single, double or solid.

A verification diagram is said to be a CHAIN diagram if its nodes are labeled by assertions $\varphi_0, \ldots, \varphi_m$, with φ_0 being the terminal node, and if it satisfies the following requirements:

- If a single edge connects node φ_i to node φ_j, then $i \geq j$.
- If a double or solid edge connects φ_i to φ_j, then $i > j$.
- Every node φ_i, $i > 0$, has at least one departing edge which is either double or solid. The transition labeling this edge is identified as helpful for assertion φ_i.
- Every transition can label at most one type of edge (single, double, or solid) departing from the same node.

The first two requirements ensure that the diagram is weakly acyclic in the same sense as defined for the WAIT-FOR diagram. The stronger second requirement ensures that the subgraph based on the double and solid edges is acyclic, forbidding self-connections by double or solid edges. The third requirement demands that every nonterminal node has at least one helpful transition associated with it.

Chain diagrams allow proving progress properties that do not require an infinite size well-founded domain. Later releases of STeP will allow the user to enter well-founded domains as part of the chain diagram.

Verification Conditions

Let φ be a nonterminal node and $\varphi_1, \ldots, \varphi_k$, $k \geq 0$, be the τ-successors of φ.

- If τ labels a single edge, we associate with φ and τ the usual verification condition

$$\{\varphi\}\tau\{\varphi \vee \varphi_1 \vee \ldots \vee \varphi_k\}.$$

- If τ labels a double or solid edge, we associate with φ and τ the verification condition

$$\{\varphi\}\tau\{\varphi_1 \vee \ldots \vee \varphi_k\}.$$

- If τ labels a double edge departing from φ, we associate with φ and τ the implication

$$\varphi \rightarrow En(\tau).$$

  i.e., when φ holds, τ is enabled.

- If τ labels a solid edge departing from φ, we associate with φ and τ the compassion requirement in the form of the response formula

$$\varphi \Rightarrow \Diamond(\neg\varphi \vee En(\tau)).$$

Valid Chain Diagrams

A CHAIN diagram is *P-valid* if the first three types of verification conditions are P-state valid and the compassion requirements associated with solid edges, are P-valid.

A P-valid CHAIN diagram establishes that the response formula

$$(\bigvee_{j=0}^{m} \varphi_j) \Rightarrow \Diamond \varphi_0$$

is P-valid. To establish the P-validity of

$$p \Rightarrow \Diamond q$$

the following two additional side verification conditions have to be established:

$$p \rightarrow \bigvee_{j=0}^{m} \varphi_j \quad \text{(FC1)}$$

$$\varphi_0 \rightarrow q \qquad \text{(FC2)}.$$

5.5 Compound Nodes

To make verification diagrams more readable, STeP supports encapsulation conventions similar to those of Statecharts [Harel, 1987]. The basic encapsulation construct is the *compound node* containing *internal nodes*. We refer to the contained nodes as the *descendants* of the compound node. Nodes that are not compound are called *basic nodes*.

Departing edges An edge departing from a compound node is interpreted as though it departed from each of its descendants.

Arriving edges An edge arriving at a compound node is interpreted as though it arrived at each of its descendants. An assertion φ labeling a compound node is interpreted as though it were a conjunct added to each of its descendants.

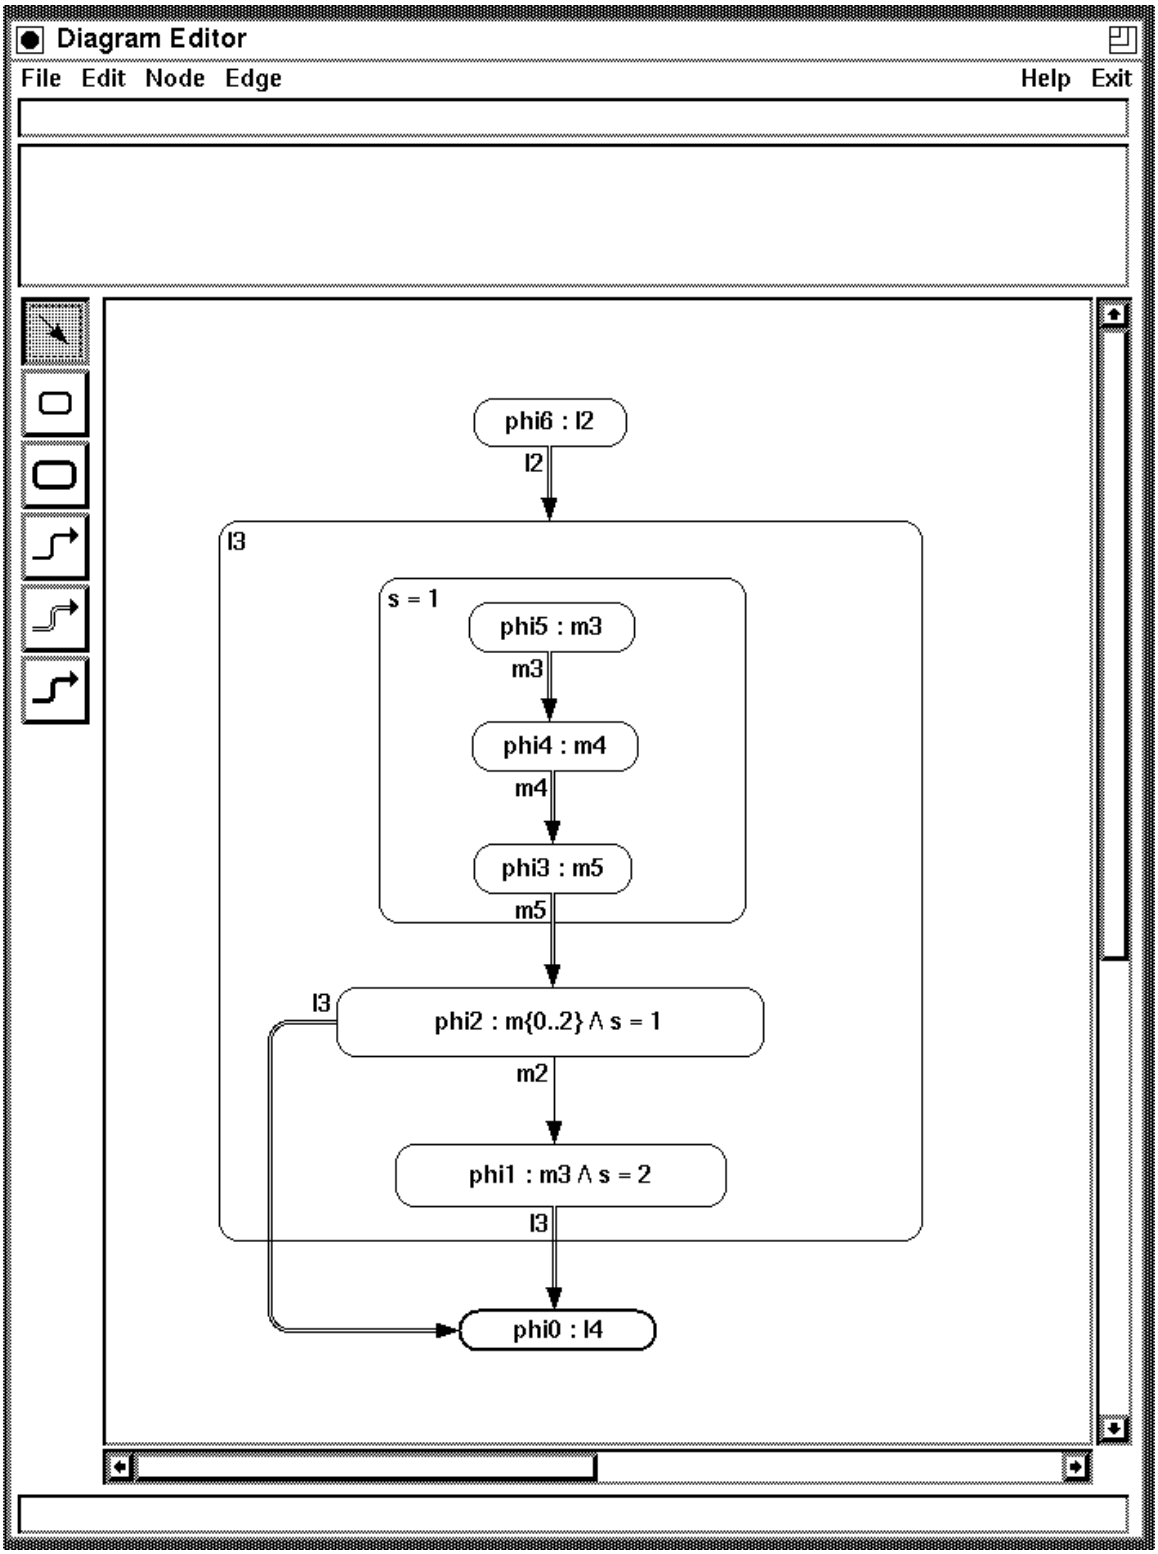

Figure 5.1: The Verification Diagram Editor

5.6 Verification Diagram Editor: Interface

The Verification Diagram Editor allows you to load and manipulate verification diagrams. The interface is shown in Figure 5.1. The two top panels are used to edit the label and the assertions on nodes and edges. The diagram itself is drawn in the large panel, using the drawing tools displayed in the column to its left.

The verification diagram does not have to be related to the current system description, and you can use the Verification Diagram Editor independently of the STeP verification session in progress. By selecting the Current-to-Verify option under the File menu, you make the current diagram the one used by the verification session. In this way, a verification diagram can be applied as a verification rule to generate verification conditions that establish a given property.

The Verification Diagram Editor does not parse the expressions labeling the nodes, or check the existence of the transitions labeling the edges. These checks are performed when the verification diagram rule is invoked. If an assertion does not parse, or there is any other problem with the diagram, a message indicating the error will appear in the message window of the Top-level Prover. The verification diagram can then be corrected within the Verification Diagram Editor, made current again, and the verification diagram rule applied again. [1]

Drawing Tools

The leftmost column contains the tools for drawing basic and compound nodes, and edges corresponding to the different justice requirements.

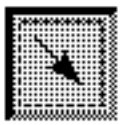

This tool sets the pointer to selection mode. It enables you to select nodes and edges by clicking on them. When an edge is selected, all the nodes to which it is attached are selected too. To select a single edge (e.g., to convert it into a different type of edge) you first have to move the edge away from the nodes. You can do this by pressing the left mouse button while pointing at the edge and then moving it quickly while keeping the button pressed. To select multiple edges and nodes, click on some location outside these nodes, and while keeping the left mouse button pressed, drag the mouse until the nodes and edges are enclosed in the region indicated by the dashed rectangle.

[1] An inquisitive user may wonder why the Verification Diagram Editor does not check for parsing errors. The reason is that to check parsing errors, one requires knowing the system for which a diagram is being applied to. However, STeP allows the Verification Diagram Editor to run independent to any system. This feature is useful when the same (or similar) diagrams uare being applied to different systems.

This tool is used to draw *non-terminal nodes*. After selecting it, move the mouse to the location where you want the upper left corner of the new node and click on the left button; while keeping it pressed, move the mouse down and to the right until the displayed shape has the proper dimensions. Releasing the button will then create the node. The node will be initially selected, and the two editable fields above the diagram will contain the default *name* and *label* for this node. You can edit the name field and enter a state assertion for this node.
To resize an existing node, select it with the left button (after first selecting the pointer tool), move the mouse to one of its corners, and resize the node while keeping the button pressed on the black square in the chosen corner.

Works exactly as the tool for non-terminal nodes, but the generated node is a terminal node.

This tool draws an edge without an associated fairness requirement. An edge is drawn using the left mouse button: press the button at the location where you want the edge to start, keep the button pressed while moving the pointer and release the button at the desired endpoint. To draw an edge with corners, release the button momentarily at the location of the turn, press again, and move the mouse in the new direction. Several consecutive turns can be made in this way.
The label of the edge will appear in the top panel and can be edited there. The label can correspond to any transition of the system being analyzed.

Works exactly as the tool for drawing a single-lined edge, but draws a double-lined edge that has an associated `Justice` requirement. The label of such a Just edge must be the name of a transition declared as Just or Compassionate.

Works exactly as the tool for drawing a single-lined edge, but draws a solid-lined edge that has an associated `Compassion` requirement. The label of a Compassionate edge must be the name of a transition declared as Compassionate.

Menu Options

The File Menu

The **File** menu is used to load and save verification diagrams, as well as to make verification diagrams accessible to the Top-level Prover.

New Resets the verification diagram editor, deleting the currently edited diagram.

Load Loads a verification diagram from a file.

Save Saves the verification diagram currently being edited to a file.

Current-to-verify Selects the currently edited verification diagram as the one used by the Top-Level Prover.

Exit Quits the Verification Diagram Editor after confirmation by the user.

The Edit Menu

The **Edit** menu is used to cut and paste parts of the diagram.

Cut Deletes one or more selected items from the editing window. The deleted item(s) (node or edge) are put on a scratchpad and can later be recovered by Paste.

Copy Copies the selected item from the editing window to the scratchpad.

Paste Takes the item most recently put on the scratchpad (either by Cut or Copy) and puts it in the editing window in the same location it was before.

Duplicate Duplicates one or more selected items. This is equivalent to doing Copy followed by Paste.

The Node Menu

The Node menu is used to group nodes into compound nodes, change terminal into nonterminal nodes, and vice-versa.

Group/Compound By selecting a large node and the nodes enclosed by the large node, Group turns the large node into a compound node, containing the enclosed and selected sub-nodes (see Figure 5.1).

Subnodes of a compound node cannot be edited. They have to be released with Ungroup before you can change their labels or move them.

The label and assertion of the compound (enclosing) node is displayed in the upper left corner of the node.

Ungroup/Uncompound Converts a selected compound node into a regular node, and releases the subnodes such that they can be edited and moved.

Terminal Converts the selected node (only one node should be selected) into a terminal node. The borders of the node are drawn with thick lines and the node is declared as terminal.

Nonterminal Converts a selected terminal node into a nonterminal node.

The Edge Menu

The **Edge** menu lets you change the fairness requirements associated with a selected edge, and check that all edges are properly connected to the nodes.

Cut Head Cuts the head hop of a selected edge. The head hop is the hop with an arrow where an edge ends.

Cut Tail Cuts the tail hop of a selected edge. The tail hop is the hop where an edge starts.

Merge Merges a selected edge with the edges chained to it. Two edges are considered chained if one's ending point is close to the other's starting point.

Unjust Converts a selected edge into an Unjust edge. Unjust edges are drawn with thin black lines.

Just Converts a selected edge into a Just edge. Just edges are drawn with double lines.

Compassionate Converts a selected edge into a Compassionate edge. Compassionate edges are drawn with thick black lines.

Connectivity Highlights those edges whose ends are not connected to nodes. This is a convenient way to check whether all edges are connected properly in the verification diagram.

Global Hookup Checks and hooks up all loose edges. An edge will be hooked up to a node if its end is close enough to or inside the node.

Help

Activation of the on-line help pages, using Mosaic or your favorite `html` (WWW) browser. The Top-level Prover help pages are accessible from the Verification Diagram Editor help pages and vice-versa, so a single help session should generally suffice. The executable used for on-line help can be changed by setting the `STEP_BROWSER` environment variable.

5.7 Hierarchical Verification Diagrams ⋆

In order to allow decomposition of verification diagrams into smaller, more manageable parts, STeP supports hierarchical verification diagrams. With hierarchical diagrams, several nodes in a large diagram can be represented by a single node, making large diagrams more readable and easier to edit.

Import Nodes

A hierarchical verification diagram has one or more import nodes. An *import node* imports a verification diagram from a file. The imported diagram may itself have import nodes that import other files. The edges going out from the import node are matched up with outgoing edges from the imported diagram, and similarly for incoming edges.

An import node is created by giving it the label `#import` and giving it as assertion the filename (with full or relative path name) within angle brackets (`<>`), for example `<mux-pet2/mux-h1.diagram>`, of the file that contains the diagram to be imported. The assertion field may also contain parameters (see next paragraph) enclosed within "%". Any text in the assertion field of an import node that is not enclosed by angle brackets or percent signs is ignored.

Parameter Passing

Hierarchical verification diagrams allow macro-like text replacement when importing a verification diagram. Parameters (in the node to be imported) are of the form "@1", "@2", "@3", etc., and "@1" is replaced by the first argument passed by the importing node, "@2" is replaced by the second argument, etc.

When the top-level diagram is processed (i.e., when it is used as a verification rule by the Top-level Prover), STeP checks whether the number of parameters each imported diagram requires is equal to the number of arguments passed. The arguments passed by the import node are included in the assertion field; each argument is enclosed within percent signs ("%"). An example of the assertion field of an import node that passes parameters is

```
<mux-pet2/mux-h1.diagram> %m3% %l3% %s=1%
```

The arguments do not have to be on the same line, and white space is ignored.

Edges

The connection of edges between the import node and the imported diagram is done by matching unconnected outgoing edges in the imported diagram (also called dangling edges) with outgoing edges of the import nodes that have the same label, and similarly for incoming edges. If multiple edges of the imported diagram match multiple edges of the import node, then edges are created for all possible combinations (taking the product over all edges). Thus, one incoming edge in an imported diagram may translate into multiple edges with the same label entering that same node in the imported diagram.

If multiple edges with the same label entering or leaving an import node need to be distinguished (e.g., because they enter different nodes in the imported diagram), the labels may be provided with a secondary label. The secondary label follows the regular transition name, separated by a colon. Thus, if `l1` is the regular transition name, then `l1:1` labels an edge for transition `l1`, with secondary label `1`. Edges connected to the import node and dangling edges in the imported diagram only match when both the primary and the secondary label are identical. White space is ignored in secondary labels. Secondary labels default to `0` (zero) if left blank, but this value is not displayed in the diagram.

Edges can also be connected to import nodes as if the import node was a compound node (i.e., an incoming edge in the import node is connected to all nodes of the imported diagram, and similarly for an outgoing edge), by adding the secondary label **All** to the label of the edge entering the import node. Thus, an edge labeled with **All** in the importing diagram is connected to all nodes of the imported diagram without any name matching. It is not meaningful to add the secondary label **All** to dangling edges in an imported diagram, and this is therefore ignored.

If import nodes are within a compound node, then edges entering or leaving the compound node are interpreted as entering or leaving each internal node of the import node. Thus, if one wants specific edges to connect with dangling edges in the import nodes one must connect the relevant edges directly to the import nodes.

Chapter 6

Automatic Theorem-Proving

A computer-aided verification system should try to make the task of proving verification conditions as easy as possible while still ensuring soundness.

To this effect, STeP provides powerful, usually fast, automatic simplification procedures. These are supplemented by the interactive Gentzen-style theorem prover, discussed in Chapter 8, used to establish properties that the automatic simplifier cannot.

6.1 The Automatic Simplifier

There is a basic trade-off between the strength and the speed of the automatic Simplifier. The more powerful the settings are, the slower the Simplifier is. The default settings are a good compromise: they are able to simplify most commonly occurring verification conditions in a reasonable time.

The impatient user may want to use the interrupt button to halt lengthy simplification operations and try alternative proof techniques.

Selecting the Simplification Flags menu option brings up the window shown in Figure 6.1. It contains the following settings:

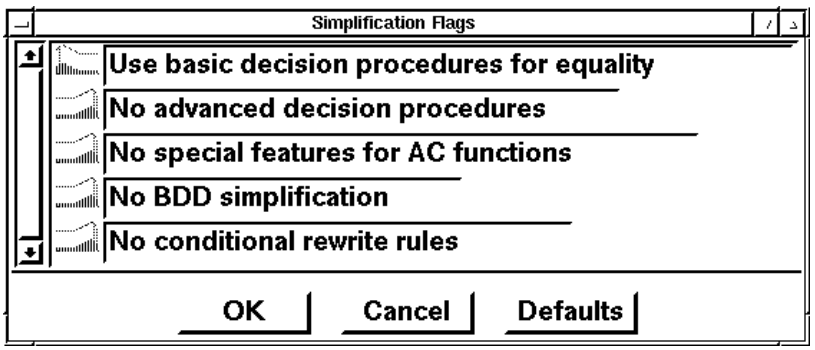

Figure 6.1: Simplifier settings

- **Equality and Inequality**
  If ON, the Simplifier will use properties of equalities and inequalities. For example it

will simplify

$$y > 1 \ \rightarrow \ y > 0$$

to *true* when this switch is ON, but not when it is OFF.

These decision procedures are quite efficient and useful in general, so they are ON by default.

- **Decision Procedures**
  If ON, the Simplifier will use decision procedures for linear arithmetic, which use properties of $>$, $+$, and $=$. It also enables decision procedures of arrays, bit-vectors, lexicographic orderings of tuples, partial orders, and inductive data-types. For example, it will simplify

  $$y \geq 0 \ \wedge \ y' = y + 1 \ \rightarrow \ y' > 0$$

  to *true* when the switch is ON, but will not be able to simplify it when the switch is OFF. The default is OFF.

  The performance of these procedures can vary widely from one formula to the next; the *interrupt button* can be used to halt the computation.

- **Associative and Commutative functions**
  If ON, STeP will take into account the associativity and commutativity of addition, multiplication, and boolean connectives, as well as those user-defined binary function symbols that have been declared as AC, ASSOCIATIVE or COMMUTATIVE (see Section 2.3).

  Note that since STeP rewrites all formulas into a normal form, it implicitly uses many instances of the associative and commutative laws even when this flag is OFF. The default is OFF.

- **BDD simplification**
  If ON, the (propositionally complete) BDD simplification method will be used as part of the simplifier, before other decision procedures are used (see Section 6.2). The default is OFF.

- **Conditional Rewrite Rules**
  If ON, the Simplifier will use a number of built-in conditional rewrite rules. For example, it will simplify

  $$(y > 0 \wedge x \cdot y > 0) \rightarrow x > 0$$

  to *true* when the switch is ON, but will not be able to simplify it when the switch is OFF.

