
**Verification of Parallel Programs with
the Owicki-Gries and Rely-Guarantee
Methods in Isabelle/HOL**

Leonor Prensa Nieto
Institut für Informatik
Technische Universität München

Institut für Informatik
der Technischen Universität München
Lehrstuhl für Software & Systems Engineering

**Verification of Parallel Programs with
the Owicki-Gries and Rely-Guarantee
Methods in Isabelle/HOL**

Leonor Prensa Nieto

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen
Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. Manfred Broy

Prüfer der Dissertation:

1. Univ.-Prof. Tobias Nipkow, Ph. D.
2. Univ.-Prof. Dr. Javier Esparza

Die Dissertation wurde am 31. Oktober 2001 bei der Technischen Univer-
sität München eingereicht und durch die Fakultät für Informatik am 31.
Januar 2002 angenommen.

Abstract

This thesis presents the first formalization of the Owicki-Gries method and its compositional version, the rely-guarantee method, in a theorem prover. These methods are widely used for correctness proofs of parallel imperative programs with shared variables. We define syntax, semantics and proof rules in Isabelle/HOL, which is the instantiation of higher-order logic in the theorem prover Isabelle. The proof rules also provide for programs parameterized in the number of parallel components. Their correctness w.r.t. the semantics is proven mechanically and the completeness proofs for both methods are extended to the new case of parameterized programs. For the automatic generation of verification conditions we define a tactic based on the proof rules. Using this tactic we verify several non-trivial examples for parameterized and non-parameterized programs.

Zusammenfassung

In dieser Arbeit wird die Owicki-Gries Methode, und ihre kompositionelle Version, die Rely-Guarantee Methode, zur Verifikation paralleler imperativer Programme mit gemeinsamen Variablen zum ersten Mal in einem Theorembeweiser formalisiert. Syntax, Semantik und Beweisregeln werden in höherstufiger Logik definiert und die Korrektheit des Beweissystems bezüglich der Semantik wird bewiesen. Zahlreiche Beispiele, darunter parametrisierte parallele Programme, werden mit Hilfe einer Taktik für die systematische Generierung der Verifikations-Bedingungen verifiziert. Außerdem wird die Vollständigkeit der formalisierten Systeme für den Fall parametrisierter paralleler Programme bewiesen.

Acknowledgements

First of all I wish to thank Tobias Nipkow and Javier Esparza for giving me the opportunity to stay in Germany and work on my Ph.D. under their supervision. During four years they have continuously given direction to my work, motivated me and given me advice whenever I had problems.

Special thanks also to David von Oheimb, Markus Wenzel and Cornelia Pusch who helped me get started with the Isabelle system and were always ready to help me solve my daily problems.

Further, I thank the (not yet mentioned) members and guests of the Isabelle working group and office mates Christine Röckl, John Harrison, Wolfgang Naraschewski, Stefan Berghofer, Bernd Grobauer, Gertrud Bauer, Gerwin Klein, Sebastian Skalberg, Martin Strecker, Norbert Schirmer, Raman Ramanujam and Giampaolo Bella as well as Manfred Broy and the remaining colleagues and staff members of our chair for the wonderful working atmosphere.

Very special thanks to the Ph.D. Program “Logic in Computer Science” for generously funding my first three years and providing an excellent environment for the realization of my work.

Finally, I would like to thank my parents Alejandro and Leonor and my friends in Germany and in Spain for their love and support during these four years.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Shared-Variable Parallel Programs	2
1.3	Hoare Logic for Parallel Programs	2
1.4	Parameterized Parallel Programs	4
1.5	Need for Machine-Supported Verification	5
1.5.1	Theorem Proving vs. Model Checking	6
1.6	Related Work	7
1.7	Formalization in Isabelle/HOL	8
1.7.1	Formalization of the Verification Methods	8
1.7.2	Application to Concrete Programs	11
1.8	The Standard Isabelle/HOL Library	12
1.9	Presentation Style	13
1.10	Overview	13
2	The Owicki-Gries Method in Isabelle/HOL	15
2.1	Abstract Syntax	16
2.1.1	Component Programs	18
2.1.2	Atomic and Parallel Programs	20
2.2	Operational Semantics	22
2.2.1	The Transition Relation	22
2.2.2	Definition of Semantics	27
2.3	Validity of Correctness Formulas	31
2.4	The Proof System	32
2.4.1	Proof System for Atomic Programs	33
2.4.2	Proof System for Component Programs	33
2.4.3	Proof System for Parallel Programs	36
2.4.4	Auxiliary Variables	40
2.5	Soundness	41

2.5.1	Soundness of the System for Atomic Programs	42
2.5.2	Soundness of the System for Component Programs . .	42
2.5.3	Soundness of the System for Parallel Programs	45
2.6	Generation of Verification Conditions	50
2.7	Concrete Syntax	50
2.7.1	Formalization of the State	51
2.7.2	Concrete Syntax for Commands and Assertions	56
2.8	Examples	58
2.8.1	Mutual Exclusion	60
2.8.2	Parallel Zero Search	65
2.8.3	Producer/Consumer	67
2.9	Summary	69
3	Case Study: Single and Multi-Mutator Garbage Collection Algorithms	71
3.1	Incremental Garbage Collection	72
3.2	Formalization of the Memory	74
3.3	The Single-Mutator Case	76
3.3.1	The Mutator	77
3.3.2	The Collector	78
3.3.3	Interference Freedom	89
3.4	The Multi-Mutator Case	91
3.4.1	The Mutators	91
3.4.2	The Collector	93
3.4.3	Interference Freedom	97
3.5	Conclusions and Related Work	98
4	The Rely-Guarantee Method in Isabelle/HOL	101
4.1	Abstract Syntax	102
4.2	Operational Semantics	103
4.2.1	Semantics of Component Programs	104
4.2.2	Semantics of Parallel Programs	105
4.2.3	Computations	106
4.2.4	Modular Definition of Computation	107
4.3	Validity of Correctness Formulas	111
4.3.1	Validity for Component Programs	111
4.3.2	Validity for Parallel Programs	112
4.4	The Proof System	113
4.4.1	Proof System for Component Programs	113
4.4.2	Proof System for Parallel Programs	116

4.5	Soundness	118
4.5.1	Soundness of the System for Component Programs . .	118
4.5.2	Compositionality of the Semantics	125
4.5.3	Soundness of the System for Parallel Programs	129
4.6	Concrete Syntax	132
4.7	Examples	134
4.7.1	Set Elements of an Array to Zero	134
4.7.2	Increment a Variable in Parallel	135
4.7.3	Find Least Element	137
4.8	Concluding Remarks	141
5	Completeness of the Proof Systems for Parameterized Parallel Programs	143
5.1	Parameterized Programs	144
5.1.1	Syntax of Parameterized Programs	144
5.1.2	The State Space	147
5.2	Completeness of the Owicki-Gries System for Parameterized Parallel Programs	148
5.2.1	Semantic Completeness	151
5.2.2	Extending the Program	153
5.2.3	Annotating the Program	154
5.3	Relative Completeness	155
5.4	Completeness of the Rely-Guarantee System for Parameterized Parallel Programs	157
5.5	Example	158
5.6	Summary	160
6	Conclusion	161
6.1	General Contributions	161
6.2	Statistics of the formalization	163
6.3	Further Work	165
6.4	Experience	166
6.5	To conclude	166
A	Automatic Generation of Verification Conditions	167
A.1	VCG for Atomic Programs	167
A.2	VCG for Component Programs	170
A.3	VCG for Parallel Programs	172
B	Formal Declaration of Concrete Syntax	181

Bibliography	185
Index	195

Chapter 1

Introduction

1.1 Motivation

Parallel programs find applications in many different areas. They offer the possibility of computing data much faster than sequential programs, which is important in applications such as weather forecasting, computer visualized brain surgery, etc. They also support real-time applications where a program (alarm system) cannot be halted to let another program run (deallocation of computer memory), or, they simply need to be synchronized (scheduling resources). However, even small parallel programs are difficult to design, and errors are more often the rule than the exception. Failure of such programs can lead to minor disruptions in daily life, like the loss of a printing job, or to endangering human lives.

Due to their complexity, parallel programs should always be proved correct. This means proving formally that a program performs its intended task for any possible input. This is not easy. Some techniques are based on testing or reasoning about the behavior of possible executions. These techniques help to find bugs and improve efficiency, but do not provide a reliable proof of correctness. In general, it is impossible to verify designs exhaustively by such methods, simply because the number of possible executions is too large (or even infinite). Especially for programs used in safety-critical situations, we need formal verification methods to prove that a program *always* satisfies its specification.

The first formal method for verifying parallel programs was invented already in 1975 by Susan Owicki and David Gries [Owicki, 1975, Owicki and Gries, 1976a, Owicki and Gries, 1976b] and Leslie Lamport [Lamport, 1977]. It was the starting point for further developments like the compositional

rely-guarantee method for shared-variable parallelism [Jones, 1981, Jones, 1983] and the assumption-commitment method for synchronous message passing [Misra and Chandy, 1981, Apt *et al.*, 1980, Levin and Gries, 1981].

This dissertation presents the first formalization of the Owicki-Gries and the rely-guarantee methods for correctness proofs of imperative parallel programs with shared variables in a theorem prover. The formalizations have been successfully applied to the verification of several (parameterized and non-parameterized) programs like mutual exclusion and garbage collection algorithms.

1.2 Shared-Variable Parallel Programs

By parallel programs we understand a collection of sequential processes which run concurrently and cooperate in accomplishing some task. This cooperation is possible through the sharing of objects. Depending on the nature of these shared objects we distinguish between shared-variable and message-passing parallelism. In the former, the different processes have access to *shared* memory cells. The latter is used when each process has its own local memory and communicates with other processes by sending messages on *shared* channels. The present work deals with parallel programs that communicate via shared variables only.

1.3 Hoare Logic for Parallel Programs

Verification means proving formally that a program satisfies its specification, which should first be written in some logical language. The approach used in this work is based on the axiomatic method initiated by Hoare for sequential programs and extended later by various researchers to parallel programs. Hoare-like techniques are based on proving that a given program together with its specification can be derived from a system of axioms and inference rules which are syntax oriented, i.e. the proof is carried out on the program's text [Hoare, 1969, Apt, 1981b].

The Owicki-Gries proof system represents the first and probably the simplest extension of Hoare logic to parallel programs with shared-variable concurrency. It provides a methodology for breaking down correctness proofs into simpler pieces. First, the sequential components of the program are annotated with suitable assertions (resulting in so-called proof outlines). Then, the proof reduces to showing that the annotation of each component is correct, and that each assertion of an annotation is invariant under

the execution of the actions of the other components (so-called interference freedom of proof outlines).

The main drawback of the Owicki-Gries method is that it is not compositional. To perform the interference-freedom tests for some component we require information about the implementation of all other components. A compositional proof method, however, should be able to infer the specification of the system from the specification of the components without knowing anything about their internal representation.

The idea leading to a compositional proof method for parallel programs is to enrich the specification of each component with additional information about the interaction with the environment during execution. Such a proof method for shared-variable concurrency was first proposed by Cliff Jones [Jones, 1981, Jones, 1983]. A complete version of this system was later designed by [Stølen, 1990].

Jones extended the traditional Hoare pre- and postcondition specification of sequential programs with two new predicates: a *rely* condition specifying what the component expects from the environment, and a *guarantee* condition expressing the task performed by that component, and how this task may influence the environment. These two conditions can be formulated independently of the actual implementation of the components. Then, the verification process consists of proving certain relations among the rely and guarantee conditions of all components (and possibly of an overall environment). As a result, the proof rule for parallel composition can be formulated in terms of the specifications without any need for additional information about the implementation of the components, i.e. the resulting proof method is compositional.

The main improvements of the rely-guarantee method over Owicki-Gries are:

1. The complexity of the verification process grows linearly with the number of components, whereas in the Owicki-Gries method the number of proof obligations grows exponentially.
2. It allows verification of *open systems*, i.e. systems whose interaction with the environment can be specified without knowing the precise implementation of the environment. This makes the method suitable for top-down design. In contrast, the Owicki-Gries method only works for verifying *closed systems*, where the environment is fully characterized in terms of concretely known processes.

Although the rely-guarantee method represents an important step forward in

the methodology of program verification, it does not make the Owicki-Gries method obsolete. In general, when it comes to verifying an algorithm, the main concern is finding a proof. In this sense, non-compositional methods are sometimes more successful than compositional ones. This is the case for systems based on shared-variable concurrency like, for example, mutual exclusion algorithms, where processes require reading from and writing to the same shared variable. In general, programs defined using invariants which cannot be easily expressed as a conjunction of local predicates [Chandy and Misra, 1984, Francez and Rodeh, 1980] are difficult for compositional verification methods. There are very few examples of non-trivial shared-variable programs which have been verified using compositional methods [Stølen, 1990, de Boer *et al.*, 1997]. In contrast, non-compositional methods have been quite successful [Gries, 1997, Prensá Nieto and Esparza, 2000, Feijen and van Gasteren, 1999] (see [de Roever *et al.*, 2000] for a discussion on this topic).

Since proving the correctness of parallel programs is indeed difficult, each particular program should be tackled with the most suitable technique. Therefore, it is advantageous to have both compositional and non-compositional systems available.

1.4 Parameterized Parallel Programs

Parameterized parallel programs have become a very important subject of research in the area of computer-aided verification. These are programs which are defined generically, depending in a regular way on a parameter that represents the number of parallel processes. Many interesting programs are of this form, for example, mutual exclusion algorithms for an arbitrary number of processes wanting to use a common resource, or, a garbage collector that interacts with an arbitrary number of user processes. The goal is to verify such systems uniformly, i.e. prove by a *single* proof that the system is correct for any value of the parameter.

The Hoare-like methods for parallel programs found in the literature present systems of rules where the rule for the parallel constructor allows us to derive the composition of some fixed number of processes. In most cases, the parallel constructor is a binary operator [Xu *et al.*, 1997, Stirling, 1988] and programs with more than two processes have to be verified by repeatedly using the rule. Other systems present the parallel composition rule for a fixed number n of known processes [Owicki and Gries, 1976a, Apt, 1981a].

In this thesis, we present a generalization of the parallel composition rule so that parameterized programs can be directly verified in the system. This is achieved by modeling the parallel constructor such that its argument is a list of component programs. Then, the length of the list can be fixed or left as a parameter. With the resulting system, we can derive parameterized parallel programs in one go, i.e. by a single derivation. Several examples have been verified using both the compositional and the non-compositional methods formalized in this thesis. In these examples, the assertions used to obtain a valid derivation in the system are, like the program instructions, parameterized in the number of components and in the particular index of each component.

This led us to the question whether it is always possible to find such assertions, i.e. whether the systems are complete for verifying parameterized parallel programs. Completeness results for the standard systems, where parallel programs have a fixed number of components, have already been presented in [Owicki, 1975, Apt, 1981a] for the Owicki-Gries system and in [Xu *et al.*, 1997] for the rely-guarantee system. These results ensure that for any correct parallel program with a fixed number of components we can find a derivation in the respective system. However, they do not solve the problem for the parameterized case. In this thesis, we present proofs of completeness of the systems for parameterized programs as a natural extension of the known completeness results. These proofs have been carried out in an informal pencil-and-paper style, i.e. not formalized in the theorem prover.

1.5 Need for Machine-Supported Verification

Whichever method we select for an application, the main difficulty lies in finding assertions that formally express the conditions of the specification and yield a derivation from the corresponding system of rules.

For the Owicki-Gries method we need to find interference free intermediate assertions that lead to the expected results. For the rely-guarantee method, we need appropriate rely and guarantee conditions. These assertions are usually difficult to find and have to be changed and tuned many times. Each time the verification process has to be essentially restarted.

It is possible to do the proof by hand. However, this is a long, tedious and error-prone process. Consider for example the Owicki-Gries method. The number of interference-freedom tests needed is $O(k^n)$, where n is the number of sequential components, and k is the maximal number of lines of a

component. This makes a complete pencil-and-paper proof very laborious, even for small examples. For this reason, many of the interference-freedom proofs, which tend to be very simple, are usually omitted. A fact that increases the possibility of a mistake.

With the rely-guarantee method the situation is not so critical because the number of conditions to be checked is considerably smaller. Nevertheless, the proofs involved in verification are very detailed and often just boring routine work. When proofs are done by hand, one tends to spend too much time checking rather simple proof steps over and over again. Therefore it is desirable to have the help of computers that automate the process and ensure that no mistakes are made.

1.5.1 Theorem Proving vs. Model Checking

There are two major approaches used for mechanizing verification: theorem proving and model checking. Theorem proving is based on using a specific deductive system and performing a formal proof in the mathematical sense. Model checking techniques are based on decision or semi-decision procedures that “check” whether a program satisfies its specification by basically exploring the possible states exhaustively. Model checking has the advantage of being essentially automatic whereas general theorem proving requires interactive input from the user. On the other hand, verification with theorem provers can be fully general whereas model checking is still only applicable to a limited class of programs.

Initially, model-checking was restricted to finite-state systems of moderate size, but thanks to the development of techniques that improve efficiency it is now possible to tackle surprisingly large examples. Unfortunately, for infinite-state systems, even the theoretical possibility of exploring the state space disappears.

A program may have an infinite state space because it operates on data structures from a potentially infinite domain (integers, queues, etc.) or because it has an infinite control part (parallel programs parameterized in the number of components). Theoretical results [Apt and Kozen, 1986] even show that the verification of parameterized parallel programs is undecidable. Nevertheless, recent work present solutions that extend the applicability of model checking to restricted cases of infinite-state systems of the first kind [Jonsson and Parrow, 1993, Alur and Dill, 1994, Henzinger, 1995], and of the second kind [German and Sistla, 1992, Clarke *et al.*, 1995, Esparza, 1995, Abdulla and Jonsson, 1998]. However, programs which are infinite in both data and control flow are out of reach for the existing model-checking

techniques.

The availability of a theorem prover allows us to reason generically about programs without restrictions on their specification. It can deal with unbounded or infinite systems and supports highly expressive specifications of properties. For example, a garbage collection algorithm that manipulates an infinite data structure representing the computer memory and interacts with a parameterized number of mutators can be naturally specified and verified [Prensa Nieto and Esparza, 2000].

In this work a powerful interactive theorem prover, Isabelle, is used to formalize two well-known axiomatic verification methods, prove their soundness and considerably automate their application to real programs. As a result, we obtain a verification tool for general parallel programs where the generation of the verification conditions and the proof of the easy cases is automatic. The user is then able to concentrate only on the most important aspects of the proof.

1.6 Related Work

The idea of embedding an imperative programming language in a theorem prover goes back at least to [Gordon, 1989], who considered the Hoare logic of a simple while-language. Gordon's idea inspired an increasingly active research area. In this section, we give a brief overview. The following list of references is by no means exhaustive.

In the line of sequential language embeddings, [Nipkow, 1996, Nipkow, 1998] formalizes the first chapters of [Winskel, 1993] in Isabelle/HOL (and even finds a mistake in the proof of completeness), [Harrison, 1998] presents a formalization in HOL of Dijkstra's classic [Dijkstra, 1976], [Homeier and Martin, 1996] and [Kleymann, 1998] deal also with recursive procedures, [von Oheimb, 2001] presents an embedding of a Hoare logic for a subset of sequential Java in Isabelle/HOL and [Filliatre, 1999] presents a formalism for the verification of imperative programs in Type Theory in Coq.

For concurrent programming languages we encounter a long list of embeddings. For example, UNITY has been formalized in the Boyer-Moore prover [Goldschlag, 1990], HOL [Andersen *et al.*, 1994], Coq [Heyd and Crégut, 1996], LP [Chetali and Heyd, 1997] and Isabelle [Paulson, 2000, Paulson, 2001]. A related framework, *action systems*, has also been formalized in HOL [Långbacka and von Wright, 1997]. CSP has been treated in HOL [Camillieri, 1990], in Isabelle/HOL [Tej and Wolff, 1997] and in PVS [Dutertre and Schneider, 1997]. CCS has been formalized in HOL [Nesi,

1994]. ACP-style process algebra has also been formalized in PVS [Basten and Hooman, 1999]. TLA is found in HOL [von Wright and Långbacka, 1993], LP [Engberg *et al.*, 1993] and Isabelle/HOL [Kalvala, 1995]. Input/Output Automata embeddings are found in Isabelle/HOL [Müller and Nipkow, 1997, Müller, 1998] and in LP [Søgaard-Andersen *et al.*, 1993]. A formalism based on a *semantical* characterization of compositional verification has been formalized in PVS [Owre *et al.*, 1995]. [Hooman, 1998] presents an embedding of an assertional compositional system for asynchronous communication in PVS. The PVS system combined with model checking techniques [Rushby, 2000, Owre *et al.*, 1996, Shankar, 1996] has been successfully applied to medium-size examples as reported for example in [Hooman, 1995, Shankar, 1998, de Roever *et al.*, 1998].

Surprisingly, it appears that there has been no work on embedding Hoare logics for shared-variable parallelism in any theorem prover.

The Owicki-Gries method marks the beginning of a vast body of literature on proof systems for concurrency which we cannot survey here. The recent book [de Roever *et al.*, 2000] presents a development of state-based verification systems and contains numerous references related to the subject. A second volume of this book is announced [Hooman *et al.*, 2000] which focuses on illustrating through examples the success of compositional techniques in correctness proofs for parallel programs as well as techniques for machine-support in the verification process using PVS.

1.7 Formalization in Isabelle/HOL

For the formalization and proofs presented in this work we use the system Isabelle/HOL. Isabelle [Paulson, 1994] is a generic interactive theorem prover and Isabelle/HOL is its instantiation for higher-order logic, which is very similar to Gordon’s HOL system [Gordon and Melham, 1993]. A recent gentle introduction to Isabelle/HOL is [Nipkow and Paulson, 2001].

In this section we briefly describe the aspects of the system required to understand this dissertation. We can divide the work done with Isabelle into two main parts: the formalization of the verification methods and their application to concrete programs.

1.7.1 Formalization of the Verification Methods

When formalizing a programming language in a theorem prover, one has to decide between using a *deep embedding*, where first the (abstract) syntax is represented via an inductive datatype and then a semantics is assigned to it,

and a *shallow embedding*, where a term in the language is essentially an abbreviation of its semantics. Deep embeddings are useful when meta-theoretic reasoning (usually by induction over the syntax) is required. Shallow embeddings on the other hand, simplify reasoning about individual programs because one may work directly with the semantics avoiding the extra syntactic level.

We use a combination of both styles that has become quite established [Nipkow, 1998, von Oheimb, 2001]. We formalize as much as possible using a shallow embedding and use a deep embedding only where needed in order to perform the meta-theoretic proofs we are interested in. For our purposes, it suffices to use a deep embedding for the programming language. Assertions, expressions and even assignments are represented semantically, i.e. as functions on *states*. Consequently, the assertion language is not restricted to first-order logic as is customary in Hoare-like frameworks. Any HOL expression, and in particular all the constants defined in the Isabelle/HOL library can be used in assertions, boolean conditions and expressions within a program.

The program syntax is defined in Isabelle/HOL via a **datatype** definition. A free datatype is defined by listing its constructors together with their argument types, separated by ‘|’. In general it has the form

$$\mathbf{datatype} (\alpha_1, \dots, \alpha_n) t = C_1 \tau_{11} \dots \tau_{1k_1} \mid \dots \mid C_m \tau_{m1} \dots \tau_{mk_m}$$

where α_i are distinct type variables, C_i are distinct constructor names and τ_{ij} are types. Type abbreviations in Isabelle are declared by the keyword **types**. They follow the syntax of ML, except that function types are denoted by \Rightarrow . Laws about datatypes, such as $C_i \neq C_j$, are automatically included in the simplification tactics for future proofs. An induction principle, namely, structural induction over the constructors of the datatype is also generated with each datatype declaration. To use it in proofs it has to be explicitly invoked. Functions about datatypes are usually defined by primitive recursion. They are introduced by the keyword **primrec**.

Constants are declared with **consts** followed by their name and type, separated by ‘::’. Non-recursive definitions are declared by the keyword **constsdefs**. The introduced constant and its definition are separated by ‘ \equiv ’. Sometimes we first declare the constant and introduce the definition later with the keyword **defs**.

The operational semantics of commands is inductively defined via a set of rules. Similarly, the set of correct specifications is defined inductively by a set of axioms and proof rules. Such inductively defined sets represent the least set which is closed under the formation rules. They are declared by

the keyword **inductive** followed by the word **intros**. From each inductive definition Isabelle generates the corresponding induction principle, called *rule induction*, which represents the most powerful proof method used in this dissertation. In particular, soundness of the system of rules for program verification is proved by rule induction.

The so-called *inductive cases* proof principle is also automatically generated by the system for any inductive definition. It can be understood as the counterpart of (structural) case distinction on inductively generated elements. Whenever we have an assumption stating that an element belongs to an inductively defined set, we can distinguish on the last rule used for its derivation. As a result, we obtain a subgoal where the given element has been replaced by the corresponding premises of each of the proof rules whose conclusion matches the element. When we speak of case analysis on an inductively generated element we refer to this proof principle.

Statements that we want to prove are preceded by **theorem** or **lemma**. There is no formal difference between them; we use one or the other depending on the importance we attach to the stated proposition. Proofs are done by applying *tactics* to the stated goals. The application of a tactic is preceded by the keyword **apply**. The basic tactics are based on *resolution*, i.e. by applying inference rules (backwards or forwards) in a natural deduction style, and *rewriting*, i.e. by applying (conditional) directed equalities. As a result, goals are reduced to simpler subgoals until they become trivial. When all subgoals are solved the proposition is proven and stored under some name given by the user.

Some tactics are based on natural deduction (forward and back-chaining of rules) where search with backtracking is automated using the so-called *classical reasoner*. Other tactics, called *simplifiers*, compose rewriting steps. More powerful tactics (like *auto*) combine both systems and are able to automatically prove complicated goals. Tactics may also be combined using control structures called *tacticals*.

In an interactive theorem prover like Isabelle, if a statement cannot be proved automatically, the user is able to direct the proof by explicitly using induction, case distinction, instantiating variables and, in general, giving hints to the prover whenever automatic tactics do not succeed. Very often, automatic tools do succeed if they are supplied with the suitable auxiliary lemmas, which can be often obtained from the Isabelle library or have to be previously proven by the user.

1.7.2 Application to Concrete Programs

The relation between the program syntax and its semantics need only be studied once in the proof of soundness of the system. Once this is done, we can forget about the semantics and just use the system of rules for the verification of programs.

In order to make the verification task easier, we enrich the formalization with two additional features: the definition of a familiar concrete syntax for the programming language, and a tactic that, given a program specification, automatically generates the verification conditions.

Isabelle offers several facilities to define concrete (or external) syntax. It is possible to declare mixfix syntax notation for types and constants. The new notation may contain mathematical symbols and user-defined precedences. They can be directly given with the constant or type declaration (by writing the mixfix notation in parenthesis) or by defining new syntactic constants under the keyword **syntax** and putting the corresponding translation equations into the internal (abstract) syntax under the keyword **translations**. The translation equations are directed. The direction from left to right is represented by the arrows \rightarrow , and the translation in both directions by \Leftrightarrow .

When the transformations cannot easily be done via translations, Isabelle offers the possibility of defining ML programs that perform the translations from concrete syntax into abstract syntax (**parse_translation**), and vice versa (**print_translation**). These facilities are used to obtain an alternative syntax that allows us to write programs and assertions essentially like they are found in the literature.

The correctness of a program specification depends upon the validity of certain conditions called *verification condition* (also called *proof obligations*). These conditions are pure higher-order logic predicates with no mention of the programming language. Their validity is thus proven using standard Isabelle proving techniques.

For the automatic generation of the verification conditions, we define a so-called *verification conditions generator* (*vcg*) as an Isabelle tactic. Isabelle allows the user to construct new tactics by programming them in ML as combinations of existing ones. Without this tactic, the verification of the larger examples presented in this thesis would have been unbearably tedious.

1.8 The Standard Isabelle/HOL Library

The formalization uses some types and constants defined in the standard Isabelle/HOL library. We briefly present the most frequently used. Others will be explained when needed in the subsequent chapters.

The product type $\alpha \times \beta$ comes with the projection functions *fst* and *snd*. Tuples are pairs nested to the right, e.g. $(a, b, c) = (a, (b, c))$. They may also be used as patterns like in $\lambda(x, y). f\ x\ y$.

List notation is similar to ML (e.g. *@* is ‘append’) except that the ‘cons’ operation is denoted by *#* (instead of *::*). The *i*th component of a list *xs* is written *xs!**i*, where the first element has the index 0, i.e. *xs!*0, also defined as the head of the list, *hd xs*. The rest of the list (or tail) can be represented by the function *tl xs*. *last xs* represents the last element of a non-empty list. The syntax *xs*[*i* := *x*] denotes *xs* with the *i*th component replaced by *x*. The functional *map* :: $(\alpha \Rightarrow \beta) \Rightarrow \alpha\ list \Rightarrow \beta\ list$ applies a function to all elements of a list. The function *length* :: $\alpha\ list \Rightarrow nat$ returns the length of a list. The conversion function *set* :: $\alpha\ list \Rightarrow \alpha\ set$ builds a set from the elements of a list. The function *filter* :: $(\alpha \Rightarrow bool) \Rightarrow \alpha\ list \Rightarrow \alpha\ list$, returns the list of elements formed from the elements of a given list for which a given predicate holds.

The datatype $\alpha\ option = None \mid Some\ \alpha$ is frequently used to add a distinguished element to some existing type. It comes with the function *the* :: $\alpha\ option \Rightarrow \alpha$ such that *the* (*Some x*) = *x*.

Set comprehension syntax is $\{x. Px\}$ expressing the set of all elements that satisfy the predicate *P*. This notation is also available for tuples $\{(x, y, z). Pxyz\}$. A more general syntax for sets is $\{e\vec{x} \mid \vec{x}. P\vec{x}\}$ which abbreviates the set $\{u. \exists \vec{x}. u = e\vec{x} \wedge P\vec{x}\}$. The image of a set *A* under a function *f* is denoted by *f* ‘ *A* and is predefined in the Isabelle library as $\{f\ x \mid x. x \in A\}$. The complement of a set *A* is denoted by $-A$.

The notation $\llbracket A_1; \dots; A_n \rrbracket \Longrightarrow A$ represents an implication with assumptions A_1, \dots, A_n and conclusion *A*. It is also important to distinguish between the object implication ‘ \longrightarrow ’ and the meta-implication ‘ \Longrightarrow ’. The first one is a normal implication as known from mathematics and the second one separates assumptions from conclusions in proofs. In other words, if we state the goal $a \longrightarrow b$ in Isabelle, then we want to prove the proposition $a \longrightarrow b$ from the empty set of assumptions. However, if we state $a \Longrightarrow b$, then we want to prove *b* by assuming *a*. The first goal can be reduced to the second one by applying the deduction rule $(P \Longrightarrow Q) \Longrightarrow P \longrightarrow Q$ backwards.

1.9 Presentation Style

Isabelle provides tools for the automatic generation of L^AT_EX documents from Isabelle theories. The formal content of this dissertation has been written with this system. Hence, all definitions and theorems presented are part of actual Isabelle theories, where the side explanations have been inserted in special “text” environments that are ignored when the theory is being processed by Isabelle. This procedure ensures the consistency of the information presented. However, it forces the presentation to be bottom-up (as in the original formal theories) where probably a top-down explanation of the contents would be more appropriate. Nevertheless, we manage to avoid showing unnecessary details of the formalization (like user-defined preferences in syntax declarations) and more technical parts that are not relevant for the general understanding (like the syntax transformation functions or the tactics programmed in ML) are placed in the appendix.

The proofs done in this work follow the “traditional” tactic style, resulting in so-called “proof scripts”, which are not easily readable for non-experienced users. Now there is an alternative based on a more developed proof language that allows us to write proofs basically as they are found in mathematics books. This new system is called Isar (Intelligible semi-automated reasoning) and has been developed by Markus Wenzel [Wenzel, 2001a]. We have adopted the definition language of Isar but have maintained the tactic-style proofs (for historical reasons). Thus, we do not show the mechanical proofs in this thesis, but simply state the lemmas and explain their proofs informally in the text. Longer proofs are announced by **Proof** and finished with \square . The complete Isabelle theories and proof scripts can be obtained from <http://isabelle.in.tum.de/hoare-parallel/>.

1.10 Overview

Chapter 2 presents the formalization of the Owicki-Gries method and its application to some typical examples.

Chapter 3 is devoted to the main case study, namely, the verification of two parallel garbage collection algorithms, the second one parametric in the number of mutators.

Chapter 4 presents the formalization of the rely-guarantee method and its application to some typical examples.

Chapter 5 is devoted to the completeness of both the Owicki-Gries and rely-guarantee systems for parameterized parallel programs.

Chapter 6 summarizes the main results and gives suggestions for further work.

All results except for the completeness theorems of chapter 5 have been obtained using the theorem prover Isabelle/HOL.

Part of the materials contained in this thesis have been previously published in [Nipkow and Prensa Nieto, 1999], [Prensa Nieto and Esparza, 2000] and [Prensa Nieto, 2001].

Chapter 2

The Owicki-Gries Method in Isabelle/HOL

In this chapter we present the first formalization in a theorem prover of the Owicki-Gries method. First published by Susan Owicki in her Ph. D. thesis under the supervision of David Gries [Owicki, 1975], this method is widely accepted as the most fundamental methodology for correctness proofs of shared-variable concurrency.

Our formalization closely follows the description in [Apt and Olderog, 1991]. Thus, the present work can also be seen as an exercise in formalizing textbooks on programming language semantics. Yet, the search for a suitable and efficient adaptation in the theorem prover yields some improvements over the original presentation. Especially interesting is the parametric nature of the parallel composition rule (allowing verification of parameterized parallel programs directly in the system) and the soundness proof (because it does not explicitly mention program locations).

The different parts of the formalization are introduced following the usual steps required by a formal system. First, the abstract syntax of the programming language presents a simple while-language with concurrent execution of commands and synchronization via an await-command. Then, the operational semantics and the proof system are inductively defined as sets of rules. The soundness of the latter w.r.t. the former represents the main meta-theoretical result of the formalization.

Our interest is focused on the practical application. Therefore we prove soundness but not completeness of the proof system. Instead, we provide an automatic procedure for the generation of the verification conditions and define a familiar concrete syntax for writing programs. Finally, some typical

examples illustrate the applicability of the formalization. Furthermore, the verification of several schematic programs, where the number of parallel components is a parameter, shows that this embedding is more than just a verification condition generator. The use of a theorem prover allows us to tackle problems outside the range of fully automatic methods like model checking.

In the next chapter the method is applied to two garbage collection algorithms, the second one parametric in the number of mutators. These nontrivial case studies demonstrate the success of the approach, mainly due to the high degree of automation.

2.1 Abstract Syntax

We follow [Apt and Olderog, 1991] in stratifying the language. Only top-level parallelism is allowed, i.e. the parallel operator (\parallel) must not be nested. Hence, each c_i in $c_1 \parallel \dots \parallel c_n$ is a sequential command, called a (*sequential component*) of the parallel composition. Nevertheless, parallelism may occur within sequential composition, conditional statements and while-loops.

The third sublanguage in this stratification is the one used in the bodies of await-commands. They are called *atomic programs* because they are executed atomically, i.e. without interruption from other components. Summarizing, the programming language combines the following three layers:

Parallel commands include parallel composition of component programs as a construct. Component programs are purely sequential programs, thus nested parallelism is excluded.

Component commands represent the language of the programs appearing within a parallel composition and can be synchronized via an await-command. They have to be annotated with assertions before each command, thus, they are also called *annotated commands*.

Atomic commands are used for the bodies of await-statements, which are executed atomically.

We illustrate by the following example the scope of each layer. The delimiters **cobegin-coend** enclose a list of programs that are to be executed in

parallel:

```
x := 0;
cobegin
{x = 0} await True then x := x + 1; x := x + 1 end {True}
||
{True} x := 0 {x = 0 ∨ x = 1 ∨ x = 2}
coend
```

The whole is a *parallel* command (non-annotated) consisting of the sequential composition of an assignment ($x:=0$) and a parallel composition of two component (annotated) commands. The first component consists of an await-statement whose body is an atomic command (non-annotated). In HOL, there are two ways to encode this stratification:

1. Define the type of all programs and require well-formedness predicates for each sublanguage, or
2. Define the syntax in layers with different types.

We have chosen a combination of both that simplifies statements and proofs about the language. Component programs and parallel programs are defined in different layers, whereas atomic programs are defined as a sublanguage of parallel programs via a simple well-formedness predicate.

Although a number of constructs appear duplicated in both the parallel and component layers, they differ in that the latter attaches annotations to them. Proofs about those constructs may have to be duplicated but this duplication is quite mechanical.

Following the established combination of shallow and deep embedding we start by defining the parameterized type abbreviations:

types

```
 $\alpha$  bexp =  $\alpha$  set
 $\alpha$  assn =  $\alpha$  set
```

representing both assertions and boolean expressions as sets (of states). The α stands for the state of a program, which is a parameter of the program type. The reason for this formalization of the state will be explained in §2.7.

The three levels of the language depend on each other. An await-command is a constructor of the language of component programs but its body must be atomic. Similarly, a parallel construct is not part of the component's language but its arguments are component programs.

Atomic commands represent the simpler layer in this hierarchy. Its constructors are also used in parallel commands so that atomic programs can be defined as a sublanguage of parallel programs. Hence, only two datatypes are required: α *ann-com* for annotated sequential programs and α *com* for atomic and parallel programs. Datatypes that depend on each other are called *mutually recursive*. They are defined in HOL under a single datatype declaration joined via the keyword **and**:

```

datatype  $\alpha$  ann-com =
  AnnBasic ( $\alpha$  assn) ( $\alpha \Rightarrow \alpha$ )
  | AnnSeq ( $\alpha$  ann-com) ( $\alpha$  ann-com)
  | AnnCond1 ( $\alpha$  assn) ( $\alpha$  bexp) ( $\alpha$  ann-com) ( $\alpha$  ann-com)
  | AnnCond2 ( $\alpha$  assn) ( $\alpha$  bexp) ( $\alpha$  ann-com)
  | AnnWhile ( $\alpha$  assn) ( $\alpha$  bexp) ( $\alpha$  assn) ( $\alpha$  ann-com)
  | AnnAwait ( $\alpha$  assn) ( $\alpha$  bexp) ( $\alpha$  com)
and  $\alpha$  com =
  Parallel ( $\alpha$  ann-com option  $\times$   $\alpha$  assn) list
  | Basic ( $\alpha \Rightarrow \alpha$ )
  | Seq ( $\alpha$  com) ( $\alpha$  com)
  | Cond ( $\alpha$  bexp) ( $\alpha$  com) ( $\alpha$  com)
  | While ( $\alpha$  bexp) ( $\alpha$  assn) ( $\alpha$  com)

```

2.1.1 Component Programs

The language of component programs has the type α *ann-com*. It is a standard sequential while-language augmented with a synchronization construct (*AnnAwait*). It departs from the usual presentation of the language by the inclusion of assertions directly in the syntax: every construct, apart from sequential composition, is annotated with a precondition, and the loop is also annotated with an invariant. Due to this special presentation they are called *annotated commands*. We emphasize that these assertions are merely annotations and do not change the semantics of the language. Next, we discuss the different constructors:

AnnBasic represents a basic atomic state transformation, for example an assignment, a multiple assignment, or even any non-constructive specification. Concrete syntax for single assignments of the form $x := e$, where x is a program variable and e is an expression of the corresponding type, is also available (cf. §2.7).

AnnSeq is the sequential composition of commands.

$AnnCond_1$ is the standard conditional. Both subprograms are themselves annotated commands, thus preceded by a precondition, subject to the so called *interference freedom* test in the verification process.

$AnnCond_2$ is the conditional without else-branch. In some cases, it is convenient to ignore the else-branch altogether because its precondition is bound to fail at the interference freedom test.

$AnnWhile$ is the loop, annotated with an invariant.

$AnnAwait$ is the synchronization construct. Notice that its body is of type $\alpha\ com$.

It might seem more natural to define $AnnCond_2$ as a special case of the constructor $AnnCond_1$, namely $AnnCond_1\ r\ b\ c\ (AnnBasic\ p\ id)$, where id represents the identity transformation. This, however, would still require proving interference freedom for the assertion p which must, in general, include the clause $\neg b$. Sometimes, e.g. in §3.4.1, this is not possible and we prefer to directly ignore the else-part. Another possible way of avoiding this assertion is to consider $AnnCond_2$ as an abbreviation of $AnnCond_1\ r\ b\ c\ (AnnAwait\ r\ True\ (Cond\ b\ c\ (Basic\ id)))$. This way we also avoid proving that $\neg b$ is interference free, but it is an unnecessarily complicated solution.

The meaning of await-commands is quite intuitive to understand. Imagine an execution of a parallel composition of programs where one component intends to execute the statement $AnnAwait\ r\ b\ c$. If b evaluates to *True*, then c is executed as an atomic region. If b evaluates to *false*, the component becomes blocked. In sequential programs this behavior does not make any sense, but in the case of parallel programs it means that other components can take over the execution. If eventually b becomes true, the blocked component can resume its execution. Otherwise, it remains blocked forever. Programs with this construct may end in a *deadlock*. This happens when some component of a parallel program is blocked and there are no other active components. A component is *active* if it is neither blocked nor finished.

To reason about parallel programs with shared variables we need to reason about each atomic step taken in the computations of its components. To this end, proofs of component programs are presented in the form of *proof outlines*, i.e. interleaved with assertions at appropriate places. Furthermore, we directly present them as a special case of proof outlines called *standard proof outlines* obtained by minimizing the number of annotations. Each command c is then preceded by an assertion, $pre\ c$, and apart from these and

loop invariants there are no other assertions¹. These annotations describe (a subset of) the set of states that are reachable at each point of control.

For purely sequential programs such a presentation is not necessary. The intermediate annotations can be derived as the weakest precondition from the postcondition and loop invariants. The so computed assertions invariably hold at their respective locations since no other action can modify the expected results. In contrast, a component in a parallel program has the ability to modify shared variables, endangering the task of other components. For this reason we explicitly annotate each point of control with an assertion, whose invariance under the actions in the other parallel components can be checked.

To obtain this special presentation we include the precondition in the syntax of component programs. Moreover, it turns out that for proof-theoretic reasons it is very helpful to define the semantics of the language directly on annotated commands. The precondition of each annotated command is extracted by the function *pre*. Its definition is:

```

consts
  pre ::  $\alpha$  ann-com  $\Rightarrow$   $\alpha$  assn
primrec
  pre (AnnBasic r f) = r
  pre (AnnSeq c1 c2) = pre c1
  pre (AnnCond1 r b c1 c2) = r
  pre (AnnCond2 r b c) = r
  pre (AnnWhile r b i c) = r
  pre (AnnAwait r b c) = r

```

2.1.2 Atomic and Parallel Programs

Atomic and parallel commands are similar enough to both be represented by the same datatype (α com) together with a simple well-formedness predicate characterizing atomic programs. We could define them in two different datatypes, but this would make the specification of the language unnecessarily long and proofs about (otherwise) identical constructors would have to be duplicated. By minimizing the number of constructors in the language we obtain a clearer and shorter presentation of the theories and proofs.

Atomic commands form the body of await-statements which are executed as atomic regions, i.e. its activation cannot be interrupted by the other components. Hence, they behave as pure sequential programs. On the other

¹Non-standard proof outlines admit two assertions after each other provided that the second one is a logical consequence of the first one.

hand, parallel commands are themselves not executed in parallel, i.e. nested parallelism is not allowed. However, if c_1 and c_2 are both parallel programs, they can be sequentially composed ($Seq\ c_1\ c_2$) or appear in conditional ($Cond\ b\ c_1\ c_2$) statements and loop constructions ($While\ b\ i\ c_1$).

The only kind of commands in this layered language that is executed in parallel with others are the annotated commands.

Consequently, atomic programs and programs containing parallel programs are similar in the sense that they are both sequential. This is characterized by two important aspects:

1. They do not contain await-statements, which are only meaningful in a parallel context. To ensure termination, it is also usual to forbid while-commands inside atomic regions. However, this condition is not necessary for the soundness of the system and so we leave it out.
2. They do not contain intermediate annotations. For purely sequential commands there is no need to record a proof outline to be checked for interference freedom.

The *Parallel* constructor encloses a list of pairs (c, q) , where c is a sequential command or the empty program if the execution has terminated, and q is the postcondition (remember that the precondition is already part of the annotated c). Strictly speaking it is not necessary to include the postcondition, but it simplifies program verification.

Although each component consists of a pair, they can be seen as Hoare triples. The three elements are the precondition, which can be extracted from the command, followed by the program itself (ignoring the precondition), and finally the postcondition.

The remaining commands are almost like their namesakes in the sequential layer, but with a slightly different concrete syntax for sequential composition, to avoid confusion.

Nevertheless, atomic commands do not contain parallel constructs. We introduce a well-formedness predicate that characterizes the subset of programs of $\alpha\ com$ that may appear inside atomic regions:

consts $atom-com :: \alpha\ com \Rightarrow bool$

primrec

$atom-com\ (Parallel\ Ts) = False$

$atom-com\ (Basic\ f) = True$

$atom-com\ (Seq\ c_1\ c_2) = (atom-com\ c_1 \wedge atom-com\ c_2)$

$atom-com\ (Cond\ b\ c_1\ c_2) = (atom-com\ c_1 \wedge atom-com\ c_2)$

$atom-com\ (While\ b\ i\ c) = atom-com\ c$

If desired while-commands could also be excluded by writing

$$\text{atom-com } (\text{While } b \text{ i } c) = \text{False}$$

instead.