  Conditional rewrite rules may slow down the Simplifier considerably, so they are switched OFF by default.

Commonly used Simplifier settings are also available as separate buttons in the action region of the Top-level Prover:

Fast-simplify Applies only the most basic simplification procedures, disregarding the background properties. This is useful to rewrite formulas into a more standard form, after which they can be compared more easily with each other.

Bool-simplify Applies a propositional satisfiability check, using the active background properties. If the given formula is a propositional tautology with respect to the active background properties, it will be simplified to *true*. This procedure is complete for propositional logic. The basic simplification procedures (see **Fast-simplify** above) are called before the propositional test.

This procedure uses STeP's BDD package (see Section 6.2).

Simplify Invokes the simplifier and decision procedures as given by the current simplification settings.

Check-Valid Invokes the validity checker with support for first-order reasoning and theories where STeP provides decision procedures. It often performs much faster than Simplify, but may occasionally engage in an exhaustive and time-consuming search. The principles behind this button are described in [Bjørner *et al.*, 1997c].

For material related to the decision procedures and indexing used in STeP, see [Bledsoe, 1975], [Shostak, 1979], [Nelson and Oppen, 1980], [Boyer and Moore, 1988], [Stickel, 1989], [McCune, 1992], [Zhang, 1993].

6.2 BDD Simplification

STeP includes a package for the creation and manipulation of propositional (i.e., boolean) formulas using Binary Decision Diagrams (BDDs). This package is used in the BDD-split proof rule as well as in boolean simplification steps.

Since the amount of memory used by the BDD package can grow indefinitely (in the form of accumulated *BDD nodes*), the package can be reset at any time with the Reset BDDs option under the Settings menu. This only incurs a (usually small) short-term efficiency loss, and does not affect the soundness of the prover.

This option should also be used if the number of BDD nodes reaches the maximum given by the `STEP_BDD_NODES` environment variable (see Appendix E). The number of existing BDD nodes can be obtained through the Sytem Information option.

A number of environment variables control the amount of memory that BDDs use:

- `STEP_BDD_NODES` The maximum number of BDD nodes kept around at one time.
- `STEP_BDD_CACHE` The size of the BDD *result cache*.
- `STEP_BDD_SPLIT` The maximum number of subgoals that can be generated by a BDD split operation in both the Top-level and Interactive Provers (see Section 6.2.1 below).

For more information on BDDs, see [Bryant, 1986, Bryant, 1992].

6.2.1 BDD split ⋆

The BDD-split proof rule uses BDDs to convert the current goal into a list of subgoals, each a disjunction of atomic formulas. The subgoals correspond to the negation of the propositional models that can falsify the formula. If all the subgoals are valid, the original goal is valid as well.

In the worst case, exponentially many children are produced using this expansion, but in most applications there is a good chance that it will work well. The acceptable maximum number of subgoals is given by the `STEP_BDD_SPLIT` environment variable.

In the Top-level Prover, the active background formulas are automatically used to reduce the number of subgoals generated; in the Interactive Prover, the background formulas must be added to the sequent before invoking BDD-split. This is analogous to the behavior of the Simplifier in the two provers.

6.3 Tactics ⋆

Tactics are a method to automate parts of the high-level proof search by encoding long or repetitive sequences of proof commands. An individual proof command is an *atomic tactic*.

The Top-level Prover and the Interactive Prover share the same tactic language, parameterized by their atomic tactics (since each has a different set of proof commands). The general syntax of a tactic is as follows:

tactic	::=	*predefined*	Predefined tactic
	|	*composite*	Composite tactic
	|	*atomic*	Atomic tactic
predefined	::=	`skip`	Identity tactic
	|	`fail`	Failure tactic
composite	::=	`try` (*tactic* , *tactic* , *tactic*)	
	|	`then` *[form] tactics*	
	|	`sequence` *tactics*	
	|	`else` *tactics*	
	|	`repeat` *[form] tactic*	
	|	`spread` *tactic tactics*	
	|	`branch` *tactic tactics* [`otherwise` *tactic*]	
	|	*tactic* ; *tactic*	
	|	(*tactic*)	
tactics	::=	`[` [*tactic* { , *tactic* }] `]`	
form	::=	`depth first`	
	|	`breadth first`	
atomic	::=	`{` *atomic-tactic* `}`	

The grammar for *atomic-tactic* for the Top-level Prover is given in Section 6.3.2; the grammar for the Interactive Prover atomic tactics is given in Section 6.3.3.

The curly brackets can be omitted if the atomic tactic can be written without spaces. In other cases, curly brackets have to enclose each atomic tactic.

6.3.1 Composition Options

When a verification rule is applied to a goal to which it does not syntactically make sense (for example B-INV applied to a response formula) it *fails*, otherwise it *succeeds*. Tactics, which are composed from verification rules can also fail and succeed as described below.

try(*tac1*,*tac2*,*tac3*)

Try tactic *tac1*, if it succeeds, apply *tac2* to each subgoal generated. If *tac1* fails, apply tactic *tac3* to the original goal. The entire tactic fails if *tac2* fails on one of the generated subgoals of *tac1*, or *tac3* fails.

Figure 6.2 gives an idea of how the proof-search looks like after *tac1* and *tac2* have been applied successfully, or *tac1* failed and *tac3* has been applied. The lower-most dot represents the original goal, the first left-most layer of dots represent the subgoals generated from *tac1*, the second left-most layer, the subgoals generated from *tac2*. To the right is a schematic representation of the subgoals generated by *tac3*.

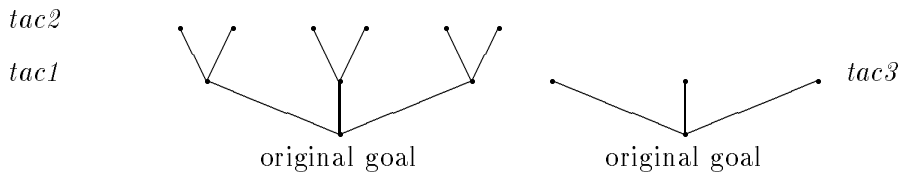

Figure 6.2: `try`(*tac1*,*tac2*,*tac3*)

then *[form] tactics*

`then []` has the same effect as `skip` - the original subgoal is reduced to itself. For non-empty lists we define inductively:

`then breadth first [` tac_1,tac_2,...,tac_n`]` applies tactic tac_1 first. If it succeeds, it applies the tactic `then breadth first [` tac_2,...,tac_n`]` to *each* of the generated subgoals. If it fails it applies the tactic `then breadth first [` tac_2,...,tac_n`]` to the original subgoal.

Figure 6.3 illustrates the levels obtained from successful completion of `then breadth first` [*tac1*,*tac2*,*tac3*]. So for instance, the nodes above the original goal in the bottom represent the result of applying *tac1*.

`then depth first [` tac_1,tac_2,...,tac_n`]` applies the tactic tac_1 first. If it succeeds, it applies `then depth first [` tac_2,...,tac_n`]` to the first subgoal generated. If it fails, it applies `then depth first [` tac_2,...,tac_n`]` to the original goal. If any left-most subgoal is closed by a tactic tac_i, where $i < n$, then $tac_{i+1} \ldots tac_n$ are ignored.

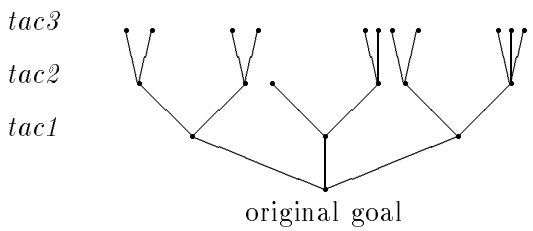

Figure 6.3: `then breadth first` [*tac1,tac2,tac3*]

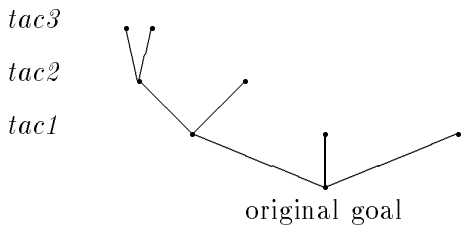

Figure 6.4: `then depth first` [*tac1,tac2,tac3*]

Figure 6.4 illustrates the effect of successful completion of `then depth first` [*tac1,tac2,tac3*].

If no *form* is given the default is `breadth first`.

sequence *tactics*

Applies the tactics in left-to-right order breadth first until the first tactic fails. The constructor `sequence` can be defined in terms of `try`, namely `sequence` [tac_1, tac_2, tac_3] has the same effect as `try`(tac_1, `try`(tac_2, tac_3,`fail`), `fail`).

tactic ; *tactic*

tac_1 ; tac_2 is shorthand for `then [` tac_1, tac_2 `]`. The semicolon ; associates to the left and has the highest precedence.

else *tactics*

Applies the first tactic that succeeds from the list of tactics supplied.

Figure 6.5 illustrates the result of `else` [*tac1,tac2,...,*tac_i*,...,*tac_n], where tac_i is the first tactic that succeeds.

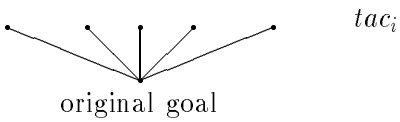

Figure 6.5: `else` [*tac1,tac2,...,*tac_i*,...,*tac_n]

repeat *[form] tactic*

`repeat breadth first` *tac* applies *tac* repeatedly until failure appears on one of the generated subgoals in a layer.

`repeat depth first` *tac* applies *tac* repeatedly on the left-most subgoal generated.

spread *tactic tactics*

`spread` *tac* `[` $tac_1, tac_2, \ldots, tac_n$ `]` applies first *tac*, and then applies tac_i to the i'th subgoal generated. The tactic fails if the number of subgoals generated from *tac* is different from n. Also if any of the tactics *tac* or tac_i fail, the entire tactic fails.

branch *tactic tactics [* **otherwise** *tactic]*

`branch` *tac* $[tac_1, \ldots, tac_n]$ works like `spread`, except that if the number of subgoals is $m \geq n$, then it applies tactic tac_n to the subgoals indexed $n, n+1, \ldots, m$.

`branch` *tac* $[tac_1, \ldots, tac_n]$ `otherwise` *tac'* first applies the tactic `branch` *tac* $[tac_1, \ldots, tac_n]$. If this fails, it applies *tac'* to the subgoal.

6.3.2 Top-level Prover Tactics

The atomic tactics for the Top-level Prover mirror the names of the verification rules available there. The identifiers are identical to the names of the rules, whereas rules requireing auxiliary expressions take a list of expressions enclosed in quotes.

atomic-tactic	::=	`B-INV`
	|	`B-WAIT`
	|	`B-BACKTO`
	|	`B-CAUS`
	|	`B-SAFE`
	|	`Simplify`
	|	`Check-Valid`
	|	`Undo`
	|	`Redo`
	|	`Postpone`
	|	`BDD`
	|	`Interactive tactic[` *interactive-tactic*`]`
	|	`Verification-Diagram diagram[` *diagram*`]`
	|	`G-INV` *expns*
	|	`G-WAIT` *expns*
	|	`G-BACKTO` *expns*
	|	`G-CAUS` *expns*
	|	`G-SAFE` *expns*
	|	`Modelcheck` *expns*
expns	::=	{ ``*expn*'' }+
interactive-tactic	::=	tactic whose atomi components are from Section 6.3.3
diagram	::=	textual format of verification diagrams

As a special feature the Top-level accepts not only one tactic for the current proof search, but a list of tactics, i.e., the correct input format for the Top-level Prover is:

$$top\text{-}level\text{-}input \quad ::= \quad \{tactic\}+$$

Given input tac_1 tac_2 ... tac_n and open searches $s_1, \ldots, s_m$, $m \geq n$, STeP will apply tac_1 to s_1, tac_2 to s_2 et.c.. Hence, when establishing 3 simple invariants one inputs the tactic:

```
B-INV;Simplify
B-INV;Simplify
B-INV;Simplify.
```

6.3.3 Interactive Prover Tactics

To appreciate this section it is helpful to have read Chapter 8.

atomic-tactic	::=	*basic-rule [positions]*	Basic rule not requiring arguments
	|	*e-rule (expns | expn-poss)*	Rule taking expression and optional position argument
	|	*p-rule [positions]*	Rule taking optional position argument
	|	*replace*	Replacement of equals for equals
	|	*skolemize*	Skolemization
basic-rule	::=	`Undo`	Basic rules behave the
	|	`Redo`	same as their
	|	`Postpone`	pushbutton counterparts
	|	`Flatten`	
	|	`Simplify`	
	|	`Check-Valid`	
	|	`Make-first-order`	
	|	`Next`	
	|	`Presburger`	
	|	`All-Propositional`	
	|	`BDD-split`	
	|	`PTL-expansion`	
e-rule	::=	`Cut`	adds argument expressions as cuts
	|	`Add-Axiom`	adds argument expressions as axioms
	|	`Rewrite`	arguments is a tuple of pairs each pair represents a rewrite rule
	|	`Free-Induction`	arguments are the integer variables in the induction
	|	`Instantiate`	Instantiate existential force quantifiers at indicated positions by the listed expressions
p-rule	::=	`Duplicate`	duplicates formulas at

		indicated positions
	`Delete`	duplicates formulas at indicated positions
	`Hide`	hides formulas at indicated positions
	`Unhide`	unhides formulas located at their hidden positions.
	`1-Step-Propositional`	
	`Induction`	
	`Sum-Monotonicity`	
	`Let-reduction`	

replace ::= `Replace at` *i-val* `left by right`
| `Replace at` *i-val* `right by left`
In both cases *i-val* must be < 0.

skolemize ::= `Skolemize`
| `Skolemize "` *expn* `" at` *position*
| `Skolemize` *position*

expns ::= { `"` *expn* `"` }+
positions ::= { *position* }+
position ::= *i-val*
| `[` *i-val* { `,` *i-val* } `]`
expn-poss ::= { `"` *expn* `" at` *position* }+

The following conventions apply:

1. Optional positions can be supplied to a *basic-rule*. This has only effect for rule `Simplify`, which will simplify the indicated subexpression.

2. If an *e-rule* is supplied only with an expression, and not a position, it will apply the rule exhaustively to all the positions where it applies. For instance { `Instantiate "0"` } instantiates `0` at all existential force quantifiers.

3. If an *e-rule* is supplied with both an expression and position it will apply the rule only at the given position. For instance { `Rewrite "(1,(x+0,x))" 2`} rewrites only the formula at index 2.

4. If a *p-rule* is given without arguments, it is applied exhaustively in a breadth-first manner.

6.3.4 Interrupting Tactics

Tactics are not guaranteed to terminate, or can otherwise run for a very long time; the *interrupt button* (the STeP logo) is useful to stop the execution of a batchmode tactic.

Two things can actually happen if the interrupt button is pressed during the execution of a tactic: if the interrupt button is activated during a simplification step, the simplification will be canceled and the system will continue with the next one; if the system is interrupted when no simplification is going on, the entire tactic will be interrupted.

6.3.5 Commonly used Tactics

B-INV; Simplify

Applies the rule B-INV followed by simplification of each subgoal. This differs from the effect of setting Automatic Simplification ON by not backtracking to the unsimplified verification condition when simplification fails.

B-INV; Simplify; Undo

If simplification does not establish the given subgoal of B-INV; the current leaf in the proof search tree is undone, getting back to the unsimplified subgoal.

repeat(B-INV; else[Check-Valid, skip ; WPC)]

Repeats application of B-INV to each unestablished subgoal that has been strengthened by the WPC rule. Check-Valid fails, when it does not establish a goal, the `else` construct then continues with the next tactic in its argument list, this is `skip`.

{ G-INV "`(l5 /\ m5) /\ (l4 /\ m5 --> s = 2 /\ y1)`" }; Simplify

Applies rule G-INV with a strengthening formula, and each subgoal is then simplified. G-INV only requires one strengthening formula; rules such as G-WAIT require more, so a list of expressions must be entered following the keyword "G-WAIT".

Skolemize; Simplify; BDD-split; Simplify

Called from the Interactive Prover, this tactic attempts to establish a subgoal by first skolemizing quantifiers of universal force, and then simplifying in the hope that simplification will reduce the boolean complexity of the verification condition. BDD-split is then invoked to split the goal into subgoals of atomic formulas; the simplifier is finally invoked to establish each of these subgoals.

1-Step-Propositional; Simplify

The Interactive Prover 1-Step-Propositional rule is used without a position. It is therefore used exhaustively until it cannot be applied. Then each generated subgoal is simplified by Simplify.

Chapter 7

Automatic Generation of Invariants

STeP includes three different methods that automatically generate *bottom-up invariants*, based solely on the system to be verified, of increasing power and complexity. Since these invariants are, by construction, guaranteed to be true of the system, they are automatically added to the list of active background properties.

Of the three methods, the first, *local invariants*, operates on an SPL program; the other two operate on the transition system directly. A description of the general principles behind STeP's invariant generation algorithms is [Bjørner *et al.*, 1997a].

To obtain invariants, select the appropriate option under the Properties menu; when the invariant generation is done, each invariant is simplified, displayed in the output window, and added to the set of background properties. By default they are active, and thus used in the simplification of goals. They may be deactivated with the Activate/Deactivate menu option.

Each invariant is simplified on its own (i.e., no background properties are used), using the basic decision procedures for equality (see Section 6.1).

When the invariant generator is finished a message will appear in the message window. Invariant generation and simplification can be halted by clicking on the interrupt button (the STeP logo).

7.1 Local Invariants

Local invariant generation produces very simple but useful invariants, based on local analysis of the program text. The generated invariants can be classified in the following three groups:

1. *Finite-domain invariants.* Invariants of the form

$$l_i \Rightarrow x = c_1 \vee x = c_2 \vee \ldots \vee x = c_n$$

   where l_i is a program location, $c_1, .., c_n$ are constants, and x is a program variable.

2. *Conditional invariants.* Invariants of the form $l_i \Rightarrow cond$, where l_i is a program location and *cond* is a formula arising from the initial condition or a branch statement,

e.g., the program fragment `if cond then l1: S1 else l2: S2` generates the conditional invariants `l1` $\Rightarrow$ `cond`, and `l2` $\Rightarrow \neg$`cond`, if no variables in `cond` are modified by parallel statements.

3. *Range invariants.* Invariants of the form $l_i \Rightarrow lb \leq x \wedge x \leq ub$, where lb and ub are integer or rational constants (or infinity).

7.2 Linear Invariants

Linear invariant generation arises from analyzing the linear relationships between system variables in the course of a program.

The invariants generated are usually linear equalities of the form

$$a_1 x_1 + a_2 x_2 + \ldots + a_n x_n = b$$

where $a_1, \ldots, a_n, b$ are constants. Each x_i can be (a) a variable, (b) the sum `Sum i:[m..M]. A[i]`, of the values of an array, or (c) a sum `Sum i:[m..M]. x[i]`, where `x` is a local variable of a parameterized process `P[i:[m..M]]`.

7.3 Polyhedral Invariants $\star$

The polyhedral invariant generator is the most powerful of the methods provided (currently in "experimental" state). The invariants generated depend on the set of variables under consideration, which can be controlled by the user. The more variables that are included, the more powerful the invariants that can be generated, but the time spent generating invariants will also increase (possibly in an exponential manner).

When the polyhedral invariant generation is selected, a dialog window appears where the user can enter the following:

1. *Variables to be ignored:* you may enter a list of program and auxiliary variables, separated by spaces or commas. These variables will be *ignored* by the invariant generator, that is, the invariant generation will not try to deduce polyhedral relationships which include them.

2. *Location variables to be considered:* By default, location variables are not considered by the polyhedral invariant generator. Here you may enter a list of location variables to be added to the vocabulary of the invariant generator.

The generator uses an external package, whose time and space limitations can be controlled by setting the following environment variables:

`STEP_POLY_SPACE` (Megabytes)
`STEP_POLY_TIME` (minutes)
`STEP_POLY_CPU` (minutes)

The invariant generation can be interrupted by clicking on the interrupt button (the STeP logo); this may leave a separate UNIX process running, which users may have to terminate independently from a UNIX prompt (see the UNIX `kill` command man pages).

Chapter 8

The Interactive Prover ⋆

The Interactive Prover is based on a *Gentzen-style proof system*, specialized to linear-time temporal logic. For more background, see, e.g., [Gallier, 1987]. This chapter explains the basic functionality of the various rules provided in the Interactive Prover and gives examples of their use.

8.1 Proof Structure

A goal or subgoal in the Interactive Prover is a *sequent*, of the form

$$\Gamma \Longrightarrow \Delta$$

where Γ and Δ are *multisets* of formulas (i.e., formulas can be repeated). A sequent stands for the formula

$$\bigwedge \gamma_i \rightarrow \bigvee \delta_j$$

where γ_i are the elements of Γ and δ_j are the elements of Δ. Free variables are implicitly universally quantified.

The Interactive Prover creates a *proof tree* similar to that in the Top-level Prover. An inference rule applies to zero or more of the formulas in a sequent $\Gamma \Longrightarrow \Delta$, and produces a number of subgoals, the branches of the tree, again in the form of sequents. The proof is finished when all subgoals are closed.

A subgoal is closed when one of the following conditions is satisfied:

- The sequent contains `true` on the right-hand side of the long doubly lined implication.

$$\Gamma \Longrightarrow \Delta, \texttt{true}$$

- The sequent contains `false` on the left-hand side of the long doubly lined implication:

$$\Gamma, \texttt{false} \Longrightarrow \Delta$$

- The sequent contains a formula `p` in both the left and right sides of the long doubly lined implication:

$$\texttt{p}, \Gamma \Longrightarrow \Delta, \texttt{p}$$

- The sequent contains a formula `(-)p` on the left side of the long doubly lined implication, and no applications of Next (see Section 8.3) separate the sequent from the root of the proof tree:

$$(-)\mathtt{p}, \Gamma \Longrightarrow \Delta$$

  This axiom reflects the well-foundedness of time. In the first time instance, the truth-value of `(-)p` is false for any formula `p`, because there is no past at this instance.

- The sequent $\Gamma \Longrightarrow \Delta$ is *duplicated* along a branch in the proof tree, i.e., the following conditions are satisfied:

  1. There is a sequent $\Gamma' \Longrightarrow \Delta'$ between the root and $\Gamma \Longrightarrow \Delta$, such that $\Gamma' \subseteq \Gamma$ and $\Delta' \subseteq \Delta$.
  2. The two sequents are separated by at least one application of the **Next**-rule.
  3. There is an *unfulfilled* formula, in the sense of the temporal tableau construction [Manna and Pnueli, 1995b], for the sequents below applications of Next and above $\Gamma' \Longrightarrow \Delta'$.

  It helps to think of duplication in the presence of unfulfilled formulas as corresponding to an induction step. The requirement also replaces inference rules formalizing induction on temporal operators.

Fortunately all these conditions are kept track of automatically and you do not have to worry about the detailed conditions, so that you can focus on the outline of the proof.

8.2 User Interface

Figure 8.1 shows the interface of the Interactive Prover. It has a menu bar on top and action buttons to the right. Most interaction in the Interactive Prover is accomplished by button clicks. If a rule is applicable to multiple formulas in the sequent, the candidates are highlighted (shown in reverse video) and you can select the one to which you intend to apply the rule. Some rules require additional information to be entered through dialog windows.

Menu Options

Refresh redraws the Interactive Prover windows.

Tactics is equivalent to the Tactics menu in the Top-level Prover (Section 3.1) except now it invokes tactics composed of Interactive Prover primitive commands (see Section 6.3).

Settings

Simplification Flags This menu option is equivalent to the Simplification Flags option in the Top-level Prover. It brings up a window used to set the power of the automatic Simplifier; see Section 6.1.

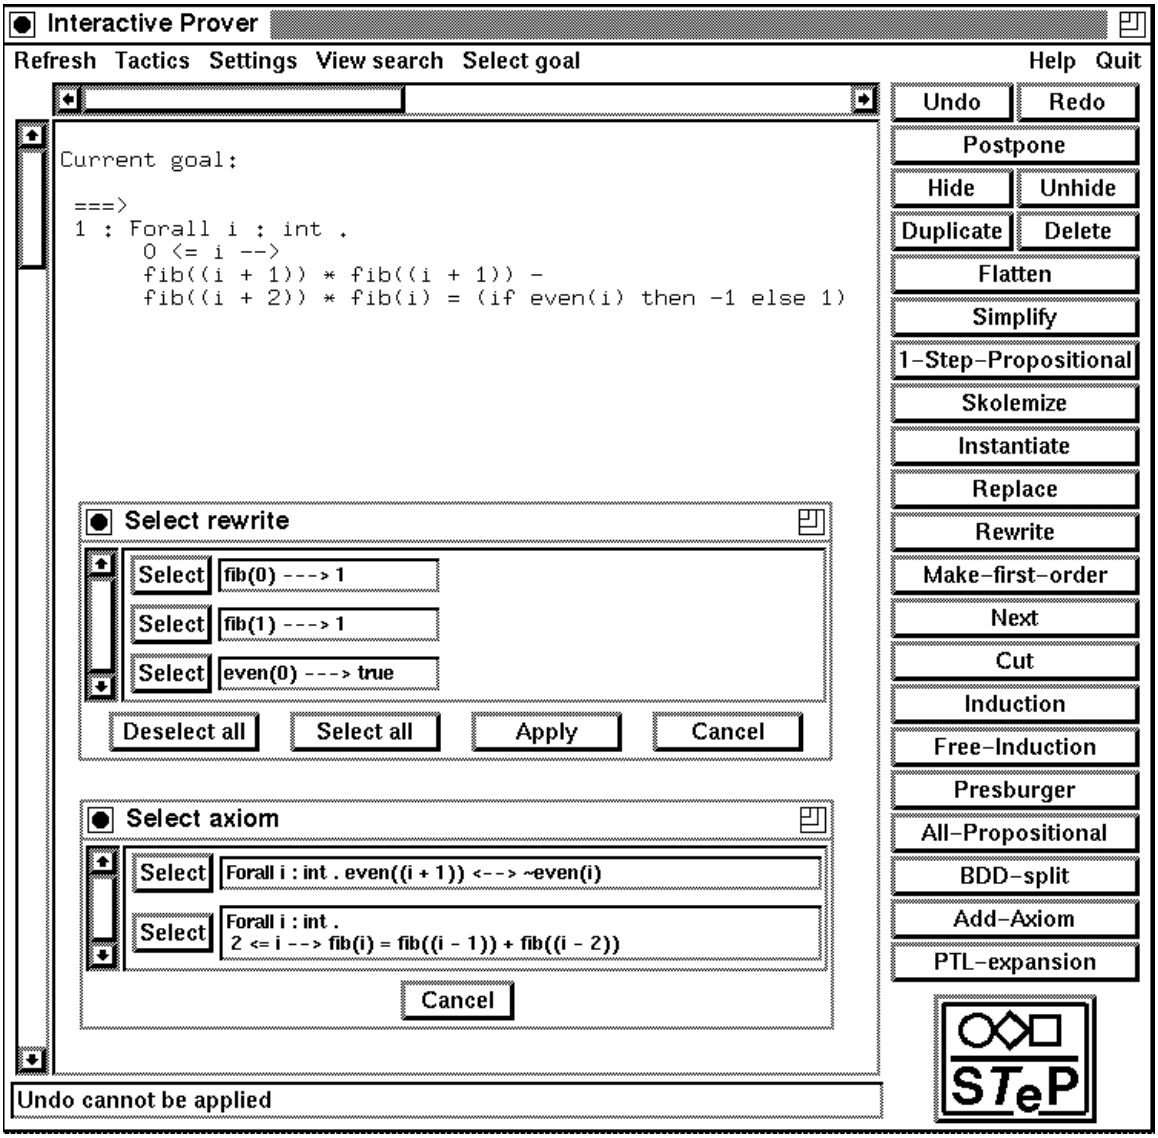

Figure 8.1: Interactive Prover Window

Select boxes This menu option allows the user to control whether boxes, i.e., the temporal operator □ known as "henceforth", in front of axioms added to the sequent during the proof should be kept or not. By default boxes are removed from axioms, when axioms are added to the sequent. See also the explanation of Add-Axiom below for an example of the usage of Select boxes.

View search Gentzen style proofs are well visualized as rooted trees. The menu option View search allows the user to navigate in the so-far developed proof search via a graphical interface of the proof tree. The current goal is drawn with a white node, other pending goals, and goals that have been backtracked from are grey, whereas intermediary nodes are black. By selecting nodes in the graphical interface one can change the current goal. The graphical display gets continually updated when navigating using the conventional verification rules.

Figure 8.2 illustrates how the graphical tool displays a proof tree with three pending goals, one white and two grey.

Select goal Brings up a dialog box containing a list of the pending goals. The user can use this utility to inspect the list of open goals as well as select a goal different from the current.

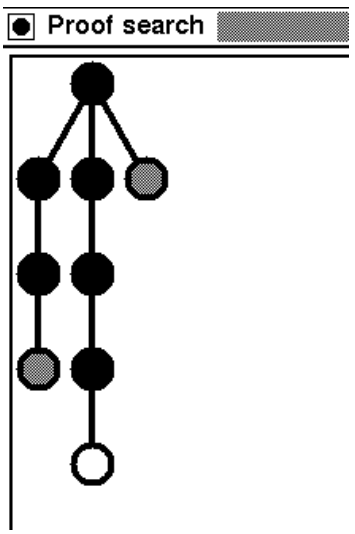

Figure 8.2: Proof search interface

Quit brings up the window shown in Figure 8.3.

If you choose to use the results obtained, any unfinished goals are returned to the Top-level Prover as a new goal with the name `Interactive`. If there are no goals left, the current goal in the Top-level Prover is closed. If you choose to ignore the effect of the Interactive Prover the proof tree in the Top-level Prover stays as it was before invoking the Interactive Prover.

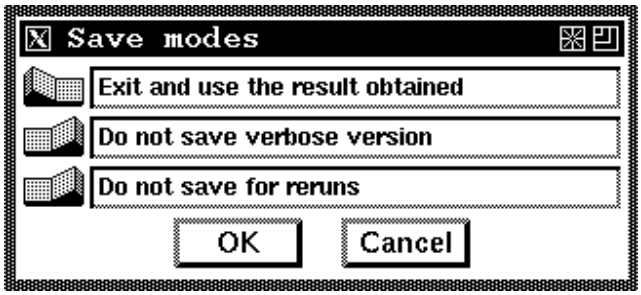

Figure 8.3: Exit Window for Interactive Prover

The button saving *verbose* versions allows you to save the entire proof in the format seen during verification. Each sequent is printed in its full length, and annotated with its position in the proof-search.

The button saving for reruns, saves a dense proof-script, which can be loaded as a tactic for later automatic recreation of the proof.

8.3 Interactive Prover Rules

The right column of the Interactive Prover contains buttons to manipulate the proof tree and invoke the Gentzen proof rules. The proof search proceeds as usual, with logical rules that reduce the current goal to a number of subgoals, and simplification steps that close or simplify the current goal.

A number of inference rules require the user to select subformulas in the sequent. The subformulas, that may be selected are highlighted when selecting the inference rule, and the user clicks on the relevant highlighted formula to invoke the rule on that formula.

Undo Makes the ancestor of the current subgoal into the current goal. The former current goal is retained to enable a Redo, unless another proof rule is applied to the new current goal.

Redo Redoes a previous Undo.

Hide Temporarily hides a formula from the current sequent. Hidden formulas do not participate in simplification steps, and cannot be acted upon by inference rules. Hidden formulas, however, can be recovered with the Unhide rule, below. Hiding formulas can make the sequent more readable and may speed up simplification.

After clicking this button all formulas are highlighted, and you can select the formula to hide. To hide several formulas you must use Hide multiple times.

Unhide Recovers a formula hidden with Hide.

Postpone The next subgoal in the subgoal queue becomes the current subgoal.

Duplicate Duplicates a formula in the sequent. This is useful if you need to instantiate a quantified variable more than once, using different values. Note that duplicated formulas can be hidden with Hide, above, until they are needed.

Delete This button implements the *weakening rule*. It allows you to delete formulas that are not needed to establish the goal. Deletion of unnecessary or redundant formulas makes the current goal easier to read and speeds up the simplification procedures.

Flatten Creates a new sequent in which all conjunctions in the antecedent and all disjunctions in the consequent are split into separate formulas.

Example: For formulas `a, b, c, d, e, f, g`, Flatten has the following effect:

$$\frac{\texttt{a, b, b \/ c} \implies \texttt{d /\ e, f, g}}{\texttt{a /\ b, b \/ c} \implies \texttt{d /\ e, f \/ g}}$$

Simplify Invokes the Simplifier on the current goal. The strength of the Simplifier is controlled via the Simplification Flags option under the Settings menu.

Check-Valid Invokes the validity checking on the current goal.

1-Step-Propositional Application of propositional inference rules is controlled uniformly with this button.

After clicking this button, formulas which contain propositional connectives that can be analyzed are highlighted. Clicking on one of the highlighted formulas applies the appropriate propositional inference rule to the selected formula.

The boolean connectives analyzed are:

$$\wedge\ ,\ \vee\ ,\ \rightarrow, \leftrightarrow, \neg, \textbf{if-then-else},$$

which are entered as:

`/\ \/, -->, <-->, !, if-then-else`

The temporal connectives analyzed are:

$$\Box, \Diamond, \mathcal{U}\ ,\ \mathcal{W}\ ,\ \boxminus, \diamondminus, \mathcal{S}\ ,\ \mathcal{B}$$

which are entered as:

`[], <>, Until, Awaits, [-], <->, Since, Backto`

Notice that $\bigcirc$ and $\ominus$ are not in this group.

Figure 8.3 shows all the rules that can be applied by 1-Step-Propositional.

Propositional Rules

Sequents and formulas in square brackets [] are only generated when the sequent below the line is above an application of the rule *Next*. This guarantees that statements such as $\Box\boxminus(\mathit{first})$ are not derivable, as opposed to $\boxminus\mathit{first}$.
(*first* is shorthand for $\neg\ominus\mathtt{true}$, which only holds at position 0.)

$$(\Rightarrow\neg)\quad\frac{\Gamma,\varphi\Rightarrow\Delta}{\Gamma\Rightarrow\Delta,\neg\varphi}\qquad(\neg\Rightarrow)\quad\frac{\Gamma\Rightarrow\Delta,\varphi}{\Gamma,\neg\varphi\Rightarrow\Delta}$$

$$(\Rightarrow\vee)\quad\frac{\Gamma\Rightarrow\Delta,\varphi,\psi}{\Gamma\Rightarrow\Delta,\varphi\vee\psi}\qquad(\vee\Rightarrow)\quad\frac{\Gamma,\psi\Rightarrow\Delta\quad\Gamma,\varphi\Rightarrow\Delta}{\Gamma,\varphi\vee\psi\Rightarrow\Delta}$$

$$(\Rightarrow\rightarrow)\quad\frac{\Gamma,\varphi\Rightarrow\Delta,\psi}{\Gamma\Rightarrow\Delta,\varphi\rightarrow\psi}\qquad(\rightarrow\Rightarrow)\quad\frac{\Gamma,\psi\Rightarrow\Delta\quad\Gamma\Rightarrow\Delta,\varphi}{\Gamma,\varphi\rightarrow\psi\Rightarrow\Delta}$$

$$(\Rightarrow\leftrightarrow)\quad\frac{\Gamma,\varphi\Rightarrow\Delta,\psi\quad\Gamma,\psi\Rightarrow\Delta,\varphi}{\Gamma\Rightarrow\Delta,\varphi\leftrightarrow\psi}\qquad(\leftrightarrow\Rightarrow)\quad\frac{\Gamma,\psi,\varphi\Rightarrow\Delta\quad\Gamma\Rightarrow\Delta,\psi,\varphi}{\Gamma,\varphi\leftrightarrow\psi\Rightarrow\Delta}$$

$$(\Rightarrow\wedge)\quad\frac{\Gamma\Rightarrow\Delta,\psi\quad\Gamma\Rightarrow\Delta,\varphi}{\Gamma\Rightarrow\Delta,\varphi\wedge\psi}\qquad(\wedge\Rightarrow)\quad\frac{\Gamma,\varphi,\psi\Rightarrow\Delta}{\Gamma\Rightarrow\Delta,\varphi\wedge\psi}$$

$$(\Rightarrow\Box)\quad\frac{\Gamma\Rightarrow\Delta,\varphi\quad\varphi,\Gamma\Rightarrow\Delta,\bigcirc\Box\varphi}{\Gamma\Rightarrow\Delta,\Box\varphi}\qquad(\Box\Rightarrow)\quad\frac{\Gamma,\varphi,\bigcirc\Box\varphi\Rightarrow\Delta}{\Gamma,\Box\varphi\Rightarrow\Delta}$$

$$(\Rightarrow\Diamond)\quad\frac{\Gamma\Rightarrow\Delta,\varphi,\bigcirc\Diamond\varphi}{\Gamma\Rightarrow\Delta,\Diamond\varphi}\qquad(\Diamond\Rightarrow)\quad\frac{\Gamma,\varphi\Rightarrow\Delta\quad\bigcirc\Diamond\varphi,\Gamma\Rightarrow\Delta,\varphi}{\Gamma,\Diamond\varphi\Rightarrow\Delta}$$

$$(\Rightarrow\mathcal{W})\quad\frac{\Gamma,\varphi,\bigcirc(\varphi\,\mathcal{W}\,\psi)\Rightarrow\Delta\quad\Gamma,\psi\Rightarrow\Delta}{\Gamma,\varphi\,\mathcal{W}\,\psi\Rightarrow\Delta}\qquad(\mathcal{W}\Rightarrow)\quad\frac{\Gamma\Rightarrow\Delta,\varphi,\psi\quad\varphi,\Gamma\Rightarrow\Delta,\bigcirc(\varphi\,\mathcal{W}\,\psi),\psi}{\Gamma\Rightarrow\Delta,\varphi\,\mathcal{W}\,\psi}$$

$$(\Rightarrow\mathcal{U})\quad\frac{\Gamma,\varphi,\bigcirc(\varphi\,\mathcal{U}\,\psi)\Rightarrow\Delta,\psi\quad\Gamma,\psi\Rightarrow\Delta}{\Gamma,\varphi\,\mathcal{U}\,\psi\Rightarrow\Delta}\qquad(\mathcal{U}\Rightarrow)\quad\frac{\Gamma\Rightarrow\Delta,\varphi,\psi\quad\Gamma\Rightarrow\Delta,\bigcirc(\varphi\,\mathcal{U}\,\psi),\psi}{\Gamma\Rightarrow\Delta,\varphi\,\mathcal{U}\,\psi}$$

$$(\Rightarrow\boxminus)\quad\frac{\Gamma\Rightarrow\Delta,\varphi\quad[\varphi,\Gamma\Rightarrow\Delta,\ominus\boxminus\varphi]}{\Gamma\Rightarrow\Delta,\boxminus\varphi}\qquad(\boxminus\Rightarrow)\quad\frac{\Gamma,\varphi[,\ominus\boxminus\varphi]\Rightarrow\Delta}{\Gamma,\boxminus\varphi\Rightarrow\Delta}$$

$$(\Rightarrow\diamondminus)\quad\frac{\Gamma\Rightarrow\Delta,\varphi[,\ominus\diamondminus\varphi]}{\Gamma\Rightarrow\Delta,\diamondminus\varphi}\qquad(\diamondminus\Rightarrow)\quad\frac{\Gamma,\varphi\Rightarrow\Delta\quad[\ominus\diamondminus\varphi,\Gamma\Rightarrow\Delta,\varphi]}{\Gamma,\diamondminus\varphi\Rightarrow\Delta}$$

$$(\Rightarrow\mathcal{B})\quad\frac{\Gamma\Rightarrow\Delta,\varphi,\psi\quad[\varphi,\Gamma\Rightarrow\Delta,\ominus(\varphi\,\mathcal{B}\,\psi),\psi]}{\Gamma\Rightarrow\Delta,\varphi\,\mathcal{B}\,\psi}\qquad(\mathcal{B}\Rightarrow)\quad\frac{\Gamma,\varphi[,\ominus(\varphi\,\mathcal{B}\,\psi)]\Rightarrow\Delta\quad\Gamma,\psi\Rightarrow\Delta}{\Gamma,\varphi\,\mathcal{B}\,\psi\Rightarrow\Delta}$$

$$(\Rightarrow\mathcal{S})\quad\frac{[\Gamma\Rightarrow\Delta,\varphi,\psi]\quad\Gamma\Rightarrow\Delta[,\ominus(\varphi\,\mathcal{S}\,\psi)],\psi}{\Gamma\Rightarrow\Delta,\varphi\,\mathcal{S}\,\psi}\qquad(\mathcal{S}\Rightarrow)\quad\frac{[\Gamma,\varphi,\ominus(\varphi\,\mathcal{S}\,\psi)\Rightarrow\Delta,\psi]\quad\Gamma,\psi\Rightarrow\Delta}{\Gamma,\varphi\,\mathcal{S}\,\psi\Rightarrow\Delta}$$

Skolemize Quantifiers of universal force can be eliminated by skolemization, which replaces the bound variable by a new variable.

If there are multiple instances of universally quantified variables all of the corresponding formulas are highlighted upon clicking on this button, and you can select the one you want to skolemize by clicking on it.

Skolemize implements the inference rule:

$$\frac{\Gamma,\ \texttt{f(b)} \quad \Longrightarrow \quad \Delta}{\Gamma,\ \texttt{Exists x . f(x)} \quad \Longrightarrow \quad \Delta}$$

where `b` is a fresh variable not occurring free in Γ, Δ or `Exists x . f(x)`.

$$\frac{\Gamma \quad \Longrightarrow \quad \Delta,\ \texttt{f(b)}}{\Gamma \quad \Longrightarrow \quad \Delta,\ \texttt{Forall x . f(x)}}$$

where again `b` is not free in either Γ, Δ or `Forall x . f(x)`.

Besides these cases, the Interactive Prover supports skolemization of quantifiers under any nesting of the boolean connectives `/\, \/, -->, ~`, and the `then` and `else` branches in `if-then-else`.

The type of the new variable `b` is that of `x`, where range types have been replaced by `int`. Range constraints are reconstructed in a type condition that is added to the formula `f`.

Example: For formulas `e,f,g, Forall x : [1..N]. h(x)`, and `y` of type integer not free in these, Skolemize eliminates the universal force quantifier by:

$$\frac{\texttt{f} \quad \Longrightarrow \quad \texttt{e --> g /\ (1 <= y /\ y <= N --> h(y))}}{\texttt{f} \quad \Longrightarrow \quad \texttt{e --> g /\ Forall x : [1..N]. h(x)}}$$

Instantiate Quantifiers of existential force can be instantiated by terms.

After clicking on this button, a data entry window appears for entering the term you want to instantiate the bound variable with.

Instantiate implements the inference rules:

$$\frac{\Gamma \quad \Longrightarrow \quad \Delta,\ \texttt{p(t)}}{\Gamma \quad \Longrightarrow \quad \Delta,\ \texttt{Exists x . p(x)}}$$

$$\frac{\Gamma,\ \texttt{p(t)} \quad \Longrightarrow \quad \Delta}{\Gamma,\ \texttt{Forall x . p(x)} \quad \Longrightarrow \quad \Delta}$$

where `t` is a term and, `p(x)` is a formula containing `x`. Complementing skolemization, instantiation applies to any existential force quantifiers under
`-->, \/, /\, ~, if-then-else`.

If the user-supplied term `t` is not obviously of the same type as `x`, Instantiate generates additional type constraint conditions along with the instantiated formula.

Example: For formulas `(N > 0)`, `Exists : [0..N-1]. h(x)`, and term `t`, Instantiate `x` with `t` eliminates the existential quantifier in the following way:

$$\frac{\texttt{N > 0} \implies \texttt{0 <= t /\ t <= N - 1 /\ h(t)}}{\texttt{N > 0} \implies \texttt{Exists x : [0..N-1]. h(x)}}$$

Replace This button allows you to replace equals by equals, that is, any term that is asserted to be equal to another term in the antecedent may be replaced with that other term throughout the sequent.