2.2 Operational Semantics

The semantics defines the input/output behavior of programs, i.e. given a program c , its semantics $Sem\ c$ defines a mapping from (initial) states to (final) states. There are two classical ways of defining this mapping:

A denotational semantics [Scott and Strachey, 1971, Gordon, 1979] defines $Sem\ c$ by induction on the structure of c , as a partial function on states. In particular, fixed point techniques are used to deal with recursion.

An operational semantics [Hennessy and Plotkin, 1979, Plotkin, 1981] defines first a transition relation between so-called *configurations* and then defines $Sem\ c$ using this relation.

Although the denotational style is more abstract and can theoretically handle all programming languages, it becomes very complicated for parallel programs. In contrast, the operational semantics remains simple and is thus preferred for assigning meaning to parallel constructors.

2.2.1 The Transition Relation

The transition relation used to define the operational semantics is inductively defined by a set of axioms and rules about transitions (or steps) between configurations. A *configuration* is a pair of a program fragment and a state, where a program fragment is either an atomic command or, if execution has come to an end, the empty program. Each transition is regarded as one step in the computation. For example, $(c, s) \rightarrow (c', s')$ means that the execution of one instruction in c from state s leads to the configuration consisting of a command c' and a state s' from which execution continues.

Basically, we need to define axioms and rules for all possible transition steps. Since the programming language constructs are defined in two different layers, two kinds of transitions are defined: *transition*, for steps of commands of type $\alpha\ com$ and, *ann-transition*, for steps of commands with type $\alpha\ ann-com$.

In order to define the rules we need some way to represent the fact that the command is empty. The language of component programs does not contain the empty program. Adjoining a new element to a type is naturally modeled by the standard Isabelle/HOL datatype $\alpha \text{ option} = \text{None} \mid \text{Some } \alpha$. In this case, *None* represents the empty program. Otherwise, the command is wrapped up by the *Some* constructor. To abbreviate we define a new type for optional annotated commands:

types $\alpha \text{ ann-com-op} = \alpha \text{ ann-com option}$

A new type abbreviation for an optional annotated command followed by its postcondition stands for the type of each component in the list of a parallel composition constructor. It can be seen as a triple since the precondition is part of the command's type:

types $\alpha \text{ ann-triple-op} = (\alpha \text{ ann-com-op} \times \alpha \text{ assn})$

Two selector functions extract the command part and the postcondition:

consts $\text{com} :: \alpha \text{ ann-triple-op} \Rightarrow \alpha \text{ ann-com-op}$

primrec $\text{com } (c, q) = c$

consts $\text{post} :: \alpha \text{ ann-triple-op} \Rightarrow \alpha \text{ assn}$

primrec $\text{post } (c, q) = q$

Equations defining a primitive recursive function are automatically added to the simplifier.

In the language of atomic and parallel programs we can consider the parallel composition with the empty list as argument as the equivalent of the empty program. This way we avoid the option type. The execution of a parallel composition terminates when all components do, i.e. when all components are *None*. The following predicate characterizes a terminated parallel composition:

constdefs

$\text{All-None} :: \alpha \text{ ann-triple-op list} \Rightarrow \text{bool}$

$\text{All-None } Ts \equiv \forall (c, q) \in \text{set } Ts. c = \text{None}$

Now we can define the transition relations. They are inductively defined as sets of relations between configurations:

consts

ann-transition :: $((\alpha \text{ ann-com-op} \times \alpha) \times (\alpha \text{ ann-com-op} \times \alpha)) \text{ set}$
transition :: $((\alpha \text{ com} \times \alpha) \times (\alpha \text{ com} \times \alpha)) \text{ set}$

Concrete syntax for a transition step as well as for its n -fold iteration and the reflexive transitive closure is provided via infix syntax annotations

syntax

-ann-transition :: $(\alpha \text{ ann-com-op} \times \alpha) \Rightarrow (\alpha \text{ ann-com-op} \times \alpha) \Rightarrow \text{bool}$
 $(- \text{ -1} \rightarrow -)$
-ann-transition-n :: $(\alpha \text{ ann-com-op} \times \alpha) \Rightarrow \text{nat} \Rightarrow (\alpha \text{ ann-com-op} \times \alpha)$
 $\Rightarrow \text{bool} \quad (- \text{ --} \rightarrow -)$
*-ann-transition-** :: $(\alpha \text{ ann-com-op} \times \alpha) \Rightarrow (\alpha \text{ ann-com-op} \times \alpha) \Rightarrow \text{bool}$
 $(- \text{ -*} \rightarrow -)$

-transition :: $(\alpha \text{ com} \times \alpha) \Rightarrow (\alpha \text{ com} \times \alpha) \Rightarrow \text{bool} \quad (- \text{ -P1} \rightarrow -)$
-transition-n :: $(\alpha \text{ com} \times \alpha) \Rightarrow \text{nat} \Rightarrow (\alpha \text{ com} \times \alpha) \Rightarrow \text{bool} \quad (- \text{ -P-} \rightarrow -)$
*-transition-** :: $(\alpha \text{ com} \times \alpha) \Rightarrow (\alpha \text{ com} \times \alpha) \Rightarrow \text{bool} \quad (- \text{ -P*} \rightarrow -)$

The corresponding syntax translations are:

translations

$con_0 \text{ -1} \rightarrow con_1 \Leftrightarrow (con_0, con_1) \in \text{ann-transition}$
 $con_0 \text{ -n} \rightarrow con_1 \Leftrightarrow (con_0, con_1) \in \text{ann-transition}^{\wedge n}$
 $con_0 \text{ -*} \rightarrow con_1 \Leftrightarrow (con_0, con_1) \in \text{ann-transition}^*$

$con_0 \text{ -P1} \rightarrow con_1 \Leftrightarrow (con_0, con_1) \in \text{transition}$
 $con_0 \text{ -Pn} \rightarrow con_1 \Leftrightarrow (con_0, con_1) \in \text{transition}^{\wedge n}$
 $con_0 \text{ -P*} \rightarrow con_1 \Leftrightarrow (con_0, con_1) \in \text{transition}^*$

The last two arrows are syntactic sugar for the n -fold and the $*$ postfix operators which are part of Isabelle/HOL's theory of relations.

The two kinds of transitions defining the semantics depend on each other. Thus, the rules for both systems are defined simultaneously and atomic programs and parallel programs “share” the rules for the common constructors.

The transition rules defined below are also called small-step rules and the semantics they define is called small-step semantics because it describes the execution of programs step by step. (In contrast to the so-called big-step semantics, where the rules represent transitions that may correspond to several steps in the execution of the program.)

inductive *ann-transition transition*

intros

AnnBasic: $(\text{Some } (\text{AnnBasic } r f), s) \text{--}1 \rightarrow (\text{None}, f s)$

AnnSeq1: $(\text{Some } c_0, s) \text{--}1 \rightarrow (\text{None}, t) \implies$
 $(\text{Some } (\text{AnnSeq } c_0 c_1), s) \text{--}1 \rightarrow (\text{Some } c_1, t)$

AnnSeq2: $(\text{Some } c_0, s) \text{--}1 \rightarrow (\text{Some } c_2, t) \implies$
 $(\text{Some } (\text{AnnSeq } c_0 c_1), s) \text{--}1 \rightarrow (\text{Some } (\text{AnnSeq } c_2 c_1), t)$

AnnCond₁T: $s \in b \implies (\text{Some } (\text{AnnCond}_1 r b c_1 c_2), s) \text{--}1 \rightarrow (\text{Some } c_1, s)$

AnnCond₁F: $s \notin b \implies (\text{Some } (\text{AnnCond}_1 r b c_1 c_2), s) \text{--}1 \rightarrow (\text{Some } c_2, s)$

AnnCond₂T: $s \in b \implies (\text{Some } (\text{AnnCond}_2 r b c), s) \text{--}1 \rightarrow (\text{Some } c, s)$

AnnCond₂F: $s \notin b \implies (\text{Some } (\text{AnnCond}_2 r b c), s) \text{--}1 \rightarrow (\text{None}, s)$

AnnWhileF: $s \notin b \implies (\text{Some } (\text{AnnWhile } r b i c), s) \text{--}1 \rightarrow (\text{None}, s)$

AnnWhileT: $s \in b \implies (\text{Some } (\text{AnnWhile } r b i c), s) \text{--}1 \rightarrow$
 $(\text{Some } (\text{AnnSeq } c (\text{AnnWhile } i b i c)), s)$

AnnAwait: $\llbracket s \in b; \text{atom-com } c; (c, s) \text{--}P^* \rightarrow (\text{Parallel } [], t) \rrbracket \implies$
 $(\text{Some } (\text{AnnAwait } r b c), s) \text{--}1 \rightarrow (\text{None}, t)$

Parallel: $\llbracket i < \text{length } Ts; Ts!i = (\text{Some } c, q); (\text{Some } c, s) \text{--}1 \rightarrow (r, t) \rrbracket$
 $\implies (\text{Parallel } Ts, s) \text{--}P1 \rightarrow (\text{Parallel } (Ts [i := (r, q)]), t)$

Basic: $(\text{Basic } f, s) \text{--}P1 \rightarrow (\text{Parallel } [], f s)$

Seq1: $\text{All-None } Ts \implies (\text{Seq } (\text{Parallel } Ts) c, s) \text{--}P1 \rightarrow (c, s)$

Seq2: $(c_0, s) \text{--}P1 \rightarrow (c_2, t) \implies (\text{Seq } c_0 c_1, s) \text{--}P1 \rightarrow (\text{Seq } c_2 c_1, t)$

CondT: $s \in b \implies (\text{Cond } b c_1 c_2, s) \text{--}P1 \rightarrow (c_1, s)$

CondF: $s \notin b \implies (\text{Cond } b c_1 c_2, s) \text{--}P1 \rightarrow (c_2, s)$

WhileF: $s \notin b \implies (\text{While } b i c, s) \text{--}P1 \rightarrow (\text{Parallel } [], s)$

WhileT: $s \in b \implies (\text{While } b i c, s) \text{--}P1 \rightarrow (\text{Seq } c (\text{While } b i c), s)$

The transition rules for the similar constructs in both, annotated and non-annotated commands, are practically identical. The only difference is that the empty program is represented by *None* in the former and by *Parallel []* in the latter.

The basic commands are executed in one step, performing the corresponding state transformation. The one-step execution of a sequential composition is determined by two rules. If the first command of the sequential composition finishes in one step, the next configuration indicates that only the second command remains to be executed. Otherwise, the first command is simply substituted by its reduction after one step.

The rules for the conditional statement lead, depending on the initial state s being an element of the set b or not, to the corresponding subprogram, leaving in both cases the state unchanged. While-statements terminate if the boolean condition is not fulfilled. Otherwise, the execution of one body followed by the original while-loop proceeds.

The transition rule *AnnAwait* formalizes the meaning of conditional atomic regions, where the body is required to be a well-formed atomic command whose execution terminates. If $s \in b$, the await-statement is executed uninterrupted in one step, provided the computation of the body terminates. If $s \notin b$, no transition is possible and the component is blocked.

Basic statements and evaluation of boolean expressions are all executed in one step. This is called a *high level* semantics, which abstracts from all the details of the evaluation of expressions.

Observe that both the preconditions denoted by r as well as the loop invariants i are merely annotations and do not play any role in the semantics. However, in the rule *AnnWhileT*, the precondition of the second *AnnWhile* has been changed to the invariant i . Although this does not influence the semantics it is important for the proof theory in §2.4.

The execution of the parallel composition of a list of annotated triples Ts proceeds by executing one non-*None* component of Ts . This form of modeling concurrency is called *interleaving*.

A terminating computation of a parallel composition of commands is a finite transition sequence starting in a state s such that in the last configuration each component program is *None*. The computation cannot be extended because there is no possible transition from *None*. For example,

$$\begin{aligned} & (Parallel [(Some (AnnBasic p f), q), (Some (AnnBasic p' g), q')], s) \\ & -P1 \rightarrow (Parallel [(None, q), (Some (AnnBasic p' g), q')], f s) \\ & -P1 \rightarrow (Parallel [(None, q), (None, q')], g (f s)) \end{aligned}$$

This *one-step* semantics is particularly appropriate for concurrent languages where different executions have to be interleaved. For simplicity, we employ this style at all levels, although strictly speaking it is only necessary for component programs.

2.2.2 Definition of Semantics

There are different definitions of the semantics of programs. They differ mainly in the amount of information they provide about the input/output behavior of programs. Two of the possible definitions are the so-called *total correctness semantics* and *partial correctness semantics*. Both approaches differ in the way they deal with divergent computations. The former considers the possibility of divergence while the latter one ignores it. Other definitions [Apt and Olderog, 1991] include information about *fairness* or *deadlocks* for example. In this work we concentrate on proving partial correctness, thus, only this interpretation will be formalized.

Following Apt and Olderog, we define the *partial correctness semantics of annotated commands* by the function *ann-SEM*:

constdefs

$$\begin{aligned} \text{ann-sem} &:: \alpha \text{ ann-com} \Rightarrow \alpha \Rightarrow \alpha \text{ set} \\ \text{ann-sem } c &\equiv \lambda s. \{t. (\text{Some } c, s) \text{ --}^* \rightarrow (\text{None}, t)\} \end{aligned}$$

$$\begin{aligned} \text{ann-SEM} &:: \alpha \text{ ann-com} \Rightarrow \alpha \text{ set} \Rightarrow \alpha \text{ set} \\ \text{ann-SEM } c \ S &\equiv \bigcup \text{ann-sem } c \ ' S \end{aligned}$$

The auxiliary function *ann-sem* returns for some annotated command *c* and some initial state *s* the set of all possible final states.

The semantics *ann-SEM* of an annotated command *c* is the union of the sets of final states that result from applying *ann-sem c* to each initial state in the set *S*. In other words, *ann-SEM c S* is the union of all possible final states of *c* executed from some state in *S*.

The image of a set *A* under a function *f* is denoted by *f ' A* and is predefined in the Isabelle library as $f \ ' \ A \equiv \{f \ x \mid x. x \in A\}$.

The definition of partial correctness semantics for non-annotated programs is slightly different because of the lack of *None* as an indicator of termination:

constdefs

$$\begin{aligned} \text{sem} &:: \alpha \text{ com} \Rightarrow \alpha \Rightarrow \alpha \text{ set} \\ \text{sem } c &\equiv \lambda s. \{t. \exists Ts. (c, s) \text{ -}P^* \rightarrow (\text{Parallel } Ts, t) \wedge \text{All-None } Ts\} \end{aligned}$$

$$\begin{aligned} \text{SEM} &:: \alpha \text{ com} \Rightarrow \alpha \text{ set} \Rightarrow \alpha \text{ set} \\ \text{SEM } c \ S &\equiv \bigcup \text{sem } c \ ' S \end{aligned}$$

The semantics *SEM* satisfies several properties that we shall need in the

sequel. For a property about the semantics of while-commands we define an auxiliary program called Ω .

syntax $\text{-Omega} :: \alpha \text{ com} \quad (\Omega)$

It is defined as an abbreviation of the following while-statement:

translations $\Omega \rightleftharpoons \text{While UNIV UNIV (Basic id)}$

where UNIV stands for the universal set of a fixed type, i.e. $\{x. \text{True}\}$ and id represents the identity transformation. This particular program enjoys the following property:

lemma SEM-Omega: $\text{SEM } \Omega \text{ } S = \{\}$

The primitive recursive function fwhile defines a sequence of deterministic programs that simulates the behavior of a while-statement:

consts $\text{fwhile} :: \alpha \text{ bexp} \Rightarrow \alpha \text{ com} \Rightarrow \text{nat} \Rightarrow \alpha \text{ com}$

primrec

$\text{fwhile } b \text{ } c \text{ } 0 = \Omega$

$\text{fwhile } b \text{ } c \text{ } (\text{Suc } n) = \text{Cond } b \text{ } (\text{Seq } c \text{ } (\text{fwhile } b \text{ } c \text{ } n)) \text{ (Basic id)}$

We prove the following lemmas about SEM as stated in [Apt and Olderog, 1991] (the proofs in the book are left as an exercise). We briefly review the main steps of our mechanized version.:

lemma SEM-mono: $X \subseteq Y \Longrightarrow \text{SEM } c \text{ } X \subseteq \text{SEM } c \text{ } Y$

The proof is automatic.

lemma SEM-Seq: $\text{SEM } (\text{Seq } c_1 \text{ } c_2) \text{ } X = \text{SEM } c_2 \text{ } (\text{SEM } c_1 \text{ } X)$

Proof. The \subseteq -inclusion is proved using the following lemma:

lemma SEM-Seq-onlyif:

$\llbracket (\text{Seq } c_1 \text{ } c_2, s) -P^* \rightarrow (\text{Parallel } Ts, t); \text{All-None } Ts \rrbracket$
 $\Longrightarrow \exists y \text{ } Rs. (c_1, s) -P^* \rightarrow (\text{Parallel } Rs, y) \wedge \text{All-None } Rs \wedge$
 $(c_2, y) -P^* \rightarrow (\text{Parallel } Ts, t)$

The proof relies on the following auxiliary lemma solved by induction on n :

lemma SEM-Seq-onlyif-aux:

$$\begin{aligned} & \llbracket (\text{Seq } c_1 \ c_2, s) -Pn \rightarrow (\text{Parallel } Ts, t); \text{ All-None } Ts \rrbracket \\ & \implies \exists y \ m \ Rs. (c_1, s) -P* \rightarrow (\text{Parallel } Rs, y) \wedge \text{ All-None } Rs \wedge \\ & \quad (c_2, y) -Pm \rightarrow (\text{Parallel } Ts, t) \wedge m \leq n \end{aligned}$$

For the \supseteq -inclusion we use the lemma:

lemma SEM-Seq-if:

$$\begin{aligned} & \llbracket (c_1, s_1) -P* \rightarrow (\text{Parallel } Ts, s_2); \text{ All-None } Ts; \\ & \quad (c_2, s_2) -P* \rightarrow (\text{Parallel } Rs, s_3); \text{ All-None } Rs \rrbracket \\ & \implies (\text{Seq } c_1 \ c_2, s_1) -P* \rightarrow (\text{Parallel } Rs, s_3) \end{aligned}$$

which is proven by induction on the length of the transition sequence given by $(c_1, s_1) -P* \rightarrow (\text{Parallel } Ts, s_2)$. \square

lemma SEM-Seq-assoc:

$$SEM (\text{Seq } (\text{Seq } c_1 \ c_2) \ c_3) X = SEM (\text{Seq } c_1 \ (\text{Seq } c_2 \ c_3)) X$$

The proof is trivial.

lemma SEM-Cond:

$$SEM (\text{Cond } b \ c_1 \ c_2) X = SEM c_1 (X \cap b) \cup SEM c_2 (X \cap \neg b)$$

where $\neg b$ represents the complement set of b . Both inclusions are easily solved by properly manipulating the transitive closure and doing case analysis on the conditional statement.

lemma SEM-While: $SEM (\text{While } b \ i \ c) = (\lambda x. \bigcup k. SEM (fwhile \ b \ c \ k) \ x)$

Proof. The \subseteq -inclusion is proved using the lemma:

lemma SEM-While-onlyif:

$$\begin{aligned} & \llbracket (\text{While } b \ i \ c, s) -Pn \rightarrow (\text{Parallel } Ts, t); \text{ All-None } Ts \rrbracket \\ & \implies \exists k. (fwhile \ b \ c \ k, s) -P* \rightarrow (\text{Parallel } Ts, t) \end{aligned}$$

which is proved by complete induction on n . With the induction hypothesis we obtain the subgoal:

$$\begin{aligned} & \llbracket \forall m < n. (\forall s. (\text{While } b \ i \ c, s) -Pm \rightarrow (\text{Parallel } Ts, t) \wedge \text{ All-None } Ts \\ & \quad \rightarrow (\exists k. (fwhile \ b \ c \ k, s) -P* \rightarrow (\text{Parallel } Ts, t))) \rrbracket \end{aligned}$$

$$\begin{aligned} & (\textit{While } b \ i \ c, s) -Pn \rightarrow (\textit{Parallel } Ts, t); \textit{All-None } Ts \] \\ \implies & \exists k. (\textit{fwhile } b \ c \ k, s) -P* \rightarrow (\textit{Parallel } Ts, t) \end{aligned}$$

We want to find an appropriate k satisfying the conclusion. If $s \notin b$, it suffices to take $k = 1$. If $s \in b$, we obtain by case analysis:

$$(\textit{Seq } c \ (\textit{While } b \ i \ c), s) -Pm \rightarrow (\textit{Parallel } Ts, t)$$

where $n = \textit{Suc } m$. By *SEM-Seq-only-if-aux* we can split the computation of the sequential composition into the computations of its two components. Then, for some y we obtain:

$$\begin{aligned} & (c, s) -P* \rightarrow (\textit{Parallel } Rs, y) \wedge \textit{All-None } Rs \\ & (\textit{While } b \ i \ c, y) -Pm' \rightarrow (\textit{Parallel } Ts, t) \wedge m' \leq m \end{aligned}$$

we apply the induction hypothesis obtaining for some k' ,

$$(\textit{fwhile } b \ c \ k', y) -P* \rightarrow (\textit{Parallel } Ts, t)$$

choosing $k = k' + 1$ the conclusion becomes

$$(\textit{Cond } b \ (\textit{Seq } c \ (\textit{fwhile } b \ c \ k)) \ \textit{Basic id}, s) -P* \rightarrow (\textit{Parallel } Ts, t)$$

but $s \in b$ holds, so it can be simplified to

$$(\textit{Seq } c \ (\textit{fwhile } b \ c \ k), s) -P* \rightarrow (\textit{Parallel } Ts, t)$$

which we prove via *SEM-Seq-if*.

The \supseteq -inclusion requires also an auxiliary lemma:

lemma *SEM-While-if*:

$$\begin{aligned} & \llbracket (\textit{fwhile } b \ c \ k, s) -P* \rightarrow (\textit{Parallel } Ts, t); \textit{All-None } Ts \rrbracket \\ \implies & (\textit{While } b \ i \ c, s) -P* \rightarrow (\textit{Parallel } Ts, t) \end{aligned}$$

proved by induction on k as follows. The base case is easy; since Ω does not terminate, there is a contradiction among the premises.

For the induction step, we first distinguish whether the transitive closure in the premise performs at least one step. If not, the proof is trivial. Otherwise, from the definition of *fwhile* we obtain a conditional statement, thus by case analysis there are two possible situations:

1. If $s \in b$, then $(\textit{Seq } c \ (\textit{fwhile } b \ c \ n), s) -P* \rightarrow (\textit{Parallel } Ts, t)$. Decomposing the computation of the sequential composition by applying

the lemma *SEM-Seq-onlyif-aux* and using the induction hypothesis we obtain for some R_s and some state y

$$\begin{aligned} (c, s) -P^* \rightarrow (Parallel\ R_s, y) \wedge All\text{-}None\ R_s \wedge \\ (While\ b\ i\ c, y) -P^* \rightarrow (Parallel\ T_s, t) \end{aligned}$$

as premises. Since we have $s \in b$, the conclusion becomes

$$(Seq\ c\ (While\ b\ i\ c), s) -P^* \rightarrow (Parallel\ T_s, t)$$

The proof follows by applying the lemma *SEM-Seq-if*.

2. If $s \notin b$, then the conclusion becomes

$$(Parallel\ [], s) -P^* \rightarrow (Parallel\ T_s, t)$$

We need to show that $s = t$ and $T_s = []$. This follows easily from the premise *(Basic id, s) -P^* → (Parallel T_s, t)* and the lemma

$$(Parallel\ [], s) -P_n \rightarrow (Parallel\ T_s, t) \implies T_s = [] \wedge n = 0 \wedge s = t$$

□

2.3 Validity of Correctness Formulas

Before introducing the proof system that shows us the rules to derive correct programs, we need to formalize what we mean by “correct”, and in particular correct in the sense of partial correctness. Informally, we say that a program is correct if it satisfies the intended input/output relation, where input and output are usually described as predicates over states, or equivalently as sets of states. Thus, a program specification consists of a triple, also called *Hoare triple*, of the form $\{p\} c \{q\}$ where c is a program and p and q are the corresponding *precondition* and *postcondition*. The precondition describes the set of initial or input states, in which the program c is started, and the postcondition describes the set of final or output states.

Then, we say that a formula $\{p\} c \{q\}$ is *valid* (or true) in the sense of partial correctness iff every terminating computation of c that starts in a state s satisfying p ends in a state satisfying q . This definition does not take diverging computations of c into account. Following [Apt and Olderog, 1991], we formalize this interpretation as set theoretic inclusions.

For a command c of type α *com* we define *validity of a partial correctness formula*, and write $\models p \ c \ q$ as follows:

constdefs

$$\begin{aligned} \text{com-validity} &:: \alpha \text{ assn} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ assn} \Rightarrow \text{bool} && (\models - - -) \\ \models p \ c \ q &\equiv \text{SEM } c \ p \subseteq q \end{aligned}$$

where p and q are sets of initial and final states, respectively.

Validity of a partial correctness formula for an annotated command is defined analogously:

constdefs

$$\begin{aligned} \text{ann-com-validity} &:: \alpha \text{ ann-com} \Rightarrow \alpha \text{ assn} \Rightarrow \text{bool} && (\models - -) \\ \models c \ q &\equiv \text{ann-SEM } c \ (\text{pre } c) \subseteq q \end{aligned}$$

The only difference with programs of type α *com* is that we do not need an extra argument for the precondition since it is included as part of the annotated command and can be extracted via the function *pre*.

2.4 The Proof System

Given a correctness formula, we can reason about its validity directly in terms of the semantics. This methodology is commonly known as *operational* or *behavioral* reasoning and basically consists on observing the effects of the computation by unfolding the steps according to the rules of the operational semantics. This procedure should be repeated for each possible initial state. Obviously, this is very tedious for non-trivial programs (or even impossible since the set of initial states may be infinite). Programmers using this procedure tend to give an informal account of the possible behaviors of the program for certain inputs, but experience has shown that this lack of structure is bound to fail due to a non-exhaustive study of the possibilities.

Hoare's approach to program verification is based on proving that a given correctness formula is derivable in a system of axioms and inference rules which are syntax oriented. The relation between the program syntax and its semantics is studied exclusively in the proof of soundness of the system. Once this is achieved, we can forget the semantics completely and just use the system of rules for program verification.

The main properties of such systems are soundness and completeness. The rules of a *sound* system inductively define a set of correct programs. If the system is also *complete* then it defines the set of *all* correct programs.

We define three systems of rules, one for each level of the language. However, due to the interdependencies among the different levels, the three sets which are inductively defined from the three systems of rules are mutually recursive. Thus, they must be declared in a single inductive set definition. Atomic and parallel programs share the rules for the common constructors.

In this section we progressively show the three systems and give the full set of rules only at the end. We start with the system for atomic programs. The definition of the rules and the soundness proof is independent of the other rules. We continue with the system for component programs which depends on the previous one and finish with the rule for parallel composition.

2.4.1 Proof System for Atomic Programs

The set of derivable correctness formulas of atomic programs is inductively defined by the set:

consts *oghoare* :: (α *assn* \times α *com* \times α *assn*) *set*

The syntax $\Vdash p \ c \ q$ denotes that the triple (p, c, q) is an element of the inductively defined set. This amounts to say that it can be derived in the system.

syntax *-oghoare* :: α *assn* \Rightarrow α *com* \Rightarrow α *assn* \Rightarrow *bool* (\Vdash - - -)

translations $\Vdash p \ c \ q \Leftrightarrow (p, c, q) \in$ *oghoare*

A complete system contains at least one axiom or inference rule for each constructor of the language. The rules defining this set are shown in table 2.1.

2.4.2 Proof System for Component Programs

Correctness formulas for component programs are not just annotated with a precondition and a postcondition. As explained in §2.1 they appear as proof outlines, where each command is preceded by its precondition. The set of derivable proof outlines is defined by the constant:

consts *ann-hoare* :: (α *ann-com* \times α *assn*) *set*

A pair (c, q) is an element of the set if it has derivation in the system, denoted $\vdash c \ q$:

<i>Basic:</i>	$\Vdash \{s. f s \in q\} \text{ (Basic } f) q$
<i>Seq:</i>	$\llbracket \Vdash p c_1 r; \Vdash r c_2 q \rrbracket \implies \Vdash p \text{ (Seq } c_1 c_2) q$
<i>Cond:</i>	$\llbracket \Vdash (p \cap b) c_1 q; \Vdash (p \cap \neg b) c_2 q \rrbracket$ $\implies \Vdash p \text{ (Cond } b c_1 c_2) q$
<i>While:</i>	$\llbracket \Vdash (p \cap b) c p \rrbracket \implies \Vdash p \text{ (While } b i c) (p \cap \neg b)$
<i>Conseq:</i>	$\llbracket p' \subseteq p; \Vdash p c q; q \subseteq q' \rrbracket \implies \Vdash p' c q'$

Table 2.1: Proof rules for atomic commands.

syntax $\text{-ann-hoare} :: \alpha \text{ ann-com} \Rightarrow \alpha \text{ assn} \Rightarrow \text{bool} \quad (\vdash - -)$

translations $\vdash c q \Leftrightarrow (c, q) \in \text{ann-hoare}$

The formation rules for proof outlines are shown in table 2.2. They look unusual because preconditions are hidden as part of the commands' syntax. The consequence rule does not permit to strengthen the precondition. However, this possibility appears embedded in each of the other rules. The following result shows that this system is equivalent to the standard presentation.

Equivalence of Proof Systems.

Let $\vdash_{st} p \kappa q$ stand for provability of the correctness formula $p \kappa q$ in some standard system. By κ we mean commands without any annotation other than loop invariants. The relation $c \sim \kappa$ means that both commands are equal except for the annotations (loop invariants must also be equal). Then,

1. $\vdash_{st} p \kappa q \implies \exists c. \vdash c q \wedge p \subseteq \text{pre } c \wedge \kappa \sim c$
2. $\vdash c q \implies \exists \kappa. \vdash_{st} (\text{pre } c) \kappa q \wedge \kappa \sim c$

Proof. Both directions are proven by rule induction on \vdash_{st} and \vdash respectively. We have not formalized them in Isabelle, but we hope to convince the reader that the system above can be used without any loss of generality:

1. To prove the first implication we have to consider six cases, one for each rule (the standard system contains only one rule for the conditional-statement). We only show two representative ones:

<i>AnnBasic</i> :	$r \subseteq \{s. f s \in q\} \implies \vdash (\text{AnnBasic } r f) q$
<i>AnnSeq</i> :	$\llbracket \vdash c_0 \text{ pre } c_1; \vdash c_1 q \rrbracket \implies \vdash (\text{AnnSeq } c_0 c_1) q$
<i>AnnCond₁</i> :	$\llbracket r \cap b \subseteq \text{pre } c_1; \vdash c_1 q; r \cap -b \subseteq \text{pre } c_2; \vdash c_2 q \rrbracket \implies \vdash (\text{AnnCond}_1 r b c_1 c_2) q$
<i>AnnCond₂</i> :	$\llbracket r \cap b \subseteq \text{pre } c; \vdash c q; r \cap -b \subseteq q \rrbracket \implies \vdash (\text{AnnCond}_2 r b c) q$
<i>AnnWhile</i> :	$\llbracket r \subseteq i; i \cap b \subseteq \text{pre } c; \vdash c i; i \cap -b \subseteq q \rrbracket \implies \vdash (\text{AnnWhile } r b i c) q$
<i>AnnAwait</i> :	$\llbracket \text{atom-com } c; \Vdash (r \cap b) c q \rrbracket \implies \vdash (\text{AnnAwait } r b c) q$
<i>AnnConseq</i> :	$\llbracket \vdash c q; q \subseteq q' \rrbracket \implies \vdash c q'$

Table 2.2: Proof rules for annotated commands.

Basic Suppose that the equivalent non-annotated basic command is called **Basic** in some standard system. We have to prove

$$\begin{aligned} & \vdash_{st} \{s. f s \in q\} (\text{Basic } f) q \implies \\ & \exists c. \vdash c q \wedge \{s. f s \in q\} \subseteq \text{pre } c \wedge (\text{Basic } f) \sim c \end{aligned}$$

The annotated command $\text{AnnBasic } \{s. f s \in q\} f$ fulfills the requirements.

Seq Suppose $\vdash_{st} p (c_1; c_2) q$. From the rule of sequential composition we have for some r , $\vdash_{st} p c_1 r$ and $\vdash_{st} r c_2 q$. By induction hypothesis we obtain for some ca and some cb the assumptions:

$$\begin{aligned} & \vdash ca r \wedge p \subseteq \text{pre } ca \wedge c_1 \sim ca, \text{ and} \\ & \vdash cb q \wedge r \subseteq \text{pre } cb \wedge c_2 \sim cb. \end{aligned}$$

We want to find a command c such that

$$\vdash c q \wedge p \subseteq \text{pre } c \wedge c_1; c_2 \sim c$$

Taking $c = \text{AnnSeq } ca cb$ we have from $c_1 \sim ca$ and $c_2 \sim cb$ that $c_1; c_2 \sim \text{AnnSeq } ca cb$.

From $p \subseteq pre\ ca$ and $pre\ AnnSeq\ ca\ cb = pre\ ca$ (by definition of pre), we obtain $p \subseteq pre\ AnnSeq\ ca\ cb$. By the rule of consequence,

$$\vdash ca\ r \wedge r \subseteq pre\ cb \implies \vdash ca\ pre\ cb$$

And finally, by the rule for sequential composition,

$$\vdash ca\ (pre\ cb) \wedge \vdash cb\ q \implies \vdash AnnSeq\ ca\ cb\ q$$

2. We illustrate the opposite direction with two cases not considered before:

AnnAwait Suppose the await-command in the standard system is called **Await**. We want to prove

$$\vdash (AnnAwait\ r\ b\ c)\ q \implies \exists \kappa. \vdash_{st}\ r\ \kappa\ q \wedge \kappa \sim (Await\ r\ b\ c)$$

By the proof rule *AnnAwait*, we know that $\Vdash (r \cap b)\ c\ q$. The system of rules for atomic programs is like the standard system and since c is an atomic command it does not contain assertions, thus we can assume $\vdash_{st}\ (r \cap b)\ c\ q$.

If we choose $\kappa = \mathbf{Await}\ b\ c$ we can derive $\vdash_{st}\ r\ (\mathbf{Await}\ b\ c)\ q$ and obviously $AnnAwait\ r\ b\ c \sim \mathbf{Await}\ b\ c$ holds.

AnnConseq Suppose $\vdash c'\ q'$. By the rule of consequence we have the premises $\vdash c'\ q$ and $q \subseteq q'$ for some q . By hypothesis there is a non-annotated command κ such that $\vdash_{st}\ pre\ c'\ \kappa\ q$ and $\kappa \sim c'$. Obviously, κ is the searched program:

$$\vdash_{st}\ pre\ c'\ \kappa\ q \wedge q \subseteq q' \implies \vdash_{st}\ pre\ c'\ \kappa\ q'$$

This concludes the proof of equivalence of our system of rules and a standard one like the one in [Apt and Olderog, 1991], where preconditions are not directly attached to commands. \square

2.4.3 Proof System for Parallel Programs

Proof of correctness of parallel programs is much more demanding than the sequential case. The problem is that different components can interfere with each other via shared variables. Unfortunately, proving that all proof outlines are correct independently of the environment is not sufficient to conclude that the input/output specification of a parallel composition is the intersection (assertions are modeled as sets) of the input/output specification of each component. We also need to guarantee that the proof outline of

any component is not falsified by the execution of the others. This property, called *interference freedom of proof outlines*, is determined by the predicate *interfree*. Its definition requires a number of auxiliary concepts:

- An assertion p is invariant under execution of an atomic command a iff $\models (p \cap \text{pre } a) \ a \ p$.
- An atomic command a *does not interfere* with a standard proof outline $c \ q$ iff the following two conditions hold:
 1. $\models (q \cap \text{pre } a) \ a \ q$,
 2. For any assertion p within c : $\models (p \cap \text{pre } a) \ a \ p$
- Standard proof outlines $c_1 \ q_1, \dots, c_n \ q_n$ are interference free if no assignment or atomic region of a program c_i interferes with the proof outline $c_j \ q_j$ of another component with $i \neq j$.

Given two component program's proof outlines $c \ q$ and $c' \ q'$, showing interference freedom means proving that all assertions in the former remain invariant under execution of all assignments or atomic regions in the latter, and vice versa.

Atomic commands are collected by the function *atomics* which, given an annotated command, returns the set of all pairs (r, a) where a is either the body of an *AnnAwait*-command, or a *Basic*-command (from an *AnnBasic*-command) and r is the corresponding precondition.

consts *atomics* :: $\alpha \ \text{ann-com} \Rightarrow (\alpha \ \text{assn} \times \alpha \ \text{com}) \ \text{set}$

primrec

atomics (*AnnBasic* $r \ f$) = $\{(r, \text{Basic } f)\}$
atomics (*AnnSeq* $c_1 \ c_2$) = *atomics* $c_1 \cup$ *atomics* c_2
atomics (*AnnCond*₁ $r \ b \ c_1 \ c_2$) = *atomics* $c_1 \cup$ *atomics* c_2
atomics (*AnnCond*₂ $r \ b \ c$) = *atomics* c
atomics (*AnnWhile* $r \ b \ i \ c$) = *atomics* c
atomics (*AnnAwait* $r \ b \ c$) = $\{(r \cap b, c)\}$

The set of all assertions of an annotated command (including loop invariants) is collected by the function *assertions*:

consts *assertions* :: $\alpha \ \text{ann-com} \Rightarrow (\alpha \ \text{assn}) \ \text{set}$

primrec

assertions (*AnnBasic* $r \ f$) = $\{r\}$
assertions (*AnnSeq* $c_1 \ c_2$) = *assertions* $c_1 \cup$ *assertions* c_2

Parallel:

$$\begin{aligned} & \llbracket \forall i < \text{length } Ts. \exists c q. Ts!i = (\text{Some } c, q) \wedge \vdash c q; \text{interfree } Ts \rrbracket \implies \\ & \quad \Vdash (\bigcap i \in \{i. i < \text{length } Ts\}. \text{pre } (\text{the } (\text{com } (Ts!i)))) \\ & \quad \quad \quad \text{Parallel } Ts \\ & \quad \quad \quad (\bigcap i \in \{i. i < \text{length } Ts\}. \text{post } (Ts!i)) \end{aligned}$$

Table 2.3: Proof rule for parallel programs.

$$\begin{aligned} \text{assertions } (\text{AnnCond}_1 r b c_1 c_2) &= \{r\} \cup \text{assertions } c_1 \cup \text{assertions } c_2 \\ \text{assertions } (\text{AnnCond}_2 r b c) &= \{r\} \cup \text{assertions } c \\ \text{assertions } (\text{AnnWhile } r b i c) &= \{r, i\} \cup \text{assertions } c \\ \text{assertions } (\text{AnnAwait } r b c) &= \{r\} \end{aligned}$$

The interference freedom test in one direction, i.e. where the assertions in (co, q) are checked for invariance against the atomic actions in co' , is realized by the function *interfree-aux*:

$$\begin{aligned} \text{constdefs } \text{interfree-aux} &:: (\alpha \text{ ann-com-op} \times \alpha \text{ assn} \times \alpha \text{ ann-com-op}) \Rightarrow \text{bool} \\ \text{interfree-aux} &\equiv \lambda(co, q, co'). co' = \text{None} \vee \\ & \quad (\forall(r, a) \in \text{atomics } (\text{the } co'). \Vdash (q \cap r) a q \wedge \\ & \quad (co = \text{None} \vee (\forall p \in \text{assertions } (\text{the } co). \Vdash (p \cap r) a p))) \end{aligned}$$

The function *interfree-aux* must be applied to all possible combinations of component programs, except for a component program with itself. Hence, the definition of *interfree* becomes:

$$\begin{aligned} \text{constdefs } \text{interfree} &:: (\alpha \text{ ann-triple-op}) \text{ list} \Rightarrow \text{bool} \\ \text{interfree } Ts &\equiv \forall i j. i < \text{length } Ts \wedge j < \text{length } Ts \wedge i \neq j \longrightarrow \\ & \quad \text{interfree-aux } (\text{com } (Ts!i), \text{post } (Ts!i), \text{com } (Ts!j)) \end{aligned}$$

The rule for parallel composition shown in table 2.3 claims that if all component programs are correct and interference free, then the parallel composition satisfies the formula where the precondition is the intersection of all the components' preconditions and the postcondition is the intersection of all the components' postconditions. Each element of Ts is a pair of an optional command $\alpha \text{ ann-com-op}$ and a postcondition $\alpha \text{ assn}$. The function *post* extracts the postcondition, *com* extracts the optional command $\alpha \text{ ann-com-op}$, the predefined function *the* extracts the command c from *Some* c (by assumption all commands are wrapped up in *Some*), and finally

pre extracts the precondition. This rule together with the rules for atomic programs constitute the system of rules for parallel programs.

Because of the interdependencies among the rules, the full system consisting of the rules that define the set *oghoare* and the rules that define the set *ann-hoare* must be declared simultaneously in a so-called *mutually inductive definition*:

inductive *oghoare ann-hoare*

intros

$$\text{AnnBasic: } r \subseteq \{s. f s \in q\} \Longrightarrow \vdash (\text{AnnBasic } r f) q$$

$$\text{AnnSeq: } \llbracket \vdash c_0 \text{ pre } c_1; \vdash c_1 q \rrbracket \Longrightarrow \vdash (\text{AnnSeq } c_0 c_1) q$$

$$\text{AnnCond}_1: \llbracket r \cap b \subseteq \text{pre } c_1; \vdash c_1 q; r \cap -b \subseteq \text{pre } c_2; \vdash c_2 q \rrbracket \\ \Longrightarrow \vdash (\text{AnnCond}_1 r b c_1 c_2) q$$

$$\text{AnnCond}_2: \llbracket r \cap b \subseteq \text{pre } c; \vdash c q; r \cap -b \subseteq q \rrbracket \Longrightarrow \vdash (\text{AnnCond}_2 r b c) q$$

$$\text{AnnWhile: } \llbracket r \subseteq i; i \cap b \subseteq \text{pre } c; \vdash c i; i \cap -b \subseteq q \rrbracket \\ \Longrightarrow \vdash (\text{AnnWhile } r b i c) q$$

$$\text{AnnAwait: } \llbracket \text{atom-com } c; \Vdash (r \cap b) c q \rrbracket \Longrightarrow \vdash (\text{AnnAwait } r b c) q$$

$$\text{AnnConseq: } \llbracket \vdash c q; q \subseteq q' \rrbracket \Longrightarrow \vdash c q'$$

$$\text{Parallel: } \llbracket \forall i < \text{length } Ts. \exists c q. Ts!i = (\text{Some } c, q) \wedge \vdash c q; \text{interfree } Ts \rrbracket \\ \Longrightarrow \Vdash (\bigcap i \in \{i. i < \text{length } Ts\}. \text{pre } (\text{the } (\text{com } (Ts!i)))) \\ \text{Parallel } Ts \\ (\bigcap i \in \{i. i < \text{length } Ts\}. \text{post } (Ts!i))$$

$$\text{Basic: } \Vdash \{s. f s \in q\} (\text{Basic } f) q$$

$$\text{Seq: } \llbracket \Vdash p c_1 r; \Vdash r c_2 q \rrbracket \Longrightarrow \Vdash p (\text{Seq } c_1 c_2) q$$

$$\text{Cond: } \llbracket \Vdash (p \cap b) c_1 q; \Vdash (p \cap -b) c_2 q \rrbracket \Longrightarrow \Vdash p (\text{Cond } b c_1 c_2) q$$

$$\text{While: } \llbracket \Vdash (p \cap b) c p \rrbracket \Longrightarrow \Vdash p (\text{While } b i c) (p \cap -b)$$

$$\text{Conseq: } \llbracket p' \subseteq p; \Vdash p c q; q \subseteq q' \rrbracket \Longrightarrow \Vdash p' c q'$$

Like in the definition of the semantics, atomic and parallel programs share the rules for the common constructors. Thus, to denote that a triple (p, c, q) , where c is a parallel program, is derivable in the system we use the same syntax as for atomic programs $\Vdash p \ c \ q$.

We can refer to a particular rule, for example *Seq*, by writing the prefix *oghoare-ann-hoare*, i.e. *oghoare-ann-hoare.Seq*.

2.4.4 Auxiliary Variables

An important aspect of the Owicki-Gries method is the use of auxiliary variables. They augment the program with additional information for proof purposes. Therefore, auxiliary variables should neither affect the control flow nor the data flow of the program. In fact, they are only allowed to appear in assignments of the form $a := t$, where a is an auxiliary variable. Since auxiliary variables may not appear in boolean expressions or in assignments to program variables, they are superfluous to the real computation and can therefore be eliminated.

Auxiliary variables record information about the course of the computation in a program which cannot in general be captured by the program variables alone. There are two main kinds of auxiliary variables:

- *History variables*: only one such auxiliary variable is required for the full parallel program. It records the values of all program variables atomically with every assignment or atomic region. At the end of the computation the history variable contains the full sequence of states that the execution has gone through.
- *Location variables*: in this case, a different auxiliary variable of this kind is introduced for each component of a parallel program. These variables keep track of the location where control flow resides at each moment of the computation by means of labels, where a different label is needed for every possible control point.

There are several well-known systematic procedures for the introduction of auxiliary variables of both kinds [Owicki, 1975, Apt and Olderog, 1991, Best, 1996, de Roever *et al.*, 2000]. Such general procedures are essential for the completeness proof of the Owicki-Gries method because they work for every possible program. However, history variables are too complicated in practice and the general procedure for location variables introduces too many auxiliary variables. A more clever proof for a particular program

can normally extract the needed information from a few suitable auxiliary variables.