Upon clicking on this button all terms that are candidates to be replaced are highlighted, and you can select the term that you want to replace with.

Example: Starting with the sequent

$$\texttt{x=3, y=4} \implies \texttt{x*x + y*y = 25}$$

clicking on Replace highlights: `x`, `y`, 3, and 4. Clicking on 3 replaces `x` with 3 in the consequent, resulting in

$$\texttt{x=3, y=4} \implies \texttt{3*3 + y*y = 25}$$

Rewrite Brings up the **Select rewrite** window, which displays all rewrite rules that were declared in the specification during the current verification session via

`REWRITE` *desc* : *expn* `--->` *expn*

Rules can be selected and de-selected by clicking on the Select buttons. When clicking on the Apply button, the current goal is simplified according to the selected rewrite rules.

Make-first-order Upon clicking on this button with current goal sequent G, a first-order subgoal G' is generated, such that the first-order validity of G' implies the first-order temporal validity of G.

Next Upon clicking on this button a new subgoal is created in which the operator $\bigcirc$ is deleted from (sub)formulas preceded by it, and in which $\ominus$ is appended to (sub)formulas not preceded by $\bigcirc$.

Cut With the **Cut** rule you can perform a case split analysis on the current goal. Upon clicking on this button, a dialog window will appear, where you can enter the formula to split on. Thus upon having entered the cut-formula `p`, Cut implements the inference rule

$$\frac{\Gamma, \texttt{p} \implies \Delta \qquad \Gamma \implies \Delta, \texttt{p}}{\Gamma \implies \Delta}$$

Induction Upon clicking on this button a dialog window will appear in which you can enter the variable you want to do induction on. The variable has to be a universally quantified integer variable. The button implements the inference rule:

$$\frac{\begin{array}{lll}\Gamma & \Longrightarrow & \Delta,\ \texttt{p(0)} \\ \Gamma & \Longrightarrow & \Delta,\ \texttt{Forall x : int . (x >= 0 /\ p(x) --> p(x+1))} \\ \Gamma & \Longrightarrow & \Delta,\ \texttt{Forall x : int . (x <= 0 /\ p(x) --> p(x-1))}\end{array}}{\Gamma \Longrightarrow \Delta,\ \texttt{Forall x : int . p(x)}}$$

where `p` is a formula containing `x`.

Free-Induction supports simple mathematical induction for free variables. Suppose that the integer variable `x` occurs free in the sequent $\Gamma \Longrightarrow \Delta$, which we recall corresponds to a formula:

$$\texttt{p(x)}: \quad \bigwedge_{\gamma(x)\in\Gamma} \gamma(x) \rightarrow \bigvee_{\delta(x)\in\Delta} \delta(x).$$

Since x is implicitly universally quantified it makes sense to formulate the induction principle:

$$\frac{\begin{array}{ll}\Longrightarrow & \texttt{p(0)} \\ \Longrightarrow & \texttt{Forall x : int . (x >= 0 /\ p(x) --> p(x+1))} \\ \Longrightarrow & \texttt{Forall x : int . (x <= 0 /\ p(x) --> p(x-1))}\end{array}}{\Gamma \Longrightarrow \Delta}$$

Presburger Decides formulas of Presburger arithmetic. Formulas belonging to Presburger arithmetic may involve integer valued variables, addition, multiplication by constants, boolean connectives, and quantifiers over integer variables.

The decision problem is super-exponential, and thus only works on small examples. It is seldom needed in practice, as the decision procedure cannot handle uninterpreted function symbols (e.g., arrays).

Example: The following are Presburger formulas:

```
Forall x : int . Exists y : int . (x + y > 0)
Exists x : int . Forall y : int . Exists z : int .
     (x > y \/ x < y \/ z < y - x /\ 5*z = x)
```

The following are not Presburger formulas:

```
Forall x : int . Exists y : int . (f(x) + y > 0)
Exists x : int . Forall y : rat . (x > y)
```

since the first contains an uninterpreted function symbol `f`, while the second includes the unsupported type `rat`.

All-Propositional Applies 1-Step-Propositional recursively until there are no new formulas to which it is applicable.

BDD-split Simplifies the purely propositional structure of the formulas in the sequent. Thus, a formula whose main connective is a temporal operator `[]`, `<>`, `Awaits`, `Until`, `Since`, `Backto`, `[-]`, `<->`, is treated as an atomic formula (as opposed to All-Propositional). See Section 6.2 for further explanations.

Add-Axiom A list of previously established properties or asserted axioms is displayed in a separate window. The user selects the previously proved property or asserted axiom to be inserted in the antecedent of the sequent.

The Select boxes menu-option under Settings allows the user to automatically remove or retain preceding boxes from the axioms added by Add-Axiom. By default boxes are removed from axioms added to the sequent.

Example: By default, adding the axiom `[](p(x) /\ q(y))` to the sequent $\Gamma \Longrightarrow \Delta$ generates

$$\texttt{(p(x) /\ q(y))}, \Gamma \Longrightarrow \Delta.$$

Selecting Select boxes keeps the boxes. Hence adding the same axiom generates the sequent:

$$\texttt{[](p(x) /\ q(y))}, \Gamma \Longrightarrow \Delta.$$

PTL-expansion This invokes a symbolic PTL decision procedure. The returned model is furthermore checked for satisfiability with respect to the ground theory of equality and arithmetic.

Chapter 9

Basic Tutorial

This chapter presents detailed verification sessions that can be followed as a tutorial for STeP and used as templates for other verifications. All input files mentioned in this chapter are available in the `examples` directory in the STeP distribution. We assume that you have successfully started STeP (see Section 1.2).

9.1 Preliminaries

The input to STeP will be an SPL program P and a temporal-logic formula φ that expresses a property of P to be verified. A verification session deals with a single program and one or more properties to be verified for that program. Thus, loading a new program reinitializes the system, and starts a new verification session.

Loading a program

To load an SPL program, select the File menu option in the upper left corner (see Figure 3.2), and select Load program. This will start up the File Browser with a Load program window, as shown in Figure 3.3. By clicking on the *.spl button, you can choose to see only the files with extension `.spl`, the standard extension for SPL program files. The right column shows the directory hierarchy, and the left column shows the files in the current directory. Clicking on a directory in the right column makes it the current directory. You can load a file by clicking on its filename in the left column. Alternatively, you may enter the full filename directly in the entry line under `File name`. After entry, click the OK button, and the selected file will be loaded.

Loading the specification

To load a specification, select the Load specification option from the File menu. Again, this activates the File Browser, which brings up a Load specification window. The default extension for specification files is `.spec`. If the specification file contains a single formula, this formula is understood to be the top-level goal, and is displayed in the top window as the new current goal. If the file contains multiple formulas to be proved, the first one

is displayed as the current top-level goal. The other top-level goals can be accessed by selecting the Next search or Select search option under the Properties menu.

Specification files can also contain declarations, axioms and definitions—see Chapter 2.

Help

At any time, you can get online help by selecting the Help option from the menu. This will bring up a WWW browser with STeP's help pages.

Quitting STeP

To leave STeP, select the Quit option from the top-level menu or select the Quit option from the File menu option. Both will bring up a confirmation window asking if you really want to exit STeP.

9.2 MUX-SEM: Mutual Exclusion

We now assume that you have started STeP and are able to load program and specification files. The program and specifications mentioned in this section are available in the directory `examples/mux-sem`.

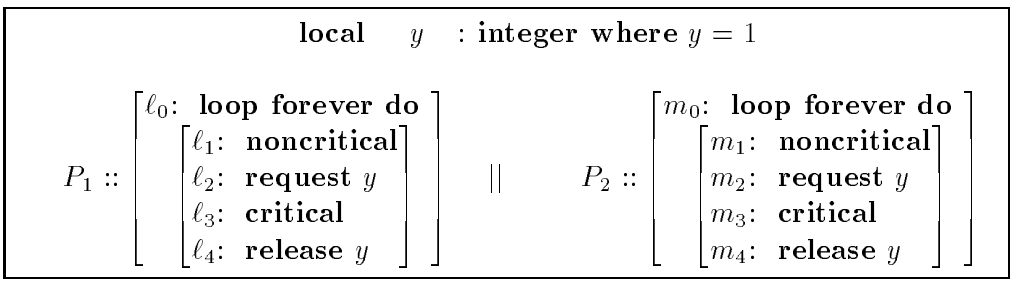

Figure 9.1: Program MUX-SEM (mutual exclusion by semaphores)

We begin with the proof of two simple safety properties for the program MUX-SEM, shown in Figure 9.1: the first, φ_1, states that y is always greater than 0, and the second, φ_2, expresses mutual exclusion. These properties are expressed by the formulas

$$\varphi_1 : \Box(y \geq 0)$$

$$\varphi_2 : \Box(\neg(\ell_3 \wedge m_3))$$

or, in STeP's input format,

```
[](y >= 0)
```

```
[]!(l3 /\ m3).
```

To illustrate the various methods available for verification we will prove these properties in a number of different ways. In this example we will load each property from a separate file.

9.2.1 Using B-INV

The first property, $\Box(y \geq 0)$, is *inductive*, so it should be easy to verify it using rule B-INV. We can do this in the following steps:

1. Load the program MUX-SEM by loading the file `mux-sem.spl`. This will bring up the Program Text window shown in Figure 9.2. Apart from the program, it contains a list of how the program's control locations correspond to STeP's internal variables. Control locations are explained in detail in Appendix C.2. When an SPL program is loaded, a fair transition system is automatically generated, which you can view by selecting View Transitions from the File menu.

2. Load the specification file `mux-sem/yge0.spec`. The property appears in the current goal window, as shown in Figure 9.3.

3. Enable automatic validity checking, by selecting Check-Valid in the Automatic Simplify / Check-Valid / Off option under the Settings menu.

4. Click on the B-INV button in the action region. One by one, the verification conditions for all transitions are displayed in the current goal window, and the result of their verification is displayed in the output window (see Figure 9.4). Since all verification conditions are simplified to *true*, the proof is complete, as indicated. Note that only the transitions `l2`, `l4`, `m2`, and `m4` are analyzed. The other transitions do not affect `y` and STeP realizes before generating verification conditions that the invariant is trivially maintained over these transitions.

The single original goal has now been proved, so there is no current (unproven) goal, as indicated in the current goal window. The proved property is saved as a background property and can be used in the proof of subsequent properties. The background properties can be viewed by selecting the Activate/Deactivate option under the Properties menu. This brings up a Property Checklist window, shown in Figure 9.5. By default, properties are active and will be used in the proof of subsequent properties. However, you may deactivate a property by clicking on the switch next to it. Inactive properties can be reactivated later on; however, you may also entirely delete a property by clicking on the skull.

To save the proof to a file in ASCII format, select the Save last proof option of the File menu. The File Browser brings up a window with the title "`Save completed proof`". As when loading a file, you can select an existing filename or enter a new one. Enter a filename with suffix `.search`. The window also allows you to enter three lines of text that will be added to the top of the file as a comment. Finally, click on OK to save your proof. Figure 9.6 shows the format of the saved proof.

9.2.2 Using G-INV

The next property we will prove, φ_2, expresses mutual exclusion. This property is not inductive, even relative to the simple property φ_1 proved before. However, we can prove it by *strengthening* the invariant being proved, as follows:

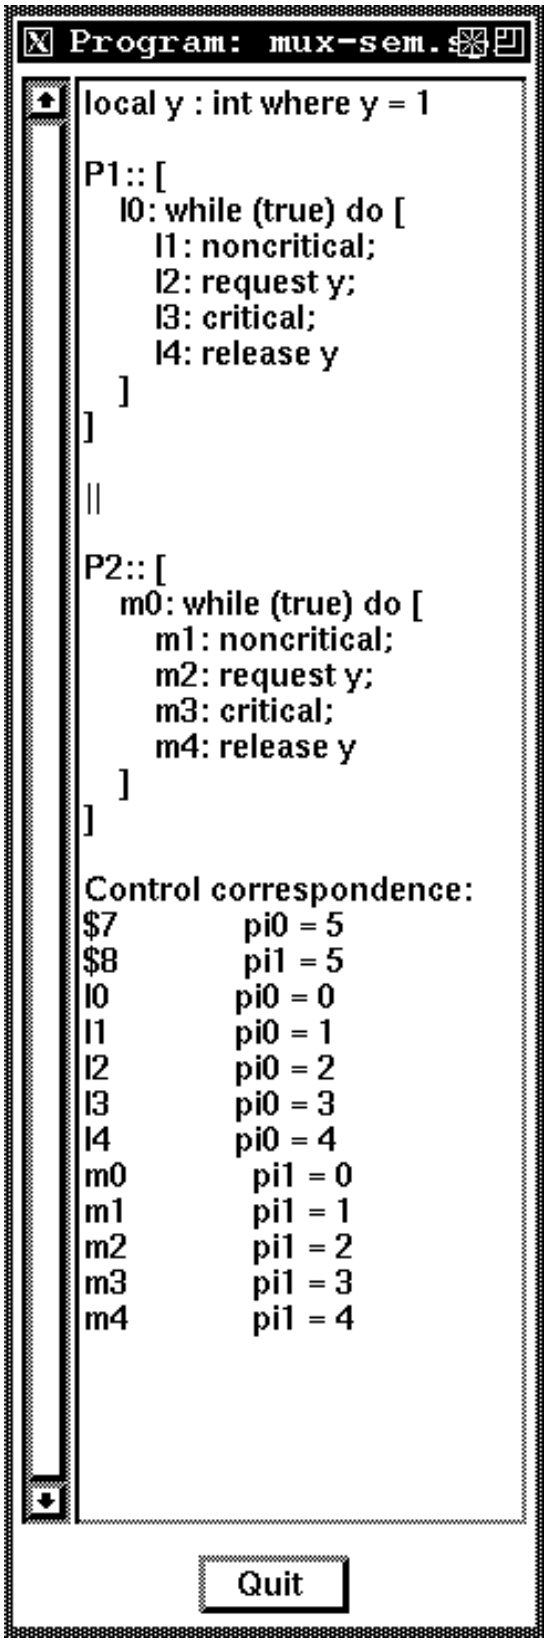

Figure 9.2: MUX-SEM program text window

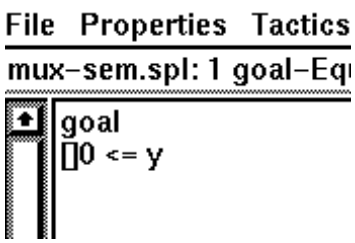

Figure 9.3: Loaded specification for MUX-SEM

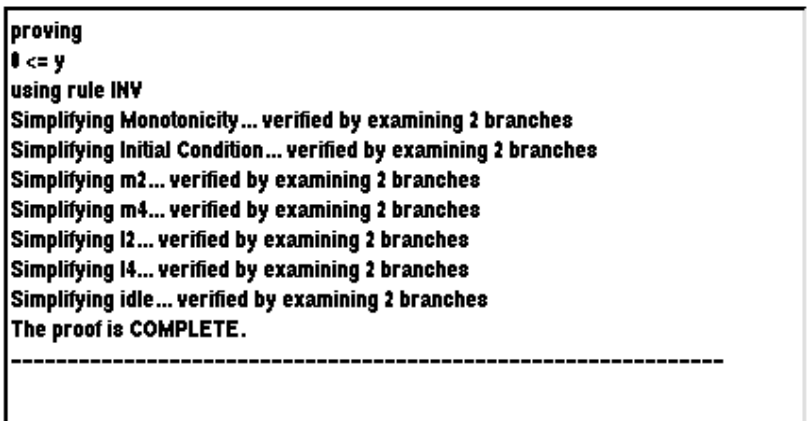

Figure 9.4: Verification of `[]y >= 0` for MUX-SEM

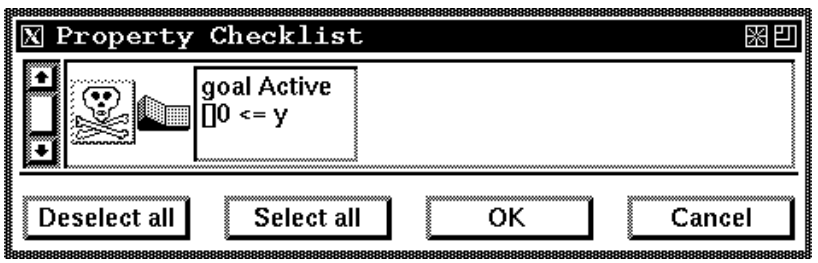

Figure 9.5: Background properties

```
% Proof of [](y>=0) for program mux-sem, using rule B-INV.
Formula abbreviations
phi0:        0 <= y

Reduced via G-INV: []0 <= y
given the strengthening formula:  phi0
1 Established via Check-Valid: true
2 Established via Check-Valid: y = 1 /\ l0 /\ m0 --> 0 <= y
3 Established via Check-Valid: {phi0}m2{phi0}
4 Established via Check-Valid: {phi0}m4{phi0}
5 Established via Check-Valid: {phi0}l2{phi0}
6 Established via Check-Valid: {phi0}l4{phi0}
7 Established via Check-Valid: {phi0}idle{phi0}
```

Figure 9.6: Transcript for proof of `[](y>=0)` over MUX-SEM

1. Load the specification file `mux-sem/mux.spec`. The property φ_2 will appear in the current goal window.

2. Click on G-INV. This will bring up a window with the title "`Input auxiliary assertions`". By default, the auxiliary assertion is the goal to be proven, so when the default is used G-INV behaves exactly as B-INV. Now enter the following formula after deleting the default one:

   ```
   (l3+l4+m3+m4+y) = 1
   ```

   and click on OK. If there is a syntax error in the entered formula, a message will be displayed at the bottom of the entry window, and you will be asked to try again.

   Using boolean terms such as `l3` and `m3` as integers is a convenient feature for writing succinct specifications, and is called *arithmetization*. The boolean value *true* is also interpreted as 1, and *false* as 0 (see Chapter 2).

   After the strengthened assertion is entered, all verification conditions are simplified to *true* and the proof is complete.

Note: The strengthened property only implies mutual exclusion based on the invariant $\varphi_1 : \Box(y \geq 0)$ that we proved before. If you do not prove this property before attempting the proof of mutual exclusion, the Monotonicity verification condition will fail to simplify to *true*.

The basic but useful property $y \geq 0$ can also be generated automatically using STeP's invariant generation tools. To see how this can be used, reset the searches with the Reset searches option under the File menu, and repeat the two steps using G-INV above. The Monotonicity property will not be simplified to *true*. Now we can do the following:

3. Select the Get local invariants option from the Properties menu. Seven invariants appear in the output window. The second invariant, in Figure 9.7, is φ_1, the property we need.

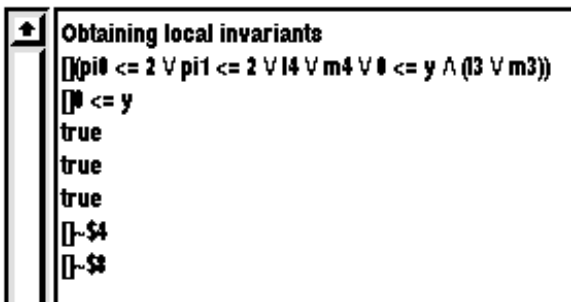

Figure 9.7: Local invariants for MUX-SEM

4. Click on the Check-Valid button. Now the Monotonicity condition should simplify to *true*, and the proof is complete.

9.2.3 Using the Model Checker

Since program MUX-SEM is a finite-state program we may use the Model Checker to prove mutual exclusion, as follows:

1. Load again the mutual-exclusion property from the `mux-sem/mux.spec` file. (The model checker does not use background properties, so you don't have to reset the current background properties at this point.)

2. Click on the Modelcheck button. This will bring up a window titled "`Input additional initial condition`". This window allows the user to enter additional assumptions under which model checking is to be done. Since we don't need any in this case, we can just click the OK button. A File Browser window titled "`Modelcheck log file`" will appear. Here, you have to specify a file in which you want to have model checker results recorded. The default extension of the model checker log file is `.mclog`.

3. After the log file is specified, the model checker will run. We will then see in the output window:

   ```
   Modelcheck succeeded. The proof is COMPLETE.
   ```

   The model checker log file provides more detailed information, such as the number of states explored. In general, however, the log file is interesting only when the model checker fails. In this case, the file will include a counterexample computation, that is, a computation of the system which violates the original property.

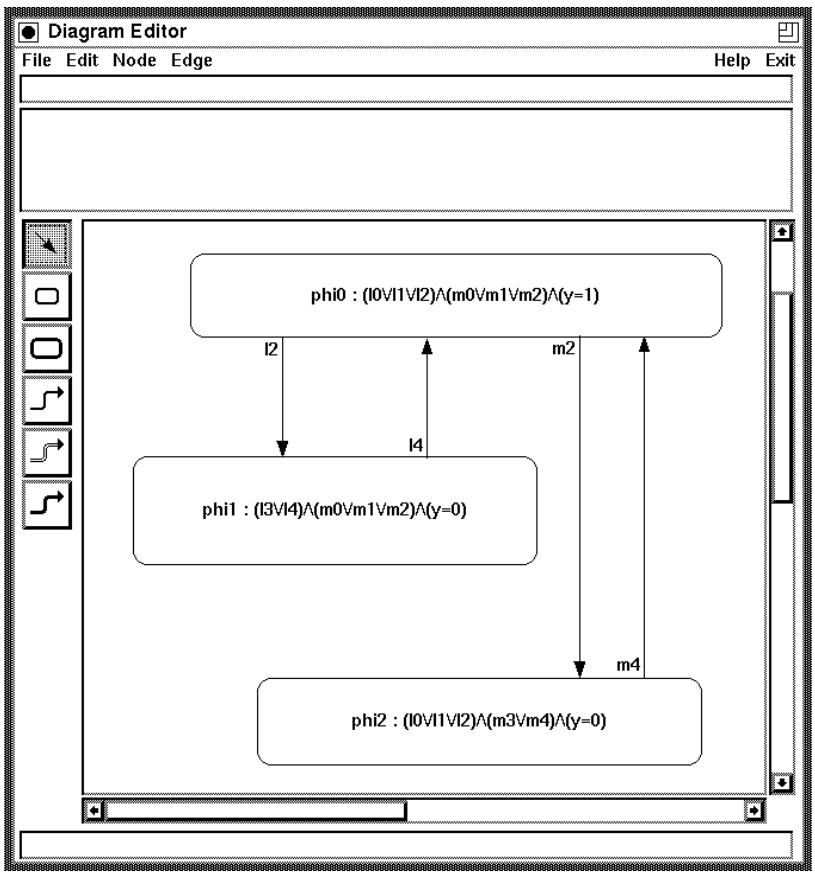

Figure 9.8: Invariance verification diagram for MUX-SEM

9.2.4 Using Verification Diagrams

In the discussion of MUX-SEM, [Manna and Pnueli, 1995b] presents a state-space partition graph. This graph can easily be converted into the invariance verification diagram shown in Figure 9.8. This diagram may be used to prove mutual exclusion for MUX-SEM as follows:

1. Reset the background properties and current searches with the Reset searches option under the File menu.

2. Load again the mutual-exclusion property from the `mux-sem/mux.spec` file.

3. Select the Edit verification diagram option under the Diagrams menu. This brings up the Verification Diagram Editor.

4. Select the Load option under the File menu in the Verification Diagram Editor window. This brings up the File Browser.

5. Select the file `mux-sem/state-partition.diagram`. This loads the verification diagram, which is displayed in the editor window, as shown in Figure 9.8.

6. Select the Current-to-Verify option of the File menu of the Verification Diagram Editor. This will make the verification diagram accessible to the Top-level Prover.

7. Select the Verification diagram rule option of the Diagrams option. This generates all the verification conditions corresponding to the given verification diagram as subgoals. If automatic validity checking is enabled, the simplifier will proceed to verify each subgoal.

   The first two subgoals are *side verification conditions*. They state that the label on each of the nodes implies the invariant to be proven, and that the initial condition implies the label of one of the nodes. Then the verification conditions corresponding to each of the edges in the diagram are simplified (including those for the implicit self-loops). In this case, all verification conditions simplify to *true* and the proof is complete.

9.2.5 Using Linear Invariants

Finally, we show how we can prove φ_2 based on φ_1 and automatically generated *linear invariants*:

1. Load the `examples/mux-sem.spl` program and prove $\varphi_1 : \Box(y \geq 0)$ using B-INV, as described in Section 9.2.1.

2. Load the goal φ_2 from `examples/mux.spec`.

3. Select the Get global linear invariants option under the Properties menu. This will automatically generate a number of linear invariants which are added to the set of background properties, including

   $$\Box(l_0 + l_1 + l_2 = m_3 + m_4 + y).$$

   Your set of background properties should now contain this formula, as well as φ_1.

4. Click on Check-Valid. As the conjunction of all background invariants implies the original goal, the property will be simplified to *true*, and the proof is complete.

9.3 Two-process Bakery

The *bakery* algorithm is based on the idea that customers, as they enter, pick numbers which form an ascending sequence. Then, a customer with a lower number has higher priority in accessing its critical section.

The bakery algorithm is shown in Figure 9.9 as a SPL-program. Two processes, `P1` and `P2`, draw their numbers at locations `l1` and `m1`. If both processes wish to access the critical region at location `l3` or `m3` only the process with the lowest value obtains access.

In the following we demonstrate the generation of invariants and verify three properties of the Bakery algorithm.

```
local y1, y2 :int where y1 = 0, y2 = 0
P1::[loop forever do [
        l0: noncritical;
        l1: y1 := y2 + 1;
        l2: await (y2 = 0 \/ y1 <= y2);
        l3: critical;
        l4: y1 := 0
        ]
]
||
P2::[loop forever do [
        m0: noncritical;
        m1: y2 := y1 + 1;
        m2: await (y1 = 0 \/ y2 < y1);
        m3: critical;
        m4: y2 := 0
        ]
]
```

Figure 9.9: Two-process Bakery.

9.3.1 Static analysis

The first step in the verification process is usually the generation of invariants since these will facilitate the verification of other properties. After the program is loaded into STeP select the Get global linear invariants and Get global polyhedral invariants options under the Properties menu. The generation of local and polyhedral invariants in gives the bottom-up properties in Figure 9.10. Once generated these properties are part of the background properties and will be used by the Simplify and Check-Valid rules. The properties are also available in the interactive prover through the Add-axiom rule.

9.3.2 Mutual exclusion

Mutual exclusion can be expressed as the following property:

```
[]!(l3 /\ m3)
```

With the generated invariants, this property is verified by applying the B-INV rule.

9.3.3 1-bounded overtaking

When process `P1` waits to access its critical section, process `P2` can only access its section at most once. This property, called *1-bounded overtaking*, is specified by:

$$\textit{1-bo:}\ \ at_\ell_2 \rightarrow \neg m3\ \mathcal{W}\ m3\ \mathcal{W}\ \neg m3\ \mathcal{W}\ l3$$

or in STeP format:

```
% local invariants
[]~$10
[]~$9
pi0 <= 2 \/ $9 ==> y1 = 0
pi1 <= 2 \/ $10 ==> y2 = 0
[]0 <= y2
[]0 <= y1

[](l5 \/
  m5 \/
  y2 < y1 /\ m4 \/
  (l1 \/ l2) \/
  (m1 \/ m2) \/
  y1 <= y2 /\ l4 \/
  l3 /\ y1 = y2 + 1 \/
  m3 /\ y2 = y1 + 1 \/
  y1 = 0 /\ (l0 \/ m4) \/ y2 = 0 /\ (l4 \/ m0))

% polyhedral invariants

[]~$10
[]~$9
l0 ==> 0 <= y2 /\ l0 /\ y1 = 0
l1 ==> 0 <= y2 /\ l1 /\ y1 = 0
l2 ==> 0 <= y2 /\ l2 /\ y1 = 0
l3 ==> 0 < y1 /\ 0 <= y2 /\ l3
l4 ==> 0 < y1 /\ 0 <= y2 /\ l4
l5 ==> 0 < y1 /\ 0 <= y2 /\ l5
m0 ==> 0 <= y1 /\ m0 /\ y2 = 0
m1 ==> 0 <= y1 /\ m1 /\ y2 = 0
m2 ==> 0 <= y1 /\ m2 /\ y2 = 0
m3 ==> 0 < y2 /\ 0 <= y1 /\ m3
m4 ==> 0 < y2 /\ 0 <= y1 /\ m4
m5 ==> 0 < y2 /\ 0 <= y1 /\ m5
```

Figure 9.10: Bottom-up invariants

```
l2 ==> ~m3 Awaits m3 Awaits ~m3 Awaits l3
```

The property has to be strengthened to enable verification. Similarly to the general invariance rule, the G-WAIT rule offers the possibility to strengthen the intermediate assertions to make the verification conditions state-valid.

In STeP we can either apply G-WAIT directly to the property, and then enter the strengthening formulas in the dialog window provided, or we can apply it in the form of a tactic as follows:

```
{ G-WAIT
  "m2 /\ l2 /\ y2 < y1"                    % strengthens ~m3
  "m3 /\ l2"                               % strengthens m3
  "~m3 /\ l2 /\ (m4 \/ (y2=0\/y1<=y2))"    % strengthens ~m3
  "l3"                                     % strengthens l3
};
Check-Valid
```

The resulting verification conditions are proved automatically (using the generated invariants), which completes the proof. Nested wait-for properties can also be verified with WAIT-FOR diagrams. Figure 9.11 shows a WAIT-FOR diagram for the one-bounded overtaking property of the bakery algorithm. We see that the assertions labeling the nodes are exactly the ones we used to strengthen the intermediate formulas when applying rule G-WAIT.

9.3.4 Response

The property of one-bounded overtaking makes sure that process P2 will enter its critical section at most once before process P1 is allowed to do so. However, it does not guarantee that P1 will eventually be allowed access to its critical section if it is interested. This is because a wait-for formula allows each interval to extend to infinity, so it is possible that one of the intermediate intervals will not end. The *accessibility* property states that P1 will eventually reach its critical section, once it is at ℓ_1:

```
l1 ==> <> l3
```

To verify it, load and apply the progress diagram shown in Figure 9.12.

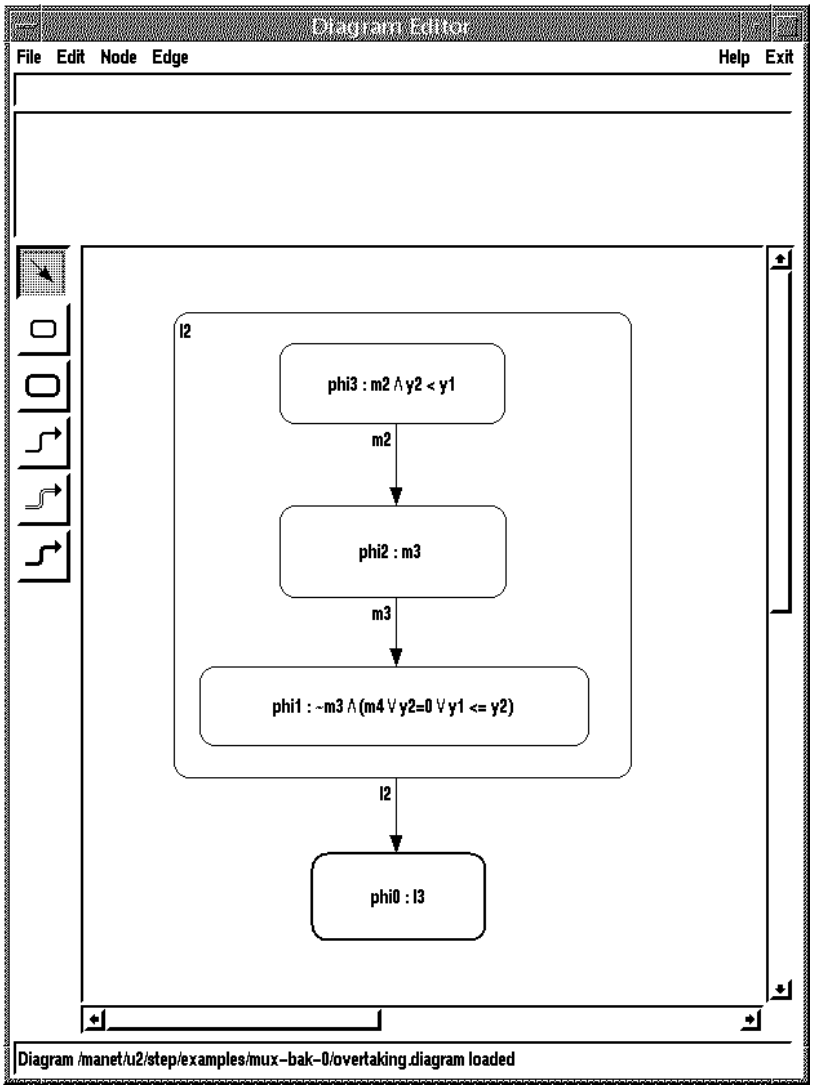

Figure 9.11: Wait-for verification diagram for bakery.

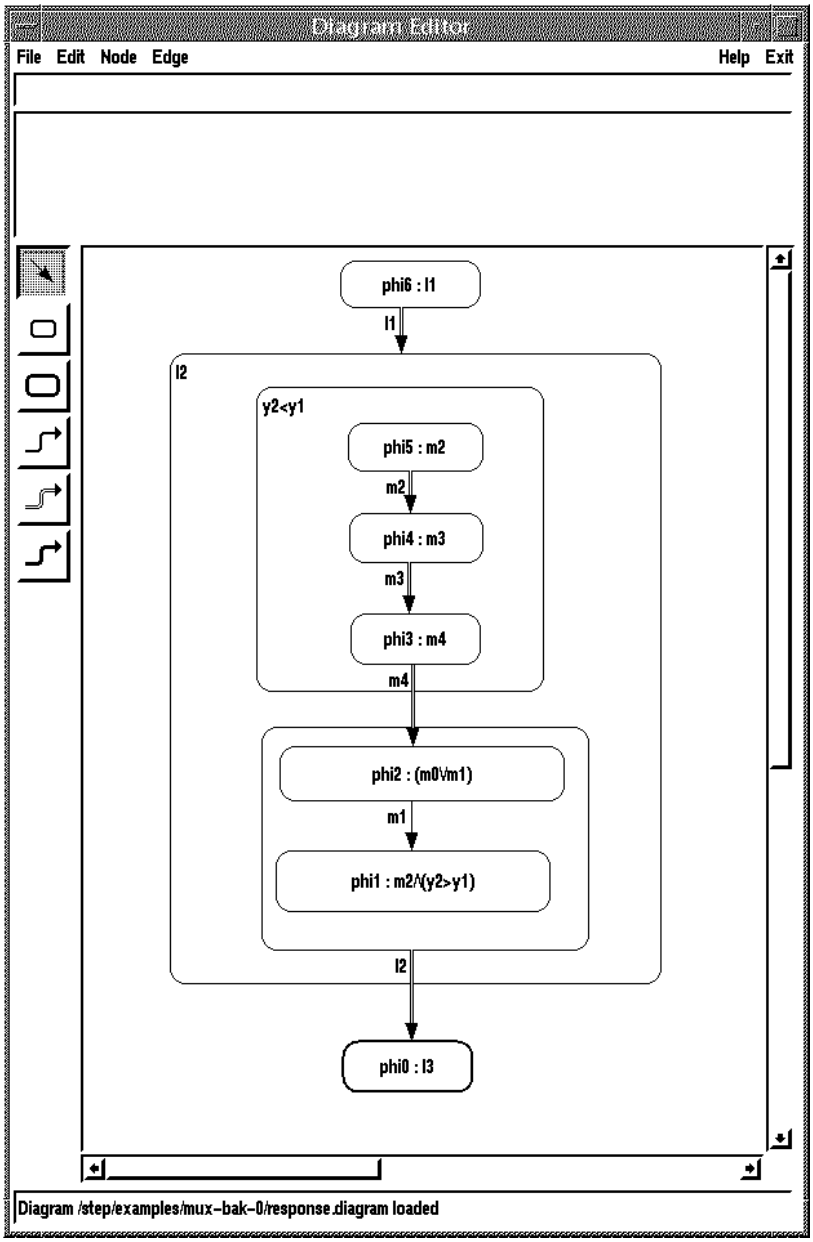

Figure 9.12: Progress verification diagram for bakery.

Appendix A

Computational Model

A.1 Fair Transition Systems

An SPL program is compiled into a fair transition system. A fair transition system $\langle V, \Theta, \mathcal{T}, \mathcal{J}, \mathcal{C} \rangle$ is given by the following components:

- V: *system variables*, including both *data variables* and *control variables*.
- Θ: *initial condition*, which is a satisfiable assertion characterizing all the initial states of a computation.
- $\mathcal{T}$: set of *transitions*. Each transition $\tau \in \mathcal{T}$ is a function

$$\tau : \Sigma \mapsto 2^{\Sigma},$$

  where Σ is the set of all states, and a state is a type-consistent interpretation of V. Each state in $\tau(s)$ is called a τ-successor of s. A transition τ is said to be *enabled* on S if $\tau(s) \neq \emptyset$. Otherwise it is said to be *disabled*.
- $\mathcal{J} \subseteq \mathcal{T}$: the set of *just* transitions (also called weakly fair transitions).
- $\mathcal{C} \subseteq \mathcal{T}$: the set of *compassionate* transitions (also called strongly fair transitions).

Each transition $\tau \in \mathcal{T}$ is represented by a first-order formula $\rho_\tau(V, V')$, called the *transition relation*, which may refer to both unprimed and primed versions of the system variables. The transition relation expresses the relation holding between a state s and any of its τ-successors $s' \in \tau(s)$. The unprimed version of a variable refers to its value in s, and the primed version refers to its value in s'.

Thus, the state s' is a τ-successor of the state s if the formula $\rho_\tau(V, V')$ evaluates to `true`, when we interpret each $x \in V$ as $s[x]$, and its primed version x' as $s'[x]$.

If not explicitly specified, STeP adds to each fair transition system the *idling* transition τ_I (also called the stuttering transition), whose transition relation is $\rho_I : (V = V')$.

A.2 Computations

An infinite sequence of states

$$\sigma : s_0, s_1, s_2, \ldots$$

is defined to be a *computation* of a fair transition system P if it satisfies the following requirements:

- *Initiality:* s_0 is initial, that is, it satisfies the initial condition.
- *Consecution:* For each $j = 0, 1, \ldots$ the state s_{j+1} is a τ-successor of the state s_j.
- *Justice:* For each transition $\tau \in \mathcal{J}$, it is not the case that τ is continually enabled beyond some position j in σ but not taken.
- *Compassion:* For each transition $\tau \in \mathcal{C}$, it is not the case that τ is enabled at infinitely many positions in σ but taken at only finitely many positions.

A.3 SPL semantics

System Variables

The system variables V consist of the program variables declared in the program, and a set of control variables (approximately one for each spawned subprocess), and a set of location variables (one for each label in the program). A control variable ranges over a subset of the integers, where each value corresponds to a location in a process (see Section C.2). A location variable is a boolean variable with the same name as a program label. Its value is true iff control currently resides at the location of the label, i.e., iff one of the control variables equals the value of the corresponding location.

Initial Condition

The initial condition Θ for program P, with n top-level processes, is defined as

$$\Theta : \pi_0 = 0 \ \wedge \ \pi_1 = 0 \ \wedge \ \ldots \ \wedge \ \pi_n = 0 \ \wedge \ \varphi,$$

where φ is the conjunction of all the assertions that appear in the *where* clause of the declaration of P. For a parameterized program, with top-level processes $P[i]$, $1 \leq i \leq N$, the initial condition is given by

$$\Theta : \forall i : [1 \ldots N].\pi_0[i] = 0 \ \wedge \ \varphi.$$

Transitions

A transition relation is defined for each of the statements in the language.

To make the definitions more succinct, we define the abbreviation

$$move_S(\ell, \widehat{\ell}) : \pi_S = val(\ell) \ \wedge \ \pi'_S = val(\widehat{\ell}) \ \wedge \ loc(\ell) \ \wedge \ loc(\widehat{\ell'}),$$

where π_S refers to the control variable of the innermost subprocess in which the statement S appears, and $val(\ell)$ refers to the integer value associated with the location of ℓ, and $loc(\ell)$ refers to the location variable associated with the location of ℓ.

Since every transition usually modifies only a few variables, we define

$$pres(U) : \bigwedge_{u \in U} (u' = u),$$

where $U \subseteq V$. It states that all variables in U preserve their value at the current step.

In the following we represent a statement S with label ℓ and post-label $\widehat{\ell}$ in the form $[\ell : S; \widehat{\ell}]$. For a statement $[\ell : S; \widehat{\ell}]$ we denote by Y_S, $V - \{\pi_S, loc(\ell), loc(\widehat{\ell})\}$. Unless otherwise specified, all transitions are `Just`.

Basic Statements

- *Skip*
  Statement: $\ell :$ **skip**; $\widehat{\ell} :$
  Transition relation:
  $$\rho_\ell : move_S(\ell, \widehat{\ell}) \wedge pres(Y_S)$$

- *Assignment*
  Statement: $\ell : \bar{u} := \bar{e}; \widehat{\ell} :$
  Transition relation:
  $$\rho_\ell : move_S(\ell, \widehat{\ell}) \wedge \bar{u}' = \bar{e} \wedge pres(Y_S - \{\bar{u}\}).$$
  The assignment is always enabled (if control is at the assignment statement), even if $\bar{e}$ is not in the range of $\bar{u}$, which may happen with e.g., range types. A proof obligation to check for this type of runtime error is generated by the **Check Runtime System** menu option of the Top-level Prover.

- *Await*
  Statement: $\ell :$ **await** c; $\widehat{\ell} :$
  Transition relation:
  $$\rho_\ell : move_S(\ell, \widehat{\ell}) \wedge c \wedge pres(Y_S)$$

- *Asynchronous Send*
  Statement: $\ell : \alpha \Leftarrow e; \widehat{\ell} :$, where α is an asynchronous channel.

  Unbounded channel Transition relation:
  $$\rho_\ell : move_S(\ell, \widehat{\ell}) \wedge \alpha' = append(\alpha, e) \wedge pres(Y_S - \{\alpha\}).$$

  Bounded channel For a channel α with bound k, the transition relation is:
  $$\rho_\ell : move_S(\ell, \widehat{\ell}) \wedge length(\alpha) < k \wedge \alpha' = append(\alpha, e) \wedge pres(Y_S - \{\alpha\}).$$
  To ensure that a process cannot be excluded from sending to a bounded channel if it is infinitely often non-full, we associate the `Compassionate` justice requirement with this transition.

- *Asynchronous Receive*
  Statement: $\ell:\ \alpha \Rightarrow u; \widehat{\ell}:$
  Transition relation:

  $$\rho_\ell:\ move(\ell, \widehat{\ell}) \wedge length(\alpha) > 0 \wedge u' = head(\alpha) \wedge \alpha' = tail(\alpha) \wedge pres(Y_S - \{u, \alpha\}).$$

  The asynchronous receive statement is enabled only if the channel is currently nonempty.

  To ensure that a process must receive something if the channel is infinitely often nonempty, the `Compassionate` fairness requirement is associated with this transition.