The general approach consists of extending the program by the assignments to auxiliary variables, proving the correctness of the extended program and then deleting the added assignments. This last step is done with the elimination rule

$$p \ c \ q \Longrightarrow p \ c^* \ q$$

where for some set of auxiliary variables A used in the program c such that $(\text{free variables in } q) \cap A = \emptyset$, the program c^* is obtained from c by deleting all assignments to the variables in A .

The need for the auxiliary variable rule is a recurrent issue in research about verification systems for parallel programs. Stirling describes the need for this elimination rule as a “major disadvantage” [Stirling, 1988]. Further studies have demonstrated that it is possible to design a complete verification calculus for parallel programs with shared variables where the auxiliary structure is only a part of the logic, so that the program text need not be modified [Soundararajan, 1984, Stølen, 1991].

Our current proof system is incomplete because there is no rule for removing auxiliary variables. We can prove the correctness of the extended program, but it is left to the user to ensure that auxiliary variables are used correctly.

2.5 Soundness

We are interested in the following soundness property:

$$\Vdash p \ c \ q \Longrightarrow \models p \ c \ q$$

that whenever a correctness formula for a parallel program is derivable in the proof system then it is also valid in the sense of partial correctness.

Properties of inductively defined sets are usually proven by rule induction. The theorem describing this proof principle is automatically generated by Isabelle for every inductively defined set. The idea is to prove that a given property is true for all axioms of the system and that it is preserved by all inference rules. Since an inductively defined set is the least set closed under the given axioms and rules, every element of the set that has a derivation in the system satisfies the property.

The proof of soundness proceeds in three stages that correspond to the three subsystems. In §2.5.1, we prove soundness of the subsystem for atomic

programs. This result is necessary for the soundness proof of the subsystem for component (annotated) programs in §2.5.2. Section §2.5.3 presents the soundness of the rule for parallel programs. With these results, we finally prove soundness of the full system.

2.5.1 Soundness of the System for Atomic Programs

The proof is done by rule induction on the set of rules defining *oghoare*.

theorem *atom-hoare-sound*: $\llbracket \Vdash p \ c \ q; \text{atom-com } c \rrbracket \implies \models p \ c \ q$

This amounts to proving the soundness of each rule separately. We require that the program be atomic, consequently the subgoal concerning the *Parallel* rule is trivially eliminated. The proofs for the other rules follow directly from the lemmas about the semantics *SEM* (cf. §2.3).

2.5.2 Soundness of the System for Component Programs

A correctness formula is valid in the sense of partial correctness iff whenever a program c started in a state satisfying the precondition terminates, then the final state satisfies the postcondition. Observe that this definition does not mention the intermediate assertions. This is fine for non-annotated programs, but in our case satisfiability of the intermediate annotations is also relevant.

Informally speaking, proof outlines fulfill the property that whenever the control of c in a given computation starting in a state $s \in p$ reaches a point annotated by an assertion, this assertion is true. This property is called *strong soundness*. The standard soundness property follows trivially from this theorem.

Strong Soundness for Component Programs

This property of proof outlines is formally proven in the following theorem:

theorem *Strong-Soundness*:

$$\begin{aligned} & \llbracket (\text{Some } c, s) \text{ -*} \rightarrow (co, t); s \in \text{pre } c; \vdash c \ q \rrbracket \\ & \implies \text{if } co = \text{None then } t \in q \text{ else } t \in \text{pre } (\text{the } co) \end{aligned}$$

where *the* is a predefined function that extracts c from *Some* c .

Proof. The proof is by induction on the length of the computation. However, the *Strong-Soundness* theorem is not yet suitably formulated. To be able to apply the induction hypothesis we need the fact that the program rest, i.e. *the co* is also derivable in the system:

lemma *Strong-Soundness-aux*: $\llbracket (Some\ c, s) \multimap (co, t); s \in pre\ c; \vdash\ c\ q \rrbracket$
 \implies if $co = None$ then $t \in q$ else $t \in pre\ (the\ co) \wedge \vdash\ (the\ co)\ q$

If the length of the computation is 0 then $co = (Some\ c)$ and $t = s$. By hypothesis we know that $s \in pre\ c$, then $t \in pre\ c$. Suppose the length is now positive. Then, for some co' and t' we have

$$(Some\ c, s) \multimap (co', t') \multimap (co, t)$$

co' cannot be *None* because there is no possible transition from *None* in the system *ann-transition*. Thus, there is a c' so that $co' = Some\ c'$. By the induction hypothesis we know that $t' \in pre\ c'$ and $\vdash\ c'\ q$. The proof follows by rule induction on the last step. This is achieved via the following auxiliary lemma:

lemma *Strong-Soundness-aux-aux*:
 $\llbracket (co, s) \multimap (co', t); co = Some\ c; s \in pre\ c; \vdash\ c\ q \rrbracket$
 \implies if $co' = None$ then $t \in q$ else $t \in pre\ (the\ co') \wedge \vdash\ (the\ co')\ q$

We discuss three representative cases from the ten that result from applying rule induction on $(co, s) \multimap (co', t)$.

Seq2: From this rule we obtain $(Some\ c_0, s) \multimap (Some\ c_2, t)$ in the premises. After applying the induction hypothesis and using $pre\ (AnnSeq\ c_0\ c_1) = pre\ c_0$ we obtain:

$$\llbracket (Some\ c_0, s) \multimap (Some\ c_2, t); s \in pre\ c_0; \vdash\ (AnnSeq\ c_0\ c_1)\ q; \\ \forall q. \vdash\ c_0\ q \longrightarrow t \in pre\ c_2 \wedge \vdash\ c_2\ q \rrbracket \\ \implies t \in pre\ c_2 \wedge \vdash\ (AnnSeq\ c_2\ c_1)\ q$$

By case analysis on $\vdash\ (AnnSeq\ c_0\ c_1)\ q$ we obtain $\vdash\ c_0\ (pre\ c_1)$ and $\vdash\ c_1\ q$ from the *Seq2* rule. Unfortunately, we also obtain $\vdash\ (AnnSeq\ c_0\ c_1)\ q'$ and $q' \subseteq q$ from the rule of consequence. This second subgoal is basically the original subgoal. This circular effect is due to the generality of the consequence rule. This rule is so general that it can always be applied. To avoid this, we prove a more appropriate version of the inductive cases principle which applies the consequence rule at most once:

lemma *ann-hoare-case-analysis*: $\vdash C q' \Longrightarrow$
 $(\forall r f. C = \text{AnnBasic } r f \longrightarrow (\exists q. r \subseteq \{s. f s \in q\} \wedge q \subseteq q')) \wedge$
 $(\forall c_0 c_1. C = \text{AnnSeq } c_0 c_1 \longrightarrow (\exists q. q \subseteq q' \wedge \vdash c_0 \text{ pre } c_1 \wedge \vdash c_1 q)) \wedge$
 $(\forall r b c_1 c_2. C = \text{AnnCond}_1 r b c_1 c_2 \longrightarrow (\exists q. q \subseteq q' \wedge$
 $r \cap b \subseteq \text{pre } c_1 \wedge \vdash c_1 q \wedge r \cap \neg b \subseteq \text{pre } c_2 \wedge \vdash c_2 q)) \wedge$
 $(\forall r b c. C = \text{AnnCond}_2 r b c \longrightarrow$
 $(\exists q. q \subseteq q' \wedge r \cap b \subseteq \text{pre } c \wedge \vdash c q \wedge r \cap \neg b \subseteq q)) \wedge$
 $(\forall r i b c. C = \text{AnnWhile } r b i c \longrightarrow$
 $(\exists q. q \subseteq q' \wedge r \subseteq i \wedge i \cap b \subseteq \text{pre } c \wedge \vdash c i \wedge i \cap \neg b \subseteq q)) \wedge$
 $(\forall r b c. C = \text{AnnAwait } r b c \longrightarrow (\exists q. q \subseteq q' \wedge \Vdash (r \cap b) c q))$

Using this theorem instead we obtain only $\vdash c_0$ ($\text{pre } c_1$) and $\vdash c_1 q$. By instantiating $\forall q. \vdash c_0 q \longrightarrow t \in \text{pre } c_2 \wedge \vdash c_2 q$ with $\text{pre } c_1$ we prove $t \in \text{pre } c_2$ and obtain $\vdash c_2$ ($\text{pre } c_1$). The remaining conclusion, $\vdash (\text{AnnSeq } c_2 c_1) q$, follows from the rule *Seq2*.

WhileT: After some simplification the corresponding subgoal is:

$$\begin{aligned} & \llbracket s \in b; s \in r; \vdash (\text{AnnWhile } r b i c) q; r \subseteq i; i \cap b \subseteq \text{pre } c; \vdash c i; i \cap \neg b \subseteq q \rrbracket \\ & \Longrightarrow \vdash (\text{AnnSeq } c (\text{AnnWhile } i b i c)) q \end{aligned}$$

After applying *AnnSeq* backwards we obtain two subgoals. The first one

$$\begin{aligned} & \llbracket s \in b; s \in r; \vdash (\text{AnnWhile } r b i c) q; r \subseteq i; i \cap b \subseteq \text{pre } c; \vdash c i; i \cap \neg b \subseteq q \rrbracket \\ & \Longrightarrow \vdash c \text{ pre } (\text{AnnWhile } i b i c) \end{aligned}$$

is proven by simplification because $\text{pre } (\text{AnnWhile } i b i c) = i$. Observe that this follows from the definition of the semantics rule *AnnWhileT*. We mentioned in §2.2 that the annotations do not play any role in the definition of the rules of the semantics. However, if we wrote r instead of i for that precondition, this subgoal would not be provable. The second subgoal is solved by applying *AnnWhile* backwards.

Await: This is an axiom of the system so there is no induction hypothesis:

$$\begin{aligned} & \llbracket s \in b; \text{atom-com } c; (c, s) -P^* \rightarrow (\text{Parallel } [], t); s \in r; \\ & \vdash (\text{AnnAwait } r b c) q \rrbracket \Longrightarrow t \in q \end{aligned}$$

By case analysis on $\vdash (\text{AnnAwait } r b c)$ we obtain $\Vdash (r \cap b) c q$. The system for atomic programs is sound, thus $\Vdash (r \cap b) c q$. From $s \in r \cap b$ and the

definition of validity for atomic programs we prove $t \in q$. □

Finally, we state the soundness theorem for component programs:

theorem *ann-hoare-sound*: $\vdash c q \implies \models c q$

The proof is immediate using the *Strong-Soundness* theorem.

2.5.3 Soundness of the System for Parallel Programs

The most interesting result of the soundness proof is the soundness of the *Parallel* rule. Like for component programs, we first show the corresponding stronger property called *strong soundness for parallel programs*.

Intuitively, if the proof outline of each component program satisfies the strong soundness property and the actions of the other components do not “interfere”, then every component is able to establish the intended postcondition. Then, if all components finish their computations the final state satisfies all postconditions simultaneously.

For example, consider the standard proof outlines (written in a standard syntax) $\{x = 0\} x := x+1 \{x = 1\}$ and $\{True\} x := 0 \{x = 0\}$. They are obviously correct, but not interference free. For instance, the postcondition $\{x = 0\}$ is not preserved under the execution of $x := x+1$.

However, if we weaken the postconditions and consider the annotations $\{x = 0\} x := x+1 \{x = 0 \vee x = 1\}$ and $\{True\} x := 0 \{x = 0 \vee x = 1\}$ we obtain both, correct and interference free proof outlines.

Finding annotations for each component that are strong enough to satisfy its specification, and yet weak enough to remain invariant under the execution of all atomic actions of other components is often a difficult task that requires perseverance.

The strong soundness theorem for parallel programs states that whenever flow of control reaches a point annotated by an assertion, this assertion is true. The difference is that in a parallel program the control resides simultaneously at several points. Thus, we have to prove that the assertions attached to those points are simultaneously true.

Strong Soundness Theorem for Parallel Programs

Let Ts be a list of pairs formed by (optional) component programs and their postcondition. Suppose that each component $Ts!i$ such that $Ts!i =$

(*Some c, q*) for some annotated command c and some postcondition q , has a derivation in the system *ann-hoare*, i.e. $\vdash c \ q$, and *interfree Ts* also holds.

Assume also that $(\text{Parallel } Ts, s) -P^* \rightarrow (\text{Parallel } Rs, t)$ for some list of component programs Rs and some states s, t such that s satisfies the precondition of all component programs $Ts!i$. Then, all component programs $Rs!j$ of *Parallel Rs* satisfy that

- if $\text{com } (Rs!j) = \text{Some } c$ for a command c , then $t \in \text{pre } c$,
- if $\text{com } (Rs!j) = \text{None}$, then $t \in \text{post } (Rs!j)$.

In particular if $\text{com } (Rs!j) = \text{None}$ for all j , we have that $t \in \text{post } (Rs!j)$ for all j such that $j < \text{length } Rs$.

The formal lemma as formulated in Isabelle is:

lemma *Parallel-Strong-Soundness*:

$$\begin{aligned} & \llbracket (\text{Parallel } Ts, s) -P^* \rightarrow (\text{Parallel } Rs, t); \text{interfree } Ts; j < \text{length } Rs; \\ & \quad \forall i < \text{length } Ts. \exists c \ q. Ts!i = (\text{Some } c, q) \wedge s \in \text{pre } c \wedge \vdash c \ q \rrbracket \\ & \implies \text{if } \text{com } (Rs!j) = \text{None} \text{ then } t \in \text{post } (Ts!j) \text{ else } t \in \text{pre } (\text{the } (\text{com } (Rs!j))) \end{aligned}$$

Proof. Like in the case of component programs, the theorem in the above form is too weak. The conclusion must establish two more properties of the reached configuration, namely, that the program fragment *the (com (Rs!j))* has a derivation in the system and that the list of component programs after the transition still satisfies the interference freedom property, i.e. *interfree Rs*. In other words, we have to prove that derivability of a component's proof outline and the interference freedom of a list of proof outlines are preserved throughout the computation:

lemma *Parallel-Strong-Soundness-aux*:

$$\begin{aligned} & \llbracket (Ts', s) -P^* \rightarrow (Rs', t); Ts' = (\text{Parallel } Ts); \\ & \quad \forall i < \text{length } Ts. \exists c \ q. Ts!i = (\text{Some } c, q) \wedge s \in \text{pre } c \wedge \vdash c \ q; \text{interfree } Ts \rrbracket \\ & \implies \forall Rs. Rs' = (\text{Parallel } Rs) \longrightarrow \\ & \quad (\forall j < \text{length } Rs. (\text{if } \text{com } (Rs!j) = \text{None} \text{ then } t \in \text{post } (Ts!j) \\ & \quad \text{else } t \in \text{pre } (\text{the } (\text{com } (Rs!j)))) \wedge \vdash \text{the } (\text{com } (Rs!j)) \ \text{post } (Ts!j))) \wedge \\ & \quad \text{interfree } Rs \end{aligned}$$

The proof is by induction on the length of the computation. If the length is 0, the proof is trivial since $Ts' = Rs'$ and $s = t$. If the length is > 0 , then for some list of component programs Ss and some state b :

$$(\text{Parallel } Ts, s) -P^* \rightarrow (\text{Parallel } Ss, b) -P1 \rightarrow (\text{Parallel } (Ss[i:=\text{(co, q)}]), t)$$

where the last step is performed by the i th component of Ss through transition $(Some\ c, b) \rightarrow (co, t)$ and $Rs' = Parallel\ (Ss[i := (co, q)])$.

The conclusion of the lemma is a conjunction of two clauses. The second one, namely $interfree\ (Ss[i := (co, q)])$ is proven with the following lemma:

lemma *interfree-lemma*:

$$\begin{aligned} & \llbracket (Some\ c, s) \rightarrow (co, t); interfree\ Ts; i < length\ Ts; Ts!i = (Some\ c, q) \rrbracket \\ & \implies interfree\ (Ts[i := (co, q)]) \end{aligned}$$

The proof of this lemma follows from two symmetric properties of the predicate $interfree\ aux$, both proven by rule induction on the *ann-transition* relation:

lemma *interfree-aux1*:

$$\llbracket (co, s) \rightarrow (co', t); interfree\ aux\ (co_1, q, co) \rrbracket \implies interfree\ aux\ (co_1, q, co')$$

lemma *interfree-aux2*:

$$\llbracket (co, s) \rightarrow (co', t); interfree\ aux\ (co, q, co_1) \rrbracket \implies interfree\ aux\ (co', q, co_1)$$

For the other clause of the conclusion, two cases arise: $i = j$ or $i \neq j$. The first one means that the transition occurred in the same component that we were observing, i.e. component j . The proof amounts to checking strong soundness of a proof outline which is exactly the *Strong-Soundness* theorem.

The case $i \neq j$ means that the last transition $(Some\ c, b) \rightarrow (co, t)$ was performed by a component i which is not the one we were observing, i.e. not the fixed j . We must prove that component j fulfills the conclusion of the theorem just the same.

The proof proceeds by rule induction on the last *ann-transition* relation. The corresponding ‘‘appropriate’’ auxiliary lemma is:

lemma *Parallel-Strong-Soundness-aux-aux*:

$$\begin{aligned} & \llbracket (Some\ c, b) \rightarrow (co, t); i < length\ Ts; com\ (Ts!i) = Some\ c; \\ & \quad \forall i < length\ Ts. \text{ if } com\ (Ts!i) = None \text{ then } b \in post\ (Ts!i) \\ & \quad \text{ else } b \in pre\ (the\ (com\ (Ts!i))) \wedge \vdash the\ (com\ (Ts!i))\ post\ (Ts!i); \\ & \quad interfree\ Ts; j < length\ Ts; i \neq j \rrbracket \\ & \implies \text{ if } com\ (Ts!j) = None \text{ then } t \in post\ (Ts!j) \\ & \quad \text{ else } t \in pre\ (the\ (com\ (Ts!j))) \wedge \vdash the\ (com\ (Ts!j))\ post\ (Ts!j) \end{aligned}$$

If the last step in the computation consists of the evaluation of a Boolean expression, then $b = t$. The proof follows by instantiating the universal quantification in the premise with j .

Otherwise, the last step consists of the execution of a basic action, an atomic region or some transition in a sequential composition of commands. We discuss two of them; the remaining two cases are analogous:

Basic: Suppose the command of the i th component is $AnnBasic\ r\ f$, then the last transition is $(Some\ (AnnBasic\ r\ f),\ b) \dashv\!\rightarrow (None,\ f\ b)$. By assumption $b \in r$. Then,

- If $com\ (Ts!j) = None$, then by assumption $b \in post\ (Ts!j)$. By the interference freedom hypothesis we have $\models (post\ (Ts!j) \cap r)\ Basic\ f\ post\ (Ts!j)$. From the definition of validity and $b \in post\ (Ts!j) \cap r$, we conclude that $f\ b \in post\ (Ts!j)$.
- If $com\ (Ts!j) = Some\ y$ for some command y , we obtain from the assumptions that $b \in pre\ y$. By the interference freedom of Ts we have:

$$\forall p \in assertions\ y. \models (p \cap r)\ Basic\ f\ p.$$

By structural induction on c we prove the lemma: $pre\ c \in assertions\ c$. Hence, we can instantiate the previous assumption with $pre\ y$ obtaining $\models (pre\ y \cap r)\ Basic\ f\ (pre\ y)$. Finally, from the definition of validity $f\ b \in pre\ y$.

Seq2: Suppose now the command of the i th component is $AnnSeq\ c_0\ c_1$, and the last transition is

$$(Some\ (AnnSeq\ c_0\ c_1),\ b) \dashv\!\rightarrow (Some\ (AnnSeq\ c_2\ c_1),\ t)$$

Then, from the *ann-transition* rule *Seq2* we know that

$$(Some\ c_0,\ b) \dashv\!\rightarrow (Some\ c_2,\ t)$$

We instantiate the universally quantified variable Ts in the induction hypothesis with $Ts[i := (Some\ c_0,\ pre\ c_1)]$. Thereby, we obtain the information we need about t but referring to the above instantiation. Since all components of Ts other than the component i remain unchanged, the conclusion of the induction hypothesis is exactly the conclusion of the subgoal. Thus, it remains to be shown that the premises required to validate the conclusion of the induction hypothesis are indeed fulfilled by the instantiation:

lemma *Parallel-Strong-Soundness-Seq:*

$$\begin{aligned} & \llbracket \forall i < length\ Ts. \text{ if } com\ (Ts!i) = None \text{ then } b \in post\ (Ts!i) \\ & \text{ else } b \in pre\ (the\ (com\ (Ts!i))) \wedge \vdash the\ (com\ (Ts!i))\ post\ (Ts!i); \end{aligned}$$

$$\begin{aligned}
& \text{com } (Ts!i) = \text{Some } (\text{AnnSeq } c_0 \ c_1); \ i < \text{length } Ts; \ \text{interfree } Ts \] \implies \\
& (\forall ia < \text{length } Ts. \ (\text{if } \text{com } (Ts[i:=\text{Some } c_0, \ \text{pre } c_1]]!ia) = \text{None} \\
& \quad \text{then } b \in \text{post } (Ts[i:=\text{Some } c_0, \ \text{pre } c_1]]!ia) \\
& \quad \text{else } b \in \text{pre } (\text{the } (\text{com } (Ts[i:=\text{Some } c_0, \ \text{pre } c_1]]!ia)))) \\
& \wedge \vdash \text{the } (\text{com } (Ts[i:=\text{Some } c_0, \ \text{pre } c_1]]!ia)) \ \text{post } (Ts[i:=\text{Some } c_0, \ \text{pre } c_1]]!ia)) \\
& \wedge \text{interfree } (Ts[i:=\text{Some } c_0, \ \text{pre } c_1])
\end{aligned}$$

The proof is fairly straightforward. The only modification concerns component i , i.e. we substitute $(\text{Some } \text{AnnSeq } c_0 \ c_1, \ q)$ for $(\text{Some } c_0, \ \text{pre } c_1)$. The postcondition remains the same, and so does the precondition since $\text{pre } (\text{AnnSeq } c_0 \ c_1) = \text{pre } c_0$.

To show $\vdash c_0 \ \text{pre } c_1$, we instantiate the assumption for component i obtaining $\vdash (\text{AnnSeq } c_0 \ c_1) \ q$. By the rule for sequential composition $\vdash c_0 \ \text{pre } c_1$ also holds.

At last showing $\text{interfree } Ts \implies \text{interfree } Ts[i:=\text{Some } c_0, \ \text{pre } c_1]$ is straightforward. This concludes the proof of the *Parallel-Strong-Soundness* theorem. \square

The final result is the soundness of the full system of rules for parallel programs:

theorem *oghoare-sound*: $\vdash \ p \ c \ q \implies \Vdash \ p \ c \ q$

Soundness of the rule for parallel composition follows directly from the *Parallel-Strong-Soundness* theorem. The proofs of soundness for the remaining inference rules have already been proven in the soundness proof for the system of atomic programs.

Our soundness proof is new with respect to those found in the literature. By including preconditions in the program's syntax we achieve a simpler and more intuitive formulation. The textbook we follow as a model for our formalization, namely [Apt and Olderog, 1991], defines the program syntax devoid of any annotation and attaches the preconditions separately. In order to refer to the precondition reached by the execution of a program S , they define a recursive function at such that, given a program S and a subprogram T , $at(T, S)$ returns the remainder of S that is to be executed when the control is at subprogram T . Using this function they are able to refer to the precondition of the remaining subprogram. Unfortunately, this function is not well-defined: a program S might contain several identical subprograms T , so that the behavior of at is unclear. To resolve these ambiguities they

informally propose to attach labels to each basic statement in the program. In contrast, the function *pre* of our formalization is trivially well-defined and devoid of such difficulties.

2.6 Generation of Verification Conditions

Due to the presence of the test for interference freedom, the proof method of Owicki and Gries is not *compositional*, i.e. it does not allow a derivation of a correctness specification of a parallel program from the specifications of its components *without* reference to their internal structure. This causes this method to be very costly in practice. For example, in the case of two component programs of length l_1 and l_2 , proving interference freedom requires proving $l_1 \times l_2$ additional correctness formulas. Most of them are trivially satisfied because they check an assignment or atomic region a against an assertion which is disjoint from the variables changed in a . By automating this tedious work the user can be sure that all cases are considered.

Fortunately, Hoare-like methods possess the necessary structure to be automated. The proof rules of the system are syntax directed and can be used to generate the necessary verification conditions by using the rules backwards. This process has been encapsulated in an Isabelle tactic. The generated verification conditions are statements of the logic of assertions devoid of any mention of the programming language. The correctness of the program specification depends upon the validity of these conditions, which can be checked using standard Isabelle proof strategies.

As far as the user is concerned, only the name of the defined tactic is relevant. We call the tactic *oghoare*. It is simply applied to a goal stating that some parallel program's specification (with full proof outlines for the component programs) is derivable in the system with the same name (*oghoare*). As a result a subgoal for each verification condition is generated. A detailed explanation of the design of the tactics can be found in appendix A.

2.7 Concrete Syntax

In the previous sections we used an abstract representation for the syntax of the programming language. In particular, the type of the state was left completely indeterminate. This is convenient for meta-theoretical reasoning about the language, however, in order to apply the method for verification, we need to write real programs. In this section we describe the particular

formalization of the state and introduce concrete syntax for commands and assertions that allow us to write programs in a familiar way.

2.7.1 Formalization of the State

The state of a program at some point during execution is usually defined as the tuple of values of program variables (or as a mapping that returns the values of the program variables) at each point during execution. In any case, imperative programs manipulate the state by referencing and assigning program variables. Thus, we need a representation for the state that allows these two operations.

Finding an adequate model for the representation of state spaces has been a tricky issue in the story of formal tools for verification of programs. We briefly describe two of the solutions that have been previously implemented in HOL, and the approach used in this thesis.

State as Tuple

The first one is a rather simple approach proposed by [von Wright *et al.*, 1993]. The state is represented as the tuple of the variables appearing in a particular program and implicitly abstract expressions involving variables over this tuple². For example, suppose we have the annotated program (written in a familiar syntax):

$$\mathbf{vars} \ x \ y. \ \{\!| \text{True} |\!\} \ x:=0; \ \{\!| x = 0 |\!\} \ y:=1$$

The state is then represented by the tuple (x, y) . The internal representation would be

$$\begin{aligned} \text{AnnSeq} \quad & (\text{AnnBasic} \ \{(x, y). \ \text{True}\} \ (\lambda(x, y). \ (0, y))) \\ & (\text{AnnBasic} \ \{(x, y). \ x = 0\} \ (\lambda(x, y). \ (x, 1))) \end{aligned}$$

The explicit declaration of variables $\mathbf{vars} \ x \ y$ is important for translation functions in order to distinguish program variables depending on the state from constants of the underlying logic.

With this approach variables can have any name and any type. Moreover, operations, syntax, etc. on their values can be directly inherited from Isabelle/HOL's theories and used in programs.

²This was the model that we adopted originally. A previous version of the examples verified with the Owicki-Gries method were carried out using this approach.

The main disadvantage is that variable names are bound, thus they have no first-class existence. In HOL there is no difference between $\lambda(x, y). x + y$ and $\lambda(s, t). s + t$. This complicates the translation functions which have to “remember” the original names as given by the explicit declaration to avoid renaming of variables throughout the transformations. Besides, abstraction over tuples is not primitive in HOL. It is achieved by suitable combinations of ordinary abstraction and an uncurrying function of type $(\alpha \Rightarrow \beta \Rightarrow \gamma) \Rightarrow \alpha \times \beta \Rightarrow \gamma$. A one-to-one translation between input and output syntax is sometimes impossible. For example, if the state of a program is (x, y) and the input program contains the dummy assignment $x := x$ the translation into internal syntax will return $\lambda(x, y). (x, y)$. The function that translates from internal into external syntax cannot distinguish whether the original input was $x := x$, or $y := y$.

Another disadvantage is the poor modularity: program fragments can be defined separately, but they can only be put together if they depend on the exact same state tuple.

State as Function

An alternative, originally used in [Gordon, 1989], consists of defining the state as a partial function from variable names to values: *name* \Rightarrow *value option*. This approach gives variables a first-class existence, however, types of variables have to be modeled explicitly. The simplest model would be to require all variables in a program to have the same type, this was the first formalization of program variables in Isabelle/HOL [Galm, 1995]. While this model is maybe enough for simple programs with variables ranging over numbers or booleans it is already insufficient when composed variables like arrays are required. A first way out of this situation is to use an enumerated type containing all the required types. However, at least theoretically, imperative programs can use an unlimited range of types (e.g. arrays of arrays of arrays ...). A better choice is to use a properly *recursive* type. This is the approach used in [Harrison, 1998]. It can be achieved by disjoint sum types or by inductive datatypes as follows:

types α *array* = *nat* \times α *list*

datatype *value* =

Bool bool
 | *Nat nat*
 | *Array value array*
 | *Pointer value*

Following [Harrison, 1998], arrays are represented as a pair consisting of a starting index and a list of elements. With this model it is possible to have an unlimited range of types built from a fixed set of constructors. The main problem is the need to cope with type structure within the logic. This basically means that type-correctness of programs has to be proved every time. This may be feasible for meta-theoretical studies of a programming language, but quite cumbersome in concrete verification tasks.

State as Record

Finally, a model that has all the advantages of the previous two models is based on a formalization of the state as an Isabelle/HOL record type [Naraschewski and Wenzel, 1998]. This type automatically supplies selecting and updating functions for each field. This model was first used by Markus Wenzel in his version of the Hoare logic for sequential programs in Isabelle/Isar [Wenzel, 2001b]. Program variables have a first-class existence and can range over any type. Operations on their value domains can be inherited directly from Isabelle/HOL theories.

Program variables must be previously declared as an Isabelle/HOL record type. Each variable is represented by a record field whose type is the value domain of the variable. For example, consider a program with a single variable x ranging over the natural numbers. Before writing the program, we declare the following record:

record $state = x :: nat$

Automatically we obtain a selector function: $x :: state \Rightarrow nat$, and an update function: $x\text{-update} :: nat \Rightarrow state \Rightarrow state$ such that the standard properties of record fields hold. This is optimal for our purposes: the selector function x is used to reference the value of x at a certain state, and the update function is used to model assignments to the variable.

As we shall see in the examples throughout this thesis, concrete syntax can be defined in a very elegant way. The basic idea is based on the *quote/antiquote* technique. A *quotation* is an expression which is implicitly abstracted, in our case over the state space. An *antiquotation* is a marked expression (for example by the antiquote symbol ‘`''`’) within a quotation that refers to the implicit argument, in our case to the state. An antiquotation would select (or even update) components from the state.

The syntax for quoted expressions and antiquoted expressions inside a quotation is:

syntax

-*quote* :: $\beta \Rightarrow (\alpha \Rightarrow \beta)$ ($\langle\langle - \rangle\rangle$)
 -*antiquote* :: $(\alpha \Rightarrow \beta) \Rightarrow \beta$ ($\langle' - \rangle$)

A quotation $\langle\langle b \rangle\rangle$ where b has type β represents a function $\lambda s. b$ with type $\alpha \Rightarrow \beta$. Antiquoted expressions appear within a quotation and are preceded by the symbol \langle' . For example, assume f is a function with type $\alpha \Rightarrow \beta$, then $\langle\langle'f\rangle\rangle$ is a quotation, i.e. an abstraction over some bound variable s where the function f is antiquoted, i.e. applied to the implicit argument s . Thus, the internal expression is $\lambda s. f s$.

For the case where a variable has been declared in a record, for example x above, then x is a selector function. If we write $\langle\langle'x = 0\rangle\rangle$ then the quoted expression $\langle'x = 0$ is delimited by an abstraction $(\lambda s. \langle'x = 0)$. The expression x appears antiquoted so that x is translated as a function that refers to the implicit argument: $(x s)$. As a result we obtain the internal expression $\lambda s. x s = 0$. Let us see now how assignments to variables are modeled with these techniques.

A basic-command represents any state transformation. However, programming languages usually use single assignments of the form $x := e$ where x is a variable and e a expression of the proper type. First, we define external syntax for both annotated and non-annotated basic-commands:

syntax

-*Assign* :: $idt \Rightarrow \beta \Rightarrow \alpha \text{ com}$ ($\langle' - := - \rangle$)
 -*AnnAssign* :: $\alpha \text{ assn} \Rightarrow idt \Rightarrow \beta \Rightarrow \alpha \text{ com}$ ($\langle - \langle' - := - \rangle \rangle$)

On the left side of the assignment we write simply an identifier which stands for the variable name. The variable appears “artificially” antiquoted in order to keep a uniform notation for variables inside the program. On the right side we write the assigned expression. Variables appearing within this expression must be antiquoted. The internal syntax uses the function *-update-name* on syntax trees which, given an argument x , returns x -*update*.

translations

$\langle'x := a \rangle \rightarrow \text{Basic} \langle\langle'(-\text{update-name } x) a \rangle\rangle$
 $r \langle'x := a \rangle := a \rightarrow \text{AnnBasic } r \langle\langle'(-\text{update-name } x) a \rangle\rangle$

As a result of this translation from external into internal syntax, if we write in our program for example $\langle'x := \langle'x + 1 \rangle$, then internally Isabelle turns it into *Basic* $(\lambda s. s(\langle'x := x s + 1 \rangle))$, where $s(\langle'x := x s + 1 \rangle)$ represents the record s where the field x has been updated to the value $x s + 1$. The defined

concrete syntax for assignments does not allow for multiple assignments, but they can obviously be expressed in the abstract syntax.

Assertions are enclosed in special brackets to avoid confusion with set notation.

syntax

-Assert :: $\alpha \Rightarrow \alpha \text{ set}$ ($\{\!\{-\}\!\}$)

An object enclosed by an assertion, say b , is a boolean expression (possibly containing antiquoted variable names), which is internally quoted, i.e. abstracted over the state. This function is then passed on as the argument of *Collect* :: $(\alpha \Rightarrow \text{bool}) \Rightarrow \alpha \text{ set}$. Internally, $\{\!\{b\}\!\}$ represents the set of states satisfying the predicate b .

translations

$\{\!\{b\}\!\} \mapsto \text{Collect } \langle\!\langle b \rangle\!\rangle$

Further advantages of this model are:

- Antiquotations mark an expression as dependent on the implicit state abstraction. This expression may be “non-atomic”, e.g. composition of functions is allowed. This is useful in chapter 4 where abstraction occurs over pairs of states (s, t) .
- Isabelle/HOL record types may be extended in a linear fashion. For example, if we verify a program that uses a variable x ranging over naturals we declare the record:

record $program_1 = x :: \text{nat}$

If later we wish to verify a second program with variables x ranging over naturals and b of boolean type, it suffices to declare the extended record:

record $program_2 = program_1 + b :: \text{bool}$

This is useful for proving derivability of a program by first proving it separately for its subprograms.

- We can also define abbreviations for parts of assertions, parts of programs, etc. as functions over the record type. Such expressions might depend on the values of the program variables in a fixed way. For

example, assume a program with variables x , y and z declared in a record called $vars$ such that the expression $'x + 'y - 'z = 'x - 'y + 'z$ appears in many assertions. We can define a predicate $P :: vars \Rightarrow bool$ over the record type as the quotation $\ll 'x + 'y - 'z = 'x - 'y + 'z \gg$ and then simply write $'P$ in the assertions.

However, as we shall see in the examples in chapter 3 some expressions over the program variables appear repeatedly but do not always depend on a fixed form of the variables. For example, consider the previous example where the expression $'x + 'y - 'z = 'x - 'y + 'z$ appears sometimes like that and sometimes as $('x + 1) + 'y - 'z = ('x + 1) - 'y + 'z$. Then, we can define the predicate P as a function over the record-state with a parameter for the value of the variable $'x$, i.e. the type of P would be $vars \Rightarrow nat \Rightarrow bool$ and its definition $P \equiv \ll \lambda x. x + 'y - 'z = x - 'y + 'z \gg$. Then, we can write $'P 'x$ or $'P ('x + 1)$ depending on the kind of occurrence in the assertions. For large programs with many variables where the predicates in the assertions are frequently repeated, these abbreviations allow us to write clear and short annotations.

With the previous method of representing the state via a tuple of bound variables [von Wright *et al.*, 1993], abbreviations could also be declared as functions over the types of the variables concerned. For example, the predicate P of the previous example would be defined as $P \equiv \lambda(x, y, z). x + y - z = x - y + z$. However, every time the predicate is used in an assertion, the arguments have to be written, i.e. $P(x, y, z)$ or $P(x + 1, y, z)$. When the predicate depends on many variables which appear always in a fixed form, the abbreviations themselves can be unnecessary long.

The many advantages of the representation of program variables used in this thesis will be clearly illustrated in the examples presented here.

2.7.2 Concrete Syntax for Commands and Assertions

In this section we introduce the concrete syntax for commands and assertions in a recipe style. For the reader interested in understanding the examples shown in this thesis and maybe also interested in using the formalization as a verification tool, it suffices to read the rest of this section. The formal specification of the syntax and the corresponding translations are shown in the appendix B.

Variables in the program code appear always marked by an antiquote symbol (`'`). For example, the variable x is written `'x` in the program text and inside assertions. In this thesis however, we automatically substitute “antiquoted” variables by the variable name in *sans serif* font (looks nicer on paper). For example, the variable `'x`, is written x in the program text and inside assertions in this thesis.

Assertions are written as boolean expressions (predicates) enclosed between the brackets `{|}` and `|}`. Thus, if we write `{|r|}` it is internally translated as the set of states satisfying the predicate r . Boolean conditions for `if`- `while`- or `await`-statements are written as normal predicates without any special marking.

Abbreviations for predicates used frequently in the assertions of a program can be given an abbreviation by stating the equality in the premises or defining it previously via **constdefs**. They are defined as abstractions over the state (quoted expressions). Inside assertions they appear antiquoted, i.e. in *sans serif* font.

Table 2.4 shows an overview of the external syntax. Each constructor of the abstract syntax is given a concrete representation. Some commands, separated by a horizontal line in the table, are simply abbreviations for special cases of the commands declared in the abstract syntax. They can be defined by declaring syntax and one-to-one translations (see appendix B). The “new” commands introduced this way are:

Skip is a basic-command whose state-transformation function is the identity.

Atomic regions are *AnnAwait*-statements where the waiting condition is *True*. They appear enclosed in angle brackets `<` and `>`.

Wait -statements are *AnnAwait*-statements where only the waiting condition is important, i.e. the body is *Skip*.

If-then -statements for commands of type α *com* can be defined by a translation with the else-part being *Skip*.

A minor problem appears if we try to define the same syntax for the sequential composition of programs at both layers. Since both have two arguments, Isabelle cannot solve the ambiguity. Thus, we define the syntax at each level slightly different. In the case of annotated programs, sequential composition of c and c' is denoted by $c;$; c' , this choice avoids clashes with the predefined `;` in Isabelle. For programs of type α *com* the sequential composition of two commands c and c' is denoted by $c., c'$.

For parallel programs there is some concrete syntax of the form

cobegin $c_0 \{q_0\} \parallel \dots \parallel c_n \{q_n\}$ **coend**

for the case where a given number n of component programs are composed in parallel. We also define concrete syntax for program schemas, where the number of components n is a parameter, such as

$A := A[0 := 0] \parallel \dots \parallel A := A[n-1 := 0]$

which sets to 0 the components 0 to $n - 1$ in the array A , where arrays are usually modeled as lists. Although the syntax of the programming language does not cater for “...”, HOL does. Using the well-known function *map* and the construct $[i..j()]$, which represents the list of natural numbers from i to $j-1$, we can express the above schematic program in HOL as follows

Parallel (*map* ($\lambda i. \{i < \text{length } A\} A := A[i:=0] \{A!i = 0\}$) $[0..n()]$)

where the necessary annotations to prove the triple

$\Vdash \{n < \text{length } A\} \textit{Parallel} \dots \{\forall i < n. A!i = 0\}$

have already been inserted.

With the defined concrete syntax for parameterized programs the example above would be written as

cobegin
scheme $[0 \leq i < n] \{i < \text{length } A\} A := A[i:=0] \{A!i = 0\}$
coend

Schematic programs can also appear in parallel with other component programs in the same **cobegin-coend** environment. The index i ranges between the limits indicated in $[- \leq i < -]$. Note that i is a bound variable.

In the next section, devoted to the verification of concrete examples, we shall see a sample of all program features presented using the concrete syntax.

2.8 Examples

We have verified all the relevant examples in [Apt and Olderog, 1991]. This section presents solutions to the mutual exclusion, parallel zero search and the producer/consumer problems. Two of the programs for mutual exclusion handle the problem for a non-fixed number of components. In the

Assertion	$\{r\}$
<i>Non-annotated commands</i>	
<i>Basic</i>	$x := e$
<i>Seq</i>	a_0, a_1
<i>Cond</i>	if b then a_0 else a_1 fi
<i>While</i>	while b inv $\{inv\}$ do a od
<i>Skip</i>	skip
<i>Cond₂</i>	if b then a_0 fi
<i>Annotated commands</i>	
<i>AnnBasic</i>	$\{r\} x := e$
<i>AnnSeq</i>	$c_0; c_1$
<i>AnnCond₁</i>	$\{r\}$ if b then c_1 else c_2 fi
<i>AnnCond₂</i>	$\{r\}$ if b then c fi
<i>AnnWhile</i>	$\{r\}$ while b inv $\{inv\}$ do c od
<i>AnnAwait</i>	$\{r\}$ await b then a end
<i>AnnSkip</i>	$\{r\}$ skip
<i>AnnAtom</i>	$\{r\} \langle a \rangle$
<i>AnnWait</i>	$\{r\}$ wait b end
<i>Parallel commands</i>	
<i>Parallel</i>	cobegin $c_0 \{q_0\} \parallel \dots \parallel c_n \{q_n\}$ coend
<i>Schematic</i>	scheme $[j \leq i < k]$ $c \{q\}$
<i>Convention</i>	
<p> x: program variable e: expression of the type of x r, b, inv, q, q_i: boolean expressions a, a_0, a_1: non-annotated commands c, c_0, c_1: annotated commands j, k: limits for indexing the component programs </p>	

Table 2.4: Concrete syntax for programs.

next chapter we present a more involved case study of a parameterized program, namely, the verification of a parallel garbage collection algorithm for n mutators. All of them have been verified using the *vcg-tactic* to generate the verification conditions and standard Isabelle automatic tactics to prove them.

The examples use array variables which are modeled as lists. Access to components of arrays, usually written $A[i]$, becomes $A!i$; assignments to components of arrays are written $A := A[i:=e]$, where $A[i:=e]$ means that the component at index i in the list A has been replaced by e .

2.8.1 Mutual Exclusion

Mutual exclusion algorithms synchronize n processes, $n \geq 2$, which share a resource. Several properties are expected to be satisfied. The mutual exclusion property guarantees that never more than one process uses the common resource at a time, i.e. only one process may be inside its critical section at each moment. Other properties like deadlock-freedom, defined in §2.1.1, or fairness, which means that all the components get “fair” turns to perform steps, cannot be directly verified in the actual formalization.

Each process S_i in a mutual exclusion algorithm is an infinite loop of the form:

```

 $S_i \equiv$  while True do
     $NC_i$ ;           (non-critical section)
     $ACQ_i$ ;         (acquire protocol)
     $CS_i$ ;           (critical section)
     $REL_i$ ;         (release protocol)
od

```

We consider a parallel program

$$S \equiv INIT, , \mathbf{cobegin} S_1 \parallel \dots \parallel S_n \mathbf{coend}$$

where $INIT$ is a loop free program in which the variables used in ACQ_i and REL_i are initialized.

We prove correctness of three solutions to this problem: a *busy wait solution*, i.e. without synchronization, for a two-process algorithm, a second solution using semaphores via synchronization constructs, and the so-called ticket algorithm. The last two examples work for n -processes waiting to access the critical region.

A Busy Wait Solution

A mutual exclusion algorithm where the “acquire protocol” part for each process ACQ_i is of the form T_i ; **while** b_i **do skip od** with T_i loop free, is called a *busy wait solution* and the loop **while** b_i **do skip od** is called a *busy wait loop*.

The following such solution to the mutual exclusion problem for two processes is due to [Peterson, 1981]:

```
record Busy-wait-mutex =
  flag1 :: bool
  flag2 :: bool
  turn  :: nat
  after1 :: bool
  after2 :: bool
```

lemma *Busy-wait-mutex*:

```
⊢ { True }
flag1 := False,, flag2 := False,,
cobegin
{ ¬flag1 }
while True
  inv { ¬flag1 }
  do { ¬flag1 } ⟨flag1 := True,, after1 := False⟩;
    { flag1 ∧ ¬after1 } ⟨turn := 1,, after1 := True⟩;
    { flag1 ∧ after1 ∧ (turn = 1 ∨ turn = 2) }
    while ¬(flag2 → turn = 2)
      inv { flag1 ∧ after1 ∧ (turn = 1 ∨ turn = 2) }
      do { flag1 ∧ after1 ∧ (turn = 1 ∨ turn = 2) } skip od;
    { flag1 ∧ after1 ∧ (flag2 ∧ after2 → turn = 2) }
    flag1 := False
  od
{ False }
||
{ ¬flag2 }
while True
  inv { ¬flag2 }
  do { ¬flag2 } ⟨flag2 := True,,after2 := False⟩;
    { flag2 ∧ ¬after2 } ⟨turn := 2,,after2 := True⟩;
    { flag2 ∧ after2 ∧ (turn = 1 ∨ turn = 2) }
```

```

while  $\neg(\text{flag}_1 \longrightarrow \text{turn} = 1)$ 
  inv  $\{ \text{flag}_2 \wedge \text{after}_2 \wedge (\text{turn} = 1 \vee \text{turn} = 2) \}$ 
  do  $\{ \text{flag}_2 \wedge \text{after}_2 \wedge (\text{turn} = 1 \vee \text{turn} = 2) \}$  skip od;
   $\{ \text{flag}_2 \wedge \text{after}_2 \wedge (\text{flag}_1 \wedge \text{after}_1 \longrightarrow \text{turn} = 1) \}$ 
   $\text{flag}_2 := \text{False}$ 
od
 $\{ \text{False} \}$ 
coend
 $\{ \text{False} \}$ 

```

The boolean variable flag_i is set to true when the component S_i intends to enter its critical section. The variable turn is used to manage conflicts, which appear when both components intend to enter their critical sections at the same time. The component which sets the variable turn first is delayed in a busy wait loop. The auxiliary variables after_i indicate whether the assignment $\text{turn} := i$ in ACQ_i has been executed.

The critical sections CS_i proceed after the busy wait loop. The preconditions of the critical sections represent the set of states reachable at those points. Observe that program control cannot be ready to enter both critical sections at the same time because no state can satisfy both assertions simultaneously.

The vcg-tactic *oghoare* generates a total of 122 verification conditions, all of them are automatically proven by the Isabelle tactic *auto*, which solves all subgoals simultaneously.

A Solution Using Semaphores

The next solution to the mutual exclusion problem is due to Dijkstra [Dijkstra, 1968]. The algorithm can be generalized to the case where n processes compete to enter their critical sections. It has the simple form:

$$S \equiv \text{out} := \text{True}, , \text{cobegin } S_1 \parallel \dots \parallel S_n \text{ coend}$$

Written with the formalized syntax for schematic (parameterized) programs (cf. §2.7) the corresponding specification is:

```

record Semaphores-parameterized-mutex =
  out :: bool
  who :: nat

```

lemma *Semaphores-parameterized-mutex*: $0 < n \implies$

```

⊢ { True }
out := True,,
cobegin
  scheme [0 ≤ i < n]
    { True }
    while True inv { True }
      do { True } await out then out := False,, who := i end;;
        { ¬out ∧ who = i } out := True
    od
  { False }
coend
{ False }

```

This algorithm uses binary semaphores as a synchronization primitive. A *binary semaphore* is a semaphore that can only take two values. The standard operations on semaphores can be implemented via the synchronization constructor *AnnAwait*. In the program the variable `out` is a binary semaphore that indicates whether all processes are out of their critical sections. The auxiliary variable `who` serves to indicate which component, if any, is inside the critical section.

The critical section would be after the assertion $\{ \neg \text{out} \wedge \text{who} = i \}$. It is easy to see that this precondition cannot hold simultaneously for two different components. Thus, the mutual exclusion property holds.

The vcg-tactic *oghoare* generates 20 verification conditions which are all solved automatically by *auto*.

The Ticket Algorithm

The next example is also a mutual exclusion algorithm for n processes. We prove its correctness based on the proof outline given in [de Roever *et al.*, 2000].

Predicates that are often used in the assertions of a program can be given an abbreviated name by stating the equality in the premises. For reasons concerning the translations between internal and external syntax, the abbreviated expressions appear enclosed between ‘ \ll ’ and ‘ \gg ’. Then the given name (*Inv* in the next example) appears in *sans serif* font in the assertions, i.e. *Inv*. These abbreviations are not to be confused with program variables, the reason why they both have the same representation in the program, namely the same font, lies in the fact that both, variables and abbreviations, depend upon the program state.

```

record Ticket-mutex =
  num :: nat
  nextv :: nat
  turn :: nat list
  ind :: nat

lemma Ticket-mutex:
  [[ 0 < n; Invariant= << n = length turn ∧ 0 < nextv ∧
    (∀ k < n. ∀ l < n. k ≠ l → turn!k < num ∧ (turn!k = 0 ∨ turn!k ≠ turn!l)) >> ]]
  ⇒
  ⊢ { n = length turn }
  ind := 0,,
  while ind < n
    inv { n = length turn ∧ (∀ i < ind. turn!i = 0) }
  do turn := turn [ind := 0],, ind := ind + 1 od,,
  num := 1,, nextv := 1 ,,
  cobegin
    scheme [0 ≤ i < n]
      { Invariant }
      while True inv { Invariant }
      do { Invariant } ⟨turn := turn [i := num],, num := num + 1⟩;;
        { Invariant } wait turn!i = nextv end;;
        { Invariant ∧ turn!i = nextv } nextv := nextv + 1
      od
      { False }
    coend
    { False }

```

The critical section would be entered before the assignment to `nextv`, i.e. at the moment of entering the assertion $\{ \text{Invariant} \wedge \text{turn!}i = \text{nextv} \}$ holds. The mutual exclusion property is guaranteed because the conjunction of two or more assertions with different values for i before entering the critical section implies *false*. These assertions represent the possible states at that point, consequently, the set of states from which more than one component could enter the critical section is empty.

The application of the `vcg-tactic` returns 35 subgoals. Simplification tactics leave 11 verification conditions unsolved. Their proofs only need to be further directed by several case distinctions as hinted in the pencil and paper proof of [de Roever *et al.*, 2000].

2.8.2 Parallel Zero Search

The next example is a program that finds a zero of a function f from naturals to naturals, searching in parallel for zeroes that are bigger or smaller than a certain natural a .

record *Zero-search* =

turn :: *nat*
found :: *bool*
x :: *nat*
y :: *nat*

lemma *Zero-search*:

$\llbracket I_1 = \ll a \leq x \wedge (\text{found} \longrightarrow (a < x \wedge f(x) = 0) \vee (y \leq a \wedge f(y) = 0))$
 $\wedge (\neg \text{found} \wedge a < x \longrightarrow f(x) \neq 0) \gg ;$
 $I_2 = \ll y \leq a + 1 \wedge (\text{found} \longrightarrow (a < x \wedge f(x) = 0) \vee (y \leq a \wedge f(y) = 0))$
 $\wedge (\neg \text{found} \wedge y \leq a \longrightarrow f(y) \neq 0) \gg \rrbracket \Longrightarrow$

$\vdash \{ \exists u. f(u) = 0 \}$

turn := 1,, *found* := *False*,,

x := *a*,, *y* := *a* + 1 ,,

cobegin $\{ I_1 \}$

while $\neg \text{found}$

inv $\{ I_1 \}$

do $\{ a \leq x \wedge (\text{found} \longrightarrow y \leq a \wedge f(y) = 0) \wedge (a < x \longrightarrow f(x) \neq 0) \}$

wait *turn* = 1 **end**;;

$\{ a \leq x \wedge (\text{found} \longrightarrow y \leq a \wedge f(y) = 0) \wedge (a < x \longrightarrow f(x) \neq 0) \}$

turn := 2;;

$\{ a \leq x \wedge (\text{found} \longrightarrow y \leq a \wedge f(y) = 0) \wedge (a < x \longrightarrow f(x) \neq 0) \}$

$\langle x := x + 1, \text{ if } f(x) = 0 \text{ then } \text{found} := \text{True} \text{ else skip fi} \rangle$

od;;

$\{ I_1 \wedge \text{found} \}$

turn := 2

$\{ I_1 \wedge \text{found} \}$

\parallel

$\{ I_2 \}$

while $\neg \text{found}$

inv $\{ I_2 \}$

do $\{ y \leq a + 1 \wedge (\text{found} \longrightarrow a < x \wedge f(x) = 0) \wedge (y \leq a \longrightarrow f(y) \neq 0) \}$

wait *turn*=2 **end**;;

$\{ y \leq a + 1 \wedge (\text{found} \longrightarrow a < x \wedge f(x) = 0) \wedge (y \leq a \longrightarrow f(y) \neq 0) \}$

```

    turn := 1;;
    { y ≤ a + 1 ∧ (found → a < x ∧ f(x) = 0) ∧ (y ≤ a → f(y) ≠ 0) }
    ⟨ y := y - 1,, if f(y) = 0 then found := True else skip fi ⟩
  od;;
  { I2 ∧ found }
  turn := 1
  { I2 ∧ found }
coend
{ f(x) = 0 ∨ f(y) = 0 }

```

The tactic generates 98 verification conditions. After applying the general automatic tactic *auto*, 40 subgoals remain³. They are all proven with the tactic *arith* that automatically solves basic arithmetical problems.

We verify a second simpler solution to this problem without using synchronization:

lemma *Zero-Search₂*:

```

[[ I1 = ⟨ a ≤ x ∧ (found → (a < x ∧ f(x) = 0) ∨ (y ≤ a ∧ f(y) = 0))
    ∧ (¬found ∧ a < x → f(x) ≠ 0) ⟩;
   I2 = ⟨ y ≤ a + 1 ∧ (found → (a < x ∧ f(x) = 0) ∨ (y ≤ a ∧ f(y) = 0))
    ∧ (¬found ∧ y ≤ a → f(y) ≠ 0) ⟩ ]] ⇒
⊢ { ∃ u. f(u) = 0 }
found := False,,
x := a,, y := a + 1,,
cobegin { I1 }
  while ¬found
    inv { I1 }
  do { a ≤ x ∧ (found → y ≤ a ∧ f(y) = 0) ∧ (a < x → f(x) ≠ 0) }
    ⟨ x := x + 1,, if f(x) = 0 then found := True else skip fi ⟩
  od
{ I1 ∧ found }
||
{ I2 }
  while ¬found
    inv { I2 }
  do { y ≤ a + 1 ∧ (found → a < x ∧ f(x) = 0) ∧ (y ≤ a → f(y) ≠ 0) }
    ⟨ y := y - 1,, if f(y) = 0 then found := True else skip fi ⟩
  od

```

³The tactic *auto* simplifies all subgoals simultaneously and might generate several simpler subgoals out of one.


```

  { l2 ∧ found }
coend
  { f(x) = 0 ∨ f(y) = 0 }

```

Only 20 verification conditions are generated in this simplified version. The application of *auto* leaves 32 arithmetical subgoals, all solved automatically by the tactic *arith*.

2.8.3 Producer/Consumer

This problem coordinates two processes, producer and consumer, that share a common, bounded buffer. The producer puts information into the buffer, the consumer takes it out. Trouble arises when the producer attempts to put a new item in a full buffer or the consumer tries to remove an item from an empty buffer. Following Owicki-Gries we express the problem as a parallel program with shared variables and await-commands. It copies the elements of an array *a* into an array variable *b*. Note that *a* is not a variable of the program, thus its value cannot be overwritten.

$$\{0 < \text{length } a \wedge 0 < \text{length } \mathbf{buffer} \wedge \text{length } \mathbf{b} = \text{length } a\}$$

```

cobegin producer || consumer coend
  {∀k < length a. a!k = b!k}

```

The precondition imposes that the length of *a* and *b* be equal, and *a*, *b* and *buffer* have non-zero length. The full program is shown below:

```

record Producer-consumer =
  ins :: nat
  outs :: nat
  i :: nat
  j :: nat
  vx :: nat
  vy :: nat
  buffer :: nat list
  b :: nat list

```

For readability we used some abbreviations that can be defined in the premises of the lemma.

lemma *Producer-consumer*:

```

[[ INIT = « 0 < length a ∧ 0 < length buffer ∧ length b = length a » ;

```

$$I = \ll (\forall k. \text{outs} \leq k \wedge k < \text{ins} \longrightarrow a!k = \text{buffer}!(k \bmod (\text{length } \text{buffer})))$$

$$\wedge \text{outs} \leq \text{ins} \wedge \text{ins} - \text{outs} \leq \text{length } \text{buffer} \gg ;$$

$$I_1 = \ll I \wedge i \leq \text{length } a \gg ;$$

$$p_1 = \ll I_1 \wedge i = \text{ins} \gg ;$$

$$I_2 = \ll I \wedge (\forall k < j. a!k = b!k) \wedge j \leq \text{length } a \gg ;$$

$$p_2 = \ll I_2 \wedge j = \text{outs} \gg \rrbracket \Longrightarrow$$

$$\Vdash \{ \text{INIT} \}$$

$$\text{ins} := 0,, \text{outs} := 0,, i := 0,, j := 0,,$$

cobegin { $p_1 \wedge \text{INIT}$ }

while $i < \text{length } a$

inv { $p_1 \wedge \text{INIT}$ }

do { $p_1 \wedge \text{INIT} \wedge i < \text{length } a$ }

$$vx := a!i;;$$

$$\{ p_1 \wedge \text{INIT} \wedge i < \text{length } a \wedge vx = a!i \}$$

wait $\text{ins} - \text{outs} < \text{length } \text{buffer}$ **end**;;

$$\{ p_1 \wedge \text{INIT} \wedge i < \text{length } a \wedge vx = a!i \wedge$$

$$\text{ins} - \text{outs} < \text{length } \text{buffer} \}$$

$$\text{buffer} := \text{buffer} [\text{ins} \bmod (\text{length } \text{buffer}) := vx];;$$

$$\{ p_1 \wedge \text{INIT} \wedge i < \text{length } a \wedge$$

$$a!i = \text{buffer}!(\text{ins} \bmod (\text{length } \text{buffer})) \wedge \text{ins} - \text{outs} < \text{length } \text{buffer} \}$$

$$\text{ins} := \text{ins} + 1;;$$

$$\{ I_1 \wedge \text{INIT} \wedge i + 1 = \text{ins} \wedge i < \text{length } a \}$$

$$i := i + 1$$

od

$$\{ p_1 \wedge \text{INIT} \wedge i = \text{length } a \}$$

$$\parallel$$

$$\{ p_2 \wedge \text{INIT} \}$$

while $j < \text{length } a$

inv { $p_2 \wedge \text{INIT}$ }

do { $p_2 \wedge j < \text{length } a \wedge \text{INIT}$ }

wait $\text{outs} < \text{ins}$ **end**;;

$$\{ p_2 \wedge j < \text{length } a \wedge \text{outs} < \text{ins} \wedge \text{INIT} \}$$

$$vy := \text{buffer}!(\text{outs} \bmod (\text{length } \text{buffer}));;$$

$$\{ p_2 \wedge j < \text{length } a \wedge \text{outs} < \text{ins} \wedge vy = a!j \wedge \text{INIT} \}$$

$$\text{outs} := \text{outs} + 1;;$$

$$\{ I_2 \wedge j + 1 = \text{outs} \wedge j < \text{length } a \wedge vy = a!j \wedge \text{INIT} \}$$

$$b := b [j := vy];;$$

$$\{ I_2 \wedge j + 1 = \text{outs} \wedge j < \text{length } a \wedge a!j = b!j \wedge \text{INIT} \}$$

$$j := j + 1$$

od

```

  { p2 ∧ j = length a ∧ INIT }
coend
  { ∀ k < length a. a!k = b!k }

```

Both components share the variables `ins` and `outs`, which count the values added to the buffer and the values removed from the buffer, respectively. Thus, the buffer contains `ins - outs` values at each moment. Expressions `ins mod (length buffer)` and `outs mod (length buffer)` determine the subscript of the buffer element where the next value is to be added or removed.

The verification of this problem involves proving a total of 138 conditions. Half of them are trivially solved since they refer to triples of the form $(A \cap pre\ r)$ where the atomic action r does not change the variables in A . The rest are automatically solved by Isabelle standard simplification tactics.

2.9 Summary

We have presented the first formalization of the Owicki-Gries method in a general purpose theorem prover. This method can be considered a classic and has been studied extensively since its introduction in 1975 (cf. [Dijkstra, 1976, Apt, 1981a, Apt and Olderog, 1991, Best, 1996, de Roever *et al.*, 2000]). Nevertheless, our formalization yields two main new contributions:

1. A simpler and more intuitive soundness proof with no explicit reference to program locations.
2. A generalized proof rule for parallel composition that allows direct verification of parameterized parallel programs in the system.

In addition, we provide the following features that turn the formalization into a tool suitable for real program verification:

1. Familiar concrete syntax for writing programs like they are usually found in the literature, and
2. A tactic that automatically generates all the verification conditions.

So far we have only verified typical examples from the literature. The next chapter presents the verification of two garbage collection algorithms. These examples better illustrate the applicability of the formalization for two reasons: first, they are larger and more involved programs and second,

no complete Owicki-Gries proof existed in the literature. We believe that the availability of the tool was decisive in the search for successful proof outlines.

Chapter 3

Case Study: Single and Multi-Mutator Garbage Collection Algorithms

In this chapter we show that the Owicki-Gries method and its mechanization can be successfully applied to larger examples than those considered in §2.8. We study two incremental garbage collection algorithms, the second one parametric in the number of mutators. These algorithms are particularly tricky and very distinguished scientists have published flawed proofs, some of which were first detected by mechanization attempts. An excellent account of these flaws can be found in [Russinoff, 1994].

We first verify Ben-Ari's classic algorithm [Ben-Ari, 1984]. A pencil and paper proof using the Owicki-Gries method plus ad-hoc reasoning was presented in [van de Snepscheut, 1987]. Our proof follows [van de Snepscheut, 1987], but it manages to formulate the extra reasoning within the Owicki-Gries method. Ben-Ari's algorithm has also been mechanically proven using the Boyer-Moore prover [Russinoff, 1994] and PVS [Havelund, 1996], but none of these proofs uses Owicki-Gries. This makes the algorithm an excellent example for comparing Owicki-Gries with other methods, and for comparing Isabelle/HOL with other theorem provers.

In §3.4 we verify a parameterized garbage collector in which an arbitrary number of mutators work in parallel. This implies that the correctness proof must be carried out for an infinite family of algorithms, which introduces an additional difficulty.

The first extension of Ben-Ari's algorithm to several mutators was published in [Pixley, 1988]. We verify an improved version from [Jonker, 1992]

which is finer-grained, uses less colors and has less overhead for the mutators. The author of [Jonker, 1992] gives a proof of correctness using an ad-hoc technique based on observing the behavior of appropriate variant functions. In the same paper it is argued that the Owicki-Gries method is not suitable for this problem. To our knowledge this is the first Owicki-Gries proof and the first mechanized proof of this algorithm.

Thanks to Isabelle’s facilities in dealing with concrete syntax, the formalization can be done in a very natural way, where the algorithms and lemmas are as readable as in the original papers.

The chapter is structured as follows: the basics of garbage collection algorithms are described in section 3.1. In section 3.2 we formalize a model for computer memory. Section 3.3 presents the proof of Ben-Ari’s algorithm in detail. Section 3.4 presents the proof of the parametric algorithm. In both cases a safety property stating that only garbage nodes are collected is verified. In section 3.5 we compare our proofs with other related works and draw conclusions.

3.1 Incremental Garbage Collection

Garbage collection is the automatic reclamation of memory space¹. User processes, called *mutators*, might produce garbage while performing their computations. The *collector*’s task is to identify this garbage and to recycle it for future use by appending it to the *free list*. *Incremental* (also called on-the-fly) garbage collection systems are those where the garbage collection work is randomly interleaved with the execution of instructions in the running programs. This is important for real-time applications where memory management operations should never halt the executing program for more than a very brief period.

The *memory* is modeled as a finite directed graph with a fixed number of nodes, where each node has a fixed set of outgoing edges. A predetermined subset of nodes, called the *roots*, is always accessible to the running program. A node is called *reachable* or *accessible* if a directed path exists along the edges from at least one root to that node, otherwise, it is called *garbage*. For marking purposes, each node is associated a color, which can be black or white. The memory structure can only be modified by one of the following three operations: redirect an edge from a reachable node towards a reachable node, append a garbage node to the free list, or change the color of a node.

¹An excellent survey about garbage collection algorithms can be found in [Wilson, 1992].

The mutators abstractly represent the changes that user programs produce on the memory structure. It is assumed that they only work on nodes that are reachable, having the ability to redirect an edge to some new target. To make garbage collection safe, the mutators cooperate with the collector by assuming the overhead of blackening the new target. Thus, a mutator repeatedly redirects some edge R to some reachable node T , and then colors the node T black.

It is customary to describe the collector's task in this way: identify the nodes that are garbage, i.e. no longer reachable, and append them to the free list, so that their space can be reused by the running program. We abstract from the particular implementation of the free list and simply assume that the collector makes garbage nodes accessible again: since the mutator has the ability to redirect arbitrary accessible edges, it may reuse these nodes. In the sequel adding a node to the free list just means making it accessible.

The collector repeatedly executes two phases, traditionally called *marking phase* and *sweep* or *appending phase*.

During the marking phase, the collector traverses the graph, starting by blackening the set of roots, and marks accessible nodes by coloring them black. This process finishes when all reachable nodes are black. During the appending phase the memory is swept, appending all unmarked (garbage) nodes to the free list. The outline of the algorithms is:

- Marking phase:
 1. Color the roots black.
 2. Visit each edge. If the source is black, color the target black.
 3. Count the black nodes.
 4. If not all reachable nodes are black, go to step 2.

- Appending phase:
 5. Visit each node. If it is white, append it to the free list; if it is black, color it white.

The safety property we prove says that *no reachable node is garbage collected*. In other words, if during the appending operation a node is white, then it is garbage. Clearly, this property holds if step 4 is correct. But how do we determine that all reachable nodes are black? In the case of one mutator, Ben Ari's solution is to keep the result of the last count, and compare it with the result of the current count. If they coincide, then all

reachable nodes are black. For n mutators, we compare the results of the last $n + 1$ counts. So the algorithms for one and several mutators differ only in step 4.

3.2 Formalization of the Memory

The memory is formalized using two lists of fixed size. The first list, represented by the variable M in the algorithms, has an index for each memory node, i.e. nodes are referred to by natural numbers that range from 0 to *length* $M - 1$; the color of node i can be consulted by accessing the contents of index i , which in Isabelle’s list notation is written as $M!i$. The datatype *node* is the color of a node, which can be black or white.

datatype *node* = *Black* | *White*

The list of nodes M has the type

types *nodes* = *node list*

The second list, which we refer to by E in the algorithms, models the edges, which are numbered by the indices of the list; each edge is a pair of natural numbers corresponding to the source and the target nodes

types

edge = *nat* \times *nat*

edges = *edge list*

and *Roots* is an arbitrary set of nodes

consts *Roots* :: *nat set*

Figure 3.1 shows an example of a memory where the set of roots is $\{1, 2\}$, the list M of nodes is [*White*, *Black*, *White*, *White*, *White*] and the list E of edges is [(0, 0), (3, 4), (1, 2), (2, 3), (4, 2)].

We define some sets and predicates that are frequently used in the annotations. *Blacks* of a list of nodes returns the set of nodes that are *Black*. *BtoW* is true of the edges that point from a *Black* node to a *White* node. Finally, given a list of edges e , *Reach e* is the set of nodes reachable from *Roots*, i.e. the *Roots* themselves and those nodes such that there exists a

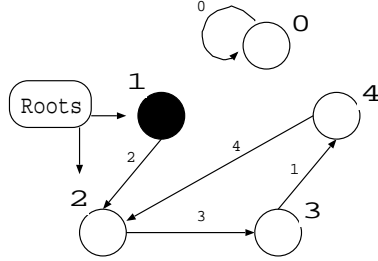


Figure 3.1: An example of the memory.

path along the edges from the node to some root. The formal definitions are:

constdefs

$BtoW :: (edge \times nodes) \Rightarrow bool$

$BtoW \equiv \lambda(e, m). (m!fst\ e) = Black \wedge (m!snd\ e) \neq Black$

$Blacks :: nodes \Rightarrow nat\ set$

$Blacks\ m \equiv \{i. i < length\ m \wedge m!i = Black\}$

$Reach :: edges \Rightarrow nat\ set$

$Reach\ e \equiv \{x. x \in Roots \vee$

$(\exists path. 1 < length\ path \wedge path!(length\ path - 1) \in Roots \wedge x = path!0$

$\wedge (\forall i < length\ path - 1. (\exists j < length\ e. e!j = (path!(i + 1), path!i))))\}$

The next predicates indicate whether a given set of roots or edges is well-formed:

$Proper-Roots :: nodes \Rightarrow bool$

$Proper-Roots\ m \equiv Roots \neq \{\} \wedge Roots \subseteq \{i. i < length\ m\}$

$Proper-Edges :: (nodes \times edges) \Rightarrow bool$

$Proper-Edges \equiv (\lambda(m, e). \forall i < length\ e. fst\ (e!i) < length\ m$
 $\wedge snd\ (e!i) < length\ m)$

Given a list of nodes, a proper set of roots is a non-empty subset of nodes. Proper edges are those that point from a node to a node, i.e. the first and second components of an edge-pair must be within the range of node indices.

The separate treatment of colors and edges in our data structure is an abstraction that considerably simplifies proofs relating to the changes in the graph. If an edge is redirected, the variable representing the memory, namely M , remains invariant, while coloring does not modify the variable representing the edges E .

Finally, we introduce a last predicate to express that all reachable nodes are black. It is called *Safe* because as we shall see this situation represents a safe state for the memory.

constdefs

Safe :: (*nodes* × *edges*) ⇒ *bool*
Safe ≡ λ(*m*, *e*). *Reach e* ⊆ *Blacks m*

3.3 The Single-Mutator Case

We verify van de Snepscheut’s version of Ben-Ari’s algorithm. We follow the ideas of [van de Snepscheut, 1987], but formulate the proof completely within the Owicki-Gries system.

The program consists of two components, namely, the *Mutator* and the *Collector*. First, we study each component separately and prove that they achieve their intended task whenever they are executed in isolation. Then, we prove that both components can indeed be executed in parallel without interfering with each other.

In order to use the defined concrete syntax for programs (see 2.7), the variables used in the collector and mutator are first declared in a record:

record *gar-coll-state* =
M :: *nodes*
E :: *edges*
bc :: *nat set*
obc :: *nat set*
Ma :: *nodes*
ind :: *nat*
k :: *nat*
z :: *bool*

In the program text and assertions, variables are printed in *sans serif* font. The variable M represents the list of nodes and E the list of edges. The role of the other variables will be explained in the following sections.

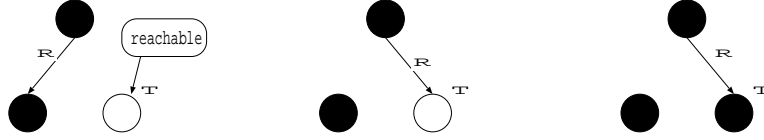


Figure 3.2: The Mutator.

3.3.1 The Mutator

The mutator first redirects an arbitrary edge R from an arbitrary accessible node towards an arbitrary accessible node T . It then colors the new target T black. A graphical description of the actions are shown in figure 3.2. We declare the arbitrarily selected node and edge as constants:

consts

$R :: nat$

$T :: nat$

The following predicate states, given a list of nodes m and a list of edges e , the conditions under which the selected edge R and node T are valid:

constdefs

$Mut-init :: gar-coll-state \Rightarrow bool$

$Mut-init \equiv \ll T \in Reach\ E \wedge R < length\ E \wedge T < length\ M \gg$

For a more structured proof we have divided the algorithms into modules. A *module* is a piece of code, consisting of one or several instructions. For the mutator we consider two modules, one for each action. An auxiliary variable z is set to false if the mutator has already redirected an edge but has not yet colored the new target. Note that the state is the previously declared record of program variables.

constdefs

$Redirect-Edge :: gar-coll-state\ ann-com$

$Redirect-Edge \equiv \{\{Mut-init \wedge z\} \langle E := E [R := (fst\ (E!R),\ T)],, z := (\neg z)\rangle\}$

$Color-Target :: gar-coll-state\ ann-com$

$Color-Target \equiv \{\{Mut-init \wedge \neg z\} \langle M := M [T := Black],, z := (\neg z)\rangle\}$

$Mutator :: gar-coll-state\ ann-com$

$Mutator \equiv$

```

    {Mut-init  $\wedge$  z}
  while True inv {Mut-init  $\wedge$  z}
  do Redirect-Edge ;; Color-Target od

```

We prove that the mutator’s proof outline is correct, by the soundness theorem it suffices to prove that it is derivable in the proof system for component programs. To obtain the full proof outline a postcondition has to be added to the annotated command. Since the program is an infinite loop, no state ever reaches the postcondition.

lemma *Mutator*: \vdash *Mutator* { *False* }

The verification conditions are generated with the tactic *annhoare*. All are trivially solved except for one which requires the following lemma:

lemma *Graph1*:

$$\llbracket t \in \text{Reach } e; r < \text{length } e \rrbracket \implies t \in \text{Reach } (e [r := (\text{fst } (e!r), t)])$$

stating that an accessible node cannot be rendered inaccessible by redirecting an edge to it. For the proof it is not necessary to require that the source of the selected edge R be reachable. However, for implementations purposes, it is important that mutators work on nodes that are reachable. Otherwise, the following scenario explained in [Pixley, 1988] could cause problems in the system: Suppose that a certain edge R has been selected by the mutator. If its source becomes unreachable before the redirection is executed, it could be garbage collected. Once it is appended to the free list another process might re-use it resulting in unexpected results when the mutator performs the pending redirection.

3.3.2 The Collector

The collector works in two phases. The first one, called the marking phase first blackens the roots and then executes a loop. The body of the loop consists of first traversing M , coloring all reachable nodes black, and then counting the number of black nodes. The loop terminates if the results of the current count and the previous one coincide. After termination of the loop, the appending phase starts. Here the collector traverses M once more, this time making all white nodes reachable and all black nodes white. Figure 3.3 shows an example of the execution of the collector in isolation.

To structure the proof outline of the collector we define four modules: *Blacken-Roots*, *Propagate-Black*, *Count-Blacks* and *Append*.

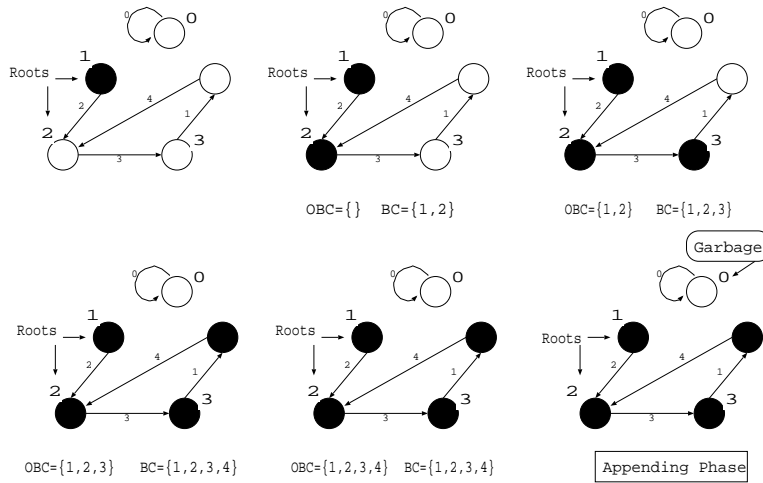


Figure 3.3: The marking phase.

The collector uses, apart from the list of nodes M and the list of edges E , five more variables. bc (black count) and obc (old black count) are used to determine if the set of black nodes has grown during the last *Propagate-Black* phase. Following [van de Snepscheut, 1987], obc is initialized to the empty set, and bc to the set $Roots$ ². A single auxiliary variable Ma is used for “recording” the value of M after the execution of *Propagate-Black*. This is just an assignment to an auxiliary variable, used exclusively for proof purposes, and therefore not part of the computation. Finally, ind is a counter for loops and k is used to achieve a finer grain of interleaving inside the *Propagate-Black* phase.

consts

Blacken-Roots :: *gar-coll-state ann-com*
Propagate-Black :: *gar-coll-state ann-com*
Count-Blacks :: *gar-coll-state ann-com*
Append :: *gar-coll-state ann-com*

A constant $M-init$ is used to give Ma a suitable first value, defined as a list of nodes where only the *Roots* are black.

² obc and bc are modeled as sets of black nodes whereas in the original algorithm they represent their cardinalities. We found the set approach more elegant but it simplifies neither the algorithm nor the proofs.

```

consts M-init :: nodes
constdefs
  Proper-M-init :: nodes  $\Rightarrow$  bool
  Proper-M-init m  $\equiv$  Blacks M-init = Roots  $\wedge$  length M-init = length m

```

For readability of the assertions we introduce the following abbreviations:

```

constdefs
  Proper :: gar-coll-state  $\Rightarrow$  bool
  Proper  $\equiv$   $\ll$  Proper-Roots M  $\wedge$  Proper-Edges (M, E)  $\wedge$  Proper-M-init M  $\gg$ 

```

The proof outline of the collector with modules is:

```

constdefs
  Collector :: gar-coll-state ann-com
  Collector  $\equiv$ 
  { Proper }
  while True inv { Proper }
  do
    Blacken-Roots;;
    { Proper  $\wedge$  Roots  $\subseteq$  Blacks M } obc := {};;
    { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc = {} } bc := Roots;;
    { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc = {}  $\wedge$  bc = Roots } Ma := M-init;;
    { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc = {}  $\wedge$  bc = Roots  $\wedge$  Ma = M-init }
    while obc  $\neq$  bc
      inv { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  length Ma = length M
           $\wedge$  obc  $\subseteq$  Blacks Ma  $\wedge$  Blacks Ma  $\subseteq$  bc  $\wedge$  bc  $\subseteq$  Blacks M
           $\wedge$  (obc  $\subset$  Blacks Ma  $\vee$  Safe (M, E)) }
      do { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M }
        obc := bc;;
        Propagate-Black;;
        { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M
           $\wedge$  bc  $\subseteq$  Blacks M  $\wedge$  (obc  $\subset$  Blacks M  $\vee$  Safe (M, E)) }
        Ma := M;;
        { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks Ma
           $\wedge$  Blacks Ma  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M  $\wedge$  length Ma = length M
           $\wedge$  (obc  $\subset$  Blacks Ma  $\vee$  Safe (M, E)) }
        bc := {};;
        Count-Blacks
    od;;

```

Append

od

Safe (M, E) states that all reachable nodes are black, i.e. $Reach\ E \subseteq Blacks\ M$. Since it holds before *Append*, all white nodes are garbage right before the appending module starts. This is almost the safety property we wish to prove. The algorithm must ensure that only garbage nodes are collected during the appending phase. We shall show later when describing the *Append* module that if a white node is garbage before *Append*, then it remains so until *Append* makes it reachable.

The key parts of the invariant are the second and third lines. The second line guarantees that after any execution of the body the cardinalities of *obc* and *bc* are a lower and upper bound, respectively, of the number of black nodes after *Propagate-Black*. It is clear that *obc* is a lower bound, because black nodes stay black until the beginning of the appending phase. That *bc* is an upper bound would be difficult to prove without the auxiliary variable *Ma* since the mutator can blacken nodes while the collector executes *Count-Blacks*. The third line of the invariant guarantees that, if an execution of the body does not establish the safety property, then *obc* is a proper lower bound, which means that some white node was colored black during the execution of *Propagate-Black*. As we shall see, the *Propagate-Black* and *Count-Blacks* modules have very clear tasks: *Propagate-Black* establishes the third line, while *Count-Blacks* establishes the second.

Let us now see that the conjunction of the invariant and the negation of the guard $obc \neq bc$ imply the safety condition. If $obc = bc$, then the upper and lower bound coincide, and so *obc* cannot be a proper lower bound. Hence, no white node was colored black during *Propagate-Black*, and we obtain *Safe* (M, E).

We prove the derivability of the collector's proof outline. Since it is an infinite loop, the postcondition is the empty set of states, i.e. the set of states that satisfy the predicate *False*:

lemma *Collector*: $\vdash Collector \{ \{ False \} \}$

Blackening Roots

In this module a loop visits all roots and colors them black. The corresponding annotated command is:

defs *Blacken-Roots-def*:
Blacken-Roots \equiv

```

{ Proper }
ind := 0;;
{ Proper ∧ ind = 0 }
while ind < length M
  inv { Proper ∧ (∀ i < ind. i ∈ Roots → M!i = Black) ∧ ind ≤ length M }
do { Proper ∧ (∀ i < ind. i ∈ Roots → M!i = Black) ∧ ind < length M }
  if ind ∈ Roots then
    { Proper ∧ (∀ i < ind. i ∈ Roots → M!i = Black) ∧ ind < length M
      ∧ ind ∈ Roots }
    M := M [ind := Black] fi;;
  { Proper ∧ (∀ i < ind + 1. i ∈ Roots → M!i = Black) ∧ ind < length M }
  ind := ind + 1
od

```

This module establishes in the postcondition that the set of roots is a subset of the set of black nodes. Its derivability in the *ann-hoare* system is easy to prove.

lemma *Blacken-Roots*: $\vdash \text{Blacken-Roots } \{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks } M \}$

Propagation of the Coloring

During this phase, the collector visits the edges in a given order, coloring the target whenever the source is *Black*. This phase establishes the third line of the invariant.

The predicate *PBInv* contains the main idea of the proof. We declare it now but explain its meaning below:

```

consts
  PBInv :: gar-coll-state ⇒ nat ⇒ bool

```

We first explain an easier version of the *Propagate-Black* module. It will be later modified to obtain a finer degree of interleaving.

```

constdefs
  Propagate-Black-aux :: gar-coll-state ann-com
  Propagate-Black-aux ≡
  { Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M ∧ bc ⊆ Blacks M }
  ind := 0;;
  { Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M ∧ bc ⊆ Blacks M
    ∧ ind = 0 }

```



```

while ind < length E
  inv { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M
         $\wedge$  bc  $\subseteq$  Blacks M  $\wedge$  PBIInv ind  $\wedge$  ind  $\leq$  length E }
  do { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M
         $\wedge$  bc  $\subseteq$  Blacks M  $\wedge$  PBIInv ind  $\wedge$  ind < length E }
  if M!fst (E!ind) = Black then
    { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
       $\wedge$  PBIInv ind  $\wedge$  ind < length E  $\wedge$  M!fst (E!ind) = Black }
    M := M [snd (E!ind) := Black];;
    { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
       $\wedge$  PBIInv (ind + 1)  $\wedge$  ind < length E }
    ind := ind + 1
  fi
od

```

If the collector is executed in isolation it suffices to define PBIInv as

$$\lambda ind. \text{obc} \subseteq \text{Blacks } M \vee (\forall i < ind. \neg \text{BtoW } (E!i, M)).$$

Upon termination, i.e. when ind = length E we would obtain

$$\text{obc} \subseteq \text{Blacks } M \vee (\forall i < \text{length } E. \neg \text{BtoW } (E!i, M))$$

in the postcondition. When all roots are black, the following lemma holds:

lemma *Graph2*:

$$\begin{aligned} & \llbracket \text{Roots} \subseteq \text{Blacks } m; \text{ Proper-Edges } (m, e); \forall i < \text{length } e. \neg \text{BtoW } (e!i, m) \rrbracket \\ & \implies \text{Reach } e \subseteq \text{Blacks } m \end{aligned}$$

Hence, upon termination we obtain $\text{obc} \subseteq \text{Blacks } M \vee \text{Safe } (M, E)$, which is the postcondition of the *Propagate-Black* phase in the proof outline of the collector.

However, it is easy to see that this definition of *BPIInv* is not invariant under the first action of the mutator: an edge that has already been visited by the collector could be redirected by the mutator to a white target. This is depicted in figure 3.4. In this example the collector is interrupted by the mutator when it reaches edge 3 (ind = 3) but before coloring. The mutator then redirects the already visited edge 2 to node 3. Thus, there is a visited edge that satisfies the black-to-white predicate, falsifying the invariant.

Fortunately, we can find a weaker predicate that is able to establish the postcondition while remaining invariant under the actions of the mutator.

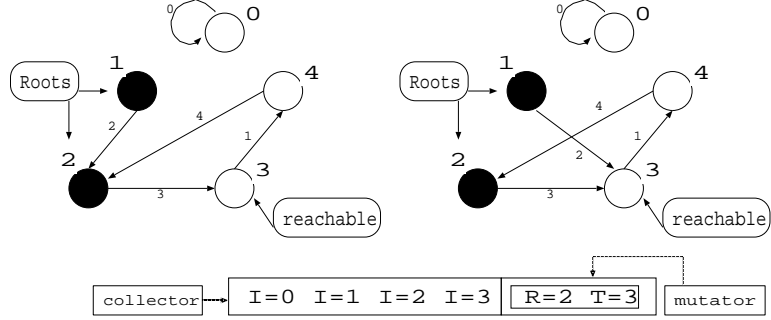


Figure 3.4: Interference with the mutator.

The following definition of BPI_{inv} is an adaptation of the one proposed in [van de Snepscheut, 1987]:

defs $BPI_{inv-def}$:

$$\begin{aligned}
 BPI_{inv} \equiv & \ll \lambda ind. \text{obc} \subset \text{Blacks } M \vee (\forall i < ind. \neg BtoW (E!i, M) \vee \\
 & (\neg z \wedge i = R \wedge (snd (E!R)) = T \wedge \\
 & (\exists r. ind \leq r \wedge r < length E \wedge BtoW (E!r, M)))) \gg
 \end{aligned}$$

Intuitively, its invariance is proved as follows.

If either the collector or the mutator blacken some white node then, after execution of the body, the predicate $\text{obc} \subset \text{Blacks } M$ holds. Otherwise, i.e. no coloring occurs, there are two situations:

1. All edges visited by the collector point to a *Black* node, i.e. $\forall i < ind. \neg BtoW (E!i, M)$ holds.
2. Some visited edge points to a white node because the mutator has redirected it. Then this edge has target node T . Ben-Ari observes that in this situation there must be another $BtoW$ edge among those that have not yet been visited by the collector. This holds because the new target T is reachable by assumption, thus, there exists a path to T from some root. Since all roots are black, some edge along this path must be a $BtoW$ edge.

Observe that upon termination of the loop this last clause cannot hold because the counter ind reaches the value of $length E$. Consequently, this invariant establishes $\text{obc} \subset \text{Blacks } M \vee \text{Safe } (M, E)$ upon termination and is interference free under the mutator's actions. We prove the derivability of the corresponding triple:

lemma *Propagate-Black-aux*:

⊢ *Propagate-Black-aux*

{ Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M ∧ bc ⊆ Blacks M
 ∧ (obc ⊆ Blacks M ∨ Safe (M, E)) }

The assignment $M := M [snd (E!ind) := Black]$ contains two references to shared variables (M and E), which leads to implementation problems. Fortunately, the loop body can be replaced by

```

if M! fst (E!ind) = Black
then k := snd (E!ind);; ⟨M := M [k := Black],, ind := ind + 1⟩
else ⟨if M! fst (E!ind) ≠ Black then ind := ind + 1 fi⟩

```

where at most one shared variable is accessed in each atomic action.

This introduces a new point of interference with the mutator. After this modification the program remains correct although the precondition of the atomic region $\langle M := M [k := Black],, ind := ind + 1 \rangle$ is non-trivial. It includes the following predicate, proposed in [van de Snepscheut, 1987]:

constdefs

Auxk :: *gar-coll-state* ⇒ *bool*

$Auxk \equiv \ll k < length\ M \wedge (M!k \neq Black \vee \neg BtoW (E!ind, M) \vee obc < Blacks\ M \vee (\neg z \wedge ind = R \wedge snd (E!R) = T \wedge (\exists r. ind < r \wedge r < length\ E \wedge BtoW (E!r, M))) \gg$

Van de Snepscheut leaves its verification as an exercise that we carried out successfully.

defs

Propagate-Black-def:

Propagate-Black ≡

{ Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M ∧ bc ⊆ Blacks M }

ind := 0;

{ Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M

∧ bc ⊆ Blacks M ∧ ind = 0 }

while ind < length E

inv { Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M

∧ bc ⊆ Blacks M ∧ PBI_{inv} ind ∧ ind ≤ length E }

do { Proper ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M

∧ bc ⊆ Blacks M ∧ PBI_{inv} ind ∧ ind < length E }

if M! fst (E!ind) = Black

```

then
  { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
     $\wedge$  PBlndv ind  $\wedge$  ind  $<$  length E  $\wedge$  M!fst (E!ind) = Black }
  k := snd (E!ind);;
  { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
     $\wedge$  PBlndv ind  $\wedge$  ind  $<$  length E  $\wedge$  M!fst (E!ind) = Black  $\wedge$  Auxk }
  (M := M [k := Black],, ind := ind + 1)
else
  { Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
     $\wedge$  PBlndv ind  $\wedge$  ind  $<$  length E }
  (if M!fst (E!ind)  $\neq$  Black then ind := ind + 1 fi)
fi
od

```

The evaluation of the condition $M!fst (E!ind) \neq Black$ and the incrementing of ind must be done atomically. If we let a point of interference in between, the mutator could blacken the node, which would falsify the assertion.

lemma *Propagate-Black*:

```

 $\vdash$  Propagate-Black
{ Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
 $\wedge$  (obc  $\subseteq$  Blacks M  $\vee$  Safe (M, E)) }

```

Counting Black Nodes

This phase finally re-establishes the invariant of the collector's outermost loop. The computed set bc must contain all nodes which were black *upon termination of Propagate-Black* or, since Ma records precisely this set, the *Count-Blacks* phase must ensure that $Blacks Ma \subseteq bc$ holds.