- *Synchronous Send-Receive*
  With each pair of matching *send* and *receive* statements:

  $$\ell:\ \alpha \Leftarrow e; \widehat{\ell}: \qquad m:\ \alpha \Rightarrow u;\ \widehat{m}:$$

  where two parallel statements are considered matching if they form an $\alpha \Leftarrow e$, $\alpha \Rightarrow u$ pair for some e and u for the same synchronous channel α.
  Transition relation:

  $$\begin{array}{lrcl} \rho_{\langle \ell, m \rangle}: & \pi_{send} & = & val(\ell) \wedge \pi'_{send} = val(\widehat{\ell}) \wedge \\ & \pi_{receive} & = & val(m) \wedge \pi'_{receive} = val(\widehat{m}) \wedge \neg loc(\ell) \wedge \\ & & & loc(\widehat{\ell}) \wedge \neg loc(m) \wedge loc(\widehat{m}) \wedge \\ & & & u' = e \wedge pres(Y_{send,receive}) \end{array}$$

  Thus, transition $\tau_{\langle \ell, m \rangle}$ is enabled if control is at ℓ and m simultaneously. When executed, the transition causes joint progress in the two processes containing the send and receive statement.

  The fairness requirement associated with the generated transitions is `Compassionate`, to ensure that a process cannot wait to transmit or receive a message indefinitely if infinitely many messages are sent on the channel it is waiting on.

- *Request*
  With the statement ℓ: **request**(r); $\widehat{\ell}$:, we associate a transition τ_ℓ, with the transition relation

  $$\rho_\ell: \quad move(\ell, \widehat{\ell}) \ \wedge \ r > 0 \ \wedge \ r' = r - 1 \ \wedge \ pres(Y_S - \{r\}).$$

  Thus, this statement is enabled when control is at ℓ and r is positive. When executed it decrements r by 1.

  To model fair scheduling of semaphores (in the form of for instance single or multi-level feedback queues), the transition associated with **request** is `Compassionate`.

- *Release*
  With the statement ℓ: **release**(r); $\widehat{\ell}$:, we associate a transition τ_ℓ, with the transition relation

  $$\rho_\ell: \quad move(\ell, \widehat{\ell}) \ \wedge \ r' = r + 1 \ \wedge \ pres(Y_S - \{r\}).$$

  This statement increments r by 1.

- *Noncritical*
  With the statement ℓ: **noncritical**; $\widehat{\ell}$:, we associate a transition τ_ℓ, with transition relation

  $$\rho_\ell: \quad move(\ell, \widehat{\ell}) \ \wedge \ pres(Y_S).$$

  Thus, the only observable action of this statement is to terminate. The situation that execution of the noncritical section does not terminate is modeled by a computation that does not take transition τ_ℓ. This is allowed by excluding τ_ℓ from the justice set.

- *Critical*
  With the statement ℓ: **critical**; $\widehat{\ell}$:, we associate a transition τ_ℓ, with transition relation

  $$\rho_\ell: \quad move(\ell, \widehat{\ell}) \ \wedge \ pres(Y_S).$$

  The only observable action of the critical statement is to terminate.

- *Produce*
  With the statement ℓ: **produce** x; $\widehat{\ell}$:, we associate a transition τ_ℓ, with transition relation

  $$\rho_\ell: \quad move(\ell, \widehat{\ell}) \ \wedge \ x' \neq 0 \ \wedge \ pres(Y_S - \{x\}).$$

  The observable action of the produce statement is to assign a nonzero value to variable x. Note that this transition is nondeterministic; that is, a state s may have more than one τ_ℓ-successors. In fact it may have infinitely many successors, one for each possible value of x'.

- *Consume*
  With the statement ℓ: **consume** y; $\widehat{\ell}$:, we associate a transition τ_ℓ, with transition relation

  $$\rho_\ell: \quad move(\ell, \widehat{\ell}) \ \wedge \ y' = 0 \ \wedge \ pres(Y_S - \{y\}).$$

  The observable action of the consume statement is to set variable y to zero.

- *Guarded Assignment*
  Statement: ℓ : **guard** c **do** $\bar{u} := \bar{e}$; $\widehat{\ell}$:
  Transition relation:

  $$\rho_\ell : \ c \ \wedge \ moves_S(\ell, \widehat{\ell}) \ \wedge \ \bar{u'} = \bar{e} \ \wedge \ pres(Y_S - \{\bar{u}\}).$$

  The assignment is enabled if control is at the assignment statement and c evaluates to *true*. Run-time type errors are not considered in the enabling condition of the transition. As with standard assignment, the possibility of run-time errors can be checked with the Check Runtime System option under the Properties menu.

Composite Statements

Standard control constructs and primitives for concurrency compose simpler program statements into more elaborate statements. A composite statement is in general of the form $\mathcal{C}[S_1, ...S_n]$, where $S_1, \ldots, S_n$ are the immediate substatements. The transitions associated with a composite statement comprise of the union of the transitions associated with

$S_1, \ldots, S_n$ and the transitions associated with the construct $\mathcal{C}$. Below we describe the transitions associated with each construct. Here Y_S refers to all system variables except the control and location variables changed by each construct.

- *Conditional*

  One-way With the statement ℓ: [**if** c **then** $\widetilde{\ell}$:S]; $\widehat{\ell}$:, we associate a transition τ_ℓ, with the transition relation

$$\rho_\ell: \quad \left(\begin{array}{ccc} c & \wedge & move(\ell, \widetilde{\ell}) \\ & \vee & \\ \neg c & \wedge & move(\ell, \widehat{\ell}) \end{array} \right) \wedge \; pres(Y_S).$$

  According to ρ_ℓ, when c evaluates to *true* control moves from ℓ to $\widetilde{\ell}$, and when c evaluates to *false* control moves from ℓ to $\widehat{\ell}$.

  Two-way With the statement ℓ: [**if** c **then** $\widetilde{\ell}_1$:S_1 **else** $\widetilde{\ell}_2$:S_2]; $\widehat{\ell}$:, we associate a transition τ_ℓ, with the transition relation

$$\rho_\ell: \quad \left(\begin{array}{ccc} c & \wedge & move(\ell, \widetilde{\ell}_1) \\ & \vee & \\ \neg c & \wedge & move(\ell, \widetilde{\ell}_2) \end{array} \right) \wedge \; pres(Y_S).$$

  According to ρ_ℓ, when c evaluates to *true* control moves from ℓ to $\widetilde{\ell}_1$, and when c evaluates to *false* control moves from ℓ to $\widetilde{\ell}_2$.

- *While*
  With the statement ℓ: [**while** c **do** [$\widetilde{\ell}$:$\widetilde{S}$]]; $\widehat{\ell}$: we associate a transition τ_ℓ, with the transition relation

$$\rho_\ell: \quad \left(\begin{array}{ccc} c & \wedge & move(\ell, \widetilde{\ell}) \\ & \vee & \\ \neg c & \wedge & move(\ell, \widehat{\ell}) \end{array} \right) \wedge \; pres(Y_S).$$

  According to ρ_ℓ, when c evaluates to *true* control moves from ℓ to $\widetilde{\ell}$, and when c evaluates to *false* control moves from ℓ to $\widehat{\ell}$. Note that, in this context, the post-location of $\widetilde{S}$ is ℓ. Note also that the enabling condition of τ_ℓ is at_ℓ, which does not depend on the value of c.

- *Loop forever*
  With the statement ℓ: [**loop forever do** [$\widetilde{\ell}$:$\widetilde{S}$]]; $\widehat{\ell}$: we do not associate any transitions. However the locations ℓ, $\widetilde{\ell}$ and post-location of $\widetilde{S}$ are equivalent.

- *Repeat until*
  With the statement ℓ: [**repeat** [$\widetilde{\ell}$: $\widetilde{S}$]$\widehat{\ell}_1$: **until** c];$\widehat{\ell}_2$: we associate a transition $\tau_{\widehat{\ell}_1}$, with the transition relation

$$\rho_{\widehat{\ell}_1}: \quad \left(\begin{array}{ccc} c & \wedge & move(\widehat{\ell}_1, \widehat{\ell}_2) \\ & \vee & \\ \neg c & \wedge & move(\widehat{\ell}_1, \ell) \end{array} \right) \wedge \; pres(Y_S).$$

According to $\rho_{\widehat{\ell}_1}$, when c evaluates to *true* control moves from $\widehat{\ell}_1$ to $\widehat{\ell}_2$, and when c evaluates to *false* control moves from $\widehat{\ell}_1$ to ℓ. The enabling condition of $\tau_{\widehat{\ell}_1}$ is $at_\widehat{\ell}_1$, which does not depend on the value of c. Note that the locations ℓ and $\widetilde{\ell}$ are equivalent.

- *When*
  With the statement ℓ: [**when** c **do** $[\widetilde{\ell}{:}\widetilde{S}\,]$]; $\widehat{\ell}$: we associate a transition τ_ℓ, with the transition relation
  $$\rho_\ell\colon \quad c \ \wedge \ move(\ell, \widetilde{\ell}) \ \wedge \ pres(Y_S).$$
  According to ρ_ℓ, when c evaluates to *true* control moves from ℓ to $\widetilde{\ell}$.

- *Or*
  No special transition is associated with the statement S_1 **or** S_2. When control resides at this statement, it is interpreted to reside simultaneously at the starting points of S_1 and S_2.

- *Cooperation*
  With the statement ℓ: $[\ell_1\colon\ S_1{:}\widehat{\ell}_1]||[\ell_2\colon\ S_2{:}\widehat{\ell}_2]||\ldots||[\ell_n\colon\ S_n{:}\widehat{\ell}_n]$ $\widehat{\ell}$: we associate two transitions, $\tau_{\ell_{Entry}}$ and $\tau_{\ell_{Exit}}$, with transition relations:
  $$\rho_{\ell_{Entry}}\colon \quad move(\ell, \{\ell_1, \ell_2, \ldots, \ell_n\}) \ \wedge \ pres(Y_S).$$
  $$\rho_{\ell_{Exit}}\colon \quad move(\{\widehat{\ell}_1, \widehat{\ell}_2, \ldots, \widehat{\ell}_n\}, \widehat{\ell}) \ \wedge \ pres(Y_S).$$
  All they do is to move control from the guarding location ℓ into the prelocation of each parallel statement, and from the post-location of each parallel statement to the post-location $\widehat{\ell}$ of the entire statement.

  STeP makes a special case for parallel composition at the top-level of the program. They do not get a designated pre- and post-location, but the initial state is instead what corresponds to $\ell_1, \ell_2, \ldots, \ell_n$ in the schema.

- *Concatenation*
  No special transition is associated with sequential composition of statements:
  $$S_1 \ ; \ S_2$$
  Since the post-location of S_1 is the pre-location of S_2 the transitions generated from each substatement S_1 and S_2 will move control appropriately from S_1 to S_2.

- *Parameterized Cooperation*
  With the statement ℓ: $i : [1..N]$. || [$\ell[i]$: $S[i]$ $\widehat{\ell}[i]$:] $\widehat{\ell}$: we associate two transitions , $\tau_{\ell_{Entry}}$ and $\tau_{\ell_{Exit}}$, with transition relations:
  $$\rho_{\ell_{Entry}}\colon \quad move(\ell, \{\ell[i] \mid i = 1, \ldots, N\}) \ \wedge \ pres(Y_S).$$
  $$\rho_{\ell_{Exit}}\colon \quad move(\{\widehat{\ell}[i] \mid i = 1, \ldots, N\}, \widehat{\ell}) \ \wedge \ pres(Y_S).$$
  They are the parameterized versions of the transitions associated with non-parameterized cooperation.

- *Parameterized Or*
  No special transition is associated with parameterized selection. This reflects that the construct is the parameterized version of standard *or*.

- *Blocks*
  Blocks are used to encapsulate variable declarations, and don't generate transitions on their own.

Grouped Statements

To compose several simple statements into one atomic entity, SPL provides the grouping construct **<< >>**. For instance the transition associated with a grouped statement of the form **<<**$S_1; S_2$**>>** is the relational composition of their transition relations.

Appendix B

Linear-Time Temporal Logic

As a requirement specification language for reactive systems, we take *linear-time temporal logic*. There is an underlying first-order assertion language $\mathcal{L}$ over interpreted symbols for expressing functions and relations over concrete domains such as the integers, arrays, and lists of integers (see Section 2.1).

We refer to a formula in the assertion language $\mathcal{L}$ as a *state formula*, or simply as an *assertion*. A *temporal formula* is constructed out of state formulas by applying the boolean operators $\neg$ and $\vee$ (the other boolean operators can be defined from these), and temporal operators. There are two classes of temporal operators, *future* and *past*. The future and past temporal operators used in STeP are presented in Tables B.1 and B.2, respectively.

Operator	Name	STeP representation
$\Box p$	Henceforth p	`[]`
$\Diamond p$	Eventually p	`<>`
$p\,\mathcal{U}\,q$	p Until q	`Until`
$p\,\mathcal{W}\,q$	p Waiting-for (Unless) q	`Awaits`
$\bigcirc p$	Next p	`()`

Table B.1: Future Temporal Operators

Operator	Name	STeP representation
$\boxminus p$	So-far p	`[-]`
$\diamondminus p$	Once p	`<->`
$p\,\mathcal{S}\,q$	p Since q	`Since`
$p\,\mathcal{B}\,q$	p Back-to q	`Backto`
$\ominus q$	Previously p	`(-)`

Table B.2: Past Temporal Operators

A *model* for a temporal formula p is an infinite sequence of states $\sigma : s_0, s_1, \ldots$, where each state s_j provides an interpretation for the vocabulary of p, i.e., the variables occurring in p. For a given state s_j and state formula q, we write $s_j \vDash q$ when q is true when evaluated over s_j. Given a model σ, we now present an inductive definition for the notion of a temporal formula p holding at a position $j \geq 0$ in σ, written as $(\sigma, j) \vDash p$:

- For a state formula p,
  $(\sigma, j) \models p \iff s_j \models p$
  That is, we evaluate p locally, using the interpretation provided by s_j.
- $(\sigma, j) \models \neg p \iff (\sigma, j) \not\models p$
- $(\sigma, j) \models p \vee q \iff (\sigma, j) \models p$ or $(\sigma, j) \models q$
- $(\sigma, j) \models \Box p \iff$ for all $k \geq j$, $(\sigma, k) \models p$
- $(\sigma, j) \models \Diamond p \iff$ for some $k \geq j$, $(\sigma, k) \models p$
- $(\sigma, j) \models p \,\mathcal{U}\, q \iff$ for some $k \geq j$, $(\sigma, k) \models q$,
  and $(\sigma, i) \models p$ for every i such that $j \leq i < k$
- $(\sigma, j) \models p \,\mathcal{W}\, q \iff (\sigma, j) \models p \,\mathcal{U}\, q$ or $(\sigma, j) \models \Box p$
- $(\sigma, j) \models \bigcirc p \iff (\sigma, j+1) \models p$

For past operators, we have:

- $(\sigma, j) \models \boxminus p \iff$ for all k, $0 \leq k \leq j$, $(\sigma, k) \models p$
- $(\sigma, j) \models \diamondminus p \iff$ for some k, $0 \leq k \leq j$, $(\sigma, k) \models p$
- $(\sigma, j) \models p \,\mathcal{S}\, q \iff$ for some k, $0 \leq k \leq j$, $(\sigma, k) \models p$
  and for every i such that $k < i \leq j$, $(\sigma, i) \models p$
- $(\sigma, j) \models p \,\mathcal{B}\, q \iff (\sigma, j) \models p \,\mathcal{S}\, q$ or $(\sigma, j) \models \boxminus p$
- $(\sigma, j) \models \ominus p \iff j > 0$ and $(\sigma, j-1) \models p$

Another useful derived operator is the *entailment* operator, defined by: $p \Rightarrow q \iff \Box(p \rightarrow q)$. For a state formula p and a state s such that p holds on s, we say that s is a *p-state*. A state formula that holds on all states is called *state-valid*.

For a temporal formula p and a position $j \geq 0$ such that $(\sigma, j) \models p$, we say that j is a *p-position* (in σ). If $(\sigma, 0) \models p$, we say that p *holds* on σ, and write it as $\sigma \models p$. A formula p is called *satisfiable* if it holds on some model. A formula is called (*temporally*) *valid* if it holds on all models.

Two formulas p and q are defined to be *equivalent*, written as $p \sim q$, if the formula $p \leftrightarrow q$ is valid, i.e., $\sigma \models p$ iff $\sigma \models q$, for all models σ. We adopt the convention by which a formula p that is claimed to be valid is state-valid if p is an assertion, and is temporally valid if p contains at least one temporal operator.

The formulas p and q are defined to be *congruent*, written as $p \approx q$, if the formula $\Box(p \leftrightarrow q)$ is valid, i.e., $(\sigma, j) \models p$ iff $(\sigma, j) \models q$, for all models σ and all positions $j \geq 0$. If $p \approx q$ then p can be replaced by q in any context, i.e., $\varphi(p) \approx \varphi(q)$ for any formula $\varphi(p)$ containing occurrences of p.

The notion of (temporal) validity requires that the formula holds over *all* models. Given a program P, we can restrict our attention to the set of models which correspond to computations of P, i.e., $\mathcal{C}omp(P)$. This leads to the notion of P-validity, by which a temporal formula p is *P-valid* (valid over program P) if it holds over all the computations of P. Obviously, any formula that is temporally valid is also P-valid for any program P. In a similar way, we obtain the notions of P-satisfiability and P-equivalence.

A state s that appears in some computation of P is called a *P-accessible* state. A state formula is called *P-state valid* if it holds over all P-accessible states. Obviously, any state formula that is state-valid is also P-state valid for any program P.

Again, we refer to a P-state valid formula simply as P-valid.

Appendix C

Differences with the Book

STeP was implemented based on [Manna and Pnueli, 1995b]. However, due to implementation considerations and limitations, there are some differences between STeP and [Manna and Pnueli, 1995b], hereafter referred to as MP95. For the convenience of those using STeP together with MP95, this appendix highlights most of the differences. These are deviations from the syntax of programs as well as differences in specification style.

C.1 SPL Programs

STeP closely follows the SPL syntax described in MP95. However, some constructs had to be translated into an ASCII equivalent. To ease transcribing programs from the book, Table C.1 presents each statement as it appears in MP95 and its representation in STeP. In the table, u refers to an arbitrary variable, $\bar{u}$ refers to a list of variables (i.e., $u_1, \ldots, u_n$), e refers to an arbitrary expression, $\bar{e}$ refers to a list of expressions, z refers to an arbitrary channel, and S refers to an arbitrary statement.

Table C.2 presents the various types of expressions as they appear in MP95 and their representation in STeP. Again, e is an arbitrary expression, u is an arbitrary variable, and $\bar{e}$ refers to a list of expressions.

The various types allowed in STeP are presented in Table C.3. In this table, e is an arbitrary expression, u is an arbitrary variable, b is an arbitrary base type, t is an arbitrary base or simple type, and T is an arbitrary type (i.e., base, simple or complex). A *range* in STeP is written $e_1..e_2$.

C.2 Control Locations in STeP

As in MP95, STeP uses control variables to indicate where the current program control resides. However, MP95 uses a single control variable π, whose value is the set of locations where control currently resides. STeP, on the other hand, uses a set of *control counters* where each control counter has its own subscript. We illustrate the differences between MP95 and STeP with an example.

Consider the program schema presented in Figure C.1, which lists the value of the control variable π (in italics) as it is used in MP95. In the figure, S_i refers to any simple

Class	In MP95	In STeP	Description
Simple	**skip**	`skip`	*skip*
	$u := e$	u `:=` e	*single assignment*
	$(\bar{u}) := (\bar{e})$	$(\bar{u})$ `:=` $(\bar{e})$	*multiple assignment*
	await e	`await` e	*await*
	$z \Leftarrow e$	z `<==` e	*send*
	$z \Rightarrow u$	z `==>` u	*receive*
	request r	`request` r	*request*
	release r	`release` r	*release*
Schematic	**noncritical**	`noncritical`	*noncritical*
	critical	`critical`	*critical*
	produce u	`produce` u	*produce*
	consume u	`consume` u	*consume*
	choose u	`choose` u	*choose*
Compound	**if** e **then** S_1	`if` e `then` S_1	*1-branch cond.*
	if e **then** S_1 **else** S_2	`if` e `then` S_1 `else` S_2	*2-branch cond.*
	$S_1; S_2$	S_1 `;` S_2	*concatenation*
	when e **do** S	`when` e `do` S	*when*
	S_1 **or** S_2	S_1 `or` S_2	*selection*
	while e **do** S	`while` e `do` S	*while loop*
	repeat S **until** e	`repeat` S `until` e	*repeat*
	loop forever do S	`loop forever do` S	*loop forever*
	$S_1 \,\|\|\, S_2$	S_1 `||` S_2	*cooperation*
	[*local decl*; S]	[*local decl*; S]	*block*
	$\mathop{\mathbf{OR}}_{j=1}^{M} P[j] :: S[j]$	`or` $P[j:[1..M]]::S[j]$	*param. selection*
	$\mathop{\mathbf{OR}}_{j=1}^{M} S[j]$	`or` $j:[1..M].S[j]$	*param. selection*
	$\mathop{\Big\|\|}_{j=1}^{M} P[j] :: S[j]$	`||` $P[j:[1..M]]::S[j]$	*param. cooperation*
	$\mathop{\Big\|\|}_{j=1}^{M} S[j]$	`||` $j:[1..M].S[j]$	*param. cooperation*
Grouped	$\langle S \rangle$	`<<` S `>>`	*group*

Table C.1: The statements of SPL

In MP95	In STeP	Description
T	`true`	*true*
F	`false`	*false*
integer	*integer*	*integer value*
u	u	*variable*
$u[e]$	u`[`e`]`	*array access*
(e)	`(`e`)`	*parenthesized expression*
$(\bar{e})$	`(`$\bar{e}$`)`	*tuple of expressions*
$e_1 = e_2$	e_1 `=` e_2	*equal*
$e_1 \neq e_2$	e_1 `!=` e_2	*not equal*
$e_1 < e_2$	e_1 `<` e_2	*less than*
$e_1 > e_2$	e_1 `>` e_2	*greater than*
$e_1 \leq e_2$	e_1 `<=` e_2	*less than or equal to*
$e_1 \geq e_2$	e_1 `>=` e_2	*greater than or equal to*
$-e$	`-`e	*unary minus*
e_1 **mod** e_2	e_1 `mod` e_2	*modulo*
e_1 **div** e_2	e_1 `div` e_2	*integer division*
$e_1 + e_2$	e_1 `+` e_2	*addition*
$e_1 - e_2$	e_1 `-` e_2	*subtraction*
$e_1 \cdot e_2$	e_1 `*` e_2	*multiplication*
e_1/e_2	e_1 `/` e_2	*rational division*
$\prod_{i=e_1}^{e_2} e$	`Prod` i`:[`e_1`..`e_2`].`e	*product*
$\sum_{i=e_1}^{e_2} e$	`Sum` i`:[`e_1`..`e_2`].`e	*summation*
$\neg e$	`~`e	*negation*
$\neg e$	`!`e	*negation*
$e_1 \wedge e_2$	e_1 `/\` e_2	*conjunction*
$e_1 \vee e_2$	e_1 `\/` e_2	*disjunction*
$e_1 \rightarrow e_2$	e_1 `-->` e_2	*implication*
$e_1 \leftrightarrow e_2$	e_1 `<-->` e_2	*equivalence*
$\forall\, u.\, e$	`Forall` $u.\ e$	*universal quantification*
$\exists\, u.\, e$	`Exists` $u.\ e$	*existential quantification*
$\exists!\, u.\, e$	`Exists!` $u.\ e$	*unique existential quant.*

Table C.2: SPL expressions

Class	In MP95	In STeP	Description

Base type	**integer**
	boolean
	rational
		`int`
		`bool`
		`rat`
			integer
			boolean
			rational

Simple type	[*range*] $t_1 \times \cdots \times t_k$ { *values* }	[*range*] t_1 `*` $\cdots$ `*` t_k { *values* }	*range type* *tuple type* *enumeration type*
Complex type	**array** [*ranges*] **of** T **channel of** b **channel** [1..] **of** b **channel** [*range*] **of** b (T)	`array` [*ranges*] `of` T `channel of` b `channel [1..] of` b `channel` [*range*] `of` b `(`T`)`	*array type* *synchronous* *unbounded asynch.* *bounded asynch.* *parenthesized type*

Table C.3: SPL types in STeP

or schematic statement (i.e., any statement that has no substatements). The STeP version of the same program schema is presented in Figure C.2, which lists all the control counters used (in italics) and their value.