The invariant contains the predicate

constdefs

```

CountInv :: gar-coll-state  $\Rightarrow$  nat  $\Rightarrow$  bool
CountInv  $\equiv$   $\ll$   $\lambda ind. \{i. i < ind \wedge Ma!i = Black\} \subseteq bc$   $\gg$ 

```

which upon termination establishes $Blacks Ma \subseteq bc$. The corresponding annotated command is:

defs

```

Count-Blacks-def:
Count-Blacks  $\equiv$ 
{ Proper  $\wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  length Ma = length M

```

$$\begin{aligned}
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{Blacks M} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \wedge \text{bc} = \{\} \} \\
\text{ind} & := 0; \\
\{ & \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks M} \wedge \text{length Ma} = \text{length M} \\
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{Blacks M} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \wedge \text{bc} = \{\} \wedge \text{ind} = 0 \} \\
\text{while} & \text{ ind} < \text{length M} \\
\text{in} & \{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks M} \wedge \text{length Ma} = \text{length M} \\
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{Blacks M} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge \text{CountInv ind} \\
& \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \wedge \text{ind} \leq \text{length M} \} \\
\text{do} & \{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks M} \wedge \text{length Ma} = \text{length M} \\
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{Blacks M} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge \text{CountInv ind} \\
& \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \wedge \text{ind} < \text{length M} \} \\
\text{if} & \text{ M!ind} = \text{Black} \\
\text{then} & \{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks M} \wedge \text{length Ma} = \text{length M} \\
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{Blacks M} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge \text{CountInv ind} \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \\
& \wedge \text{ind} < \text{length M} \wedge \text{M!ind} = \text{Black} \} \\
& \text{bc} := \text{insert ind bc} \\
\text{fi}; \\
\{ & \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks M} \wedge \text{length Ma} = \text{length M} \\
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{Blacks M} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge \text{CountInv (ind} + 1) \\
& \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \wedge \text{ind} < \text{length M} \} \\
\text{ind} & := \text{ind} + 1 \\
\text{od}
\end{aligned}$$

The mutator cannot access the auxiliary variable Ma . Thus, the set of black nodes of Ma remains invariant under the mutator's blackening action so that these annotations are invariant against the mutator's actions.

The postcondition is exactly the outermost invariant of the collector's loop, which must hold at the beginning and at the end of the loop's body.

lemma *Count-Blacks*:

$$\begin{aligned}
& \vdash \text{Count-Blacks} \\
& \{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks M} \wedge \text{length Ma} = \text{length M} \\
& \wedge \text{obc} \subseteq \text{Blacks Ma} \wedge \text{Blacks Ma} \subseteq \text{bc} \wedge \text{bc} \subseteq \text{Blacks M} \\
& \wedge (\text{obc} \subset \text{Blacks Ma} \vee \text{Safe (M, E)}) \}
\end{aligned}$$

With all reachable nodes marked we can proceed to the appending phase where all unmarked nodes are appended to the free list.

Appending to the Free List

Following our predecessors, the operation of appending a garbage node ind to the free list, i.e. making ind reachable, is modeled abstractly by the function:

consts $\text{Append-to-free} :: \text{nat} \times \text{edges} \Rightarrow \text{edges}$

satisfying the following axioms:

axioms

$\text{Append-to-free}_0: \text{length} (\text{Append-to-free} (n, e)) = \text{length } e$

$\text{Append-to-free}_1: \text{Proper-Edges} (m, e)$

$\implies \text{Proper-Edges} (m, \text{Append-to-free} (n, e))$

$\text{Append-to-free}_2: n \notin \text{Reach } e$

$\implies n' \in \text{Reach} (\text{Append-to-free} (n, e)) = (n' = n \vee n' \in \text{Reach } e)$

In the annotated code, **AppendInv** ind states that all white nodes with index ind or larger are garbage, i.e. the safety property is maintained throughout the appending loop.

constdefs

$\text{AppendInv} :: \text{gar-coll-state} \Rightarrow \text{nat} \Rightarrow \text{bool}$

$\text{AppendInv} \equiv \ll \lambda \text{ind}. \forall i < \text{length } M. \text{ind} \leq i \longrightarrow i \in \text{Reach } E \longrightarrow M!i = \text{Black} \gg$

defs

$\text{Append-def}:$

$\text{Append} \equiv$

$\{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{Safe } (M, E) \}$

$\text{ind} := 0;$

$\{ \text{Proper} \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{Safe } (M, E) \wedge \text{ind} = 0 \}$

while $\text{ind} < \text{length } M$

inv $\{ \text{Proper} \wedge \text{AppendInv } \text{ind} \wedge \text{ind} \leq \text{length } M \}$

do $\{ \text{Proper} \wedge \text{AppendInv } \text{ind} \wedge \text{ind} < \text{length } M \}$

if $M!\text{ind} = \text{Black}$ **then**

$\{ \text{Proper} \wedge \text{AppendInv } \text{ind} \wedge \text{ind} < \text{length } M \wedge M!\text{ind} = \text{Black} \}$

$M := M [\text{ind} := \text{White}]$

else

```

    { Proper  $\wedge$  AppendInv ind  $\wedge$  ind < length M  $\wedge$  ind  $\notin$  Reach E }
    E := Append-to-free (ind, E)
  fi;;
  { Proper  $\wedge$  AppendInv (ind + 1)  $\wedge$  ind < length M }
  ind := ind + 1
od

```

The precondition of the assignment to E guarantees that only garbage nodes are collected.

3.3.3 Interference Freedom

The proof outline for the mutator has a total of 5 assertions and 2 atomic actions. The proof outline of the collector has 36 assertions and 20 atomic actions. Hence, the number of interference freedom tests that have to be checked is 172. Obviously many of them are trivial, but in many cases what seemed to be a perfect proof outline revealed a bug only after attempting the proof with the theorem prover.

We carry out part of the interference freeness tests using the modules used to structure the code of the collector and mutator. We prove lemmas about the invariance of the assertions in each module of the collector against the atomic actions in each module of the mutator, and vice versa. These are in total 16 lemmas of the form:

lemma *interfree-Blacken-Roots--Redirect-Edge*:
interfree-aux (Some Blacken-Roots, {}, Some Redirect-Edge)

lemma *interfree-Redirect-Edge--Blacken-Roots*:
interfree-aux (Some Redirect-Edge, {}, Some Blacken-Roots)

The verification conditions that result from the interference-freedom tests represented by these lemmas can be automatically generated. First, the definitions of the modules are unfolded, and then we apply a special tactic called *interfree-aux* which is a “subtactic” of the one used for full parallel programs (see appendix A.3).

To prove several verification conditions that result from the interference freedom tests, we need some auxiliary lemmas about graphs. The first one states that the set of reachable nodes is not increased by the first action of the mutator:

lemma *Graph3*:
 $\llbracket t \in \text{Reach } e; r < \text{length } e \rrbracket \implies \text{Reach } (e [r := (\text{fst } (e!r), t)]) \subseteq \text{Reach } e$

In the proof of *interfree-Propagate-Black--Redirect-Edge* we need the following lemma:

lemma *Graph4*:

$$\begin{aligned} & \llbracket t \in \text{Reach } e; \text{Roots} \subseteq \text{Blacks } m; \text{index} \leq \text{length } e; t < \text{length } m; r < \text{length } e; \\ & \quad \forall i < \text{index}. \neg \text{BtoW } (e!i, m); r < \text{index}; m!\text{fst } (e!r) = \text{Black}; m!t \neq \text{Black} \rrbracket \\ & \implies \exists r. \text{index} \leq r \wedge r < \text{length } e \wedge \text{BtoW } (e [r := (\text{fst } (e!r), t)]!r, m) \end{aligned}$$

establishing that whenever a visited edge $r < \text{index}$ is redirected to a white target t then there exists a not visited edge in the modified list of edges such that it satisfies the predicate *BtoW* as well. A slight variation is needed to prove the invariance of the predicate *Auxk*, namely,

lemma *Graph5*:

$$\begin{aligned} & \llbracket t \in \text{Reach } e; \text{Roots} \subseteq \text{Blacks } m; \forall i < r. \neg \text{BtoW } (e!i, m); t < \text{length } m; \\ & \quad r < \text{length } e; m!\text{fst } (e!r) = \text{Black}; m!\text{snd } (e!r) = \text{Black}; m!t \neq \text{Black} \rrbracket \\ & \implies \exists r'. r < r' \wedge r' < \text{length } e \wedge \text{BtoW } (e [r := (\text{fst } (e!r), t)]!r', m) \end{aligned}$$

Next, we prove the interference freedom of the collector against the mutator and vice versa:

lemma *interfree-Collector--Mutator*:

$$\text{interfree-aux } (\text{Some Collector}, \{\}, \text{Some Mutator})$$

lemma *interfree-Mutator--Collector*:

$$\text{interfree-aux } (\text{Some Mutator}, \{\}, \text{Some Collector})$$

Finally, we prove the derivability of the full program:

lemma *Gar-Coll*:

$$\begin{aligned} & \Vdash \{ \text{Proper} \wedge \text{Mut-init} \wedge z \} \\ & \text{cobegin} \\ & \quad \text{Collector } \{ \text{False} \} \parallel \text{Mutator } \{ \text{False} \} \\ & \text{coend} \\ & \{ \text{False} \} \end{aligned}$$

The tactic *oghoare* is applied without unfolding the definitions of the modules. As a result, the derivability of the components is directly proven by the lemmas *Collector* and *Mutator*, and the interference freedom test consists only of two subgoals, which correspond exactly to the lemmas about interference shown above. By the soundness theorem the algorithm is correct in the sense of partial correctness. The validity of the postcondition is, however,

not interesting since both the mutator and the collector are infinite cycles. The interesting property is the validity of the intermediate annotations. In particular, the precondition of the action which appends nodes to the free list ensures that these nodes are garbage.

3.4 The Multi-Mutator Case

If we allow the interaction with several mutators, new difficulties come into play. We consider a solution, first presented in [Jonker, 1992], where the collector proceeds to the appending phase only after $n + 1$ consecutive executions of the *Propagate-Black* phase where the set of black nodes is not increased. Observe that, for one mutator, this algorithm checks *twice* whether $obc = bc$. [Jonker, 1992] also shows that n consecutive executions suffice, but we do not consider this version here.

The program consists of a fixed, finite and nonempty set of mutator processes and one collector process. The external syntax for parameterized programs is shown in the table 2.4.

3.4.1 The Mutators

A mutator can only redirect an edge when its target is a reachable node. Redirecting an edge may make its old target inaccessible. If several mutators are active, then one of them may select a reachable node T as a new target. Before the edge has been redirected, however, another mutator may render T inaccessible. To solve this problem, selecting the new target and redirecting the edge is modeled as a single atomic action.

Each mutator m selects an edge R_m and a target node T_m . As in the previous section each mutator uses an auxiliary variable Z_m , that indicates if it is pending before the blackening of a node. These three objects are put together as fields of a record:

```
record mut =
  Z :: bool
  R :: nat
  T :: nat
```

Isabelle's syntax for accessing the field Z of a variable Mut of type mut is $Z Mut$. Record update is written $Mut (Z := True)$, meaning that the field Z of the record Mut is updated to the value $True$.

In the algorithm the variable `Muts` is a list of length n (the number of mutators) whose components are records of type `mut`. For example, to access the selected edge of mutator j we write $R(\text{Muts}!j)$.

The variables of the program are the same as in the case for one mutator, except for the list `Muts` used by the mutator and two new variables, `Qa` and `L`, of the collector which we explain in the following sections:

```
record mul-gar-coll-state =
  M :: nodes
  E :: edges
  bc :: nat set
  obc :: nat set
  Ma :: nodes
  ind :: nat
  k :: nat
  Qa :: nat
  L :: nat
  Muts :: mut list
```

In the assertions of the mutator we use the following predicate:

```
constdefs
  Mul-mut-init :: mul-gar-coll-state  $\Rightarrow$  nat  $\Rightarrow$  bool
  Mul-mut-init  $\equiv$   $\ll \lambda n. n = \text{length } \text{Muts} \wedge (\forall i < n. R(\text{Muts}!i) < \text{length } E$ 
     $\wedge T(\text{Muts}!i) < \text{length } M) \gg$ 
```

indicating that the selected edges and targets are within the range of edges and nodes, respectively, and that the list of records `Muts` has an entry for each of the n mutators.

The modules of the mutator's code are functions of the number of mutators n and the particular mutator's index j with $0 \leq j < n$.

```
constdefs
  Mul-Redirect-Edge :: nat  $\Rightarrow$  nat  $\Rightarrow$  mul-gar-coll-state ann-com
  Mul-Redirect-Edge j n  $\equiv$ 
  { Mul-mut-init n  $\wedge Z(\text{Muts}!j)$  }
  <if  $T(\text{Muts}!j) \in \text{Reach } E$  then
  E :=  $E[R(\text{Muts}!j) := (\text{fst } (E!R(\text{Muts}!j)), T(\text{Muts}!j))]$  fi,
  Muts :=  $\text{Muts}[j := (\text{Muts}!j) \langle Z := \text{False} \rangle]$ 

  Mul-Color-Target :: nat  $\Rightarrow$  nat  $\Rightarrow$  mul-gar-coll-state ann-com
```

Mul-Color-Target $j\ n \equiv$
 $\{ \text{Mul-mut-init } n \wedge \neg Z\ (\text{Muts!}j) \}$
 $\langle M := M [T\ (\text{Muts!}j) := \text{Black}],\ \text{Muts} := \text{Muts } [j := (\text{Muts!}j)\ (\{Z := \text{True}\})] \rangle$

Mul-Mutator $::\ \text{nat} \Rightarrow \text{nat} \Rightarrow \text{mul-gar-coll-state ann-com}$

Mul-Mutator $j\ n \equiv$
 $\{ \text{Mul-mut-init } n \wedge Z\ (\text{Muts!}j) \}$
while *True*
 inv $\{ \text{Mul-mut-init } n \wedge Z\ (\text{Muts!}j) \}$
 do *Mul-Redirect-Edge* $j\ n$;;
 Mul-Color-Target $j\ n$
 od

The annotations of the proof outline of the mutators are, like the instructions, parameterized by the number of mutators n and the index j . In chapter 5, we shall show that parameterized annotations for correct specifications of parameterized programs can always be found. We prove the derivability of the parameterized proof outline of a generic mutator:

lemma *Mul-Mutator*: $\llbracket 0 \leq j; j < n \rrbracket \Longrightarrow \vdash \text{Mul-Mutator } j\ n \{ \text{False} \}$

3.4.2 The Collector

In the case of one mutator, if an execution of the body does not establish the safety property, the reason is that some white node was colored black during the execution of *Propagate-Black*. When several mutators are present, there may be other reasons. To describe them we need a new value which represents the number of mutators that are *queueing* to blacken a white node. This value is computed by the function *Queue*, which returns the length of the list that results from filtering the queueing mutators from the list of mutator variables:

constdefs

Queue $::\ \text{mul-gar-coll-state} \Rightarrow \text{nat}$
Queue $\equiv \ll \text{length } (\text{filter } (\lambda i. \neg Z\ i \wedge M!(T\ i) \neq \text{Black})\ \text{Muts}) \gg$

The auxiliary variable *Qa* “records” this value upon termination of the *Propagate-Black* phase. The definition of the predicate *Mul-Proper* requires, besides proper nodes and proper edges, that the length of the variable *Muts* be the number of mutators n :

constdefs

Mul-Propser :: *mul-gar-coll-state* \Rightarrow *nat* \Rightarrow *bool*
Mul-Propser $\equiv \ll \lambda n. \text{ Proper-Roots } M \wedge \text{ Proper-Edges } (M, E)$
 $\wedge \text{ Proper-M-init } M \wedge n = \text{ length Muts } \gg$

We declare the modules used for the collector:

consts

Mul-Blacken-Roots :: *nat* \Rightarrow *mul-gar-coll-state ann-com*
Mul-Propagate-Black :: *nat* \Rightarrow *mul-gar-coll-state ann-com*
Mul-Count-Blacks :: *nat* \Rightarrow *mul-gar-coll-state ann-com*
Mul-Append :: *nat* \Rightarrow *mul-gar-coll-state ann-com*

The variable *L* is a counter that keeps track of how many consecutive times the values of *obc* and *bc* coincide. When it reaches the value $n + 1$ the conditions satisfy the safety requirement and the collector proceeds to collect the unmarked nodes. The proof outline of the collector is:

constdefs

Mul-Collector :: *nat* \Rightarrow *mul-gar-coll-state ann-com*
Mul-Collector *n* \equiv
 $\{ \text{ Mul-Propser } n \}$
while *True* **inv** $\{ \text{ Mul-Propser } n \}$
do
Mul-Blacken-Roots *n* ;;
 $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \} \text{ obc} := \{ \};;$
 $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \wedge \text{ obc} = \{ \} \} \text{ bc} := \text{ Roots};;$
 $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \wedge \text{ obc} = \{ \} \wedge \text{ bc} = \text{ Roots} \} \text{ L} := 0;;$
 $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \wedge \text{ obc} = \{ \} \wedge \text{ bc} = \text{ Roots} \wedge \text{ L} = 0 \}$
while $\text{ L} < n + 1$
inv $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \wedge \text{ bc} \subseteq \text{ Blacks } M \wedge$
 $(\text{ Safe } (M, E) \vee (\text{ L} \leq \text{ Queue} \vee \text{ bc} \subset \text{ Blacks } M)) \wedge \text{ L} < n + 1 \}$
do
 $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \wedge \text{ bc} \subseteq \text{ Blacks } M$
 $\wedge (\text{ Safe } (M, E) \vee \text{ L} \leq \text{ Queue} \vee \text{ bc} \subset \text{ Blacks } M) \}$
 $\text{ obc} := \text{ bc};;$
Mul-Propagate-Black *n*;;
 $\{ \text{ Mul-Propser } n \wedge \text{ Roots } \subseteq \text{ Blacks } M \wedge \text{ obc} \subseteq \text{ Blacks } M \wedge \text{ bc} \subseteq \text{ Blacks } M$
 $\wedge (\text{ Safe } (M, E) \vee \text{ obc} \subset \text{ Blacks } M \vee \text{ L} < \text{ Queue}$
 $\wedge (\text{ L} \leq \text{ Queue} \vee \text{ obc} \subset \text{ Blacks } M)) \}$

```

bc := {};;
{ Mul-Propser n ∧ Roots ⊆ Blacks M ∧ obc ⊆ Blacks M ∧ bc ⊆ Blacks M
  ∧ (Safe (M, E) ∨ obc ⊂ Blacks M ∨ L < Queue
  ∧ (L ≤ Queue ∨ obc ⊂ Blacks M)) ∧ bc = {} }
⟨Ma := M,, Qa := Queue⟩;;
Mul-Count-Blacks n;;
{ Mul-Propser n ∧ Roots ⊆ Blacks M ∧ length Ma = length M
  ∧ obc ⊆ Blacks Ma ∧ Blacks Ma ⊆ Blacks M ∧ bc ⊆ Blacks M
  ∧ Blacks Ma ⊆ bc ∧ (Safe (M, E) ∨ obc ⊂ Blacks Ma ∨
  L < Qa ∧ (Qa ≤ Queue ∨ obc ⊂ Blacks M)) ∧ Qa < n + 1 }
if obc = bc
then
  { Mul-Propser n ∧ Roots ⊆ Blacks M ∧ length Ma = length M
    ∧ obc ⊆ Blacks Ma ∧ Blacks Ma ⊆ Blacks M ∧ bc ⊆ Blacks M
    ∧ Blacks Ma ⊆ bc ∧ (Safe (M, E) ∨ obc ⊂ Blacks Ma ∨
    L < Qa ∧ (Qa ≤ Queue ∨ obc ⊂ Blacks M)) ∧ Qa < n + 1
    ∧ obc = bc }
  L := L + 1
else
  { Mul-Propser n ∧ Roots ⊆ Blacks M ∧ length Ma = length M
    ∧ obc ⊆ Blacks Ma ∧ Blacks Ma ⊆ Blacks M ∧ bc ⊆ Blacks M
    ∧ Blacks Ma ⊆ bc ∧ (Safe (M, E) ∨ obc ⊂ Blacks Ma ∨
    L < Qa ∧ (Qa ≤ Queue ∨ obc ⊂ Blacks M)) ∧ Qa < n + 1
    ∧ obc ≠ bc }
  L := 0
fi
od;;
Mul-Append n
od

```

The invariant of the one-mutator case must be compared with the precondition of the **if–then–else** instruction, because both correspond to the assertion established by the phase that counts the black nodes. The assertion $\text{Safe}(M, E) \vee \text{obc} \subset \text{Blacks } Ma$ has been weakened with a new disjunct, corresponding to the new situation which can prevent $\text{Safe}(M, E)$ from holding. The new disjunct corresponds to the case in which at least one mutator joins the queue during the *Mul-Propagate-Black* phase, i.e. $L < Qa$. In this case it is also necessary to distinguish whether some mutator leaves the queue, i.e. colors its white target ($\text{obc} \subset \text{Blacks } M$), or none leaves the queue, i.e. the queue is not decreased ($Qa \leq \text{Queue}$). Intuitively, after $n + 1$

non-blackening *Mul-Propagate-Black* iterations, the property *Safe* (M, E) must hold, since the number of queueing mutators cannot exceed n .

The codes of the modules are the same as in §3.3 except for the annotations in the *Mul-Propagate-Black* and *Count-Blacks* phases, which have to be adapted to the new invariant.

We just show the *Mul-Propagate-Black* phase.

constdefs

Mul-PBInv :: *mul-gar-coll-state* \Rightarrow *bool*

Mul-PBInv $\equiv \ll$ *Safe* (M, E) \vee *obc* \subseteq *Blacks* M \vee L < Queue
 \vee ($\forall i < \text{ind. } \neg \text{BtoW } (E!i, M)) \wedge$ L \leq Queue \gg

Mul-Auxk :: *mul-gar-coll-state* \Rightarrow *bool*

Mul-Auxk $\equiv \ll$ L < Queue \vee M!k \neq Black \vee $\neg \text{BtoW } (E!\text{ind}, M)$
 \vee *obc* \subseteq *Blacks* M \gg

defs

Mul-Propagate-Black-def:

Mul-Propagate-Black $n \equiv$

$\{ \text{Mul-Prop} } n \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{obc} \subseteq \text{Blacks } M \wedge \text{bc} \subseteq \text{Blacks } M$
 $\wedge (\text{Safe } (M, E) \vee L \leq \text{Queue} \vee \text{obc} \subseteq \text{Blacks } M) \}$

ind := 0;;

$\{ \text{Mul-Prop} } n \wedge \text{Roots} \subseteq \text{Blacks } M$
 $\wedge \text{obc} \subseteq \text{Blacks } M \wedge \text{Blacks } M \subseteq \text{Blacks } M \wedge \text{bc} \subseteq \text{Blacks } M$
 $\wedge (\text{Safe } (M, E) \vee L \leq \text{Queue} \vee \text{obc} \subseteq \text{Blacks } M) \wedge \text{ind} = 0 \}$

while *ind* < *length* E

inv $\{ \text{Mul-Prop} } n \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{obc} \subseteq \text{Blacks } M \wedge \text{bc} \subseteq \text{Blacks } M$
 $\wedge \text{Mul-PBInv} \wedge \text{ind} \leq \text{length } E \}$

do $\{ \text{Mul-Prop} } n \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{obc} \subseteq \text{Blacks } M \wedge \text{bc} \subseteq \text{Blacks } M$
 $\wedge \text{Mul-PBInv} \wedge \text{ind} < \text{length } E \}$

if M!*fst* (E!*ind*) = Black

then

$\{ \text{Mul-Prop} } n \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{obc} \subseteq \text{Blacks } M \wedge \text{bc} \subseteq \text{Blacks } M$
 $\wedge \text{Mul-PBInv} \wedge \text{M!fst } (E!\text{ind}) = \text{Black} \wedge \text{ind} < \text{length } E \}$

k := *snd* (E!*ind*);;

$\{ \text{Mul-Prop} } n \wedge \text{Roots} \subseteq \text{Blacks } M \wedge \text{obc} \subseteq \text{Blacks } M \wedge \text{bc} \subseteq \text{Blacks } M$
 $\wedge (\text{Safe } (M, E) \vee \text{obc} \subseteq \text{Blacks } M \vee L < \text{Queue} \vee$
 $(\forall i < \text{ind. } \neg \text{BtoW } (E!i, M)) \wedge L \leq \text{Queue} \wedge \text{Mul-Auxk })$
 $\wedge k < \text{length } M \wedge \text{M!fst } (E!\text{ind}) = \text{Black} \wedge \text{ind} < \text{length } E \}$

(M := M [k := Black],, ind := ind + 1)

```

else
   $\{\!|$  Mul-Proper  $n \wedge$  Roots  $\subseteq$  Blacks M  $\wedge$  obc  $\subseteq$  Blacks M  $\wedge$  bc  $\subseteq$  Blacks M
   $\wedge$  Mul-PBInv  $\wedge$  ind  $<$  length E  $\!\}$ 
   $\langle$  if M!fst (E!ind)  $\neq$  Black then ind := ind + 1 fi  $\rangle$ 
fi
od

```

If we expand the predicate Mul-PBInv in the invariant we obtain

$$\begin{aligned}
& Safe (M, E) \vee obc \subset Blacks\ M \vee L < Queue \\
& \vee (\forall i < ind. \neg BtoW (E!i, M)) \wedge L \leq Queue.
\end{aligned}$$

Any coloring establishes $obc \subset Blacks\ M$. (Observe that only coloring can make the queue shorter.) If no coloring occurs, then, either all the visited edges point to a black node, or some mutator has redirected an edge to a white source, but has not yet colored the target, which amounts to saying that the queue grows, i.e. $L < Queue$.

The next lemma proves derivability of the collector's proof outline in the system:

lemma *Mul-Collector*: \vdash *Mul-Collector* n $\{\!|$ False $\!\}$

3.4.3 Interference Freedom

The collector has a total of 40 assertions and 21 atomic actions. One mutator has 5 assertions and 2 atomic actions. For the interference freedom test it suffices to consider the following combinations:

1. Invariance of the assertions in the collector against the actions of a generic mutator with some index j such that $0 \leq j < n$, and vice versa.
2. Invariance of the assertions of one mutator j against the actions of another mutator i such that $i \neq j$.

This results in a total of 195 interference freedom proofs. Like for the one-mutator case we perform combinations of the modules in separate lemmas that are then applied directly to the verification of the parallel composition. The interference freedom among the mutators is proven in the following lemma:

lemma *Mul-interfree-Mutator--Mutator*: $\llbracket i < n; j < n; i \neq j \rrbracket \implies$
interfree-aux (Some (Mul-Mutator i n), {}, Some (Mul-Mutator j n))

The derivability of the parallel composition of the mutators, i.e. the lemma

lemma *Mul-Parameterized-Mutators*: $0 < n \implies$
 $\Vdash \{ \text{Mul-mut-init } n \wedge (\forall i < n. Z (\text{Muts!}i)) \}$
cobegin
 scheme $[0 \leq j < n]$ *Mul-Mutator* j n $\{ \text{False} \}$
coend
 $\{ \text{False} \}$

is proven by applying the tactic *oghoare* without unfolding the definition of *Mul-Mutator*. The tactic generates four subgoals. Two of them correspond to the two verification conditions about the logical implications of the overall precondition and postcondition and those of the components. A third subgoal stating the derivability of one generic mutator (proven by the lemma *Mul-Mutator*), and the fourth subgoal stating the interference between two generic mutators (proven by the lemma *Mul-interfree-Mutator--Mutator*). Notice that there is no need for induction or any other proof method; parameterized programs are directly handled by the proof rules of the system.

The rest of the interference freedom tests are similar to those of the one-mutator case:

lemma *Mul-interfree-Collector--Mutator*: $j < n \implies$
 $\text{interfree-aux } (\text{Some } (\text{Mul-Collector } n), \{\}, \text{Some } (\text{Mul-Mutator } j \ n))$
lemma *Mul-interfree-Mutator--Collector*: $j < n \implies$
 $\text{interfree-aux } (\text{Some } (\text{Mul-Mutator } j \ n), \{\}, \text{Some } (\text{Mul-Collector } n))$

Finally, we prove the derivability in the Owicki-Gries system of the full program:

lemma *Mul-Gar-Coll*:
 $\Vdash \{ \text{Mul-Proper } n \wedge \text{Mul-mut-init } n \wedge (\forall i < n. Z (\text{Muts!}i)) \}$
cobegin
 Mul-Collector n $\{ \text{False} \}$
 ||
 scheme $[0 \leq j < n]$ *Mul-Mutator* j n $\{ \text{False} \}$
coend
 $\{ \text{False} \}$

3.5 Conclusions and Related Work

We have provided mechanically checked Owicki-Gries proofs for two garbage collection algorithms. The Owicki-Gries method splits the proof into a large

number of simple interference freedom subproofs. These are very tedious to prove by hand, and so avoided by humans, who prefer to concentrate on the few difficult cases. By applying the formalized Owicki-Gries system most of the interference freedom proofs for the final annotations were automatically carried out by Isabelle/HOL. For the remaining cases, five non-trivial lemmas about graphs had to be supplied. The proofs of these lemmas, however, were very interactive.

We do not know of any complete Owicki-Gries proof for any of the two algorithms. In his proof of Ben-Ari's algorithm, [van de Snepscheut, 1987] mixes the Owicki-Gries method with ad-hoc reasoning; in particular, he does not provide an invariant for the outermost loop, implicitly claiming that doing so would be complicated. However, the invariant turns out to be simple (3 clauses), and has a clear intuitive interpretation.

For the n -mutators algorithm, [Jonker, 1992] argues that

A proof according to the Owicki-Gries theory would require the introduction of a satisfactory number of ghost variables. In an earlier version of this paper the invariant we constructed was rather unwieldy and the proof of invariance almost unreadable.

However, our proof only uses two auxiliary variables (Ma and Qa), plus a trivial auxiliary variable for each mutator. Jonker considers in his paper several variations of the algorithm. We believe that Owicki-Gries proofs for these variations should be possible to obtain from the proof presented here with reasonable effort.

We know of two other mechanized proofs of Ben-Ari's algorithm, carried out using the Boyer-Moore theorem prover [Russinoff, 1994] and the PVS theorem prover [Havelund, 1996, Havelund and Shankar, 1997]. A main advantage of our approach is the closeness to the original program text, which simplifies the interaction with the prover. Annotated programs are fairly readable by humans, and they are also directly accepted as input by Isabelle. In other approaches the program must first be translated into a different language (e.g. LISP in [Russinoff, 1994]).

Another aspect of our formalization is that we only had to prove 13 lemmas (7 of them trivial) about graph functions, whereas 100 lemmas were required in [Russinoff, 1994], and about 55 in [Havelund, 1996, Havelund and Shankar, 1997]. The reason for this is probably that many trivial lemmas about sets or lists could be automatically proven using Isabelle's built-in tactics (rewriting, classical reasoning, etc.) and Isabelle's standard libraries. The proof effort, however, took two months for the one-mutator algorithm (similar to our predecessors) and another two months for the n -mutator case.

Most of the time was consumed in finding and improving the annotations.

A disadvantage of the Owicki-Gries method (in its classical version) is that it can only be applied to safety properties. A liveness property, namely *every garbage node is eventually collected*, is also very important for garbage collection algorithms. Mechanical proofs of this liveness property are found in [Russinoff, 1994] for Ben Ari's algorithm and in [Jackson, 1998], where the safety and liveness property of a predecessor of Ben Ari's algorithm [Dijkstra *et al.*, 1978] is proven using PVS.

None of our two algorithms has been proven correct using fully automatic methods. In [Bruns, 1997], there is a proof of Ben Ari's algorithm for 1 mutator and 4 memory cells. In [Das *et al.*, 1999], a predecessor of Ben Ari's algorithm is proved correct using automatic tools for generating and proving invariants. The key invariants, however, require intelligent input from the user. The paper suggests using predicate abstraction for checking or strengthening invariants in a larger verification effort involving interactive theorem provers, which is a promising idea.

Our overall conclusion is that the application of a theorem prover greatly enhances the applicability of the Owicki-Gries method. The closeness to the original program is preserved, and the large number of routine proofs is considerably automated.

Chapter 4

The Rely-Guarantee Method in Isabelle/HOL

This chapter presents the formalization of the rely-guarantee method for correctness proofs of parallel imperative programs with shared variables in the theorem prover Isabelle/HOL. This method was first proposed by Jones [Jones, 1981, Jones, 1983] and can be seen as the compositional version of Owicki-Gries. We closely follow the presentation in [Xu *et al.*, 1997], where a sound and complete version of the system is presented in a conventional pencil and paper style. However, some aspects of our formalization differ from this model. The reasons for these modifications will be addressed as we encounter them.

The rely-guarantee system provides compositional proof rules for the verification of parallel programs. This is accomplished by enriching the specification of each component with conditions concerning the interaction with the environment. A rely-guarantee specification defines four sets: the sets of initial and final states (pre and postcondition) and the sets characterizing the effect of actions performed by the environment or by the component itself (rely and guarantee conditions).

It is important to observe the strong connection to Owicki-Gries. In the latter, programs were annotated at every point of interference and the verification process required proving interference freedom of the annotations of each component against the atomic actions of the other components. The idea of the rely-guarantee method is to record the interference information in the specification of each component. This information consists of the rely condition stating what the component expects from the environment, and the guarantee condition, stating the effect of the component itself on

the environment. By using the information given by the rely and guarantee conditions of each component, the verification of a parallel program can be carried out in a compositional way.

The compositionality of this method has two major advantages over the non-compositionality of Owicki-Gries. First of all, it drastically reduces the complexity of the verification process. The number of correctness proofs in the rely-guarantee formalism increases only linearly with the number of parallel components, while exponentially in the Owicki-Gries method. Secondly, it allows the verification of so-called *open systems*, i.e. systems where a specified margin of interaction from an arbitrary environment is permitted. Then, new components, whose actions respect this margin, can be added a posteriori. The verification of the new system can be done without looking into the structure of the previously verified program. This makes the method adequate for top-down development of parallel systems. In contrast, the Owicki-Gries method only works for *closed systems*, where all component programs must be simultaneously known. If a new component program is added to the parallel composition, the verification process must be restarted due to the non-compositionality of the interference freedom test.

The chapter is organized as follows: section 4.1 presents the abstract syntax and section 4.2 the operational semantics of the language. Section 4.3 defines the notion of validity of a specification in the rely-guarantee formalism and section 4.4 presents the rules of the proof system. The soundness of the proof system with relation to the underlying semantics is proven in section 4.5. Section 4.6 defines a user-friendly concrete syntax for using the method and section 4.7 shows the applicability of the formalization by verifying several examples. Section 4.8 summarizes the main results.

4.1 Abstract Syntax

Like in the previous formalization, the languages of sequential and parallel programs are defined in different layers.

types

α *bexp* = α *set*

datatype α *com* =

Basic ($\alpha \Rightarrow \alpha$)
 | *Seq* (α *com*) (α *com*)
 | *Cond* (α *bexp*) (α *com*) (α *com*)
 | *While* (α *bexp*) (α *com*)

| *Await* (α *be x p*) (α *com*)
types α *par-com* = α *com option list*

The language of component programs α *com* is a standard sequential while-language augmented with the known synchronization construct *Await*. The only difference with the language of component programs of section 2.1 is the lack of assertions in the syntax. Component programs are not presented as proof outlines, so intermediate assertions can be omitted.

Parallelism is defined in a separate layer with type α *par-com*; it is simply a list of optional component programs. This is analogous to the constructor *Parallel* defined in the Owicki-Gries language. For simplicity, we consider parallelism only at the top level, i.e. no nested parallelism. Moreover, contrary to the formalization in 2.1 parallelism appears as a single construction, i.e. there are neither sequential nor if- nor while-constructions for parallel programs. This way, we reduce the number of constructors and avoid having to duplicate proofs.

Like in chapter 2, this model of parallelism allows composition of any number of sequential component programs by grouping them in a *list*. Representation of parameterized parallel programs may be achieved by means of the function *map*.

4.2 Operational Semantics

The execution of a component program is characterized by an operational semantics that distinguishes between two kinds of transitions: *program* (or *component*) *transitions*, performed by the component itself, and *environment transitions*, performed by another component or an arbitrary environment. The latter affects the state but leaves the program unchanged.

The set of rules defining program transitions is analogous to the rules presented in §2.2. Execution of programs is described via *computations*, which record the sequence of transitions of both kinds. This semantics allows us to define the computation of parallel programs in terms of the computations of the components via a special operator described in section 4.5.2.

Next, we introduce the rules of the semantics and the definition of computation for each layer of the language.

4.2.1 Semantics of Component Programs

A configuration is a pair (P, σ) , where P is some program or *None* standing for a terminated program, and σ is a state.

types $\alpha \text{ conf} = \alpha \text{ com option} \times \alpha$

A transition is represented by a labelled arrow connecting the beginning and ending configurations. There are two kinds of transitions:

Environment transitions, labelled with e , represent a step from the environment and can only change the state.

consts $\text{etran} :: (\alpha \text{ conf} \times \alpha \text{ conf}) \text{ set}$
syntax $\text{-etran} :: \alpha \text{ conf} \Rightarrow \alpha \text{ conf} \Rightarrow \text{bool}$ $(- \text{-}e\rightarrow -)$
translations $P \text{-}e\rightarrow Q \Leftrightarrow (P, Q) \in \text{etran}$
inductive etran
intros
Env: $(P, s) \text{-}e\rightarrow (P, t)$

Component transitions, labelled with c , represent a step of a sequential component program

consts $\text{ctran} :: (\alpha \text{ conf} \times \alpha \text{ conf}) \text{ set}$
syntax
 $\text{-ctran} :: \alpha \text{ conf} \Rightarrow \alpha \text{ conf} \Rightarrow \text{bool}$ $(- \text{-}c\rightarrow -)$
 $\text{-ctran}^* :: \alpha \text{ conf} \Rightarrow \alpha \text{ conf} \Rightarrow \text{bool}$ $(- \text{-}c^*\rightarrow -)$
translations
 $P \text{-}c\rightarrow Q \Leftrightarrow (P, Q) \in \text{ctran}$
 $P \text{-}c^*\rightarrow Q \Leftrightarrow (P, Q) \in \text{ctran}^*$

where $P \text{-}c^*\rightarrow Q$ is the reflexive transitive closure of $P \text{-}c\rightarrow Q$.

inductive ctran

intros

Basic: $(\text{Some } (\text{Basic } f), s) \text{-}c\rightarrow (\text{None}, f s)$

Seq1: $(\text{Some } P_0, s) \text{-}c\rightarrow (\text{None}, t)$
 $\Rightarrow (\text{Some } (\text{Seq } P_0 P_1), s) \text{-}c\rightarrow (\text{Some } P_1, t)$

Seq2: $(\text{Some } P_0, s) \text{-}c\rightarrow (\text{Some } P_2, t)$
 $\Rightarrow (\text{Some } (\text{Seq } P_0 P_1), s) \text{-}c\rightarrow (\text{Some } (\text{Seq } P_2 P_1), t)$

CondT: $s \in b \implies (\text{Some } (\text{Cond } b P_1 P_2), s) -c \rightarrow (\text{Some } P_1, s)$

CondF: $s \notin b \implies (\text{Some } (\text{Cond } b P_1 P_2), s) -c \rightarrow (\text{Some } P_2, s)$

WhileF: $s \notin b \implies (\text{Some } (\text{While } b P), s) -c \rightarrow (\text{None}, s)$

WhileT: $s \in b \implies (\text{Some } (\text{While } b P), s) -c \rightarrow (\text{Some } (\text{Seq } P (\text{While } b P)), s)$

Await: $\llbracket s \in b; (\text{Some } P, s) -c^* \rightarrow (\text{None}, t) \rrbracket$

$\implies (\text{Some } (\text{Await } b P), s) -c \rightarrow (\text{None}, t)$

Basic actions and evaluation of boolean conditions are atomic. In both conditional and iteration statements, the evaluation of the boolean tests are atomic, but a step of the environment can interrupt between the boolean test and the first action from the corresponding program body. The body of an await-statement is executed atomically, thus no environment transitions can occur.

It is usual to ensure that await-statements always terminate by disallowing iteration and await-statements in the body, however, this restriction is not necessary for the soundness proofs of this formalization and is thus not required.

4.2.2 Semantics of Parallel Programs

The semantics of parallel programs is also defined by transition rules between configurations. A configuration for a parallel program is a pair formed by a program (α *par-com*) and a state. A parallel program has terminated if so have all its components, i.e. when all component programs are *None*, but this is also of type α *par-com*. Thus, we do not need to wrap the program part into an *option* type.

types

α *par-conf* = α *par-com* \times α

Transitions may be from the environment, labelled with *pe*, or from the parallel program, labelled with *pc*.

consts

par-etran :: $(\alpha$ *par-conf* \times α *par-conf*) *set*

par-ctran :: $(\alpha$ *par-conf* \times α *par-conf*) *set*

syntax

$-par-etrans:: \alpha \text{ par-conf} \Rightarrow \alpha \text{ par-conf} \Rightarrow \text{bool} \quad (- -pe \rightarrow -)$
 $-par-ctran:: \alpha \text{ par-conf} \Rightarrow \alpha \text{ par-conf} \Rightarrow \text{bool} \quad (- -pc \rightarrow -)$

translations

$P -pe \rightarrow Q \Leftrightarrow (P, Q) \in \text{par-etrans}$
 $P -pc \rightarrow Q \Leftrightarrow (P, Q) \in \text{par-ctran}$

The transition rule for environment transitions is as expected.

inductive *par-etrans***intros**

ParEnv: $(Ps, s) -pe \rightarrow (Ps, t)$

The execution of a parallel program is modeled by a nondeterministic interleaving of the atomic actions of the components. In other words, a parallel program performs a component step when one of its non-terminated components performs a component step.

inductive *par-ctran***intros**

ParComp: $\llbracket i < \text{length } Ps; (Ps!i, s) -c \rightarrow (r, t) \rrbracket$
 $\implies (Ps, s) -pc \rightarrow (Ps[i:=r], t)$

$Ps[i:=r]$ is the list of programs Ps with the program i replaced by r . This is the only transition rule. If we extend the syntax with other constructors at the parallel level, the set of rules defining the semantics should be augmented with the corresponding rules.

4.2.3 Computations

A computation is defined in [Xu *et al.*, 1997] as any sequence of the form

$$(P_0, \sigma_0) \xrightarrow{\delta_1} (P_1, \sigma_1) \xrightarrow{\delta_2} \dots \xrightarrow{\delta_n} (P_n, \sigma_n) \xrightarrow{\delta_{n+1}} \dots, \delta_i \in \{e, c\}$$

There are many ways of formalizing this concept. Given a definition of computation, the main requirement is to be able to access the program fragment and the state of each configuration, and also the kind of transition between two configurations.

The solution we adopted is to model computations as an inductive set of lists of configurations. The one-element list is always a computation, and

two inference rules, one for each kind of transition, determine which lists belong to the inductive set.

types α *confs* = α *conf list*

consts *cptn* :: α *confs set*

inductive *cptn*

intros

CptnOne: $[(P, s)] \in \textit{cptn}$

CptnEnv: $(P, t)\#xs \in \textit{cptn} \implies (P, s)\#(P, t)\#xs \in \textit{cptn}$

CptnComp: $\llbracket (P, s) -c \rightarrow (Q, t); (Q, t)\#xs \in \textit{cptn} \rrbracket$
 $\implies (P, s)\#(Q, t)\#xs \in \textit{cptn}$

Given two consecutive configurations in a computation it is always possible to determine the kind of transition between them by comparing both program fragments: environment transitions leave the program unchanged while component transitions always change it. Computations of parallel programs are defined analogously.

types α *par-confs* = α *par-conf list*

consts *par-cptn* :: α *par-confs set*

inductive *par-cptn*

intros

ParCptnOne: $[(P, s)] \in \textit{par-cptn}$

ParCptnEnv: $(P, t)\#xs \in \textit{par-cptn} \implies (P, s)\#(P, t)\#xs \in \textit{par-cptn}$

ParCptnComp: $\llbracket (P, s) -pc \rightarrow (Q, t); (Q, t)\#xs \in \textit{par-cptn} \rrbracket$
 $\implies (P, s)\#(Q, t)\#xs \in \textit{par-cptn}$

The set of computations of a program P starting from some initial state s is defined as the set of lists of configurations with first element the pair (P, s) which are a computation.

constdefs

cp :: α *com option* \Rightarrow $\alpha \Rightarrow \alpha$ *confs set*

cp P s $\equiv \{l. !!0 = (P, s) \wedge l \in \textit{cptn}\}$

par-cp :: α *par-com* \Rightarrow $\alpha \Rightarrow \alpha$ *par-confs set*

par-cp P s $\equiv \{l. !!0 = (P, s) \wedge l \in \textit{par-cptn}\}$

4.2.4 Modular Definition of Computation

The definition of computation of sequential programs presented in the previous section follows the one proposed in [Xu *et al.*, 1997] and is probably

the most natural and intuitive approach. However, it represents the execution of a program in a simplified linear way without taking into account the inherent structure of the development of a computation.

In the definition of the programming language, however, we observe a well defined structure. For example, the sequential composition is formed from two programs, and the body of a while or an await constructor is itself a program. This structure is automatically reflected in the corresponding computations.

For the proof of some properties, this modular structure is very important. Trying to retrieve this information out of the linear representation of the computation results in tedious and illegible proofs. Such proofs are not appropriate for being carried out in a theorem prover and can often be avoided by redefining concepts. The alternative definition for computations proposed in this section explicitly shows the structure of the program, thus considerably simplifying some proofs, especially those concerning properties of while-programs.

First, we define the auxiliary function *lift* that returns, given a configuration and a program Q , the same configuration where the program has been sequentially composed with Q . If the concerned program is finished, i.e. *None*, the returned program is just Q .

constdefs

$lift :: \alpha\ com \Rightarrow \alpha\ conf \Rightarrow \alpha\ conf$
 $lift\ Q \equiv \lambda(P,s). (if\ P=None\ then\ (Some\ Q,\ s)\ else\ (Some(Seq\ (the\ P)\ Q),\ s))$

The set of computations can be defined respecting the modular structure by the following rules:

consts $cptn-mod :: \alpha\ confs\ set$

inductive $cptn-mod$

intros

$CptnModOne: [(P, s)] \in\ cptn-mod$
 $CptnModEnv: (P, t)\#xs \in\ cptn-mod \implies (P, s)\#(P, t)\#xs \in\ cptn-mod$
 $CptnModNone: [(Some\ P, s) -c\to (None, t); (None, t)\#xs \in\ cptn-mod]$
 $\implies (Some\ P, s)\#(None, t)\#xs \in\ cptn-mod$

$CptnModCondT: [(Some\ P_0, s)\#ys \in\ cptn-mod; s \in\ b]$
 $\implies (Some\ (Cond\ b\ P_0\ P_1), s)\#(Some\ P_0, s)\#ys \in\ cptn-mod$
 $CptnModCondF: [(Some\ P_1, s)\#ys \in\ cptn-mod; s \notin\ b]$
 $\implies (Some\ (Cond\ b\ P_0\ P_1), s)\#(Some\ P_1, s)\#ys \in\ cptn-mod$

CptnModSeq1: $\llbracket (Some\ P_0, s)\#xs \in\ cptn\text{-}mod; zs = map\ (lift\ P_1)\ xs \rrbracket$
 $\implies (Some\ (Seq\ P_0\ P_1), s)\#zs \in\ cptn\text{-}mod$

CptnModSeq2:

$\llbracket (Some\ P_0, s)\#xs \in\ cptn\text{-}mod; fst\ (last\ ((Some\ P_0, s)\#xs)) = None;$
 $(Some\ P_1, snd\ (last\ ((Some\ P_0, s)\#xs)))\#ys \in\ cptn\text{-}mod;$
 $zs = (map\ (lift\ P_1)\ xs)\@ys \rrbracket \implies (Some\ (Seq\ P_0\ P_1), s)\#zs \in\ cptn\text{-}mod$

CptnModWhile1:

$\llbracket (Some\ P, s)\#xs \in\ cptn\text{-}mod; s \in\ b; zs = map\ (lift\ (While\ b\ P))\ xs \rrbracket$
 $\implies (Some\ (While\ b\ P), s)\#(Some\ (Seq\ P\ (While\ b\ P)), s)\#zs \in\ cptn\text{-}mod$

CptnModWhile2:

$\llbracket (Some\ P, s)\#xs \in\ cptn\text{-}mod; fst\ (last\ ((Some\ P, s)\#xs)) = None;$
 $s \in\ b; zs = (map\ (lift\ (While\ b\ P))\ xs)\@ys;$
 $(Some\ (While\ b\ P), snd\ (last\ ((Some\ P, s)\#xs)))\#ys \in\ cptn\text{-}mod \rrbracket$
 $\implies (Some\ (While\ b\ P), s)\#(Some\ (Seq\ P\ (While\ b\ P)), s)\#zs \in\ cptn\text{-}mod$

The first two rules are the same as in the set of rules defining *cptn*. The third rule of *cptn*, namely *CptnComp*, is now replaced by seven rules which not only take into account that the first step is performed by the component program but also consider the kind of program performing the step.

The rule *CptnModNone* summarizes the three possible steps where the program terminates: *Basic*, *WhileF* and *Await*. The two rules for the conditional are obvious. Observe that for these five cases the new definition does not provide any richer information than the *CptnComp* rule with case analysis on the corresponding *c*-step.

The rule *CptnModSeq1* represents the computations of a sequential composition where execution does not enter the second program, and *CptnModSeq2* those who at least finish the first program. For while-programs a computation might enter the body but not finish it (*CptnModWhile1*) or finish it at least once (*CptnModWhile2*).

The new definition is useful for proofs about computations of while-programs because, in general, we do not know how often the body is executed. By using rule induction on *cptn-mod* we directly obtain the three following cases:

1. *CptnModNone*: the while-body is not entered.
2. *CptnModWhile1*: the execution of the body is at least started.
3. *CptnModWhile2*: the body is executed completely at least once followed by a new computation of the same while-program, on which the

induction hypothesis holds.

The proof power of applying rule induction to *cptn-mod* is, at least for the while-case, decisive for the proof of soundness in §4.5. In contrast, the information obtained by using the same proof method on *cptn* was almost useless.

Equivalence of both Definitions

The new definition of computation does not represent the intuitive idea of a computation as obviously as the previous one does. The reader might not be convinced that the set generated by these rules contains all computations defined by the set *cptn* and vice versa. For this reason, and also because we still want to use the previous definition when it is convenient, we prove their equivalence in the following theorem:

theorem *cptn-iff-cptn-mod*: $(c \in \textit{cptn}) = (c \in \textit{cptn-mod})$

Proof. The if-direction is fairly easy.

lemma *cptn-if-cptn-mod*: $c \in \textit{cptn-mod} \implies c \in \textit{cptn}$

It is proven by rule induction on *cptn-mod*. The only-if-direction is more complicated since it requires recovering the missing structure.

lemma *cptn-onlyif-cptn-mod*: $c \in \textit{cptn} \implies c \in \textit{cptn-mod}$

It is proved by rule induction on *cptn*, with a nested structural induction on the program for the case where the first step is made by the component, i.e. we need to prove the following auxiliary lemma by structural induction on *a*:

lemma *cptn-onlyif-cptn-mod-aux*:

$$\begin{aligned} & \llbracket (\textit{Some } a, s) -c \rightarrow (Q, t); (Q, t) \# xs \in \textit{cptn-mod} \rrbracket \\ & \implies (\textit{Some } a, s) \# (Q, t) \# xs \in \textit{cptn-mod} \end{aligned}$$

In the proof of this lemma we need an important property stating that the computation of a sequential composition of programs can be divided into a computation of the first program and a computation of the second one.

lemma *div-seq*:

$$(\textit{Some } (\textit{Seq } P \ Q), s) \# zs \in \textit{cptn-mod} \implies$$

$$\begin{aligned}
& \exists xs. (Some P, s) \# xs \in cptn-mod \wedge \\
& (zs = map (lift Q) xs \vee fst (last ((Some P, s) \# xs)) = None \wedge \\
& (\exists ys. (Some Q, snd (last ((Some P, s) \# xs))) \# ys \in cptn-mod \\
& \quad \wedge zs = map (lift Q) xs @ ys))
\end{aligned}$$

The proof is by rule induction on *cptn-mod*. □

4.3 Validity of Correctness Formulas

A rely-guarantee correctness formula (or specification) of a program P consists of the quadruple $(pre, rely, guar, post)$. These four conditions can be classified in two parts:

- *Assumptions*, represented by the pre- and rely condition, describe the conditions under which the program runs, and
- *Commitments*, composed by the guarantee and postcondition, describe the expected behaviors of the program when it is run under the assumptions.

The pre- and postcondition are, like in the traditional Hoare logic, sets of states. They impose conditions upon the initial and final states of a computation, respectively. The rely and guarantee conditions describe properties of environment transitions and transitions of the program, respectively. Thus, they describe sets of pairs of states, formed by the state before and after the transition.

Jones first suggested in [Jones, 1981] that the rely and guarantee conditions be reflexive and transitive. However, for the soundness proof only the reflexivity of the guarantee condition is necessary.

4.3.1 Validity for Component Programs

Specifications of sequential programs are written with the syntax “ $P \text{ sat } [pre, rely, guar, post]$ ” where *sat* stands for “satisfies”. The type of these tuples is:

$$\mathbf{types} \ \alpha \text{ rformula} = \alpha \text{ com} \times \alpha \text{ set} \times (\alpha \times \alpha) \text{ set} \times (\alpha \times \alpha) \text{ set} \times \alpha \text{ set}$$

Informally, we say that P satisfies its specification if under the assumptions that

- 1 P is started in a state that satisfies pre , and
 - 2 any environment transition in the computation satisfies $rely$,
- then P ensures the following commitments:
- 3 any component transition satisfies $guar$, and
 - 4 if the computation terminates, the final state satisfies $post$.

The formal definitions are given by the functions:

constdefs

$$\begin{aligned}
assum &:: (\alpha \text{ set} \times (\alpha \times \alpha) \text{ set}) \Rightarrow \alpha \text{ confs set} \\
assum &\equiv \lambda(pre, rely) . \{c. snd (c!0) \in pre \wedge (\forall i. Suc i < length c \longrightarrow \\
&\quad c!i -e \longrightarrow c!Suc i \longrightarrow (snd (c!i), snd (c!Suc i)) \in rely)\}
\end{aligned}$$

$$\begin{aligned}
comm &:: ((\alpha \times \alpha) \text{ set} \times \alpha \text{ set}) \Rightarrow \alpha \text{ confs set} \\
comm &\equiv \lambda(guar, post) . \{c. (\forall i. Suc i < length c \longrightarrow \\
&\quad c!i -c \longrightarrow c!Suc i \longrightarrow (snd (c!i), snd (c!Suc i)) \in guar) \wedge \\
&\quad (fst (last c) = None \longrightarrow snd (last c) \in post)\}
\end{aligned}$$

A rely-guarantee specification of a sequential component program P is valid, and we use the usual syntax $\models P \text{ sat } [pre, rely, guar, post]$, iff for any initial state, all computations of P that satisfy the assumptions satisfy the commitments.

constdefs

$$\begin{aligned}
com\text{-}validity &:: \alpha \text{ com} \Rightarrow \alpha \text{ set} \Rightarrow (\alpha \times \alpha) \text{ set} \Rightarrow (\alpha \times \alpha) \text{ set} \Rightarrow \alpha \text{ set} \Rightarrow bool \\
&\quad (\models - \text{ sat } [-, -, -, -]) \\
\models P \text{ sat } [pre, rely, guar, post] &\equiv \\
&\quad \forall s. cp (Some P) s \cap assum (pre, rely) \subseteq comm (guar, post)
\end{aligned}$$

4.3.2 Validity for Parallel Programs

Parallel programs can be seen as a unit executed in a possibly interfering environment. For this reason, we include a rely and a guarantee condition in the specification of parallel programs as well. They have the form $P \text{ SAT } [pre, rely, guar, post]$ where P has the type $\alpha \text{ par-com}$.

A parallel program has finished when all its components are *None*. To abbreviate this we introduce the following definition:

constdefs

$$\begin{aligned}
All\text{-}None &:: \alpha \text{ com option list} \Rightarrow bool \\
All\text{-}None \text{ xs} &\equiv \forall c \in set \text{ xs}. c = None
\end{aligned}$$

The definition of assumptions, commitments and validity are analogous to the previous section:

constdefs

$$\begin{aligned} \text{par-assum} &:: (\alpha \text{ set} \times (\alpha \times \alpha) \text{ set}) \Rightarrow \alpha \text{ par-confs set} \\ \text{par-assum} &\equiv \lambda(\text{pre}, \text{rely}). \{c. \text{snd } (c!0) \in \text{pre} \wedge (\forall i. \text{Suc } i < \text{length } c \longrightarrow \\ &\quad c!i \text{ -pc} \rightarrow c!\text{Suc } i \longrightarrow (\text{snd } (c!i), \text{snd } (c!\text{Suc } i)) \in \text{rely})\} \end{aligned}$$

$$\begin{aligned} \text{par-comm} &:: ((\alpha \times \alpha) \text{ set} \times \alpha \text{ set}) \Rightarrow \alpha \text{ par-confs set} \\ \text{par-comm} &\equiv \lambda(\text{guar}, \text{post}). \{c. (\forall i. \text{Suc } i < \text{length } c \longrightarrow \\ &\quad c!i \text{ -pc} \rightarrow c!\text{Suc } i \longrightarrow (\text{snd } (c!i), \text{snd } (c!\text{Suc } i)) \in \text{guar}) \wedge \\ &\quad (\text{All-None } (\text{fst } (\text{last } c)) \longrightarrow \text{snd } (\text{last } c) \in \text{post})\} \end{aligned}$$

$$\begin{aligned} \text{par-com-validity} &:: \alpha \text{ par-com} \Rightarrow \alpha \text{ set} \Rightarrow (\alpha \times \alpha) \text{ set} \Rightarrow (\alpha \times \alpha) \text{ set} \Rightarrow \alpha \text{ set} \\ &\quad \Rightarrow \text{bool} \quad (\models \text{ - SAT } [-, -, -, -]) \\ \models P \text{ SAT } [\text{pre}, \text{rely}, \text{guar}, \text{post}] &\equiv \\ \forall s. \text{par-cp } P \text{ s} \cap \text{par-assum } (\text{pre}, \text{rely}) &\subseteq \text{par-comm } (\text{guar}, \text{post}) \end{aligned}$$

4.4 The Proof System

The system of axioms and inference rules for deriving partial correctness formulas of parallel programs in the rely-guarantee formalism can be regarded as a compositional reformulation of the Owicki-Gries system. Due to the layered definition of the syntax, the set of all derivable specifications is defined using two sets:

1. The set of all derivable specifications of sequential programs: *rghoare*.
2. The set of all derivable specifications of parallel programs: *par-rghoare*.

The definition of the second set uses the previous one. Thus, the declarations must follow the previous order.

4.4.1 Proof System for Component Programs

We first define a predicate about stability needed in the rules.

constdefs

$$\begin{aligned} \text{stable} &:: \alpha \text{ set} \Rightarrow (\alpha \times \alpha) \text{ set} \Rightarrow \text{bool} \\ \text{stable} &\equiv \lambda f g. \forall x y. x \in f \longrightarrow (x, y) \in g \longrightarrow y \in f \end{aligned}$$

For example, *stable pre rely* means that if a state belongs to the precondition and some transition satisfies the rely condition, then the reached state still belongs to the precondition.

The set of all derivable specifications is defined by the constant

consts *rghoare* :: α *rgformula* *set*

where α *rgformula* is the type of a specification of a sequential component program (see §4.3.1). A derivable specification is denoted with the usual syntax:

syntax

-rghoare :: α *com* \Rightarrow α *set* \Rightarrow $(\alpha \times \alpha)$ *set* \Rightarrow $(\alpha \times \alpha)$ *set* \Rightarrow α *set* \Rightarrow *bool*
 $(\vdash$ *- sat* [*-*, *-*, *-*, *-*])

translations

$\vdash P$ *sat* [*pre*, *rely*, *guar*, *post*] $\Leftrightarrow (P, pre, rely, guar, post) \in$ *rghoare*

We follow [Xu *et al.*, 1997] in the definition of the rules, but differ mainly in the representation of the variables. In [Xu *et al.*, 1997], rules are expressed in terms of a variable y representing the vector of program variables and the corresponding primed variable y' referring to the same vector after a transformation. In our formalization we describe properties of states or of pairs of states by directly describing the set of tuples of values, i.e. we do not refer to program variables. For example, the set of pairs of states representing the identity transformation is $\{(s, t). s = t\}$, whereas in [Xu *et al.*, 1997] it would be $y = y'$.

inductive *rghoare*

intros

Basic: $\llbracket pre \subseteq \{s. f s \in post\}; \{(s, t). s \in pre \wedge (t = f s \vee t = s)\} \subseteq guar;$
 $stable\ pre\ rely; stable\ post\ rely \rrbracket$
 $\Rightarrow \vdash Basic\ f\ sat\ [pre, rely, guar, post]$

In the computation of a *Basic* command there is exactly one component transition that updates the state. Before and after this component transition there can be a number of environment transitions. The initial state satisfies *pre*, thus from *stable pre rely* it follows that *pre* holds immediately before the component transition takes place. From $pre \subseteq \{s. f s \in post\}$ it follows that *post* holds immediately after the component transition, and because *post* is stable when *rely* holds, *post* holds after any number of environment transitions.

The rules for the sequential composition and conditional statements are standard:

$$\begin{aligned} \text{Seq: } & \llbracket \vdash P \text{ sat } [pre, rely, guar, mid]; \vdash Q \text{ sat } [mid, rely, guar, post] \rrbracket \\ & \implies \vdash \text{Seq } P \ Q \text{ sat } [pre, rely, guar, post] \end{aligned}$$

$$\begin{aligned} \text{Cond: } & \llbracket \text{stable } pre \text{ rely}; \vdash P_1 \text{ sat } [pre \cap b, rely, guar, post]; \\ & \vdash P_2 \text{ sat } [pre \cap \neg b, rely, guar, post]; \forall s. (s, s) \in guar \rrbracket \\ & \implies \vdash \text{Cond } b \ P_1 \ P_2 \text{ sat } [pre, rely, guar, post] \end{aligned}$$

In the while-rule the precondition plays the role of the invariant; it must hold before and after execution of the body at every iteration:

$$\begin{aligned} \text{While: } & \llbracket \text{stable } pre \text{ rely}; pre \cap \neg b \subseteq post; \text{stable } post \text{ rely}; \\ & \vdash P \text{ sat } [pre \cap b, rely, guar, pre]; \forall s. (s, s) \in guar \rrbracket \\ & \implies \vdash \text{While } b \ P \text{ sat } [pre, rely, guar, post] \end{aligned}$$

The rule for the await-statement is less obvious:

$$\begin{aligned} \text{Await: } & \llbracket \forall V. \vdash P \text{ sat } [pre \cap b \cap \{V\}, \{(s, t). s = t\}, UNIV, \\ & \{s. (V, s) \in guar\} \cap post]; \text{stable } pre \text{ rely}; \text{stable } post \text{ rely} \rrbracket \\ & \implies \vdash \text{Await } b \ P \text{ sat } [pre, rely, guar, post] \end{aligned}$$

By the semantics of the await-command, a positive evaluation of the condition and the execution of the body is done atomically. Thus, the state transition caused by the complete execution of P must satisfy the guarantee condition. This is reflected in the precondition and postcondition of P in the assumptions; since these are sets of single states, the relation between the state before and after the transformation is established by fixing the values of the first via a universally quantified variable V . The intermediate state changes during the execution of P must not guarantee anything, thus the guarantee condition is the universal set $UNIV$. However, since they are executed atomically, the environment cannot change their values. This is reflected by the rely condition $\{(s, t). s = t\}$. To ensure that the postcondition holds at the end of the computation, regardless of possible environment transitions, we require *stable post rely*.

Finally, the rule of consequence allows to strengthen the assumptions and weaken the commitments:

$$\begin{aligned} \text{Conseq: } & \llbracket pre \subseteq pre'; rely \subseteq rely'; guar' \subseteq guar; post' \subseteq post; \\ & \vdash P \text{ sat } [pre', rely', guar', post'] \rrbracket \\ & \implies \vdash P \text{ sat } [pre, rely, guar, post] \end{aligned}$$

These six rules inductively define the set of derivable specifications of sequential component programs. In §4.5 we prove that only valid specifications can be derived, i.e. we prove the soundness of the system.

The functions defined below extract the parts of a specification of a sequential parallel program and will be used in the following section:

constdefs

```

Pre ::  $\alpha$  rgformula  $\Rightarrow$   $\alpha$  set
Pre x  $\equiv$  fst (snd x)
Post ::  $\alpha$  rgformula  $\Rightarrow$   $\alpha$  set
Post x  $\equiv$  snd (snd (snd (snd x)))
Rely ::  $\alpha$  rgformula  $\Rightarrow$  ( $\alpha \times \alpha$ ) set
Rely x  $\equiv$  fst (snd (snd x))
Guar ::  $\alpha$  rgformula  $\Rightarrow$  ( $\alpha \times \alpha$ ) set
Guar x  $\equiv$  fst (snd (snd (snd x)))
Com ::  $\alpha$  rgformula  $\Rightarrow$   $\alpha$  com
Com x  $\equiv$  fst x

```

4.4.2 Proof System for Parallel Programs

This section presents the rule for deriving parallel programs whose components are sequential. Observe that in the definition of validity for parallel programs (see §4.3.2) no information about the pre, post, rely and guarantee conditions of the component programs was included. This was not important for the definition of validity, however, at the level of concrete verification of programs with the system of rules, we want to apply the rules backwards. Therefore, the conclusion should include all the information needed in the premises of the rule. For this reason, we include it as part of the elements of the set of derivable formulas. Their type is

```

types  $\alpha$  par-rgformula =  $\alpha$  rgformula list  $\times$   $\alpha$  set  $\times$  ( $\alpha \times \alpha$ ) set  $\times$  ( $\alpha \times \alpha$ ) set
 $\times$   $\alpha$  set

```

The type α rgformula corresponds to a full specification of a component program (see §4.3.1). The constant defining the corresponding set of derivations, called *par-rghoare*, and a familiar syntax for membership of an element are shown below.

```

consts par-rghoare ::  $\alpha$  par-rgformula set

```

syntax

```

-par-rghoare::  $\alpha$  rgformula list  $\Rightarrow$   $\alpha$  set  $\Rightarrow$  ( $\alpha \times \alpha$ ) set  $\Rightarrow$  ( $\alpha \times \alpha$ ) set  $\Rightarrow$   $\alpha$  set
 $\Rightarrow$  bool (| - SAT [-, -, -, -] )

```

translations

$$\vdash Ps \text{ SAT } [pre, rely, guar, post] \Leftrightarrow (Ps, pre, rely, guar, post) \in \text{par-rghoare}$$

Do not confuse the type of the first argument. Here it is a list of specifications of sequential component programs, while in $\models P \text{ SAT } [pre, rely, guar, post]$, P is just a program of type $\alpha \text{ par-com}$.

The rule for parallel composition is new in the sense that it generalizes the case of composing two programs, as given in [Xu *et al.*, 1997, de Roever *et al.*, 2000], to the generic case of composing any number of programs. This rule allows the verification of parameterized parallel programs directly in the system.

inductive *par-rghoare*

intros

Parallel:

$$\begin{aligned} & \llbracket \forall i < \text{length } Ps. \text{rely} \cup (\bigcup_{j \in \{j. j < \text{length } Ps \wedge j \neq i\}}. \text{Guar } (Ps!j)) \subseteq \text{Rely } (Ps!i); \\ & (\bigcup_{j \in \{j. j < \text{length } Ps\}}. \text{Guar } (Ps!j)) \subseteq \text{guar}; \\ & pre \subseteq (\bigcap_{i \in \{i. i < \text{length } Ps\}}. \text{Pre } (Ps!i)); \\ & (\bigcap_{i \in \{i. i < \text{length } Ps\}}. \text{Post } (Ps!i)) \subseteq \text{post}; \\ & \forall i < \text{length } Ps. \vdash \text{Com}(Ps!i) \text{ sat } [\text{Pre}(Ps!i), \text{Rely}(Ps!i), \text{Guar}(Ps!i), \text{Post}(Ps!i)] \\ & \implies \vdash Ps \text{ SAT } [pre, rely, guar, post] \end{aligned}$$

An environment transition for the component specified by $P_s!i$ consists of a component transition from any of the other processes $P_s!j$ where $i \neq j$, or of a transition from the overall environment. Hence, the strongest rely condition that component i can assume is $\text{rely} \cup (\bigcup_{j \in \{j. j < \text{length } Ps \wedge j \neq i\}}. \text{Guar } (Ps!j))$.

A component transition of the parallel program is performed by one of its components, hence they all have to satisfy the overall guarantee condition *guar*. Like in the Owicki-Gries system the precondition for the parallel composition must imply all the component's preconditions and the overall postcondition must be a logical consequence of all postconditions.

Finally, the specifications of the components have to be derivable in their system. Consequently, the soundness of the system of rules for sequential component programs is a necessary previous result for the proof of soundness of this rule.

4.5 Soundness

In this section we prove soundness of the rule for parallel composition. For this result the proof of soundness of the system for sequential programs is required.

As explained in §2.5, soundness of a system of rules that inductively define a set of correctness formulas can be shown by rule induction. This proof principle works by proving that a certain property of the elements of the set is true of all axioms and is preserved by each inference rule. Then by construction of the set, we obtain the property for any element in the set. When the property concerned is validity of a correctness formula, then any derivable formula is correct and we say that the system of rules is sound.

Using rule induction soundness of the system amounts to proving soundness of each rule. Thus, for each rule we assume that the formulas that appear in the premises are valid and we must prove that the formula in the conclusion is also valid.

4.5.1 Soundness of the System for Component Programs

We want to prove the theorem:

theorem *rgsound*:

$$\vdash P \text{ sat } [pre, rely, guar, post] \implies \models P \text{ sat } [pre, rely, guar, post]$$

The proof proceeds by rule induction. This results in six subgoals, one for each rule.

Soundness of the Basic Rule

The rule for a *Basic* command is an axiom. Hence, validity must follow directly from the premises of the rule without any induction hypothesis.

lemma *Basic-sound*:

$$\llbracket pre \subseteq \{s. f s \in post\}; \{(s, t). s \in pre \wedge (t = f s \vee t = s)\} \subseteq guar; \\ stable\ pre\ rely; stable\ post\ rely \rrbracket \implies \models Basic\ f\ sat\ [pre, rely, guar, post]$$

For this and some of the following proofs we need a lemma about the stability predicate.

Assume a computation whose environment transitions satisfy the rely condition and where all transitions of the subcomputation between two indices j and k are made by the environment. If the state of the configuration

at index j satisfies some condition p such that *stable p rely*, then the state at configuration k also satisfies the condition p and the program part is left unchanged:

lemma stability:

$$\begin{aligned} & \llbracket x \in \text{cptn}; \text{stable } p \text{ rely}; j \leq k; k < \text{length } x; \text{snd } (x!j) \in p; \\ & \forall i. \text{Suc } i < \text{length } x \longrightarrow x!i -e \rightarrow x! \text{Suc } i \longrightarrow (\text{snd } (x!i), \text{snd } (x! \text{Suc } i)) \in \text{rely}; \\ & \forall i. j \leq i \wedge i < k \longrightarrow x!i -e \rightarrow x! \text{Suc } i \rrbracket \\ & \implies \text{snd } (x!k) \in p \wedge \text{fst } (x!j) = \text{fst } (x!k) \end{aligned}$$

The soundness of the Basic rule is easy to prove with help of *stability* and two more lemmas. The first one states that if there is a component transition in the computation of a *Basic*-command, then it is the only one:

lemma unique-ctran-Basic:

$$\begin{aligned} & \llbracket x \in \text{cptn}; x!0 = (\text{Some } (\text{Basic } f), s); \text{Suc } i < \text{length } x; \\ & x!i -c \rightarrow x! \text{Suc } i; \text{Suc } j < \text{length } x; i \neq j \rrbracket \implies x!j -e \rightarrow x! \text{Suc } j \end{aligned}$$

The second one ensures that if the empty program appears in a computation of a *Basic*-command at some point, then there must be a component transition before:

lemma exists-ctran-Basic-None:

$$\begin{aligned} & \llbracket x \in \text{cptn}; x!0 = (\text{Some } (\text{Basic } f), s); i < \text{length } x; \text{fst } (x!i) = \text{None} \rrbracket \\ & \implies \exists j. j < i \wedge x!j -c \rightarrow x! \text{Suc } j \end{aligned}$$

Both lemmas are proven by induction on the length of the computation.

Soundness of the Await Rule

The induction hypothesis is applied to the derivability of the await-body, thus we can assume that its specification is valid:

lemma Await-sound:

$$\begin{aligned} & \llbracket \text{stable } \text{pre } \text{rely}; \text{stable } \text{post } \text{rely}; \\ & \forall V. \models P \text{ sat } [\text{pre} \cap b \cap \{s. s = V\}, \{(s, t). s = t\}, \text{UNIV}, \\ & \quad \{s. (V, s) \in \text{guar}\} \cap \text{post}] \rrbracket \\ & \implies \models \text{Await } b \text{ } P \text{ sat } [\text{pre}, \text{rely}, \text{guar}, \text{post}] \end{aligned}$$

The proof is similar to the previous one and requires the analogous lemmas:

lemma *unique-ctran-Await*:

$$\llbracket x \in \text{cptn}; x!0 = (\text{Some } (\text{Await } b \ c), \ s); \text{Suc } i < \text{length } x; \\ x!i \text{ -c} \rightarrow x \ ! \ \text{Suc } i; \text{Suc } j < \text{length } x; i \neq j \rrbracket \Longrightarrow x!j \text{ -e} \rightarrow x \ ! \ \text{Suc } j$$

lemma *exists-ctran-Await-None*:

$$\llbracket x \in \text{cptn}; x!0 = (\text{Some } (\text{Await } b \ c), \ s); i < \text{length } x; \text{fst } (x!i) = \text{None} \rrbracket \\ \Longrightarrow \exists j. j < i \wedge x!j \text{ -c} \rightarrow x \ ! \ \text{Suc } j$$

We also need to prove that there is a computation of the body of the await-statement that satisfies the specification given in the premises. However, the only information we obtain from the semantics about the execution of the body is that it terminates in some number of component transitions. We prove that a sequence of component transitions can also be described as a computation:

lemma *Star-imp-cptn*: $(P, \ s) \text{ -c}^* \rightarrow (R, \ t) \Longrightarrow \exists l \in \text{cp } P \ s. \text{last } l = (R, \ t)$

Soundness of the Conditional Rule

Given valid subspecifications of the if- and else-branches we prove that the correctness formula for the conditional statement is also valid:

lemma *Cond-sound*:

$$\llbracket \text{stable } \text{pre } \text{rely}; \models P_1 \ \text{sat } [\text{pre} \cap b, \ \text{rely}, \ \text{guar}, \ \text{post}]; \\ \models P_2 \ \text{sat } [\text{pre} \cap \neg b, \ \text{rely}, \ \text{guar}, \ \text{post}]; \forall s. (s, \ s) \in \text{guar} \rrbracket \\ \Longrightarrow \models (\text{Cond } b \ P_1 \ P_2) \ \text{sat } [\text{pre}, \ \text{rely}, \ \text{guar}, \ \text{post}]$$

In the proof we first distinguish whether a computation of $(\text{Cond } b \ P_1 \ P_2)$ contains a component transition or not. If not, by the following lemma all transitions are performed by the environment:

lemma *etran-or-ctran*:

$$\llbracket x \in \text{cptn}; m \leq \text{length } x; \forall i. \text{Suc } i < m \longrightarrow \neg x!i \text{ -c} \rightarrow x \ ! \ \text{Suc } i; \text{Suc } i < m \rrbracket \\ \Longrightarrow x!i \text{ -e} \rightarrow x \ ! \ \text{Suc } i$$

Thus the first part of the commitments is trivially fulfilled. By using the *stability* lemma the last program of the computation cannot be *None*, so the second part of the commitments holds too.

If there is a component transition, then by the lemma

lemma *Ex-first-occurrence*: $P \ n \Longrightarrow \exists m. P \ m \wedge (\forall i < m. \neg P \ i)$

there is a first one. This component transition satisfies the guarantee condition because of the required reflexivity property. By the stability lemma, the precondition holds after this step. Depending on whether the boolean condition of the if-statement holds or not, the fulfillment of the commitments for the rest of the computation follows from the induction hypothesis on the corresponding program.

Soundness of the Sequential Rule

Validity of a specification for sequential composition follows from the validity of appropriate subspecifications of the programs that are sequentially composed:

lemma *Seq-sound*:

$$\begin{aligned} & \llbracket \models P \text{ sat } [pre, rely, guar, mid]; \models Q \text{ sat } [mid, rely, guar, post] \rrbracket \\ & \implies \models Seq P Q \text{ sat } [pre, rely, guar, post] \end{aligned}$$

In the proof we distinguish whether a computation of the sequential composition finishes computing P or not:

1. If not, we have $\forall i < length\ x. fst\ (x!i) \neq Some\ Q$. In this case, we can find a computation of P such that c is just the corresponding “lifted” computation.
2. Otherwise, we have $\exists i < length\ x. fst\ (x!i) = Some\ Q$. Such a configuration can occur several times but we are only interested in the first occurrence. The computation c in this case can be split into a “lifted” terminated computation of P followed by a computation of Q .

The following lemmas establish these properties.

lemma *Seq-sound1*:

$$\begin{aligned} & \llbracket x \in \text{cptn-mod}; x!0 = (Some\ (Seq\ P\ Q),\ s); \forall i < length\ x. fst\ (x!i) \neq Some\ Q \rrbracket \\ & \implies \exists xs \in \text{cp}\ (Some\ P)\ s. x = \text{map}\ (\text{lift}\ Q)\ xs \end{aligned}$$

We prove it by rule induction on *cptn-mod*. The induction hypothesis is only needed at the *CptnModEnv* rule.

lemma *Seq-sound2*:

$$\begin{aligned} & \llbracket x \in \text{cptn}; x!0 = (Some\ (Seq\ P\ Q),\ s); i < length\ x; fst\ (x!i) = Some\ Q; \\ & \forall j < i. fst\ (x!j) \neq Some\ Q \rrbracket \end{aligned}$$

$$\begin{aligned} \implies & \exists xs\ ys. xs \in cp\ (Some\ P)\ s \wedge length\ xs = Suc\ i \wedge \\ & ys \in cp\ (Some\ Q)\ (snd\ (xs\ !i)) \wedge x = (map\ (lift\ Q)\ xs)@tl\ ys \end{aligned}$$

This second lemma is easier to prove by rule induction on *cptn*. When the computation results from the rule *CptnComp*, we perform case analysis on the *c*-step. After simplifying, only the cases corresponding to *Seq1* and *Seq2* remain. The first one is solved easily without need of the induction hypothesis. The proof for the second one is also straightforward because the induction hypothesis can be directly used.

Back to the main lemma, suppose that a computation of *Seq P Q*, say *c*, satisfies the assumptions *assum (pre, rely)*. If the computation does not finish computing *P*, then it does not terminate, so we only have to prove that all component transitions satisfy *guar*. This follows from $\models P\ sat\ [pre, rely, guar, mid]$ and the lemma *Seq-sound1*. In the second case, we first prove the same thing for the first part of the computation. Then, since it terminates we obtain from $\models P\ sat\ [pre, rely, guar, mid]$ that the last state satisfies *mid*. From $\models Q\ sat\ [mid, rely, guar, post]$ the rest of the computation also satisfies the commitments.

Soundness of the While Rule

The subsequent proof is the most interesting one so far.

lemma *While-sound*:

$$\begin{aligned} & \llbracket\ stable\ pre\ rely; pre \cap \neg b \subseteq post; stable\ post\ rely; \\ & \quad \models P\ sat\ [pre \cap b, rely, guar, pre]; \forall s. (s, s) \in guar \rrbracket \\ & \implies \models While\ b\ P\ sat\ [pre, rely, guar, post] \end{aligned}$$

First attempts at proving this lemma using the “flat” definition of computation turn out to be unnecessarily cumbersome. Finding a suitable definition of computation together with the proof of equivalence (§4.2.4) and finally the soundness proof of the rule took about a month of work. This seems unnecessary considering that the proof on paper from [Xu *et al.*, 1995] is barely a page and very intuitive. However, some intuitive ideas like the recursive structure of a while-computation are easy to state on paper but turn out to be difficult to formalize in a theorem prover if the definitions are inadequate.

The proof of soundness directly follows from a lemma which can be proved by rule induction on *cptn-mod*:

lemma *While-sound-aux*:

$$\begin{aligned} & \llbracket \text{stable pre rely}; \text{pre} \cap - b \subseteq \text{post}; \text{stable post rely}; \\ & \quad \models P \text{ sat } [\text{pre} \cap b, \text{rely}, \text{guar}, \text{pre}]; \forall s. (s, s) \in \text{guar}; x \in \text{cptn-mod} \rrbracket \\ & \implies \forall s \text{ xs}. x = (\text{Some } (\text{While } b P), s) \# \text{xs} \longrightarrow x \in \text{assum } (\text{pre}, \text{rely}) \\ & \quad \longrightarrow x \in \text{comm } (\text{guar}, \text{post}) \end{aligned}$$

After some simplification four subgoals remain. We briefly explain their proofs.

The first one corresponds to the *CptnModEnv* rule. After an environment step the program fragment is still a while-program so the induction hypothesis can be applied:

lemma *WhileEnv*:

$$\begin{aligned} & \llbracket \text{stable pre rely}; \text{pre} \cap - b \subseteq \text{post}; \text{stable post rely}; \forall s. (s, s) \in \text{guar}; \\ & \quad \models P \text{ sat } [\text{pre} \cap b, \text{rely}, \text{guar}, \text{pre}]; (\text{Some } (\text{While } b P), t) \# \text{xs} \in \text{cptn-mod}; \\ & \quad (\text{Some } (\text{While } b P), t) \# \text{xs} \in \text{assum } (\text{pre}, \text{rely}) \longrightarrow \\ & \quad (\text{Some } (\text{While } b P), t) \# \text{xs} \in \text{comm } (\text{guar}, \text{post}); \\ & \quad (\text{Some } (\text{While } b P), s) \# (\text{Some } (\text{While } b P), t) \# \text{xs} \in \text{assum } (\text{pre}, \text{rely}) \rrbracket \\ & \implies (\text{Some } (\text{While } b P), s) \# (\text{Some } (\text{While } b P), t) \# \text{xs} \in \text{comm } (\text{guar}, \text{post}) \end{aligned}$$

The proof follows directly using the following two lemmas:

$$\begin{aligned} \text{lemma } \text{etran-in-comm}: & (P, t) \# \text{xs} \in \text{comm } (\text{guar}, \text{post}) \\ & \implies (P, s) \# (P, t) \# \text{xs} \in \text{comm } (\text{guar}, \text{post}) \end{aligned}$$

lemma *tl-of-assum-in-assum*:

$$\begin{aligned} & \llbracket (P, s) \# (P, t) \# \text{xs} \in \text{assum } (\text{pre}, \text{rely}); \text{stable pre rely} \rrbracket \\ & \implies (P, t) \# \text{xs} \in \text{assum } (\text{pre}, \text{rely}) \end{aligned}$$

The second subgoal corresponds to the *CptnModNone* rule, i.e. the while-program terminates:

lemma *WhileNone*:

$$\begin{aligned} & \llbracket \text{stable pre rely}; \text{pre} \cap - b \subseteq \text{post}; \text{stable post rely}; \forall s. (s, s) \in \text{guar}; \\ & \quad \models P \text{ sat } [\text{pre} \cap b, \text{rely}, \text{guar}, \text{pre}]; (\text{Some } (\text{While } b P), s) -c\rightarrow (\text{None}, t); \\ & \quad (\text{None}, t) \# \text{xs} \in \text{cptn-mod}; \\ & \quad (\text{Some } (\text{While } b P), s) \# (\text{None}, t) \# \text{xs} \in \text{assum } (\text{pre}, \text{rely}) \rrbracket \\ & \implies (\text{Some } (\text{While } b P), s) \# (\text{None}, t) \# \text{xs} \in \text{comm } (\text{guar}, \text{post}) \end{aligned}$$

By rule inversion on the *c*-transition we obtain $s \notin b$ and $s = t$. The first transition of the computation in the conclusion satisfies the guarantee

condition by the reflexivity assumption. The rest of the transitions are environmental. By the assumption $pre \cap - b \subseteq post$ the state s satisfies the postcondition, since $stable\ post\ rely$ it follows by the *stability* lemma that the final state satisfies the postcondition.

The third subgoal corresponds to the rule *CptnModWhile1*. There are not subcomputations of while-programs in the premises, thus the induction hypothesis cannot be used:

lemma *While1*:

$$\begin{aligned} & \llbracket stable\ pre\ rely; pre \cap - b \subseteq post; stable\ post\ rely; \forall s. (s, s) \in guar; \\ & \models P\ sat [pre \cap b, rely, guar, pre]; (Some\ P, s) \# xs \in cptn-mod; s \in b; \\ & (Some\ (While\ b\ P), s) \# (Some\ (Seq\ P\ (While\ b\ P)), s) \# \\ & \quad map\ (lift\ (While\ b\ P))\ xs \in assum\ (pre, rely) \rrbracket \\ & \implies (Some\ (While\ b\ P), s) \# (Some\ (Seq\ P\ (While\ b\ P)), s) \# \\ & \quad map\ (lift\ (While\ b\ P))\ xs \in comm\ (guar, post) \end{aligned}$$

This kind of computation does not terminate, thus we only have to prove that all component transitions satisfy the guarantee condition. The first component transition is easy due to the reflexivity property. The rest of the computation can be reduced to a computation of P where the initial state satisfies $pre \cap b$, thus the proof follows from $\models P\ sat [pre \cap b, rely, guar, pre]$.

The interesting case is the fourth subgoal. The computation contains a full computation of the body followed by a computation of the same while-command, where the induction hypothesis is applied:

lemma *While2*:

$$\begin{aligned} & \llbracket stable\ pre\ rely; pre \cap - b \subseteq post; stable\ post\ rely; \forall s. (s, s) \in guar; \\ & \models P\ sat [pre \cap b, rely, guar, pre]; fst\ (last\ ((Some\ P, s) \# xs)) = None; \\ & (Some\ P, s) \# xs \in cptn-mod; s \in b; \\ & (Some\ (While\ b\ P), snd\ (last\ ((Some\ P, s) \# xs))) \# ys \in cptn-mod; \\ & (Some\ (While\ b\ P), snd\ (last\ ((Some\ P, s) \# xs))) \# ys \in assum\ (pre, rely) \\ & \longrightarrow (Some\ (While\ b\ P), snd\ (last\ ((Some\ P, s) \# xs))) \# ys \\ & \quad \in comm\ (guar, post); \\ & (Some\ (While\ b\ P), s) \# (Some\ (Seq\ P\ (While\ b\ P)), s) \# \\ & \quad map\ (lift\ (While\ b\ P))\ xs @ ys \in assum\ (pre, rely) \rrbracket \\ & \implies (Some\ (While\ b\ P), s) \# (Some\ (Seq\ P\ (While\ b\ P)), s) \# \\ & \quad map\ (lift\ (While\ b\ P))\ xs @ ys \in comm\ (guar, post) \end{aligned}$$

The first part of the computation, i.e. $(Some\ (While\ b\ P), s) \# (Some\ (Seq\ P\ (While\ b\ P)), s) \# map\ (lift\ (While\ b\ P))\ xs$ represents the first entire

execution of the body. Like for the proof of the last subgoal it follows that the commitments are satisfied by this subcomputation. For the rest of the computation it suffices to apply the induction hypothesis. Thus, we need to prove that it satisfies the assumptions:

lemma *assum-after-body*:

$$\begin{aligned} & \llbracket \models P \text{ sat } [pre \cap b, \text{rely}, \text{guar}, \text{pre}]; (Some\ P, s) \# xs \in \text{cptn-mod}; \\ & \quad \text{fst } (\text{last } ((Some\ P, s) \# xs)) = \text{None}; s \in b; \\ & \quad (Some\ (\text{While } b\ P), s) \# (Some\ (\text{Seq } P\ (\text{While } b\ P)), s) \# \\ & \quad \quad \text{map } (\text{lift } (\text{While } b\ P))\ xs\ @\ ys \in \text{assum } (pre, \text{rely}) \rrbracket \\ & \implies (Some\ (\text{While } b\ P), \text{snd } (\text{last } ((Some\ P, s) \# xs))) \# ys \\ & \quad \in \text{assum } (pre, \text{rely}) \end{aligned}$$

From $\llbracket \models P \text{ sat } [pre \cap b, \text{rely}, \text{guar}, \text{pre}]$ the precondition (which plays the role of the loop-invariant) holds after the full execution of the body. It is also easy to prove that if all environment transitions in a computation satisfy the rely condition, so do those of a subcomputation.

Soundness of the Rule of Consequence

The proof of the soundness of the consequence rule is trivial.

lemma *Conseq-sound*:

$$\begin{aligned} & \llbracket pre \subseteq pre'; \text{rely} \subseteq \text{rely}'; \text{guar}' \subseteq \text{guar}; \text{post}' \subseteq \text{post}; \\ & \quad \llbracket \models P \text{ sat } [pre', \text{rely}', \text{guar}', \text{post}'] \rrbracket \implies \llbracket \models P \text{ sat } [pre, \text{rely}, \text{guar}, \text{post}] \rrbracket \end{aligned}$$

The soundness of the system follows from the soundness of each rule. This concludes the proof of the theorem *rgsound*.

The next step is to prove soundness of the system for deriving correct parallel programs. This result relies on an important property of the semantics which we show in the following section.

4.5.2 Compositionality of the Semantics

The most important virtue of the semantics presented in §4.2, where we distinguish between component and environment transitions, is that it allows us to define computations of parallel programs in terms of the computations of the components. In this sense, we say that the semantics is compositional.

A computation c of a parallel program can be described in terms of a list of computations of component programs $clist$ if they *conjoin*, and we write it $c \propto clist$.

Before giving the formal definition of the conjoin-operator, we explain its intuitive meaning by means of a parallel program consisting of two components P and Q that run in parallel with an overall environment R . Then, P 's environment consists of Q and R . Analogously, Q 's environment consists of P and R .

If P and Q are executed in parallel, their respective computations should have the same sequence of states:

$$\begin{aligned} (P_0, \sigma_0) &\xrightarrow{\delta_1} (P_1, \sigma_1) \xrightarrow{\delta_2} \dots \xrightarrow{\delta_n} (P_n, \sigma_n) \xrightarrow{\delta_{n+1}} \dots, & \delta_i \in \{e, c\} \\ (Q_0, \sigma_0) &\xrightarrow{\delta'_1} (Q_1, \sigma_1) \xrightarrow{\delta'_2} \dots \xrightarrow{\delta'_n} (Q_n, \sigma_n) \xrightarrow{\delta'_{n+1}} \dots, & \delta'_i \in \{e, c\} \end{aligned}$$

All components of a parallel composition have computations of the same length. If one component terminates before the others, its computation is extended by environment transitions, which are also allowed when the program has terminated.

Another requirement for separate computations of components to make part of a parallel composition is to have compatible simultaneous transitions. This means that they do not have component transitions at the same time, i.e. δ_i and δ'_i cannot be both c .

Moreover, when some component performs a component transition, then the transition of the full parallel composition is also a component transition. The parallel composition executes an environment step only when all components simultaneously perform an environment step.