The following highlights the similarities and differences between STeP's interpretation of control counters and the MP95 use of the control variable π:

1. STeP: The ordering of the labels does not affect the value of the control counters. For example, if we renamed label `l2` to be `l25`, the control counters would still have the same values. In general, the program counter *pi-j* associated with the process P_j starts at 0 and increases by 1 as control proceeds sequentially through the process.

   MP95: The ordering of the labels does affect the value of the control variable π. So, in the above example, if we rename ℓ_2 to ℓ_{25} then $\pi = \{\ell_2\}$ would change to $\pi = \{\ell_{25}\}$.

2. STeP: A cooperation statement among N processes employs N control counters. If the cooperation is a top-level cooperation statement (i.e., no entry and exit transitions) then the program has only N control counters. This is not the case in Figure C.2, where the cooperation statement has entry and exit transitions. In this case, we have two control counters, `pi1` and `pi2`, in addition to `pi0`, whose value within the cooperation statement is constant.

   MP95: There is always only one control variable, called π.

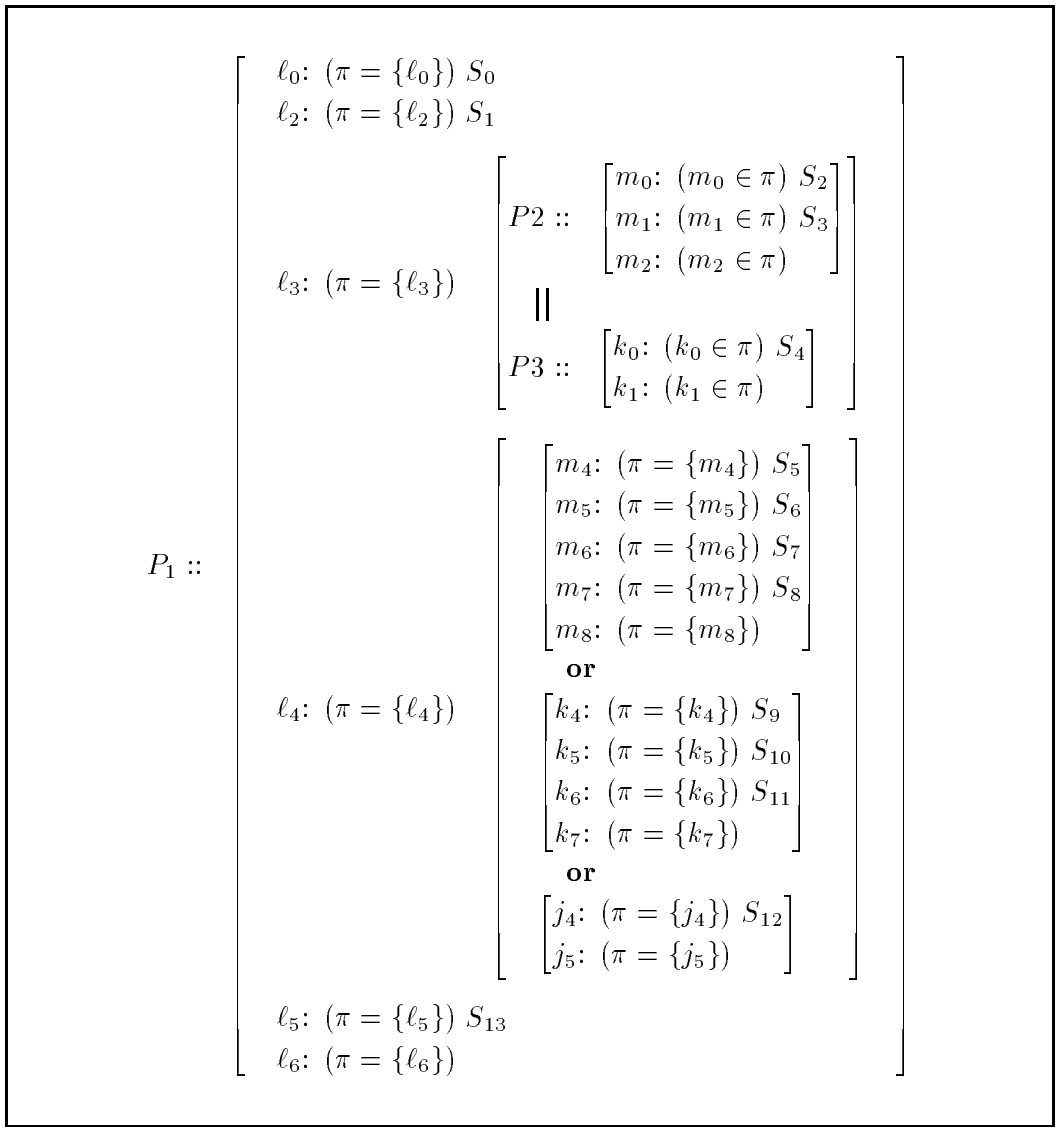

Figure C.1: Program schema illustrating MP95's use of the control variable π

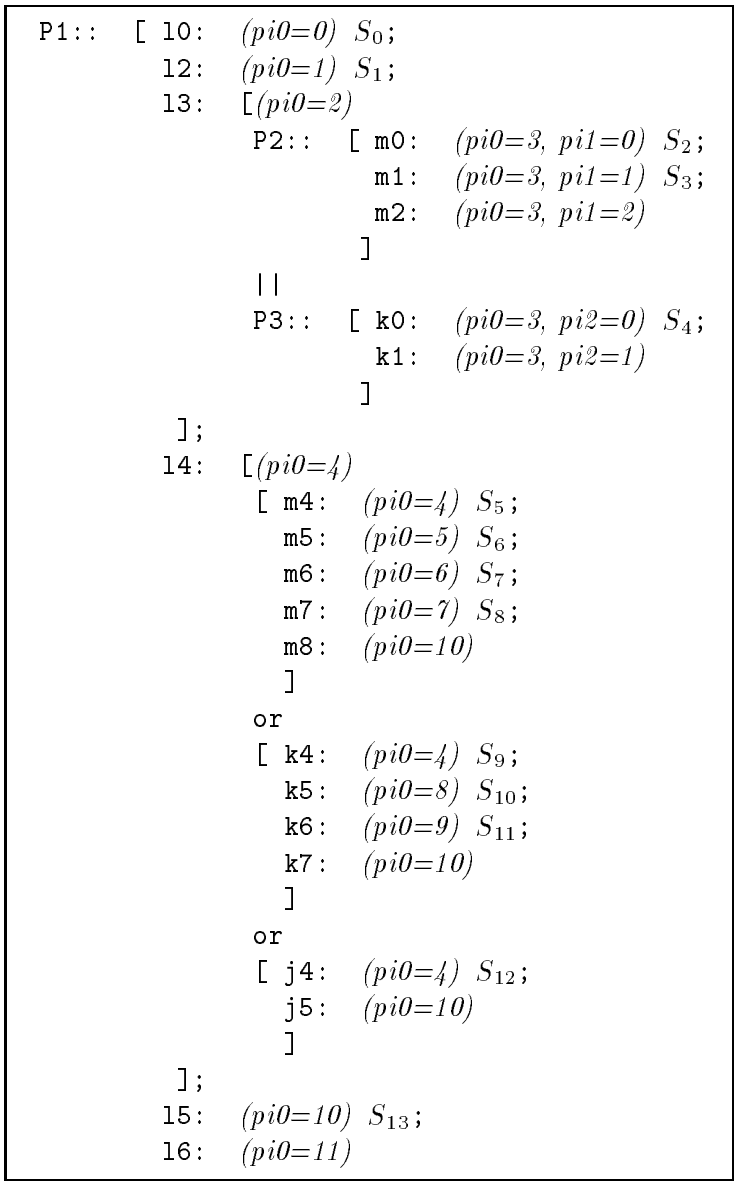

Figure C.2: Program schema illustrating STeP's use of control counters

3. In both STeP and MP95, the same rules govern the notion of label equivalence.

   STeP: This is achieved by assuming label-equivalent locations the same control-counter values.

   MP95: This is achieved by using the notational convenience that a label ℓ_i really refers to its (label) equivalence class $[\ell_i]$. Consequently, $\ell_4 \sim m_4 \sim k_4 \sim j_4$ and $\ell_5 \sim m_8 \sim k_7 \sim j_5$ (see page 13-14 in MP95).

Location equivalences

In STeP, a **loop forever do** statement does not generate a transition, so the pre-location and post-location of this statement are equivalent. However, both locations are entered as location variables in the symbol table and can thus be referred to in the specification.

C.3 Type Declarations

In several programs MP95 uses multiple lines of declarations as follows:

in	x, y, z:	**integer**
	a, b, c:	**bool**
local	d, e, f:	**integer**
	g, h, i:	**bool**

In STeP each line of declarations must be preceded by the mode of the variable, so the above has to be entered as:

```
in     x,y,z: integer
in     a,b,c: bool
local d,e,f: integer
local g,h,i: bool
```

C.4 Initialization

Array Initialization

In several programs MP95 uses the following construct to initialize an array:

in	M:	**integer**	**where** $M \geq 1$
	y:	**array**$[1..M]$ **of integer**	**where** $y = \mathrm{T}$

In STeP this causes an error due to the incompatible types of `y` and `T`. In STeP this has to be written as:

```
in M: int where M >= 1
in x: array[1..M] of int where Forall i:[1..M] . y[i] = true
```

Channel Initialization

In MP95 a channel may be initialized as the empty channel with

$$send = \Lambda$$

In STeP this is accomplished with

```
length(send) = 0
```

In MP95 a channel may be initialized with values as follows:

$$ack = \underbrace{[1 \ldots 1]}_{N}$$

STeP does not provide a construct to initialize a channel in the **where** clause. If the initial values matter, the channel has to be initialized explicitly within the program with, for example, a `while` statement. If the initial values do not matter, which is the case for program PROD-CONS-C of Figure 1.18 in MP95, the channel may be initialized with the following `where` clause:

```
in      N :int where N > 0
local   send, ack: channel [1..] of int where length(ack) = N /\
                                              length(send) = 0
```

C.5 Parameterized Programs

In MP95 (page 174), terms such as N_3 are often used to denote the number of processes residing at a particular location. This abbreviation is not available in STeP, so there are two ways to express such formulas: In a formula of the form

$$N_3 \leq 1$$

(i.e., the number of processes at location 3 is less than 1), the recommended STeP representation is the following:

```
Forall i:[1..N] . Forall j:[1..N] . (l3[i] /\ l3[j] --> i=j)
```

This may be generalized to small numbers greater than 1 using disjunction in the consequent. When this representation is not possible, e.g., in a formula of the form

$$N_3 + y = z$$

you can use `Sum` as follows:

```
(Sum i:[1..N] . l3[i]) + y = z
```

We discourage the use of `Sum` for the simpler cases because reasoning with `Sum` is more difficult than reasoning with quantifiers only.

Appendix D

Distribution and Installation

Installation

STeP is available for platforms that are supported by Standard ML of New Jersey version 1.08, including:

- Sun SPARC workstations (SunOS and Solaris)
- DEC Alpha workstations (OSF1)

You should have copied the distribution file specific to your machine. The only difference is in the binaries included with each.

Files associated with running STeP are located in the `bin` directory. The executable files are found in the `.bin` and `.heap` directories, and accessed through the `STeP`, `polyserv`, and `model_checker` scripts. These scripts determine the machine type and operating system, using the `.arch-n-opsys` command, and then invoke the appropriate binary file (from the `.run` or `.bin` directory) on the given arguments. Thus, none of these files should be moved or renamed.

You can edit the `STeP` script to set the `STEP_DIR` environment variable once and for all.

Note: These scripts use the `/bin/sh` shell; if this shell is not found, you should edit them and change the first line to indicate the appropriate location (actually, `ksh` is preferred).

STeP Distribution

When the distribution has been untarred it will have created three directories:

- The `bin` directory containing binaries `step`, `model_checker`, and `polyserv` for executing STeP, and from within it the model checker and utilities for generating polyhedral invariants.
- The `doc` directory. It contains:
    1. The STeP manual (postscript).

2. The STeP Technical Report (postscript), STAN-CS-TR-94-1518, Stanford Computer Science Department, June 1994. This report complements the STeP manual by giving a high-level description of the system and its goals and some extra technical details.
3. The list of known problems and bugs `known-bugs.txt`.
4. A quick installation guide `installation.txt`.
5. The text of the licence agreement `license-agreement.txt`.
6. The `WWW` subdirectory, which contains the help pages.

- The `examples` directory containing a selection of simple example programs and specifications that can be verified using STeP.

New releases will include updated help pages; please let us know how those pages, and this manual, can be improved.

D.1 Example Programs

The examples directory adheres to the default filename extensions used by the File Browser. These conventions are shown in Table D.1.

Filename extension	contents of file
`.spl`	SPL program
`.fts`	fair transition system
`.cts`	clocked transition system
`.hts`	hybrid transition system
`.spec`	specification
`.search`	proof search
`.diagram`	verification diagram
`.tactic`	tactic
`.mclog`	modelchecker logfile

Table D.1: Filename conventions

The example programs and specification files listed in Tables D.2 and D.3 are included in the distribution. They correspond to simple examples and exercises from [Manna and Pnueli, 1995b]. In most cases the reader is expected to complete the specifications for the exercises. Thus, possible answers to the exercises are not included in the distribution. The directory `mux-fisher` contains a simple real-time protocol implemented as a Clocked Transition System.

	directory	program file	specification files	other files
ADD-TWO	add-two	add-two.spl	yge0.spec l1x=2.spec	
ANY-NAT	any-nat	any-nat.spl		
BINOM	binom	binom.spl	part-corr.spec str-part-corr.spec	
BINOM-C	binom	binom-c.spl	part-corr.spec str-part-corr.spec	
CUBE	cube	cube.spl	part-corr.spec	
DIFF-INC	diff-inc	diff-inc.spl	psi.spec	
DINE DINE-CONTR DINE-EXCL	dine dine-contr dine-excl	dine.spl dine-contr.spl dine-excl.spl		
DOUBLE	double	double.spl		
EUCLID	euclid	euclid.fts		
FACT	fact	fact.spl	part-corr.spec	
GCD GCDM	gcd gcdm	gcd.spl gcdm.spl	part-corr.spec part-corr.spec	
INC	inc	inc.spl	psi.spec	psi.search
LOOP	loop	loop.fts	phi3.spec psi3.spec	
LOOP+	loop	loop+.fts	phi3.spec psi3.spec	

Table D.2: Example programs and specifications

	directory	program file	specification files	other files
MAX-ARRAY	max-array	max-array.spl	part-corr.spec	phi3.search phi4.search phi5.search
MPX-SEM	mpx-sem	mpx-sem.spl		
MUX-BAK-A	mux-bak-a	mux-bak-a.spl	mux.spec	
MUX-BAK-C	mux-bak-c	mux-bak-c.spl		
MUX-BAK-D	mux-bak-d	mux-bak-d.spl		
MUX-DEK	mux-dek	mux-dek.spl	mux.spec	
MUX-DEK-A	mux-dek-a	mux-dek-a.spl	mux.spec	
MUX-DEK-B	mux-dek-b	mux-dek-b.spl	mux.spec	
MUX-FISHER	mux-fisher	mux-fisher.cts	mux-fisher.spec	
MUX-PET1	mux-pet1	mux-pet1.spl	mux.spec mux-vd.spec simple-prec.spec 1bo.spec 1bo-vd.spec	mux.tactic mux.diagram 1bo.diagram
MUX-PET2	mux-pet2	mux-pet2.spl	mux.spec mux-vd.spec 1bo-back.spec 1bo.spec 1bo-vd.spec	mux.tactic mux.diagram 1bo.diagram
MUX-PET3	mux-pet3	mux-pet3.spl	mux.spec accessibility.spec 1bo.spec	mux.tactic accessibility.diagram
MUX-SEM	mux-sem	mux-sem.spl	mux.spec yge0.spec	mux.search state-partition.diagram
MUX-VAL3	mux-val3	mux-val3.spl	mux.spec	mux.diagram mux-compound.diagram
PROD-CONS	prod-cons	prod-cons.spl	no-overflow.spec	
PROD-CONS-C	prod-cons	prod-cons-c.spl	no-overflow-c.spec	
PROD-CONS-SV	prod-cons-sv	prod-cons-sv.spl	part-corr.spec	
RES-ALLOC	res-alloc	res-alloc.spl		
RES-ND	res-nd	res-nd.spl	mux.spec	
RES3	res3	res3.spl	mux.spec	
SEM3	sem3	sem3.spl	mux.spec	
SQRT	sqrt	sqrt.spl	phi-hat.spec	
SQRT-C	sqrt	sqrt-c.spl	str-part-corr.spec	

Table D.3: More example programs and specifications.

Appendix E

STeP Settings

STeP depends on two mandatory environment variables, `STEP_DIR`, which indicates the directory where STeP resides, and the `DISPLAY` environment variable, which should be set to the IP-address of your screen. If these variables are not set, you will not be able to run STeP.

To set a UNIX environment variable, enter

`setenv` *variable-name new-value*

from the UNIX prompt. You can also add these commands to your `.cshrc` file, to make them the defaults when you log in.

A number of other setttings can be chosen before running STeP by creating a `.step` file in your home directory. If there is no `.step` file the first time STeP is invoked, it will create one automatically with a number of default settings that can be edited later. These settings are described in more detail below.

General Settings

- `STEP_LARGE_FONT`: The font used to display buttons and menues.
  Default: `*helvetica-*-r-*-14-*`.

- `STEP_SMALL_FONT`: The font used to display text in the action regions
  Default: `*helvetica-*-r-*-12-*`.

- `STEP_BROWSER`: The location of the `WWW/http` browser executable used for on-line help.
  Default `/local/bin/netscape`.

- `STEP_AUTO_SIMPLIFY`: If this flag is set to anything other than "OFF," this will cause STeP to start with the automatic simplification flag (see Section 4.1) ON.

- `STEP_SHOW_LOADED`: Set this to "OFF," if you do not want loaded programs and transition systems to be automatically displayed.

Simplification Setttings

- `STEP_SIMPLIFY_EQUALITY`: Use equality reasoning in the simplifier.
  Default: `true`.

- `STEP_SIMPLIFY_ARITHMETIC`: Use decision procedures for linear arithmetic in the simplifier.
  Default: `false`.

- `STEP_SIMPLIFY_AC_FUNCTIONS`: Use information about associative/commutative function symbols while simplifying.
  Default: `false`.

- `STEP_SIMPLIFY_BDD`: Detect propositional validities using BDDs.
  Default: `false`.

- `STEP_SIMPLIFY_CONDITIONAL`: Use conditional rewrite rules in the simplification.
  Default: `false`.

Verification Diagram settings

The following values can be increased when very large diagrams need to be drawn.

- `STEP_DIAGRAM_WIDTH`: Maximum width, in pixels, of the verification diagram canvas.
  Default: 1200.

- `STEP_DIAGRAM_HEIGHT`: Maximum height, in pixels, of the verification diagram canvas.
  Default: 1200.

BDD-package settings

These variables control the amount of memory used by the BDD's.

- `STEP_BDD_NODES`: Maximum number of BDD nodes that are kept around by the BDD package. After this maximum is reached, the BDD packages has to be reset for BDD-based operations to succeed.
  Default: 100,000.

- `STEP_BDD_SPLIT`: Maximum number of subgoals that are generated by the BDD-split rule in the Top-level Prover and Interactive Prover. If this maximum is reached, the BDD-split rule fails. (BDD split can generate exponentially many subgoals in the worst case.)
  Default: 1,000.

- `STEP_BDD_CACHE`: Size of the BDD result cache, used for efficient BDD operations.
  Default: 32,000.

Model checker and polyhedral invariant resources

The following variables limit the space and time used by the external Model Checker and polyhedral invariant packages. The _MC_ variables control the Model Checker, and the _POLY_ variables control the polyhedral invariant generator. Both of them will fail if any of the limits is exceeded.