For example, consider the above transitions with the following labels:

$$\begin{aligned} (P_0, \sigma_0) &\xrightarrow{c} (P_1, \sigma_1) \xrightarrow{c} \dots \xrightarrow{e} (P_n, \sigma_n) \xrightarrow{e} \dots \\ (Q_0, \sigma_0) &\xrightarrow{e} (Q_1, \sigma_1) \xrightarrow{e} \dots \xrightarrow{e} (Q_n, \sigma_n) \xrightarrow{e} \dots \end{aligned}$$

Then, both computations could be composed in the following computation of the parallel program formed by P and Q (denoted $P \parallel Q$):

$$(P_0 \parallel Q_0, \sigma_0) \xrightarrow{c} (P_1 \parallel Q_1, \sigma_1) \xrightarrow{c} \dots \xrightarrow{e} (P_n \parallel Q_n, \sigma_n) \xrightarrow{e} \dots$$

The formal definitions of the properties involved for a list of computations to conjoin with the computation of a parallel composition are listed below.

constdefs

$same_length :: \alpha \text{ par-confs} \Rightarrow \alpha \text{ confs list} \Rightarrow bool$

$same_length \ c \ clist \equiv \forall i < length \ clist. \ length \ (clist!i) = length \ c$

All computations have the same length and the same state sequence.

$$\begin{aligned} \text{same-state} &:: \alpha \text{ par-confs} \Rightarrow \alpha \text{ confs list} \Rightarrow \text{bool} \\ \text{same-state } c \text{ clist} &\equiv \forall i < \text{length clist}. \forall j < \text{length } c. \text{snd } (c!j) = \text{snd } ((\text{clist}!i)!j) \end{aligned}$$

The parallel program must at each stage of the computation be formed by combining the program fragments of *clist*:

$$\begin{aligned} \text{same-program} &:: \alpha \text{ par-confs} \Rightarrow \alpha \text{ confs list} \Rightarrow \text{bool} \\ \text{same-program } c \text{ clist} &\equiv \forall j < \text{length } c. \text{fst } (c!j) = \text{map } (\lambda x. \text{fst } (x!j)) \text{ clist} \end{aligned}$$

And finally, the labels must be compatible. A transition is labelled as *pc* in the parallel computation if one of the transitions in *clist* at the corresponding position is a *c*-transition, and *pe* if all transitions in *clist* are also made by the environment.

$$\begin{aligned} \text{compat-label} &:: \alpha \text{ par-confs} \Rightarrow \alpha \text{ confs list} \Rightarrow \text{bool} \\ \text{compat-label } c \text{ clist} &\equiv \forall j. \text{Suc } j < \text{length } c \longrightarrow \\ &(c!j \text{ -pc} \rightarrow c!\text{Suc } j \wedge (\exists i < \text{length clist}. (\text{clist}!i)!j \text{ -c} \rightarrow (\text{clist}!i)! \text{Suc } j \wedge \\ &\quad (\forall l < \text{length clist}. l \neq i \longrightarrow (\text{clist}!l)!j \text{ -e} \rightarrow (\text{clist}!l)! \text{Suc } j))) \vee \\ &(c!j \text{ -pe} \rightarrow c!\text{Suc } j \wedge (\forall i < \text{length clist}. (\text{clist}!i)!j \text{ -e} \rightarrow (\text{clist}!i)! \text{Suc } j)) \end{aligned}$$

A parallel program conjoins with a list of components if the four properties hold:

$$\begin{aligned} \text{conjoin} &:: \alpha \text{ par-confs} \Rightarrow \alpha \text{ confs list} \Rightarrow \text{bool} \quad (- \times -) \\ c \times \text{clist} &\equiv \text{same-length } c \text{ clist} \wedge \text{same-state } c \text{ clist} \wedge \\ &\quad \text{same-program } c \text{ clist} \wedge \text{compat-label } c \text{ clist} \end{aligned}$$

We now prove a lemma stating that the set of computations of a (non-empty) parallel program consists of the computations that conjoin with some list of computations of the components.

theorem one: $xs \neq [] \implies$

$$\begin{aligned} \text{par-cp } xs \text{ } s &= \{c. \exists \text{clist}. \text{length clist} = \text{length } xs \wedge c \times \text{clist} \\ &\quad \wedge (\forall i < \text{length clist}. (\text{clist}!i) \in \text{cp } (xs!i) \text{ } s)\} \end{aligned}$$

Hence, the computation of a parallel program can be described in terms of the computations of its components, revealing the compositionality of the semantics¹.

¹The “numbered” names of some lemmas follow the terminology in [Xu *et al.*, 1997].

Proof. The if-implication

lemma one-if:

$$\begin{aligned} & \llbracket \text{length } clist = \text{length } xs; \forall i < \text{length } clist. (clist!i) \in cp (xs!i) s; c \propto clist \rrbracket \\ & \implies c \in \text{par-cp } xs \ s \end{aligned}$$

is proved by means of the following auxiliary (and equivalent) lemma:

lemma aux-if:

$$\begin{aligned} & \llbracket \text{length } clist = \text{length } xs; \forall i < \text{length } xs. (xs!i, s) \# clist!i \in \text{cptn}; \\ & \quad (xs, s) \# ys \propto \text{map } (\lambda i. (fst i, s) \# snd i) (\text{zip } xs \ clist) \rrbracket \\ & \implies (xs, s) \# ys \in \text{par-cptn} \end{aligned}$$

The proof of *aux-if* proceeds by induction on the list *ys*. The base case is solved by the *ParCptnOne* rule. The induction step is reduced by working out the first step distinguishing whether it is a step made by the environment or by the component. Then we can use the induction hypothesis on the rest of the list.

The only-if-direction is analogously proven by means of an auxiliary lemma:

lemma aux-onlyif:

$$\begin{aligned} (xs, s) \# ys \in \text{par-cptn} & \implies \exists clist. \text{length } clist = \text{length } xs \wedge \\ & (xs, s) \# ys \propto \text{map } (\lambda i. (fst i, s) \# snd i) (\text{zip } xs \ clist) \wedge \\ & (\forall i < \text{length } xs. (xs!i, s) \# (clist!i) \in \text{cptn}) \end{aligned}$$

The proof is by induction on *ys*. The base case is solved by instantiating *clist* with the empty list. The proof of the inductive step follows by case analysis on the inductive definition of *par-cptn*, and instantiating *clist* with the appropriate list in each case.

Finally, the equivalence lemma

lemma one-iff-aux: $xs \neq [] \implies$

$$\begin{aligned} & (\forall ys. (xs, s) \# ys \in \text{par-cptn} = \\ & \quad (\exists clist. \text{length } clist = \text{length } xs \wedge \\ & \quad (xs, s) \# ys \propto \text{map } (\lambda i. (fst i, s) \# snd i) (\text{zip } xs \ clist) \wedge \\ & \quad (\forall i < \text{length } xs. (xs!i, s) \# clist!i \in \text{cptn}))) = \\ & (\text{par-cp } (xs) \ s = \\ & \quad \{c. \exists clist. \text{length } clist = \text{length } xs \wedge c \propto clist \wedge \\ & \quad (\forall i < \text{length } clist. clist!i \in cp (xs!i) s)\}) \end{aligned}$$

proves that theorem *one* follows from the two auxiliary lemmas. \square

The compositionality of the semantics is necessary for the proof of soundness of the rule for parallel programs, subject of the next section.

4.5.3 Soundness of the System for Parallel Programs

This section is devoted to the soundness of the system of rules that define the set *par-rghoare*. The type of c in a derivable formula $\vdash c \text{ SAT } [pre, rely, guar, post]$ is a list of the complete specifications of the program components. However, for the validity formula we just require the corresponding parallel program. We obtain it from c with the function

constdefs

$$\begin{aligned} \text{ParallelCom} &:: \alpha \text{ rgformula list} \Rightarrow \alpha \text{ par-com} \\ \text{ParallelCom } Ps &\equiv \text{map } (\text{Some} \circ \text{fst}) Ps \end{aligned}$$

The soundness theorem is formulated using this function as follows:

theorem *par-rgsound*:

$$\begin{aligned} &\vdash c \text{ SAT } [pre, rely, guar, post] \implies \\ &\models \text{ParallelCom } c \text{ SAT } [pre, rely, guar, post] \end{aligned}$$

Proof. The proof proceeds by rule induction. The system *par-rghoare* consists of a single rule: *Parallel*. The soundness of the system is thus reduced to the soundness proof of this rule, namely

lemma *Parallel-sound*:

$$\begin{aligned} &\llbracket \forall i < \text{length } xs. \text{rely} \cup (\bigcup j \in \{j. j < \text{length } xs \wedge j \neq i\}. \text{Guar } (xs!j)) \subseteq \text{Rely } (xs!i); \\ &\quad (\bigcup j \in \{j. j < \text{length } xs\}. \text{Guar } (xs!j)) \subseteq \text{guar}; \\ &\quad pre \subseteq (\bigcap i \in \{i. i < \text{length } xs\}. \text{Pre } (xs!i)); \\ &\quad (\bigcap i \in \{i. i < \text{length } xs\}. \text{Post } (xs!i)) \subseteq \text{post}; \\ &\quad \forall i < \text{length } xs. \models \text{Com}(xs!i) \text{ sat } [\text{Pre}(xs!i), \text{Rely}(xs!i), \text{Guar}(xs!i), \text{Post}(xs!i)] \rrbracket \\ &\implies \models \text{ParallelCom } xs \text{ SAT } [pre, rely, guar, post] \end{aligned}$$

By the soundness of the system for component programs we can assume that all component programs are valid. Our proof follows the one presented in [Xu *et al.*, 1997]. It relies on the compositionality of the semantics and four other lemmas that need this property for their proofs.

The second lemma states the following: Given the assumptions of the lemma *Parallel-sound*, if a computation x of a parallel program satisfies the

assumptions and conjoins with a list of computations of component programs $clist$, then all component transitions in each of the component computations satisfy their corresponding guarantee conditions.

lemma two:

$$\begin{aligned}
& \llbracket \forall i < \text{length } xs. \text{ rely } \cup (\bigcup_{j \in \{j. j < \text{length } xs \wedge j \neq i\}}. \text{ Guar } (xs!j)) \subseteq \text{ Rely } (xs!i); \\
& \quad \text{pre} \subseteq (\bigcap_{i \in \{i. i < \text{length } xs\}}. \text{ Pre } (xs!i)); \text{ length } xs = \text{ length } clist; \\
& \forall i < \text{length } xs. \models \text{ Com } (xs!i) \text{ sat } [\text{ Pre } (xs!i), \text{ Rely } (xs!i), \text{ Guar } (xs!i), \text{ Post } (xs!i)]; \\
& x \in \text{ par-cp } (\text{ ParallelCom } xs) \text{ s}; x \in \text{ par-assum } (\text{ pre}, \text{ rely}); \\
& \forall i < \text{length } clist. clist!i \in \text{ cp } (\text{ Some } (\text{ Com } (xs!i))) \text{ s}; x \propto clist \rrbracket \\
& \implies \forall i < \text{length } clist. \forall j. \text{ Suc } j < \text{length } x \longrightarrow clist!i!j \text{ -c} \longrightarrow clist!i!\text{Suc } j \longrightarrow \\
& \quad (\text{ snd } (clist!i!j), \text{ snd } (clist!i!\text{Suc } j)) \in \text{ Guar } (xs!i)
\end{aligned}$$

The proof proceeds by contradiction. Assume that the first c -transition which does not satisfy the guarantee condition is from $xs!i$ at step m . From the compositionality of the semantics, each e -transition in the subcomputation $\text{take } (\text{Suc } (\text{Suc } m)) (clist!i)$ corresponds to a c -transition in one of the other components or to an e -transition of x , therefore it satisfies $\text{rely} \cup (\bigcup_{j \in \{j. j < \text{length } xs \wedge j \neq i\}}. \text{ Guar } (xs!j))$. Hence, we can prove

$$\text{take } (\text{Suc } (\text{Suc } m)) (clist!i) \in \text{ assum } (\text{ Pre } (xs!i), \text{ Rely } (xs!i))$$

But this contradicts

$$\models \text{ Com } (xs!i) \text{ sat } [\text{ Pre } (xs!i), \text{ Rely } (xs!i), \text{ Guar } (xs!i), \text{ Post } (xs!i)]$$

because one c -transition within the first m transitions does not satisfy the guarantee condition.

Given the assumptions of the previous lemma, the third lemma states that all e -transitions of each of the component computations $xs!i$ satisfy

$$\text{rely} \cup (\bigcup_{j \in \{j. j < \text{length } xs \wedge j \neq i\}}. \text{ Guar } (xs!j)).$$

lemma three:

$$\begin{aligned}
& \llbracket xs \neq []; \text{ pre} \subseteq (\bigcap_{i \in \{i. i < \text{length } xs\}}. \text{ Pre } (xs!i)); \text{ length } xs = \text{ length } clist; \\
& \forall i < \text{length } xs. \text{ rely } \cup (\bigcup_{j \in \{j. j < \text{length } xs \wedge j \neq i\}}. \text{ Guar } (xs!j)) \subseteq \text{ Rely } (xs!i); \\
& \forall i < \text{length } xs. \models \text{ Com } (xs!i) \text{ sat } [\text{ Pre } (xs!i), \text{ Rely } (xs!i), \text{ Guar } (xs!i), \text{ Post } (xs!i)]; \\
& x \in \text{ par-cp } (\text{ ParallelCom } xs) \text{ s}; x \in \text{ par-assum } (\text{ pre}, \text{ rely}); \\
& \forall i < \text{length } clist. clist!i \in \text{ cp } (\text{ Some } (\text{ Com } (xs!i))) \text{ s}; x \propto clist \rrbracket \\
& \implies \forall i < \text{length } clist. \forall j. \text{ Suc } j < \text{length } x \longrightarrow clist!i!j \text{ -e} \longrightarrow clist!i!\text{Suc } j
\end{aligned}$$

$$\begin{aligned} &\longrightarrow (snd (clist!i!j), snd (clist!i!Suc j)) \in \\ &\quad rely \cup (\bigcup_{j \in \{j. j < length\ xs \wedge j \neq i\}}. Guar (xs!j)) \end{aligned}$$

The proof follows directly from lemma *two*.

The fourth lemma says that each *pc*-transition satisfies *guar*:

lemma four:

$$\begin{aligned} &\llbracket xs \neq []; x \in par\text{-}cp (ParallelCom\ xs)\ s; x \in par\text{-}assum (pre, rely); \\ &\quad \forall i < length\ xs. rely \cup (\bigcup_{j \in \{j. j < length\ xs \wedge j \neq i\}}. Guar (xs!j)) \subseteq Rely (xs!i); \\ &\quad (\bigcup_{j \in \{j. j < length\ xs\}}. Guar (xs!j)) \subseteq guar; \\ &\quad pre \subseteq (\bigcap_{i \in \{i. i < length\ xs\}}. Pre (xs!i)); \\ &\quad \forall i < length\ xs. \models Com (xs!i)\ sat [Pre (xs!i), Rely (xs!i), Guar (xs!i), Post (xs!i)]; \\ &\quad Suc\ i < length\ x; x!i \text{-}pc \rightarrow x!Suc\ i \rrbracket \\ &\implies (snd (x!i), snd (x!Suc\ i)) \in guar \end{aligned}$$

The proof follows from lemma *two*.

The last lemma states that if the computation terminates, the final state satisfies all postconditions of the components and thus the postcondition of the parallel program:

lemma five:

$$\begin{aligned} &\llbracket xs \neq []; x \in par\text{-}cp (ParallelCom\ xs)\ s; x \in par\text{-}assum (pre, rely); \\ &\quad \forall i < length\ xs. rely \cup (\bigcup_{j \in \{j. j < length\ xs \wedge j \neq i\}}. Guar (xs!j)) \subseteq Rely (xs!i); \\ &\quad \forall i < length\ xs. \models Com (xs!i)\ sat [Pre (xs!i), Rely (xs!i), Guar (xs!i), Post (xs!i)]; \\ &\quad pre \subseteq (\bigcap_{i \in \{i. i < length\ xs\}}. Pre (xs!i)); \\ &\quad (\bigcap_{i \in \{i. i < length\ xs\}}. Post (xs!i)) \subseteq post; \\ &\quad fst (last\ x) = ys; All\None\ ys \rrbracket \implies snd (last\ x) \in post \end{aligned}$$

From lemma *one* there exists a list of computations *clist* such that $x \propto clist$. From lemma *three* and the hypothesis it follows that

$$\forall i < length\ clist. clist!i \in assum (Pre (xs!i), Rely (xs!i))$$

By validity of the component programs the last state of each one satisfies *Post (xs!i)*. From the definition of conjoining computations the last state of all component computations is the same as the last state of *x*. Thus, it satisfies *post*.

The soundness of the parallel composition rule *Parallel-sound* follows directly from lemmas *four* and *five*. This concludes the proof of *par-rsound*. \square

The soundness proofs of both systems represent the theoretical part of the formalization of the rely-guarantee system. The next two sections are concerned with the practical application. Section 4.6 defines concrete syntax for programs in order to facilitate verification of real programs and section 4.7 presents some examples.

4.6 Concrete Syntax

This section presents an alternative external representation for programs and assertions. It approaches the standard syntax used in the literature providing a user-friendlier interface for the application of the formalized rely-guarantee method on concrete programs.

The concrete syntax defined here is similar to that of §2.7 where we presented the concrete syntax for the language of the Owicki-Gries formalization. In particular, the representation of program variables follows the quote/antiquote technique. We refer to section 2.7.1 for detailed explanations. Here, we concentrate on the particularities of the concrete syntax for the language of the present chapter. An overview of the different elements and their external representation is shown in table 4.1.

In a rely-guarantee specification the pre- and postcondition are sets of single states. Like in 2.7 variables within such assertions are represented in sans serif font, e.g. $\{x = 0\}$. Internally, x is a selector function that refers to the value of the variable x at some state.

The rely and guarantee conditions are, however, sets of pairs of states. We need to refer to the values of the variables of the first and second states, i.e. the values before and after the transition. In this case, using simply a selector function is not enough. Antiquoted entities refer now to a pair. The solution is to “antiquote” the selector function that refers to a variable, e.g. x , composed with the predefined functions on pairs fst and snd , for the value of the variable x before or after the transition, respectively. For example, the set $\{(s_1, s_2) \mid s_1 \ s_2. x \ s_1 = 0 \wedge x \ s_2 = 1\}$ is represented by the expression $\{ \prime(x \circ fst) = 0 \wedge \prime(x \circ snd) = 1 \}$ or equivalently $\{ \{ x \ \prime fst = 0 \wedge x \ \prime snd = 1 \}$, in the concrete syntax. These expressions are, however, quite long, so we introduce new syntax for antiquotations referring to the *before* and *after* the transition.

syntax

-before $:: (\alpha \Rightarrow \beta) \Rightarrow \beta \ (^{\circ}-)$
 -after $:: (\alpha \Rightarrow \beta) \Rightarrow \beta \ (^{\text{a}}-)$

Assertion	$\{r\}$
<i>Sequential commands</i>	
Basic	$x := e$
Seq	$c_0;; c_1$
Cond	if b then c_0 else c_1 fi
While	while b do c od
Await	await b then c end
Skip	skip
Cond2	if b then c_0 fi
Atom	$\langle c \rangle$
Wait	wait b end
<i>Parallel commands</i>	
Parallel	cobegin $s_0 \parallel \dots \parallel s_n$ coend
Schematic	scheme $[j \leq i < k]$ s
<i>Notation</i>	
x : program variable \bar{x} : value of program variable after a transition e : expression of the type of x r, b : boolean expressions c, c_0, c_1 : sequential commands s, s_0, \dots, s_n : specifications for component programs j, k : limits for indexing parameterized programs	

Table 4.1: Concrete syntax for programs.

Note that these syntax declarations are similar to that of the syntax constant *-quote* of §2.7.1. The corresponding translations into internal syntax are:

translations

$$\overset{\circ}{x} \rightarrow x \text{ 'fst}$$

$$\overset{\text{a}}{x} \rightarrow x \text{ 'snd}$$

In the examples shown in this thesis we tune the output by printing entities preceded by an antiquote symbol ' in sans serif font. To maintain this nice output for all variables we print entities preceded by $\overset{\circ}$ also in sans serif and those preceded by $\overset{\text{a}}$ are printed in sans serif and overlined, e.g. the set $\{ \overset{\circ}{x} = \overset{\text{a}}{x} \}$ is printed $\{ x = \bar{x} \}$.

The declaration of syntax constants, the equational and the parse and print translations that translate the external concrete syntax into the internal abstract syntax and vice versa are very similar to those defined for the language of the Owicki-Gries formalization (see appendix B).

4.7 Examples

We show the application of the method on three known examples from the literature [Xu *et al.*, 1997, Stirling, 1988].

4.7.1 Set Elements of an Array to Zero

The first example is a very simple program. It sets the first n components of an array A to zero in parallel. There are no shared variables and no auxiliary variables are required for the verification.

```
record Example1 =
  A :: nat list
```

```
lemma Example1:
```

```
  ⊢ cobegin
```

```
    scheme [0 ≤ i < n]
```

```
    (A := A [i := 0],
```

```
    { n < length A },
```

```
    { length A = length  $\bar{A}$  ∧ A ! i =  $\bar{A}$  ! i },
```

```
    { length A = length  $\bar{A}$  ∧ (∀ j < n. i ≠ j → A ! j =  $\bar{A}$  ! j) },
```

```
    { A ! i = 0 })
```

```
  coend
```

```
  SAT [{ n < length A }, { A =  $\bar{A}$  }, { True }, { ∀ i < n. A ! i = 0 }]
```

The array (modeled as a list) must have at least as many elements as the number of parallel components; this is the only requirement in the precondition. The rely condition of component i requires that the environment does not change the value of the array A at index i . On the other hand, it guarantees not to change the values of the other indices. The length of the array is invariant, thus this holds in both the rely and the guarantee condition.

We consider this program as closed with respect to the environment. This is reflected in the overall rely condition which requires that the environment does not affect the variables used in the program. The overall guarantee condition can be set to *True*, because no other program relies on the behavior of this one. If there was an influencing environment, its effect could be stated in the overall rely and guarantee conditions

The proof just requires us to first apply the *Parallel* rule backwards. The rule *Basic* is used for the proof of derivability of the parameterized component program. The generated verification conditions are easily proven with standard Isabelle techniques.

4.7.2 Increment a Variable in Parallel

Consider the program $x := x+1 \parallel x := x+1$. The goal is to prove that if the precondition $x = 0$ holds, then the postcondition $x = 2$ is satisfied by the final states. This parallel program is a classic in the literature because, despite its simplicity, auxiliary variables are unavoidable for its verification.

We declare the shared variable x and two “private” auxiliary variables c_0 and c_1 , one for each component:

```
record Example2 =
  x :: nat
  c0 :: nat
  c1 :: nat
```

The program must be extended with assignments to the auxiliary variables. The first and second components satisfy complementary specifications, where the rely and guarantee conditions are switched. We verify the parallel composition as a closed system:

```
lemma Example2:
  ⊢ cobegin
  (⟨x := x+1;; c0 := c0 + 1⟩,
  { x = c0 + c1 ∧ c0 = 0 }},
  { c0 = c̄0 ∧ (x = c0 + c1 ⟶ x̄ = c̄0 + c̄1) },
  { c1 = c̄1 ∧ (x = c0 + c1 ⟶ x̄ = c̄0 + c̄1) },
  { x = c0 + c1 ∧ c0 = 1 }})
  ||
  (⟨x := x+1;; c1 := c1+1⟩,
```

$$\begin{aligned}
& \{ \! \{ x = c_0 + c_1 \wedge c_1 = 0 \! \} \}, \\
& \{ \! \{ c_1 = \bar{c}_1 \wedge (x = c_0 + c_1 \longrightarrow \bar{x} = \bar{c}_0 + \bar{c}_1) \! \} \}, \\
& \{ \! \{ c_0 = \bar{c}_0 \wedge (x = c_0 + c_1 \longrightarrow \bar{x} = \bar{c}_0 + \bar{c}_1) \! \} \}, \\
& \{ \! \{ x = c_0 + c_1 \wedge c_1 = 1 \! \} \} \\
\mathbf{coend} \\
& SAT [\{ \! \{ x = 0 \wedge c_0 = 0 \wedge c_1 = 0 \! \} \}, \{ \! \{ x = \bar{x} \wedge c_0 = \bar{c}_0 \wedge c_1 = \bar{c}_1 \! \} \}, \\
& \quad \{ \! \{ True \! \} \}, \{ \! \{ x = 2 \! \} \}]
\end{aligned}$$

The proof uses the rule for parallel composition *Parallel* and the system of rules for sequential component programs.

Parameterized Version

We take advantage of the possibility of proving the derivability of parameterized programs in our system and show in the next lemma how to prove the correctness of the last example for any number of components.

We introduce a composed variable c , such that the component i of the parallel program atomically updates the shared variable x and the i th component of c .

There are several possibilities of modeling such composed variables in HOL. One of them is as lists of length n (like in the previous example), however, lists cause in some cases unnecessary trouble, for example the invariance of the length has to be explicitly required in all assertions. Another more abstract possibility is to use functions from naturals to the value domain of the components. The syntax for updating the value of a function f on argument i is $f(i := t)$, where t is of the corresponding type.

The declaration of the variables is:

```

record Example2-parameterized =
  x :: nat
  c :: nat ⇒ nat

```

We want to prove that the extended parameterized program

$$\langle x := x+1, c := c(i := 1) \rangle$$

satisfies the rely-guarantee specification

$$(x = 0 \wedge (\sum i < n. c i) = 0, x = \bar{x} \wedge c = \bar{c}, True, x = n).$$

For the assertions, we use the summation function predefined in the Isabelle library. In the precondition of the program we require the summation of the values of the composed variable c on the first n natural numbers

to be 0. The rely and guarantee conditions indicate that the program is to be executed in a closed environment.

To establish the specification above, we require that each component i satisfy suitable (parameterized) local specifications.

The lemma stating the derivability of the full specification is:

lemma *Example2-parameterized*: $0 < n \implies$

\vdash **cobegin**

scheme $[0 \leq i < n]$

$\langle\langle x := x+1;; c := c (i := 1) \rangle\rangle,$

$\{\{ x = (\sum i < n. c i) \wedge c i = 0 \}\},$

$\{\{ c i = \bar{c} i \wedge (x = (\sum i < n. c i) \longrightarrow \bar{x} = (\sum i < n. \bar{c} i)) \}\},$

$\{\{ (\forall j < n. i \neq j \longrightarrow c j = \bar{c} j) \wedge$

$(x = (\sum i < n. c i) \longrightarrow \bar{x} = (\sum i < n. \bar{c} i)) \}\},$

$\{\{ x = (\sum i < n. c i) \wedge c i = 1 \}\}$

coend

SAT $[\{\{ x = 0 \wedge (\sum i < n. c i) = 0 \}\}, \{\{ x = \bar{x} \wedge c = \bar{c} \}\}, \{\{ True \}\}, \{\{ x = n \}\}]$

It is fairly easy to prove with the help of some basic lemmas about the summation function.

In section 5.5 of the next chapter we reconsider this example and study the relation with its proof in the Owicki-Gries system.

4.7.3 Find Least Element

The next example was already used in [Owicki and Gries, 1976a] and has ever since constituted a standard example of parallel program verification. We reproduce the more general version for parameterized parallel programs of [Stirling, 1988].

Let B be an array, modeled as a list of length m . We want a program that finds the first element, if there is one, satisfying the predicate P . If so, it is saved in the variable x , otherwise $x = m$. Then, the program *FINDP* should satisfy the specification

$$\{\{ True \}\} \textit{FINDP} \{\{ x < m+1 \wedge (\forall i < x. \neg P (B!i)) \wedge x < m \longrightarrow P (B!x) \}\}$$

where we already use the notation of Isabelle for access to components of a list. The program *FINDP* is of the form: *INIT*; *SEARCH*; *END*, where *INIT* initializes, *SEARCH* searches in parallel and *END* performs the final assignment to x .

In *SEARCH* we use a number of concurrent programs, *SEARCH* (0) $\parallel \dots \parallel$ *SEARCH* ($n-1$) with $n \leq m$; for simplicity let n divide m . Each *SEARCH* (i) scans the array squares $i, n+i, 2*n+i, \dots, m+i$ looking for x .

Assume that each *SEARCH* (i) has a private variable x_i for searching. Each *SEARCH* (i) should terminate when

1. $P(B ! x_i)$ or
2. $x_i > m$ or
3. *SEARCH* (k), with $k \neq i$, has found that $P(B ! x_k)$ for $x_k < x_i$.

We introduce a further private variable y_i for each *SEARCH* (i) which initially is set to $m+i$ and if $P(B ! x_i)$ holds, then y_i is set to x_i . Hence, the termination condition for *SEARCH* (i) is $\exists j < n. y_j \leq x_i$. Finally, *END* sets x to the value of $\min(y_0, \dots, y_{n-1})$ when each *SEARCH* (i) terminates.

Consider *SEARCH* (i). Then,

1. As x_i, y_i are private, the environment cannot affect their values or increase $y_k, k \neq i$.
2. The program *SEARCH* (i) cannot affect the variables x_k and y_k for $k \neq i$ and it does not increase the initial value of y_i .

SEARCH (i) will be a loop with invariant:

$$\mathbf{inv} : x_i \bmod n = i \wedge (\forall j < x_i. j \bmod n = i \longrightarrow \neg P(B ! j)) \wedge (y_i < m \longrightarrow P(B ! y_i) \wedge y_i \leq m + i)$$

The full specification is shown below. Like in the previous example we represent the private variables x_i and y_i as functions x, y from naturals to naturals, so that $x \ i$ corresponds to the private variable x_i of component i :

record *Example3* =

$x :: \text{nat} \Rightarrow \text{nat}$

$y :: \text{nat} \Rightarrow \text{nat}$

lemma *Example3*: $m \bmod n = 0 \implies$

\vdash **cobegin**

scheme $[0 \leq i < n]$

(while $(\forall j < n. x \ i < y \ j)$ **do**


```

if  $P (B ! x i)$  then  $y := y (i := x i)$ 
else  $x := x (i := (x i) + n)$  fi
od,
 $\{ (x i) \bmod n = i \wedge (\forall j < x i. j \bmod n = i \longrightarrow \neg P (B ! j)) \wedge$ 
 $(y i < m \longrightarrow P (B ! y i) \wedge y i \leq m+i) \}$ ,
 $\{ (\forall j < n. i \neq j \longrightarrow \bar{y} j \leq y j) \wedge x i = \bar{x} i \wedge y i = \bar{y} i \}$ ,
 $\{ (\forall j < n. i \neq j \longrightarrow x j = \bar{x} j \wedge y j = \bar{y} j) \wedge \bar{y} i \leq y i \}$ ,
 $\{ (x i) \bmod n = i \wedge (\forall j < x i. j \bmod n = i \longrightarrow \neg P (B ! j)) \wedge$ 
 $(y i < m \longrightarrow P (B ! y i) \wedge y i \leq m+i) \wedge (\exists j < n. y j \leq x i) \}$ 
coend
SAT [ $\{ \forall i < n. x i = i \wedge y i = m+i \}$ ,  $\{ x = \bar{x} \wedge y = \bar{y} \}$ ,  $\{ True \}$ ,
 $\{ \forall i < n. (x i) \bmod n = i \wedge (\forall j < x i. j \bmod n = i \longrightarrow \neg P (B ! j)) \wedge$ 
 $(y i < m \longrightarrow P (B ! y i) \wedge y i \leq m+i) \wedge (\exists j < n. y j \leq x i) \}$ ]

```

The initialization part of the program (*INIT*) as well as the final assignment $x := \min(y_0, \dots, y_{n-1})$ of *END* are not included because the sequential composition with parallel programs is not defined in the programming language. If so wanted, the language should be extended and so the semantics and the proof rules. Here we concentrate on the verification of the parallel part *SEARCH*.

We consider the parallel program closed and thus define the overall rely and guarantee conditions as such.

The interactive proof is easy but needs explicit hints about the assertions that should be used for the final verification conditions. This reveals an important aspect which was not obvious while studying the theory: an automatic verification generation tactic would need intermediate annotations.

It is known that sequential while-programs need only be annotated with the corresponding loop invariants. This is because invariants are not, in general, derivable from the postcondition and the program itself. Loop invariants, precondition and postcondition are the only annotations that an automatic procedure requires to extract the verification conditions for a sequential program.

For the Owicki-Gries method, more annotations than loop-invariants were required. In fact, programs had to be fully annotated as proof outlines. The automatic tactic could thus extract the verification conditions out of the intermediate annotations supplied by the user.

In the verification of this example we observe that an automatic tactic to generate the verification conditions is not possible by just annotating programs with the four conditions (pre-, rely-, guar- and postcondition).

While-invariants and even explicit intermediate annotations are also needed. To illustrate the insufficiency of the rely-guarantee specification we follow the verification of this example and show that extra information must be supplied by the user.

An automatic vcg-tactic that uses the specified rules in the theory would proceed in the following order:

- 1 First, the *Parallel* rule is applied backwards. We obtain four verification conditions and a fifth goal concerning the derivability of the component programs.
- 2 The component programs are all the same up to indexing, so the goal reduces to the following formula:

lemma $i < n \implies$

\vdash **while** $(\forall j < n. x\ i < y\ j)$
do if $P\ (B\ !\ x\ i)$ **then** $y := y\ (i := x\ i)$
else $x := x\ (i := x\ i + n)$ **fi**
od

sat

$[\{\{ x\ i \bmod n = i \wedge (\forall j < x\ i. j \bmod n = i \longrightarrow \neg P\ (B\ !\ j)) \wedge$
 $(y\ i < m \longrightarrow P\ (B\ !\ y\ i) \wedge y\ i \leq m+i) \}\},$
 $\{\{ (\forall j < n. i \neq j \longrightarrow \bar{y}\ j \leq y\ j) \wedge x\ i = \bar{x}\ i \wedge y\ i = \bar{y}\ i \}\},$
 $\{\{ (\forall j < n. i \neq j \longrightarrow x\ j = \bar{x}\ j \wedge y\ j = \bar{y}\ j) \wedge \bar{y}\ i \leq y\ i \}\},$
 $\{\{ x\ i \bmod n = i \wedge (\forall j < x\ i. j \bmod n = i \longrightarrow \neg P\ (B\ !\ j)) \wedge$
 $(y\ i < m \longrightarrow P\ (B\ !\ y\ i) \wedge y\ i \leq m+i) \wedge (\exists j < n. y\ j \leq x\ i) \}\}]$

- 3 The precondition in this case has been carefully chosen to be the invariant of the loop, thus we can apply the *While* rule directly. From this rule we obtain five subgoals. Four of them are solvable verification conditions. Only the derivability of the while-body remains:

lemma $i < n \implies$

\vdash **if** $P\ (B\ !\ x\ i)$ **then** $y := y\ (i := x\ i)$
else $x := x\ (i := x\ i + n)$ **fi**

sat

$[\{\{ x\ i \bmod n = i \wedge (\forall j < x\ i. j \bmod n = i \longrightarrow \neg P\ (B\ !\ j))$
 $\wedge (y\ i < n*q \longrightarrow P\ (B\ !\ y\ i)) \}\} \cap \{\{ \forall j < n. x\ i < y\ j \}\},$

$$\begin{aligned}
& \{ (\forall j < n. i \neq j \longrightarrow \bar{y} j \leq y j) \wedge x i = \bar{x} i \wedge y i = \bar{y} i \}, \\
& \{ (\forall j < n. i \neq j \longrightarrow x j = \bar{x} j \wedge y j = \bar{y} j) \wedge \bar{y} i \leq y i \}, \\
& \{ x i \bmod n = i \wedge (\forall j < x i. j \bmod n = i \longrightarrow \neg P (B ! j)) \wedge \\
& \quad (y i < n * q \longrightarrow P (B ! y i)) \}]
\end{aligned}$$

The problem we observe is that the precondition of this subgoal is automatically $pre \cap b$ (see rule *Cond*), which is the strongest precondition. When we attempt to verify the derivability of this subprogram, we fail at condition *stable pre rely* which results from applying the rule *Cond*.

The rely condition only ensures that the values of $x i$ and $y i$ remain invariant for the component i , but we cannot prove that $\forall j < n. x i < y j$, because the environment can decrease the values of some $y k$ with $k \neq i$.

However, we can verify the program if we first apply the *Conseq* rule and choose a weaker precondition, namely, the same one but where the condition b only refers to the variables with index i , i.e. $x i < y i$. Obviously $pre \cap b$ implies the weaker precondition. If we apply the rule *Cond* backwards with this new weaker precondition we can prove *stable pre rely*.

For an automatic tactic it is not possible to guess when the intermediate assertions should be weakened or strengthened. Even more difficult would be to guess the new pre-, or postconditions that yield a successful derivation. The solution lies in defining a new set of inference rules for annotated programs and designing the automatic tactic in terms of these rules.

4.8 Concluding Remarks

The rely-guarantee method represents the logical successor of Owicki-Gries as its compositional reformulation. We have formalized the system and a soundness proof following the presentation given in [Xu *et al.*, 1997]. Often, the definitions and proofs found in the literature turn out to be too inefficient for theorem proving techniques. In this formalization, for example, we give an alternative (but equivalent) definition for the semantics in order to carry out the proofs successfully.

So far, we have applied the formalization to the verification of simple programs. The verifications have been done by interactively applying the proof rules and using the standard Isabelle techniques for proving the generated verification conditions. This works fine for small programs, but otherwise it becomes quite tedious. To handle the verification of larger programs a tactic for the automatic generation of the verification conditions should be designed. Such a tactic requires information about intermediate annota-

tions, which means that programs have to be presented as proof outlines. Another interesting extension is the inclusion of nested parallelism in the language. This would allow us to use the compositionality of the method for top-down design and verification of large systems. Both extensions involve changes in the definition of the programming language, which would influence the whole formalization. However, the proof of soundness is essentially the same as the one presented in [Xu *et al.*, 1997] where nested parallelism is part of the language. Thus, the formalized proofs would only need minor modifications.

Chapter 5

Completeness of the Proof Systems for Parameterized Parallel Programs

In the previous chapters we have formalized the Owicki-Gries and the rely-guarantee systems in Isabelle/HOL and proved soundness of these systems w.r.t. the corresponding underlying semantics.

The complementary property of soundness is completeness. A system for verification of programs is *complete* if all correct programs can be deduced in the system. Soundness is a minimal requirement of such systems. One starts with a trivial sound system, for example, the empty one, and adds sound axioms and rules which preserve soundness until a complete system is achieved.

Our current proof systems are incomplete because there is no rule for removing auxiliary variables. This rule is fundamental for a completeness proof. However, for the practical purposes of this work, i.e. mechanical verification of parallel programs, it is sufficient to find derivations for programs augmented with assignments to auxiliary variables. When these variables are used correctly they influence neither control nor data flow. Consequently, removing all assignments to auxiliary variables does not affect the correctness of the program.

The systems with the rule for removing auxiliary variables have been proved complete in previous works. Completeness proofs for the Owicki-Gries system can be found in [Owicki, 1975], [Apt, 1981a] or in [de Roever *et al.*, 2000]. For the rely-guarantee system see for example [Xu *et al.*, 1997], [Stirling, 1988] or again [de Roever *et al.*, 2000]. Both systems are proven

to be relatively complete with respect to the standard interpretation in the natural numbers. Moreover, in [Apt, 1981a] the author proves that the assertions for the Owicki-Gries system are recursive when history variables are used as auxiliary variables but only recursively enumerable when program counters are used.

We assume these results to be known and concentrate on extending the completeness property to a kind of programs not considered in the completeness proofs mentioned above, namely, parallel programs where the number of parallel components is represented by a parameter n . Contrary to the traditional systems found in the literature, these programs can be directly derived in our systems. The reason for this lies in the representation of parallel programs as lists of component programs.

The chapter is organized as follows. In section 5.1 we give a formal characterization of parameterized programs. Section 5.2 presents the completeness proof of the extended Owicki-Gries system. Section 5.4 extends this result to the rely-guarantee method. In section 5.5 we illustrate by an example the relevance of these completeness results.

The results presented in this chapter are, apart from the examples of section 5.5, not carried out with Isabelle and thus independent from the details of the formalization. In the sequel, we use a standard notation (even for program syntax) which is not necessarily the one required by the theorem prover.

5.1 Parameterized Programs

Parameterized programs represent a family of programs by a single syntactic object. In order to represent, understand and manipulate these objects, several aspects of the programming language have to be adapted to deal with parameters. In this section we give a formal description of the syntax of parameterized programs and other aspects relevant for the following sections.

5.1.1 Syntax of Parameterized Programs

Let n and i be parameters, where n ranges over the natural numbers and i ranges over the subset $\{0, \dots, n - 1\}$. A *parameterized component program* $S(i, n)$, parameterized by n and i , is described by the following syntactic sets:

1. A finite set of simple variables $V = \{x_0, \dots, x_k\}$, where we assume without loss of generality that each x_j for $j \leq k$ ranges over the natural numbers, and a finite set of composed variables $V_c = \{y_0, \dots, y_l\}$. (Simple variables could also be considered a degenerated case of composed variables.) Composed variables are usually implemented as arrays, where the value of an array is a function from some “range” to some “set” of values [Best, 1996]. Access to components outside the range of an array are not allowed. We interpret composed variables directly on \mathcal{N} in a way that we explain in §5.2.1.
2. Parameterized arithmetic and boolean expressions given by the following BNF grammars:

$$\begin{aligned}
 a &::= j \mid m \mid N \mid x \mid y[a_0] \mid a_0 + a_1 \mid a_0 - a_1 \mid a_0 * a_1 \\
 b &::= \text{True} \mid b_0 \wedge b_1 \mid \neg b_0 \mid \exists x. b_0 \mid a_0 \leq a_1
 \end{aligned}$$

where j and m are special variables that take the values given by the parameters i and n , N is a natural number, $x \in V$ and $y \in V_c$.

Given an arithmetic expression a , we define $a(i, n)$ inductively as follows:

$$\begin{aligned}
 j(i, n) &\stackrel{\text{def}}{=} i, \quad m(i, n) \stackrel{\text{def}}{=} n, \quad N(i, n) \stackrel{\text{def}}{=} N, \\
 x(i, n) &\stackrel{\text{def}}{=} x, \quad y[a_0](i, n) \stackrel{\text{def}}{=} y[a_0(i, n)], \\
 (a_0 + a_1)(i, n) &\stackrel{\text{def}}{=} a_0(i, n) + a_1(i, n), \\
 (a_0 - a_1)(i, n) &\stackrel{\text{def}}{=} a_0(i, n) - a_1(i, n), \\
 (a_0 * a_1)(i, n) &\stackrel{\text{def}}{=} a_0(i, n) * a_1(i, n)
 \end{aligned}$$

Parameterized boolean expressions $b(i, n)$ are defined analogously.

3. Finally, the syntax of parameterized component programs:

$$\begin{aligned}
 S &::= x := a \mid y[a_0] := a_1 \mid S_1; S_2 \mid \\
 &\quad \mathbf{if} \ b \ \mathbf{then} \ S_1 \ \mathbf{else} \ S_2 \ \mathbf{fi} \mid \\
 &\quad \mathbf{while} \ b \ \mathbf{do} \ S \ \mathbf{od} \mid \mathbf{await} \ b \ \mathbf{then} \ T \ \mathbf{end}
 \end{aligned}$$

We define $S(i, n)$ inductively as follows:

$$\begin{aligned}
(x := a)(i, n) &\stackrel{\text{def}}{=} x := a(i, n), \\
(y[a_0] := a_1)(i, n) &\stackrel{\text{def}}{=} y[a_0(i, n)] := a_1(i, n), \\
(S_1; S_2)(i, n) &\stackrel{\text{def}}{=} S_1(i, n); S_2(i, n), \\
(\mathbf{if } b \mathbf{ then } S_1 \mathbf{ else } S_2 \mathbf{ fi})(i, n) &\stackrel{\text{def}}{=} \\
\mathbf{if } b(i, n) \mathbf{ then } S_1(i, n) \mathbf{ else } S_2(i, n) \mathbf{ fi} \\
&\dots
\end{aligned}$$

The family of parallel programs built from $S(i, n)$ is defined by the set $\{\|_{i=0}^{n-1} S(i, n) \mid n \in \mathbb{N}\}$. The notation $\|_{i=0}^{n-1} S(i, n)$ represents the program $S(1, n) \parallel \dots \parallel S(n, n)$, where \parallel is the parallel composition construct based on interleaving. For each n , the resulting member is a concrete (non-parameterized) parallel program like the programs considered in the original Owicki-Gries and rely-guarantee systems.

A typical example for a parameterized parallel program is the ticket mutual exclusion algorithm for distributed processes that we already considered in §2.8

$$\mathbf{ticket} \equiv num := 1; next := 1; turn := 0; \|_{i=0}^{n-1} S(i, n)$$

where $S(i, n)$ is shown in Figure 5.1. By the assignment $turn := 0$ we mean that 0 is assigned to each component of the array $turn$.

$$\begin{aligned}
S(i, n) \equiv & \mathbf{while } true \mathbf{ do} \\
& NCS(i, n); \text{ (noncritical section)} \\
& \langle turn[i] := num; num := num + 1 \rangle; \\
& \mathbf{wait } turn[i] = next \mathbf{ end}; \\
& CS(i, n); \text{ (critical section)} \\
& next := next + 1 \\
& \mathbf{od}
\end{aligned}$$

Figure 5.1: Ticket algorithm.

Syntax of assignments to composed variables

The concrete syntax of the programming language in the Isabelle/HOL formalizations of the previous chapters does not support the above notations

for assignments to components of composed variables. For example, an assignment to component i of a list a is written $a := a [i:=e]$. Similarly, an assignment to the argument i of a function a is written $a := a (i:=e)$ meaning that the function is assigned the new updated function. In the literature, such assignments are usually written $a[i]:= e$, however, both notations have the same semantics, which we explain in the next section.