- STEP_MC_SPACE, STEP_POLY_SPACE: Maximum amount of memory, in Megabytes. Default: unlimited.
- STEP_MC_TIME, STEP_POLY_TIME: Maximum amount of user time, in minutes. Default: unlimited.
- STEP_MC_CPU, STEP_POLY_CPU: Maximum amount of CPU time, in minutes. Default: unlimited.

Finally, remember that the *interrupt button* can be used to interrupt lengthy computations in STeP.

Bibliography

[Alur *et al.*, 1993] R. Alur, C. Courcoubetis, T.A. Henzinger, and P.-H. Ho. Hybrid automata: An algorithmic approach to the specification and analysis of hybrid systems. In R.L. Grossman, A. Nerode, A. Ravn, and H. Rischel, editors, *Hybrid Systems*, volume 736 of *LNCS*, pages 209–229. Springer-Verlag, 1993.

[Bjørner *et al.*, 1997a] N.S. Bjørner, A. Browne, and Z. Manna. Automatic generation of invariants and intermediate assertions. *Theor. Comp. Sci.*, 173(1):49–87, February 1997. Preliminary version appeared in 1st *Intl. Conf. on Principles and Practice of Constraint Programming*, vol. 976 of LNCS, pp. 589–623, Springer-Verlag, 1995.

[Bjørner *et al.*, 1997b] N.S. Bjørner, Z. Manna, H.B. Sipma, and T.E. Uribe. Deductive verification of real-time systems using STeP. In *4th Intl. AMAST Workshop on Real-Time Systems*, volume 1231 of *LNCS*, pages 22–43. Springer-Verlag, May 1997.

[Bjørner *et al.*, 1997c] N.S. Bjørner, M.E. Stickel, and T.E. Uribe. A practical combination of first-order reasoning and decision procedures. In *14th Intl. Conference on Automated Deduction*, volume 1249 of *LNCS*, pages 101–115. Springer-Verlag, July 1997.

[Bledsoe, 1975] W. W. Bledsoe. A new method for proving certain Presburger formulas. In *Proc. of the* 4^{th} *Intl. Joint Conference on Artificial Intelligence*, pages 15–21, Tbilisi, Georgia, USSR, September 1975.

[Boyer and Moore, 1988] R.S. Boyer and J S. Moore. Integrating decision procedures into heuristic theorem provers: A case study with linear arithmetic. *Machine Intelligence*, 11:83–124, 1988.

[Bryant, 1986] R.E. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, August 1986.

[Bryant, 1992] R.E. Bryant. Symbolic boolean manipulation with ordered binary-decision diagrams. *ACM Computing Surveys*, 24(3):293–318, September 1992.

[Dershowitz and Jouannaud, 1990] N. Dershowitz and J-P. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, chapter 6, pages 244–320. Elsevier Science Publishers, 1990.

[Gallier, 1987] J.H. Gallier. *Logic for Computer Science—Foundations for Automatic Theorem Proving*. Wiley, New York, 1987.

[Harel, 1987] D. Harel. Statecharts: A visual formalism for complex systems. *Sci. Comp. Prog.*, 8:231–274, 1987.

[Hojati *et al.*, 1993] R. Hojati, V. Singhal, and R.K. Brayton. Edge-Street/Edge-Rabin automata environment for formal verification using language containment. SRC report, University of California, Berkeley, 1993.

[Kesten *et al.*, 1996] Y. Kesten, Z. Manna, and A. Pnueli. Verifying clocked transition systems. In *Hybrid Systems III*, volume 1066 of *LNCS*, pages 13–40. Springer-Verlag, 1996.

[Manna and Pnueli, 1991] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification.* Springer-Verlag, New York, 1991.

[Manna and Pnueli, 1994] Z. Manna and A. Pnueli. Temporal verification diagrams. In *Proc. Int. Symp. on Theoretical Aspects of Computer Software*, volume 789 of *LNCS*, pages 726–765. Springer-Verlag, 1994.

[Manna and Pnueli, 1995a] Z. Manna and A. Pnueli. Clocked transition systems. In *Logic and Software Workshop*, August 1995. Beijing, China.

[Manna and Pnueli, 1995b] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety.* Springer-Verlag, New York, 1995.

[Manna and Waldinger, 1993] Z. Manna and R. Waldinger. *The Deductive Foundations of Computer Programming.* Addison-Wesley, Reading, MA, 1993.

[McCune, 1992] W.W. McCune. Experiments with discrimination-tree indexing and path indexing for term retrieval. *J. Automated Reasoning*, 9(2):147–67, October 1992.

[Nelson and Oppen, 1980] G. Nelson and D.C. Oppen. Fast decision procedures based on congruence closure. *J. ACM*, 27(2):356–364, April 1980.

[Shostak, 1979] R.E. Shostak. A practical decision procedure for arithmetic with function symbols. *J. ACM*, 26(2):351–360, April 1979.

[Stickel, 1989] M.E. Stickel. The path indexing method for indexing terms. Technical Note 473, SRI International, Menlo Park, CA, October 1989.

[Zhang, 1993] H. Zhang. Contextual rewriting in automated reasoning. Technical Report Technical Report 93-07, Department of Computer Science, University of Iowa, August 1993.

Index

Θ [initial condition], 61
τ-successor, 115
$move(\ell, \widehat{\ell})$, 117
$pres(U)$, 117
$\mathcal{T}$ [set of transitions], 61
* [type constructor], 13
1-Step-Propositional [ip-button], 94

abandon search [menu option], 50
AC: associative and commutative, 22
accessible state, 124
action region (top-level prover), 43
activate/deactivate [menu option], 49
activity, 38
activity (in hybrid system), 37
activity relation, 38
Add-axiom [ip-button], 99
All-Propositional [ip-button], 98
append [syntax], 16, 25, 30
arithmetical operators [syntax], 16
arithmetization, 14, 41
array [type declaration], 13
array initialization
    difference with book, 131
array reference [syntax], 16
array update [syntax], 16
arriving edge, 70
assertion, 123
assign [syntax], 16, 30, 35
assign field, 31
assignable expression, 29
assignment [statement]
    semantics, 117
assignment relation, 31
ASSOCIATIVE, 22
associative functions
    simplification setting, 78
asynchronous receive [statement]
    semantics, 118
asynchronous send [statement]
    semantics, 117
atomic tactic, 80
automatic simplification, 60
automatic simplification [menu option], 55
auxiliary variable, 10
auxiliary variables, 10, 23
await [statement]
    semantics, 117
await [syntax], 28
Awaits [operator], 15
Awaits [syntax], 16
axiom, 9, 18
    entering, 53
    in interactive prover, 99

B-BACKTO [tlp-button], 62
B-CAUS [tlp-button], 62
B-INV [tlp-button], 61
    tutorial, 103
B-WAIT [tlp-button], 61
back-to rule
    basic: B-BACKTO, 62
    general: B-BACKTO, 62
background properties, 59
    activating, 49
    deactivating, 49
    deselecting, 49
    selecting, 49
background property, 9
Backto [operator], 15
Backto [syntax], 16
basic back-to rule, 62
basic causality rule, 62
basic invariance rule, 61

basic node, 70, 72
basic type, 13
basic wait-for rule, 61
BDD
    settings, 138
    simplification setting, 78
BDD-split, 80
BDD-split [ip-button], 99
binding operators, 17
Bool-simplify [top-level button], 79
bound variables, 15
bound variables [syntax], 16
branch-otherwise [tactic], 83

causality monotonicity rule, 64
causality rule
    basic: B-CAUS, 62
    general: G-CAUS, 62
chain diagram, 69
channel [type declaration], 13
channel initialization
    difference with book, 132
channel operations, 25
check runtime system [menu option], 53
Check-Valid [ip-button], 94
Check-Valid [top-level button], 79
clear text window [menu option], 57
clock [syntax], 35
clock variable, 23, 37
    master, 38
clocked transition system, 36
    syntax, 35
command-line options, 4
comments, 13
COMMUTATIVE, 22
commutative functions
    simplification setting, 78
comparison operators [syntax], 16
Compassionate [syntax], 35
compassionate edge, 73
compassionate transition, 115
Compound [menu option], 74
compound node, 70
computation, 116
computational model, 115
conditional rewrite rules
    simplification setting, 78
    user-defined, 24
congruent formulas, 124
constructor, 20
consume [statement]
    semantics, 119
consume [syntax], 28
CONT [menu option], 64
continuous variable, 37
contradiction rule, 64
control locations, 125
critical [statement]
    semantics, 119
critical [syntax], 28
Cut [ip-button], 97

datatype declaration, 20
decision procedures
    simplification setting, 78
declarations, 19
    datatype, 20
    in specification file, 19
    macro, 22
    type, 19
    value, 22
    variable, 22
deconstructor, 20
default settings, 137
define field, 31
Delete [ip-button], 94
departing edge, 70
Deriv, 38
Deriv [syntax], 16
derivative, 38
diagrams menu, 55
differential inclusion, 38
dining philosophers, 26
disabled transition, 115
div [operator], 15
div [syntax], 16
documentation, 134
double edge, 69
Duplicate [ip-button], 94

edit diagram [menu option], 55

else [tactic], 82
enable [syntax], 35
enable field, 31
enabled transition, 115
entailment (temporal operator), 124
enter axiom [menu option], 53
enter batch tactic [menu option], 55
enter interactive tactic [menu option], 55
enter new goal [menu option], 53
enumeration type, 20
environment variables
    model checker, 139
    polyhedral invariants, 139
equality, 77
    simplification setting, 77
equational axiom, 24
equivalent formulas, 124
Euclid's algorithm, 32
example
    Euclid's algorithm, 32
    MUX-SEM, 102
Exists [syntax], 16, 30
expressions, 14
    in SPL programs, 29
    syntax in specifications, 15

fair transition system, 10
    definition, 115
fairness requirement
    in transition, 32
Fast-simplify [top-level button], 79
feedback on STeP, 11
file menu [top-level prover], 45
filebrowser, 45
filename conventions, 134
first (temporal logic abbreviation), 95
Flatten [ip-button], 94
flexible variable, 10, 24
Forall [syntax], 16, 30
forest [example], 20
formulas
    P-state valid, 124
    congruent, 124
    state-valid, 124
Free-Induction [ip-button], 98
Future temporal operators, 123

G-BACKTO [tlp-button], 62
G-CAUS [tlp-button], 62
G-INV [tlp-button], 61
    tutorial, 103
G-WAIT [tlp-button], 61
general back-to rule, 62
general causality rule, 62
general invariance rule, 61
general wait-for rule, 61
get [...] invariants [menu option], 53
global time, 37
goal, 9
    entering, 53
    session, 9
Group [menu option], 74
grouped statements, 26
guard [statement]
    semantics, 119
guard [syntax], 28
guarded assignment, 119

head [syntax], 16, 25, 30
help, 102
help menu [top-level menu], 57
helpful transition, 69
Hide [ip-button], 93
hierarchical verification diagrams, 75
    edges, 76
    parameter passing, 76
hybrid system, 37
hybrid transition system, 37

idling transition, 115
if-then [syntax], 28
if-then-else [syntax], 28
import node, 75
in [mode], 26
in [syntax], 29, 35
Induction [ip-button], 98
inequality
    simplification setting, 77
initial condition, 115, 116
Initially [syntax], 35
installing STeP, 133

Instantiate [ip-button], 96
Integrate, 17
Integrate [syntax], 16
Interactive Prover
  rules, 93
  tlp-button, 58
Interactive Prover button
  1-Step-Propositional, 94
  Add-axiom, 99
  All-Propositional, 98
  BDD-split, 99
  Check-Valid, 94
  Cut, 97
  Delete, 94
  Duplicate, 94
  Flatten, 94
  Free-Induction, 98
  Hide, 93
  Induction, 98
  Instantiate, 96
  Make first order, 97
  Next, 97
  Postpone, 93
  Presburger, 98
  PTL-expansion, 99
  Redo, 93
  Replace, 97
  Rewrite, 97
  Simplify, 94
  Skolemize, 96
  Undo, 93
  Unhide, 93
internal node, 70
interrupt button, 85, 139
  invariant generation, 87
  model checker, 66
  simplification, 78
  tactics, 54
invariance diagram, 68
invariance monotonicity rule, 63
invariance rule
  basic: B-INV, 61
  general: G-INV, 61
invariants, 9
  obtaining linear, 53
  tutorial, 109
  obtaining local, 53
  polyhedral, 53

just edge, 73
just transition, 115
Justice [syntax], 35

Lambda, 17
Lambda [syntax], 16
length [syntax], 16, 25, 30
linear invariants [menu option], 53
load batch tactic [menu option], 55
load interactive tactic [menu option], 55
load program [menu option], 46
load proof state [menu option], 48
load specification [menu option], 48
load transitions [menu option], 46
loading
  batch tactic, 55
  interactive tactic, 55
  specification, 48, 101
  SPL programs, 46, 101
  tactic, 55
  transition system, 46
local [mode], 26
local [syntax], 29, 35
local invariants [menu option], 53
location disjunction [syntax], 16
location equivalence, 131
logical rules menu, 55
loop forever [syntax], 28

macro declaration, 22
mailing list, 11
Make first order [ip-button], 97
master clock, 36, 38
menu option
  abandon search, 50
  activate/deactivate, 49
  automatic simplification, 55
  check runtime system, 53
  clear text window, 57
  CONT, 64
  edit diagram, 55
  enter axiom, 53

enter batch tactic, 55
enter interactive tactic, 55
enter new goal, 53
get [...] invariants, 53
load batch tactic, 55
load interactive tactic, 55
load program, 46
load proof state, 48
load specification, 48
load transitions, 46
MON-C, 64
MON-I, 63
MON-W, 64
next search, 49
Pre/Post-conditions, 54
quit, 49
reset all, 48
reset BDDs, 57
reset searches, 48
save goal, 48
save properties, 48
save search, 48
save transitions, 48
select axiom, 49
select boxes, 92
select goal, 49
select search, 50
simplification flags, 55
system information, 57
TRN-C, 64
verification diagram rule, 55
view program text, 49
view transitions, 49
mod [operator], 15
mod [syntax], 16
mode (in declaration), 131
mode of variables, 26
model (of temporal formula), 123
model checker, 64
environment variables, 139
tutorial, 107
Modelcheck [tlp-button], 58
modifying relation, 31
modrel
[syntax], 35
field, 31
modvar field, 30
MON-C [menu option], 64
MON-I [menu option], 63
MON-W [menu option], 64
monotonicity rule
causality, 64
invariance, 63
wait-for, 64
MUX-SEM (mutual exclusion), 102

Next [ip-button], 97
next [tlp-button], 57
next search [menu option], 49
NoFairness [syntax], 35
noncritical [statement]
semantics, 119
noncritical [syntax], 28
Nonterminal [menu option], 74
nonterminal node, 68

operator associativity, 17
operator precedence, 18
ORDER [syntax], 24
ordering relation, 24
out [mode], 26
out [syntax], 29, 35

P-valid
formula, 124
verification diagram, 68
parameter
in hierarchical verification diagram, 76
parameterized programs, 26
parser help, 39
past formula, 10
past operators, 10
Past temporal operators, 123
polyhedral invariants
environment variables, 139
menu option, 53
Postpone [ip-button], 93
Pre/Post-conditions [menu option], 54
precedence, 17, 18
Presburger [ip-button], 98
Presburger decision procedure, 98

previous [tlp-button], 57
primed variable, 17, 115
Prod [operator], 29
produce [statement]
    semantics, 119
produce [syntax], 28
Product [syntax], 16
program
    loading, 101
    SPL, 26
    syntax, 26
program variables, 26
program verification session
    terminating, 48
Progress [syntax], 35
progress condition, 37
proof search, 59
proof tree
    in Interactive Prover, 89
    navigating, 60
properties menu, 49
PTL-expansion [ip-button], 99

quantifiers, 17
quit [menu option], 49
quitting STeP, 102

range [type declaration], 13
real-time systems, 36
receive statement [syntax], 28
record projection, 17
Redo [ip-button], 93
redo [tlp-button], 57
release [statement]
    semantics, 118
release [syntax], 28
repeat (SPL construct) [syntax], 28
repeat [tactic], 82
Replace [ip-button], 97
request [statement]
    semantics, 118
request [syntax], 28
reset all [menu option], 48
reset BDDs [menu option], 57
reset searches [menu option], 48
Rewrite [ip-button], 97
REWRITE [syntax], 24
rewrite rule, 24
rigid variable, 24
runtime errors, 53

save goal [menu option], 48
save properties [menu option], 48
save search [menu option], 48
save transitions [menu option], 48
saving
    goal, 48
    proof, 48, 103
    properties, 48
    transitions, 48
select axiom [menu option], 49
select boxes [menu option], 92
select goal [menu option], 49
select search [menu option], 50
send statement [syntax], 28
sequence [tactic], 82
settings
    BDD settings, 138
    simplification, 138
    verification diagram, 138
settings menu [top-level prover], 55
side verification condition, 68
simplification
    automatic, 60
    use of background properties, 59
simplification flags [menu option], 55
simplification rule, 24
simplification rule (user-defined), 24
simplification settings, 138
    associative functions, 78
    BDD simplification, 78
    commutative functions, 78
    conditional rewrite rules, 78
    decision procedures, 78
    equality, 77
    inequality, 77
Simplify [ip-button], 94
SIMPLIFY [syntax], 24
Simplify [top-level button], 79
Since [operator], 15
Since [syntax], 16

skip [statement]
  semantics, 117
  syntax, 28
Skolemize [ip-button], 96
solid edge, 69
specification, 10, 18
  declarations, 19
  loading, 48, 101
SPL basic statements, 117
SPL composite statements, 119
SPL expressions, 29
SPL programs, 10
  loading, 46, 101
  syntax, 26
spread [tactic], 83
stand-alone STeP components, 4
starting STeP, 3
state formula (assertion), 123
state, accessible, 124
state-valid formulas, 124
Statecharts, 70
STeP
  distribution, 133
  documentation, 134
  installation, 133
  mailing list, 11
  session, 8
  stand-alone components, 4
STEP_AUTO_SIMPLIFY, 137
STEP_BDD_CACHE, 138
STEP_BDD_NODES, 138
STEP_BDD_SPLIT, 138
STEP_BROWSER, 137
STEP_DIAGRAM_HEIGHT, 138
STEP_DIAGRAM_WIDTH, 138
STEP_DIR, 3
STEP_LARGE_FONT, 137
STEP_MCTIME, 66
STEP_MC_CPU, 66, 139
STEP_MC_SPACE, 66, 139
STEP_MC_TIME, 139
STEP_POLY_CPU, 139
STEP_POLY_SPACE, 139
STEP_POLY_TIME, 139
STEP_SHOW_LOADED, 137
STEP_SMALL_FONT, 137
Strengthen [tlp-button], 63
strengthening rule, 63
stuttering transition, 115
subgoal, 9
Sum [operator], 29
Sum [syntax], 16
synchronous send-receive [statement]
  semantics, 118
system description, 10
system information [menu option], 57
system variables, 10, 115, 116
system verification session, 8

tactic, 80
  atomic, 80
  branch-otherwise, 83
  else, 82
  entering, 55
  examples, 86
  loading, 55
  repeat, 82
  sequence, 82
  spread, 83
  try, 81
tactics menu, 54
tail [syntax], 16, 25, 30
temporal operators
  future, 123
  past, 123
temporal operators [syntax], 16
Terminal [menu option], 74
terminal node, 68
theorem-proving session, 8
  terminating, 48
tick, 37
top-level prover, 43
  action region, 43
  current goal window, 43
  help, 43
  interrupt button, 45
  output window, 43
  status line, 43
top-level prover button
  B-BACKTO, 62

B-CAUS, 62
B-INV, 61
B-WAIT, 61
G-BACKTO, 62
G-CAUS, 62
G-INV, 61
G-WAIT, 61
Interactive prover, 58
Modelcheck, 58
next, 57
previous, 57
redo, 57
Strengthen, 63
undo, 57
WPC, 63
transition, 31, 115
compassionate, 115
disabled, 115
enabled, 115
Idling, 115
just, 115
Transition [syntax], 35
transition relation, 30, 115
transition system, 30
hybrid, 37
loading, 46
syntax, 35
transitivity rule, 64
tree [example], 20
TRN-C [menu option], 64
try [tactic], 81
tuple projection, 16, 17
type constructor, 13
type declaration, 13, 19
types, 13

Uncompound [menu option], 74
Undo [ip-button], 93
undo [tlp-button], 57
Ungroup [menu option], 74
Unhide [ip-button], 93
UNIX environment variables, 137
unprimed variable, 115
Until [operator], 15
Until [syntax], 16
update [example], 16
update [syntax], 16

validity
P-, 124
general, 8
state-, 124
value declaration, 22
variables
auxiliary, 23
bound, 15
declaration, 19, 22
flexible, 10, 24
primed and unprimed, 17, 31, 38, 115
quantified, 24
rigid, 10, 24
verification condition, 9
side- (in diagram), 68
verification diagram, 67
basic node, 72
definition, 67
double edge, 69
hierarchical, 75
import node, 75
invariance, 68
nonterminal node, 68
P-valid, 68
response, 69
settings, 138
solid edge, 69
terminal node, 73
tutorial, 108
verification rule, 55
wait-for, 68
verification diagram editor, 67
activating, 55
drawing tools, 72
Edge menu, 74
Edit menu, 74
File menu, 73
Help menu, 75
interface, 72
menu options, 73
Node menu, 74
verification diagram rule [menu option], 55

verification rules, 60
    invoking, 57
view program text [menu option], 49
view transitions [menu option], 49

wait-for diagram, 68
wait-for monotonicity rule, 64
wait-for rule
    basic: B-WAIT, 61
    general: G-WAIT, 61
water level example, 39
weakest precondition [rule], 63
when [syntax], 28
where clause, 116
while [syntax], 28
WPC [tlp-button], 63

X-windows, 3
xhost (UNIX command), 3