5.1.2 The State Space

The meaning of a program S is usually defined as a partial function $\mathcal{M}(S)$ from states to states. States are represented as tuples of values corresponding to the variables occurring in the program. For the purpose of our completeness proof it is convenient to require that these values be natural numbers. Thus, given a state s and a simple variable x , we refer to the value of the variable x in state s by $s(x)$. For a composed variable y , $s(y)$ returns the corresponding encoding of the sequence of values given by its components. This means that the value domain of a composed variable is also \mathbb{N} . As we shall see in §5.3, the reason for this is that free variables in the language of elementary arithmetic must range over numbers and not over functions.

There are many different ways of encoding a finite sequence of values as a single natural number. Any such encoding will do as long as there is an effective procedure for encoding and an effective procedure for extracting any of the original components from the coded sequence. Following the syntax in [Apt, 1981a], the encoding of the sequence a_1, \dots, a_k is represented by $[a_1, \dots, a_k]$. If $a = [a_1, \dots, a_k]$ then $a \frown c = [a_1, \dots, a_k, c]$. We denote the code of the empty sequence by $[\]$. We assume some properties of the functions “[...]” and “ \frown ”, in particular, their definability in the language of elementary arithmetic. The proofs of these properties can be found in [Schoenfield, 1967]. The state space Σ of a given parameterized program is then \mathbb{N}^k , where k is the total number of variables appearing in the program.

An assignment to a component i of a composed variable should be understood as an assignment to the composed variable as a whole. Thus, the new value is the composed variable with the component i updated. Due to the special treatment of the value domain for composed variables we informally describe the semantics of assignments to components of composed variables by the following transition rule:

$$(y[a_0] := a_1, s) \rightarrow (None, s[y \leftarrow update(y, s(a_0), s(a_1))])$$

None denotes the empty program. The notation $s[y \leftarrow t]$ stands for the usual operation of substitution in s of the value of the variable y by the new value t . The syntax $s(a_0)$ denotes the value resulting from the evaluation of the arithmetic expression a_0 by substituting the free variables by the values they have in state s . Finally, $update(y, j, l)$ is a function that replaces the j^{th} component of the sequence encoded by y with a new value l . It is possible to prove that there is a partial recursive procedure computing this function. Thus, the transformations performed on the state by a program are all computable.

The rules of the semantics for the rest of the constructors can be consulted in sections 2.2 and 4.2.

The next section deals with the completeness result for parameterized parallel programs of the Owicki-Gries system. As will be explained in section 5.4, due to the strong connection between the completeness proofs of the Owicki-Gries and the rely-guarantee systems, and to the independence of the proof in section 5.2 from the details of the Owicki-Gries system, this result can be easily extended to the rely-guarantee system.

5.2 Completeness of the Owicki-Gries System for Parameterized Parallel Programs

The proof rule for parallel composition as formulated in the original Owicki-Gries system is¹

$$\frac{\begin{array}{l} \vdash \{P_0\} S_0 \{Q_0\}, \dots, \vdash \{P_k\} S_k \{Q_k\} \\ \text{and the proof outlines are interference free} \end{array}}{\vdash \{P_0 \cap \dots \cap P_k\} S_0 \parallel \dots \parallel S_k \{Q_0 \cap \dots \cap Q_k\}}$$

This rule is concerned with the verification of parallel programs consisting of a fixed number k of (possibly different) components S_0, \dots, S_k . We refer to the Owicki-Gries system by O .

However, as encountered already in the previous chapters, many interesting parallel programs are given schematically in terms of a parameter n , representing the number of parallel components. In order to be able to verify these programs directly we consider the extension of the system O by

¹Assertions are modeled as sets of states here.

the following rule:

$$\frac{\begin{array}{l} \forall i < n. \vdash \{P(i, n)\} S(i, n) \{Q(i, n)\} \wedge \\ \forall i < n. \forall j < n. i \neq j \longrightarrow \text{the proof outlines of} \\ \{P(i, n)\} S(i, n) \{Q(i, n)\} \wedge \{P(j, n)\} S(j, n) \{Q(j, n)\} \\ \text{are interference free} \end{array}}{\vdash \{\bigcap_{i=0}^{n-1} P(i, n)\} \parallel_{i=0}^{n-1} S(i, n) \{\bigcap_{i=0}^{n-1} Q(i, n)\}}$$

where $P(i, n)$ and $Q(i, n)$ denote for each i and n the precondition and the postcondition, respectively.

This rule represents a particular case of the more general rule presented in the formalization of the Owicki-Gries system in §2.4.3. Hence, we can easily derive it from the system²:

lemma *ParamParallelRule*:

$$\begin{array}{l} \llbracket \forall i < n. \vdash (c\ i) (Q\ i); \\ \forall i < n. \forall j < n. i \neq j \longrightarrow \text{interfree-aux } (Some\ (c\ i),\ Q\ i,\ Some\ (c\ j)) \rrbracket \\ \implies \vdash (\bigcap i \in \{i. i < n\}. \text{pre}(c\ i)) \\ \quad \mathbf{cobegin\ scheme} [0 \leq i < n] (c\ i) (Q\ i) \mathbf{coend} \\ \quad (\bigcap i \in \{i. i < n\}. Q\ i) \end{array}$$

This rule allows us to prove partial correctness of a parallel program parameterized by the number of component processes n by showing it for an arbitrary but fixed value of n . We refer to the system of rules that results by adding this new rule to the original Owicki-Gries system by F (F stands for *family*). We write $\vdash_F \{P\} S \{Q\}$ to denote that $\{P\} S \{Q\}$ can be derived from F . The soundness of this system has been already proven in chapter 2. Here we prove the relative completeness of F .

The proof proceeds by induction on the structure of the programs. For all non-parameterized constructors the proofs can be found in, for example [Owicki, 1975]. The only interesting case here is that of parameterized parallel programs, i.e. given a family of valid partial correctness formulas of the form

$$\models \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$$

we want to study if we can always find a proof outline for S so that

$$\vdash_F \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$$

²This rule is specific for the derivation of one parameterized parallel program. However, the theoretical completeness result presented here can be generalized to the more general rule of the formalization.

By the completeness of the system O we know that for each value of n we can always find proof outlines for each of the n components, so that the program can be derived in O . This result, however, does not give any information about what these proof outlines look like for different values of n and i . Do they have a uniform pattern (probably parameterized by n and i) or are the proof outlines of the program with three components completely different from those of the program with four components? And, for a given value of n , is the annotation of the component i different from the annotation of the component j for some values of i and j in the set $\{0, \dots, n - 1\}$?

Previous works on parallel program verification present proofs of parameterized programs in the style of Owicki and Gries [Stirling, 1988, de Roever *et al.*, 2000]. However, they abstract from the fact that the Owicki-Gries system does not directly support reasoning on parameterized programs. An extension of the system and a proof of completeness was missing. The authors of [de Roever *et al.*, 2000] say

Since Szymanski’s algorithm³ is parameterized, it is quite natural to use assertion networks parameterized by a process index i .

In the present work we provide a formal proof that supports this “natural-ity.” The main conclusion is that, for any valid specification of a parameterized program, it is always possible to find a single (parameterized) proof outline that can be derived in the system F for all values of n .

The proof is organized as follows. In §5.2.1 we prove the semantic existence of the intermediate assertions that form a proof outline. This means that we show in our meta-language (English and mathematics) that these assertions exist (so-called *semantic completeness*). Further, we prove that these so constructed proof outlines yield valid F -derivations since they satisfy the expected properties for using the method. In §5.3 we show that these (semantic) assertions can be expressed in a language containing at least first order arithmetic over the standard model of the natural numbers (elementary arithmetic). That is, we prove that syntactic expressions corresponding to the semantic ones exist. This result is called *expressiveness*, which together with the semantic completeness implies *relative completeness*.

³A mutual exclusion algorithm for distributed processes.

5.2.1 Semantic Completeness

The proof proceeds by induction on the structure of the programs. The case for parallel programs of the form $S_0 \parallel \dots \parallel S_k$ where the number of components is fixed was first proven in [Owicki, 1975].

The remaining case is that of parameterized parallel programs. Given a valid partial correctness formula of the form

$$\models \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$$

we prove that we can always find a proof outline for S such that

$$\vdash_F \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$$

Given a parameterized (non-annotated) component program $S(i, n)$, we construct its proof outline by associating an assertion with every point of interference. For this purpose, it suffices to annotate $S(i, n)$ with a postcondition $q(i, n)$ and every normal subprogram $R(i, n)$ of $S(i, n)$ with a precondition $pre(R(i, n))$, where a *normal subprogram* of S is a subprogram which is not a proper subprogram of an await-statement. The result is an annotated program $A(S(i, n))$. We define $pre(A(S(i, n)))$ and $post(A(S(i, n)))$ to be the initial precondition and the final postcondition, respectively.

To prove correctness of parameterized parallel programs using the system F , we need to show local correctness and interference freedom of the associated proof outlines. To this end, one must in general augment actions and boolean guards with assignments to auxiliary variables. These record the progress of computation in the other processes allowing us to express the effects of parallel execution.

There are two well-known canonical methods of introducing auxiliary variables (see 2.4.4). One makes use of a single so-called *history* variable z that records, throughout execution, the index number of the component performing the atomic action and the values of the program variables before that action is executed. In a terminating state, z contains the history associated with the executions of assignments or await-statements. The second possibility consists of using auxiliary variables as labels indicating the control points that the program passes through. Thus, these special auxiliary variables (which are also called counters) are updated with *every* transition. One auxiliary variable of this form for each component suffices for any Owicki-Gries style proof [Lamport, 1977, Best, 1996].

Our programming language does not allow for the atomic evaluation of a boolean condition together with the updating of an auxiliary variable. As a

consequence, not all transitions, but only assignments and await-statements, can be recorded by auxiliary variables. For this reason we use a single history variable z , respecting this syntactic restriction, in the style of Apt [Apt, 1981a]. The set of program variables is augmented with a single variable z , containing the value corresponding to the encoding of the substring of the history sequence, where transitions corresponding to the evaluation of boolean expressions are not registered.

Given a program S with state space Σ , we denote the state space of the extended program S^* by Σ^* , where we assume that the first value of each tuple corresponds to the value of the history variable. For instance, if $s^* \in \Sigma^*$ is a state of S^* , then, by convention, s denotes the corresponding state of Σ , where the first element of the tuple has been removed.

Given $\models \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$, the proof of semantic completeness is structured as follows:

1. Extend the program by initializing the history variable:

$$z := \square; \parallel_{i=0}^{n-1} S(i, n)$$

2. Extend $S(i, n)$ to $S^*(i, n)$ by adding assignments to the history variable.
3. Annotate every point of interference in $S^*(i, n)$ with an assertion, obtaining an annotated program $A(S^*(i, n))$ such that the following conditions hold:

- (a) $\forall s \in \Sigma. s \in P(n) \longrightarrow (\square, s) \in \bigcap_{i=0}^{n-1} \text{pre}(A(S^*(i, n)))$
- (b) *Local correctness of proof outlines:* $\forall i < n. \vdash A(S^*(i, n))$
- (c) *Interference freedom:* $\forall i, j \in \{0, \dots, n-1\}$. If $i \neq j$ then, for every assertion $r_k(i, n)$ in $A(S^*(i, n))$, and for every atomic action $a(j, n)$ with precondition $\text{pre}(a(j, n))$ in $A(S^*(j, n))$, the formula

$$\vdash \{r_k(i, n) \cap \text{pre}(a(j, n))\} a(j, n) \{r_k(i, n)\}$$

holds.

- (d) $\forall s^* \in \Sigma^*. s^* \in \bigcap_{i=0}^{n-1} \text{post}(A(S^*(i, n))) \longrightarrow s \in Q(n)$.

Then, using the rule for parameterized parallel programs, we can conclude that $\vdash_F \{P(n)\} z := \square; \parallel_{i=0}^{n-1} A(S^*(i, n)) \{Q(n)\}$ holds. By deleting the

assignments to auxiliary variables with the elimination rule (see 2.4.4) and omitting the intermediate annotations we obtain

$$\vdash_F \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$$

Note that with the system O we could prove the following statement

$$\begin{aligned} \forall n. (\models \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\} \longrightarrow \\ \exists A^n. \vdash_O \{P(n)\} \parallel_{i=0}^{n-1} A^n(S^*(i, n)) \{Q(n)\}) \end{aligned}$$

where for different values m and l of n , A^m can be very different from A^l . In this paper, however, we prove the existence of a uniform annotation (parameterized by n and i) for all values of n , i.e. we prove

$$\begin{aligned} \models \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\} \longrightarrow \\ \exists A. \vdash_F \{P(n)\} \parallel_{i=0}^{n-1} A(S^*(i, n)) \{Q(n)\} \end{aligned}$$

5.2.2 Extending the Program

Let $V = \{x_0, \dots, x_k\}$ be the set of variables occurring in $S(i, n)$. Let \bar{x} denote the coding $[x_0, \dots, x_k]$. Let z be a new variable. Then, we transform $S(i, n)$ into $S^*(i, n)$ by replacing

- (i) every await-statement **await** b **then** R **end**, where R is not *skip*, by **await** b **then** $z := z \cap [i, \bar{x}]; R$ **end**,
- (ii) every assignment $y := t$ not inside an **await**-statement by **await** *true* **then** $z := z \cap [i, \bar{x}]; y := t$ **end**.

If the language allows for assignments to variables at every transition (including those consisting of evaluation of boolean conditions), completeness can be proven by extending the program with assignments to counter variables. In a parameterized program the counter variable would be a composed variable c . The i th component program would update this variable by updating its i th component, i.e. $c[i]$, at every transition. As we shall see in §5.5, parallel programs are in practice verified using auxiliary variables as counters because finding assertions in terms of history variables is too difficult.

5.2.3 Annotating the Program

Given $S^*(i, n)$, our goal is to construct intermediate assertions⁴ that yield a proof outline such that conditions (a), (b), (c) and (d) of §5.2.1 can be proven.

By the completeness of O we know that given

$$\models \{P(n)\} \parallel_{i=0}^{n-1} S(i, n) \{Q(n)\}$$

we can find for every value m of n , i.e. for the particular valid triple

$$\models \{P(m)\} S(1, m) \parallel \dots \parallel S(m, m) \{Q(m)\},$$

an extended program

$$z := \square; (S^*(1, m) \parallel \dots \parallel S^*(m, m))$$

and an annotation A^m so that

$$\vdash \{P(m)\} z := \square; \parallel_{i=0}^{m-1} A^m(S^*(i, m)) \{Q(m)\}$$

A^m associates every point of interference j inside a component i with an assertion (set of states) r_{ji}^m whose construction is described in the completeness proof for the system O [Owicki, 1975, Apt, 1981a, de Roever *et al.*, 2000]. Informally, these assertions are defined as the sets of states that are reachable by some computation starting in a state in the precondition, where the steps of the computation are determined by the rules of the operational semantics.

The assertion associated with a location j in a parameterized component program $S^*(i, n)$, represented by $r_j(i, n)$, is defined as follows:

$$r_j(i, n) = r_{ji}^n$$

where for each value m of n and for each value l of i , the set r_{jl}^m is the one constructed in the completeness proof of the Owicki-Gries system for the concrete parallel program $\parallel_{i=0}^{m-1} S(i, m)$ at the location j inside component l .

It remains to be proven that with this annotation the conditions necessary to derive the corresponding triple using the system F are satisfied. For this purpose, it suffices to prove that conditions (a), (b), (c) and (d) mentioned above hold for this annotation. This is easy, since for any arbitrary value m of n , the proof is reduced to the case for the concrete program with m parallel components. We just show the proof of (a), the others are similar:

⁴By an assertion we mean here a set of states as a semantic object.

- (a) Denote the first location (the one corresponding to the precondition) of a component by 0. Assume that there is a state s such that $s \in P(m)$. By the completeness of the system O we have $(\square, s) \in \bigcap_{i=0}^{m-1} r_{0i}^m$. By definition of $r_0(i, m)$, we have $(\square, s) \in \bigcap_{i=0}^{m-1} r_0(i, m)$.

This closes our proof of semantic completeness of the system F . By defining the intermediate assertions as parameterized functions that for each n and each i return the set of states that characterize the assertions of the particular program with those values of n and i , we have proven that these sets of states exist for every n and every i . The main goal, however, is to prove that these sets not only exist but can also be concretely described using some concrete assertion language.

5.3 Relative Completeness

In this section we prove that the semantic sets described in §5.2.1 can be syntactically expressed if we use an assertion language which contains at least the logical system of elementary arithmetic and consider programs whose boolean conditions and state transformations can be expressed in the assertion language. For this purpose we use an important result of recursion theory [Rogers, 1987]:

For any relation R , if R defines a recursively enumerable set, then R is definable in elementary arithmetic.

Thus, it suffices to prove that the sets of states $r_j(i, n)$ that we constructed in §5.2.1 are recursively enumerable sets. Intuitively, a set S is recursively enumerable (or checkable) if there exists a procedure M such that given an element t as input, “checks” whether t is in S and stops when its checking procedure succeeds. If t is not in S the procedure continues checking forever. We show that such a procedure for checking the sets $r_j(i, n)$ exists.

Let $\parallel_{i=0}^{n-1} S(i, n)$ be a parameterized parallel program and let $P(n)$ be a set of initial states. For every value m of n and for every control point j in component i , the set r_{ji}^m is recursively enumerable. We denote the corresponding checking procedure by M_{ji}^m . Informally, it works as follows:

Given an input state s , $M_{ji}^m(s)$ succeeds iff it is possible to start an execution of $S^*(1, m) \parallel \dots \parallel S^*(m, m)$ from a state in $P(m)$ and reach the control point j of component i with the values of the variables as given by the state s .

We now define the procedure M_j that checks if a state s is in $r_j(i, n)$:

It takes as arguments n, i and the state s . With the values of n and i , it builds the concrete program $S^*(1, n) \parallel \dots \parallel S^*(n, n)$ and simulates the procedure M_{ji}^n on state s . Then, $M_j(n, i, s)$ is defined such that it succeeds iff $M_{ji}^n(s)$ succeeds.

Hence, the domain of the procedure M_j (the values for which it succeeds) is a recursively enumerable set. By the theorem mentioned above, the relation that defines this set is definable in elementary arithmetic. Let us recall what this means [Rogers, 1987]:

An n -ary relation R is *definable in elementary arithmetic* if there is a formula $F a_1 \cdots a_n$ with free variables a_1, \dots, a_n such that

$$R = \{(x_1, \dots, x_n) \mid F x_1 \cdots x_n\}$$

where, for any integers x_1, \dots, x_n , $F x_1 \cdots x_n$ is the result of substituting the numeral expressions of x_1, \dots, x_n for a_1, \dots, a_n in $F a_1 \cdots a_n$.

The tuples of our set are formed by the values of n, i and those of the program variables that form the state. Consequently, the formula defining this set has as free variables not only the program variables (such as for concrete programs) but also n and i , as intuitively expected.

In the ticket algorithm, for example, the arithmetical expression defining some assertion would have as free variables, $n, i, num, next$ and $turn$ (all ranging over \mathbb{N}). Composed variables are considered simple variables with value the encoding of the values of the components. Thus, in the first-order logic expressions of the assertions there are no references to the components of $turn$. However, to reason practically in Hoare style, more expressive logics are used (e.g. HOL). If these logics allow us to express composed variables as functions, then expressions like $turn[i]$ can be written directly.

This completeness proof does not provide an effective methodology for finding the expressions for assertions. The important result is that these expressions always exist. This means that it is always possible to find parameterized proof outlines for valid parameterized parallel programs.

5.4 Completeness of the Rely-Guarantee System for Parameterized Parallel Programs

The completeness proof of the rely-guarantee system for this particular kind of programs follows by the same reasoning as in the previous section. The completeness proof for the standard rely-guarantee system is, like for the Owicki-Gries system, based on the definition of intermediate annotations as strongest postconditions, but where the environment is also taken into account.

Given a rely-guarantee specification $\vdash P \text{ SAT } [pre, rely, guar, post]$ for a parallel program P , we associate with every point of interference l in P the corresponding strongest postcondition SP_l . Following [de Roever *et al.*, 2000], a state s belongs to SP_l if there is a computation of P together with its environment that reaches location l of P , starting in a state satisfying pre , such that all environment steps satisfy $rely$. Similarly, the so-called strongest guarantee condition SG is defined as the set of pairs of states defining transitions of P which are executed by P in some computation, provided pre is satisfied initially, and every environment transition satisfies $rely$.

The proofs of completeness for the rely-guarantee system presented in [Xu *et al.*, 1997] and [de Roever *et al.*, 2000] are essentially based on these constructions. History variables as auxiliary variables are also needed for the proofs.

In [Stirling, 1988], the author presents another proof of completeness for his version of the rely-guarantee system based on a very elegant theory of invariants. The idea is to prove that the rely-guarantee system is complete with respect to the Owicki-Gries system. Intuitively, whenever we have a derivation for a correct program in the Owicki-Gries system, we can find a derivation in the rely-guarantee system by considering the *rely*-condition of a component program P_i to be the predicate describing the set of pairs of states which preserve the annotations of its own proof outline. Thus, no environment transition can falsify the local correctness of the proof outline. For the *guar*-condition it suffices to take the predicate characterizing the set of pairs of states which preserve the annotations in the proof outlines of the other components P_j with $i \neq j$. Basically, a rely-guarantee specification can be obtained from an Owicki-Gries one by registering in the rely and guarantee conditions the information provided by the interference freedom test. We illustrate this idea with a simple example in §5.5.

In any case, proofs of completeness have already been studied in previous

works and we are just interested in assuming that given any valid rely-guarantee specification of a parallel program it is possible to construct the subspecifications of its components so that the program can be derived in the system.

The introduction of auxiliary variables for a parameterized parallel program as well as the proof of semantic and relative completeness presented in the previous section for the Owicki-Gries system are independent of the particularities of the Owicki-Gries system itself and can be directly used for the rely-guarantee method.

Consequently, for any valid rely-guarantee specification of a parameterized parallel program it is possible to find a derivation in the rely-guarantee system. Moreover, the assertions needed for it are by construction recursively enumerable and can thus be expressed in elementary arithmetic.

5.5 Example

There are several examples of assertional verification of parameterized parallel programs in the literature [Stirling, 1988, de Roever *et al.*, 2000, Best, 1996]. Some of them have already been studied in the examples of the previous chapters. Here we consider a very simple program that illustrates the practical meaning of the completeness result presented in this chapter.

The example we choose is the classic program $\{x = 0\} \parallel_{i=0}^{n-1} x := x + 1 \{x = n\}$. Its verification with the rely-guarantee method has already been studied in §4.7.2. We show here its verification in the Owicki-Gries system and compare both results.

We introduce a composed variable c , such that the component i of the parallel program atomically updates the shared variable x and the i th component of c . We model the composed variable c as a function from naturals to naturals. The syntax for updating the value of a function f on argument i is $f(i:=t)$, where t is of the corresponding type. The summation function used in the assertions is predefined in the Isabelle library.

```
record Schema =
  x :: nat
  c :: nat ⇒ nat
```

```
lemma Schema: 0 < n ⇒
  ⊢ {x=0 ∧ (∑ i < n. c i)=0}
cobegin
```

```

scheme  $[0 \leq i < n]$ 
 $\{x = (\sum_{i < n} c\ i) \wedge c\ i = 0\}$ 
 $\langle x := x + 1, c := c\ (i := 1) \rangle$ 
 $\{x = (\sum_{i < n} c\ i) \wedge c\ i = 1\}$ 
coend
 $\{x = n\}$ 

```

The completeness result for parameterized parallel programs claims that the assertions can be expressed in the language of elementary arithmetic with free variables n , i , x and c . For obvious practical reasons, we prefer to use a higher order language. However, projection and finite summation are both partial recursive functions and as such can be expressed in a first order arithmetic.

As proved in §4.7.2, the correctness of the same program with respect to the rely-guarantee specification

$$(x = 0 \wedge (\sum_{i < n} c\ i) = 0, x = \bar{x} \wedge c = \bar{c}, \text{True}, x = n)$$

where the overall rely and guarantee conditions specify the program as being closed, requires that each component i satisfy the following (parameterized) local specification:

```

pre :  $x = \sum_{i < n} c\ i \wedge c\ i = 0$ 
rely :  $c\ i = \bar{c}\ i \wedge$ 
 $(x = \sum_{i < n} c\ i \longrightarrow \bar{x} = \sum_{i < n} \bar{c}\ i)$ 
guar :  $(\forall j < n. i \neq j \longrightarrow c\ j = \bar{c}\ j) \wedge$ 
 $(x = \sum_{i < n} c\ i \longrightarrow \bar{x} = \sum_{i < n} \bar{c}\ i)$ 
post :  $x = \sum_{i < n} c\ i \wedge c\ i = 1$ 

```

As observed by [Stirling, 1988], there is a direct connection between the interference freedom tests required in the Owicki-Gries proof of the same program and the rely and guarantee conditions required here.

For component i the interference freedom test is successful if

1. The truth of the equality $x = \sum_{i < n} c\ i$ is preserved and the local variable $c\ i$ is not modified by the actions of the other components.
2. The atomic actions in i do not affect the truth of the assertions in the other components.

The first part is exactly expressed in the rely condition, which represents what the component assumes from the environment. And the second one

is reflected on the guarantee condition that represents what the component ensures the environment. In this case, since the program is closed, the environment is represented by the rest of the components.

Observe that for the formulation of the rely and guarantee conditions, no assumption on the concrete implementation of the environment is made. In the Owicki-Gries formalism, however, we need to know the concrete form of the environment to be able to carry out the interference freedom tests.

This example illustrates the intuitive idea that the rely-guarantee method is complete with respect to the Owicki-Gries method, and that this also applies to parameterized programs.

5.6 Summary

The Owicki-Gries and the rely-guarantee systems are the most fundamental assertional methods for reasoning about the correctness of parallel programs with shared variables. However, they were not designed for the verification of parameterized parallel programs, but for programs with a given number of components only. In the previous chapters we have shown that the systems can be generalized to handle parameterized parallel programs and that the resulting systems are sound.

In this chapter we have studied the completeness of the extended systems. The main difficulty, however, lies in a satisfactory formalization of the problem. To this end, we have defined the syntax and meaning of parameterized parallel programs and have formalized what it means for these programs to be verified in a Hoare logic framework.

The proofs and ideas presented in this chapter have been developed in an abstract mathematical style. Hence, they are independent from the particular formalization of the systems in Isabelle. Moreover, the starting point is the assumption that the proof systems found in the literature are complete. This assumption allows us to state that for each instance of a valid specification of a parameterized parallel program there exists a proof in the system. The main result presented in this chapter, namely, the completeness of the extended systems, states that it is always possible to find a *single* proof that works for any number of components. Not surprisingly (but not obviously) the assertions that allow us to derive such proofs are themselves parameterized. Then, the proof for some particular program can be obtained by instantiating the parameters in the assertions.

Chapter 6

Conclusion

This thesis describes the first formalization of the Owicki-Gries and rely-guarantee methods in a theorem prover. Both methods are used to prove partial correctness of parallel imperative programs with shared variables. For each method, the programming language, the operational semantics and the proof rules have been defined in Isabelle/HOL. The proof of soundness of the rules w.r.t. the underlying semantics has also been carried out with the theorem prover. As a result, we obtain a verified framework for proving the correctness of parallel programs with shared variables having the choice between a compositional and a non-compositional axiomatic method.

The theorem prover Isabelle/HOL has provided a user-friendly, reliable and powerful tool that ensures a systematic, correct and understandable development of the theory underlying both verification systems. We have shown that a theorem prover like Isabelle also provides a powerful tool for the application of these methods to real programs. In particular, by considering programs like the parallel garbage collection algorithms we have gone beyond typical “toy” examples. Moreover, since we had no previous (complete) pencil and paper Owicki-Gries proofs for these examples, the tool was useful not only for the a “a posteriori” verification, but also as a checking friend in the search for a proof.

6.1 General Contributions

Apart from this concrete verification environment this thesis provides contributions which are of general interest for the areas of parallel program verification and mechanical formalization of programming languages.

Parameterized Parallel Programs

The formalized systems differ from those found in the literature mainly in that the rule for parallel composition has been generalized to handle parameterized parallel programs. Thus, it is possible to prove correctness of such systems for any instance of the parameter by finding just one derivation. This naturally led to the question whether finding such a derivation was always possible, i.e. whether the systems are complete for these kind of programs.

Although proofs of completeness for the traditional Owicki-Gries and rely-guarantee systems already exist in the literature, none of them considers the question whether parameterized parallel programs can always be verified in a schematic way. From the completeness results of the original systems we know that there is a derivation in the system for each particular instance of the parameter, however, a machine that checks derivations of the systems for all instances would never stop. In contrast, the completeness of our system demonstrates that it is always possible to find a single parameterized derivation of these programs, so that *one* correctness proof is valid for any instance of the parameter.

Concrete Representation of Programs

Isabelle also provides concrete syntax facilities which makes it possible to present programs in a familiar syntax. To achieve this, the main concern is the representation of program variables. In this thesis we have used the *quote/antiquote* technique. This technique is widely used in the functional programming world, but the idea of using it for the purpose of encoding program variables was first proposed by Markus Wenzel in [Wenzel, 2001b]. These techniques represent the most advanced “technology” in the state of the art of program variables representation in HOL: they provide a combination of the advantages of previously used methods without introducing any drawbacks. The examples shown in this dissertation represent the largest application so far and have contributed to reveal many of the advantages of this model.

Automation of the verification process

With the theorem prover we can apply the verification methods systematically ensuring that no mistakes are made and no details ignored. The verification process is made considerably less tedious because the boring and in general easy routine steps involved in such endeavors are taken care

of. This allows the user to concentrate on the fundamental aspects of the proof which require intelligent input.

The main disadvantage of using a theorem prover for the verification of programs is that they generally require extensive human guidance, and that this guidance is expressed in terms of the particular theorem prover. This problem can be partially solved by designing a tactic that automatically generates the verification conditions. Such an automatic tactic has been defined for the Owicki-Gries method and it has been successfully applied to the verification of several non-trivial examples. Depending on the complexity of these verifications conditions, proving them could require certain knowledge of the proving techniques. Nevertheless, once the verification conditions are generated and simplified it is easier to see whether they hold or not.

6.2 Statistics of the formalization

We give an overview of the amount of work carried out with Isabelle. Table 6.1 shows the number of specification lines (types, definitions, inductive sets, etc.), the number of lemmas and the total number of proof steps needed for the formalization of the theory underlying the verification methods.

Theory	Spec. lines	Lemmas	Proof steps
Owicki-Gries	220	49	340
Rely-guarantee	330	93	2240
VCG-tactic	190	42	80
Concrete syntax	260	0	0
Total	1000	184	2660

Table 6.1: Statistics of the formalizations.

We observe that the formalization of the rely-guarantee method is more involved and, in particular, the proofs are longer. This is the price we pay in obtaining a compositional method: the underlying theory requires more work, but yields a simpler proof system. The soundness proof only needs to be done once, however, we enjoy the advantages of a compositional system with every verification exercise.

The tactic for the automatic generation of the verification conditions requires a number of rules derived as lemmas from the Owicki-Gries theory. The 190 specification lines correspond to the ML code implementing the tactic.

The concrete syntax for both formalizations involves 260 lines of specification which include the declaration of syntactic constants, the translation equations and the parse and print translation functions programmed in ML. Many of the features are duplicated in both formalizations.

One of the main goals of this work is to demonstrate the applicability of the formalizations on the verification of real programs. We briefly review the results of our experience.

Table 6.2 shows the statistics of the verification of the typical examples done with the Owicki-Gries formalization.

Algorithm	Verif. cond.	Automatic	Lemmas
Peterson	122	122	0
Dijkstra	20	20	0
Ticket	35	24	0
Zero search (with sync.)	98	98	0
Zero search (without sync.)	20	20	0
Producer/Consumer	138	125	3
Total	433	409	3

Table 6.2: Examples verified with Owicki-Gries.

The next table shows the statistics concerning the verification of the two garbage collection (gc) algorithms and the theory of graphs common to both algorithms.

Theory	Spec. lines	Lemmas	Verif. cond	Proof steps
Graph	23	17	-	289
Single-mut. gc	35	28	289	408
Multi-mut. gc	35	36	328	756
Total	93	81	617	1453

Table 6.3: Verification of garbage collection algorithms.

Finally, table 6.4 gives an overview of the effort involved in verifying the examples considered for the rely-guarantee method. There is no tactic for the automatic generation of verification conditions, thus the proof steps also contain the number of steps needed in order to apply the rules from the rely-guarantee system.

Algorithm	Verif. cond.	Lemmas	Proof steps
Array to zero	8	3	40
Increment variable	14	3	23
Find least element	22	1	30
Total	44	7	93

Table 6.4: Examples verified with rely-guarantee.

6.3 Further Work

The present work represents a first basic study of verification methods for parallel programs with shared variables in a theorem prover. Several possible extensions from both the mechanical and the theoretical sides of the formalization can be addressed:

1. The completeness proofs for both systems are known theoretical results. It would be interesting to formalize them in the theorem prover especially because our systems also handle the verification of parameterized parallel programs. To obtain a complete system, however, a missing rule in the present formalization, namely the rule for elimination of auxiliary variables, should also be formalized.
2. The programming language considered for both formalizations is fairly simple. The enrichment with nested parallelism would be a meaningful extension, especially for the rely-guarantee system. For the Owicki-Gries method the lack of this feature is not a major disadvantage because all the processes involved in a parallel program have to be known prior to verification. However, the compositionality of the rely-guarantee system makes the method suitable for top-down development of programs. This is important to cope with the design and verification of large programs. Also, to automate the application of the rely-guarantee method, a tactic for the automatic generation of verification conditions similar to the one presented for the Owicki-Gries method should be designed. However, the presentation of programs should be extended to include intermediate annotations in order to avoid the use of the consequence rule at intermediate steps.
3. It would also be interesting to investigate the relation between proofs in both the compositional and non-compositional methods. Especially, to study whether Owicki-Gries proofs can be translated into rely-guarantee proofs following some systematic procedure.

6.4 Experience

By formalizing these methods in Isabelle/HOL we have gained a full understanding of the difficulty involved in designing correct proof methods for the verification of (parallel) programs. The level of detail required for such a formal work naturally leads to approach each step in the formalization with a critical eye, considering first what might come next and studying alternatives that could simplify the formalization. This often leads to improvements over the known definitions and methods found in the literature. However, to understand the theorem proving techniques and learn to optimize the effort involved in any formalization takes a great deal of time, effort and mistakes.

Verifying parallel programs using the methods formalized here can be difficult, but it is mostly tedious work. The interesting part of the verification process is to understand the program and find an intuitive proof of their correctness. However, the projection of this intuition into assertions that have to satisfy a number of requirements (like interference freedom!) implies, in general, changing and tuning the assertions too many times and a great deal of effort is expended in getting details right.

The tool presented here does not directly help in finding the right annotations but at least automates the iterative process of changing assertions and checking the proof again. One thing is sure, the theorem prover will never let the user get away with wrong or incomplete annotations. We believed to have found the perfect annotations many times only to then have the theorem prover reveal subtle (and sometimes obvious) mistakes. From our experience, we do not recommend implicitly trusting a pencil and paper proof correctness for parallel programs.

6.5 To conclude

The main point of this thesis is that, by formalizing (generalized versions) of the two well-known Owicki-Gries and rely-guarantee methods in a theorem prover, it is now possible to obtain mechanized proofs of correctness for general (parameterized and non-parameterized) parallel programs with these methods.

Appendix A

Automatic Generation of Verification Conditions

We construct a tactic for the automatic generation of the verification conditions (vcg) for full specifications of parallel programs. The construction is presented in a bottom-up style. First, a vcg-tactic is designed for atomic programs. This makes understanding the construction easier, since atomic programs are the simplest ones used here. Using this tactic another vcg-tactic is developed for annotated programs. This second vcg-tactic is employed in the construction of the final vcg-tactic, which is used to generate the verification conditions of parallel programs.

A.1 VCG for Atomic Programs

Atomic programs are sequential programs and do not contain intermediate annotations. The vcg-tactic works by deriving all intermediate assertions from the postcondition and loop invariants and establishing the corresponding verification conditions. These are expressed simply as set theoretic inclusions because assertions and boolean conditions are formalized as sets of states.

The generated intermediate assertions are generally called *weakest liberal preconditions* [Winskel, 1993, Dijkstra, 1976]. Informally, given a command c and its postcondition q , the weakest liberal precondition $wlp(c, q)$ is defined as the set of states from which the execution of c either diverges, or ends up in a final state satisfying q .

Given a program and its specification, the verification conditions are

obtained by recursively applying the proof rules backwards until only statements of the logic of assertions, without any mention of the programming language, remain. To this end, the proof rules have to be generalized because the rules defined in the proof theory cannot always be applied backwards directly. The pre- and postconditions in most of the conclusions are too specific and need to be previously manipulated via the rule of consequence. Consider for example the following specification of an atomic program:

lemma $\vdash \{x. x = 0\} \text{ Basic } (\lambda x. x+1) \{x. x = 1\}$

The rule *Basic* is too specific to be applied. The consequence rule comes to the rescue

apply (*rule Conseq*)

generating the following three subgoals:

1. $\{x. x = 0\} \subseteq ?p$
2. $\vdash ?p \text{ Basic } (\lambda x. x + 1) ?q$
3. $?q \subseteq \{x. x = 1\}$

where *?p* and *?q* are schematic variables that can be instantiated with anything of the corresponding type. By applying the rule *Basic* backwards on the second subgoal, the schematic variables *?p* and *?q* are automatically instantiated with the pre- and postcondition required by the rule:

prefer 2

apply (*rule Basic*)

Two verification conditions are left, namely

1. $\{x. x = 0\} \subseteq \{s. s + 1 \in ?q\}$
2. $?q \subseteq \{x. x = 1\}$

which are automatically solved by the tactic *auto*:

apply *auto*

done

In general, it is impossible for an automatic tactic to guess when and how many times the rule of consequence should be used to obtain suitable assertions. Instead, we derive more general rules by combining the original proof rules for each command with the rule of consequence:

lemma *SkipRule*: $p \subseteq q \implies \Vdash p \text{ (Basic id) } q$

lemma *BasicRule*: $p \subseteq \{s. (f s) \in q\} \implies \Vdash p \text{ (Basic f) } q$

lemma *SeqRule*: $\llbracket \Vdash p \ c_1 \ r; \Vdash r \ c_2 \ q \rrbracket \implies \Vdash p \text{ (Seq } c_1 \ c_2) \ q$

lemma *CondRule*:

$\llbracket p \subseteq \{s. (s \in b \longrightarrow s \in w) \wedge (s \notin b \longrightarrow s \in w')\}; \Vdash w \ c_1 \ q; \Vdash w' \ c_2 \ q \rrbracket$
 $\implies \Vdash p \text{ (Cond } b \ c_1 \ c_2) \ q$

lemma *WhileRule*: $\llbracket p \subseteq i; \Vdash (i \cap b) \ c \ i; (i \cap -b) \subseteq q \rrbracket$

$\implies \Vdash p \text{ (While } b \ i \ c) \ q$

These derived rules can always be applied. Only the rule for sequential composition remains the same. Maybe surprising is the case of the conditional statement because the rule of consequence affects the premises and not the conclusion. The original rule

$$\llbracket \Vdash (p \cap b) \ c_1 \ q; \Vdash (p \cap -b) \ c_2 \ q \rrbracket \implies \Vdash p \text{ (Cond } b \ c_1 \ c_2) \ q$$

matches any pre- and postcondition, however, the generated subgoals easily fail. For example, consider the triple

$$p \text{ (Cond True (Basic id) (Basic id)) } p.$$

Although it is trivially valid, the original rule *Cond* would fail.

In [Winskel, 1993], this problem is solved by generating the verification conditions from programs that have been previously annotated, not only with loop invariants, but also with suitable assertions before any if- or while-statement. Our approach, however, reduces the required human guidance to loop invariants.

To verify sequential programs we proceed “bottom-up”, i.e. starting by the last non-sequential command. Hence, a function called *WlpTac* recursively decomposes sequential constructions¹ until it reaches the last non-sequential command whose postcondition is known. Using the derived proof rules, the tactic generates the weakest (liberal) precondition.

The tactic controls via the parameter *precond* if the precondition is unknown. If so, the generated verification condition has the form $?p \subseteq \dots$, where $?p$ represents the unknown precondition which is subsequently instantiated by reflexivity, i.e. using the theorem $A \subseteq A$. Initially, the only known preconditions are those supplied by the user, namely the overall precondition and loop invariants. Thus, when the tactic is invoked *precond* is true; it becomes false whenever the precondition must be worked out from the postcondition.

¹The sequential operator associates to the right, i.e. $c_0;; (c_1;; c_2)$.

The vcg-tactic for atomic programs is called *HoareRuleTac* and is programmed in ML. The tacticals *THEN*, *ORELSE*, *EVERY* and *FIRST* provide control structures for combining tactics. *THEN* executes one after the other and *ORELSE* tries first one and, if it fails, it tries the second one. *EVERY* and *FIRST* are the corresponding block versions of *THEN* and *ORELSE*, respectively. The tactic *rtac* takes two arguments, a theorem and a (subgoal) number, then it applies the given theorem backwards to the subgoal at the given number. To access Isabelle theorems inside an **ML** environment, they have to be preceded by the word *thm*.

```

ML {*
fun WlpTac i = rtac (thm SeqRule) i THEN HoareRuleTac false (i+1)
and HoareRuleTac precondition i =
  WlpTac i THEN HoareRuleTac precondition i
  ORELSE
  FIRST[rtac (thm SkipRule) i,
        rtac (thm BasicRule) i,
        EVERY[rtac (thm CondRule) i,
              HoareRuleTac false (i+2),
              HoareRuleTac false (i+1)],
        EVERY[rtac (thm WhileRule) i,
              HoareRuleTac true (i+1)]]
  THEN (if precondition then (K all-tac i) else rtac (thm subset-refl) i)
*}

```

The tactic (*K all-tac*) leaves the subgoal unchanged.

A.2 VCG for Component Programs

The vcg-tactic for component programs is simpler because the rules of the system are already generic enough to be directly applied. Moreover, since all annotations must be provided in advance there is no need for finding weakest preconditions.

We only derive three new proof rules for special instances of the *AnnBasic* and the *AnnAwait* commands when the transformation performed on the state is the identity,

lemma *AnnSkipRule*: $r \subseteq q \implies \vdash (\text{AnnBasic } r \text{ id}) q$

lemma *AnnWaitRule*: $\llbracket r \cap b \subseteq q \rrbracket \implies \vdash (\text{AnnAwait } r \text{ b } (\text{Basic id})) q$

and for an *AnnAwait* command where the boolean condition is $\{s. \text{True}\}$:

lemma *AnnAtomRule*:

$\llbracket \text{atom-com } c; \vdash r \ c \ q \rrbracket \implies \vdash (\text{AnnAwait } r \ \{s. \text{True}\} \ c) \ q$

To refer to a rule of the proof system, for example *AnnBasic*, we have to write *oghoare-ann-hoare.AnnBasic*. This is quite long, so we introduce abbreviated names with the keyword *lemmas*:

lemmas *AnnBasic* = *oghoare-ann-hoare.AnnBasic*
lemmas *AnnSeq* = *oghoare-ann-hoare.AnnSeq*
lemmas *AnnCond₁* = *oghoare-ann-hoare.AnnCond₁*
lemmas *AnnCond₂* = *oghoare-ann-hoare.AnnCond₂*
lemmas *AnnWhile* = *oghoare-ann-hoare.AnnWhile*
lemmas *AnnAwait* = *oghoare-ann-hoare.AnnAwait*
lemmas *AnnConseq* = *oghoare-ann-hoare.AnnConseq*

The *vcg-tactic* is otherwise similar to the previous one.

ML $\{*$

fun AnnWlpTac i = rtac (thm AnnSeq) i THEN AnnHoareRuleTac (i+1)

and AnnHoareRuleTac i =

AnnWlpTac i THEN AnnHoareRuleTac i

ORELSE

FIRST[*rtac (thm AnnSkipRule) i,*
EVERY[*rtac (thm AnnAtomRule) i,*
HoareRuleTac true (i+1)],
rtac (thm AnnWaitRule) i,
rtac (thm AnnBasic) i,
EVERY[*rtac (thm AnnCond₁) i,*
AnnHoareRuleTac (i+3),
AnnHoareRuleTac (i+1)],
EVERY[*rtac (thm AnnCond₂) i,*
AnnHoareRuleTac (i+1)],
EVERY[*rtac (thm AnnWhile) i,*
AnnHoareRuleTac (i+2)],
EVERY[*rtac (thm AnnAwait) i,*
HoareRuleTac true (i+1)]]]

**}*

Although it is not necessary to generate the intermediate preconditions from the last postcondition, the tactic also processes the commands “bottom-up”,

i.e. the higher subgoal number first. This is necessary because, in general, the application of the proof rules backwards generates more subgoals.

Let us see what happens if the tactic starts by processing the lowest subgoal number. We illustrate the problem by an example,

lemma

$$\begin{aligned} &\vdash \text{AnnSeq} \\ &\quad (\text{AnnCond}_1 \{x. x = 0\} \{x. x = 0\} \\ &\quad\quad (\text{AnnBasic} \{x. x = 0\} (\lambda x. x+1)) \\ &\quad\quad (\text{AnnBasic} \{x. x = 0\} (\lambda x. 0))) \\ &\quad (\text{AnnBasic} \{x. x = 1\} (\lambda x. x+2)) \\ &\quad \{x. x = \text{Suc } 2\} \end{aligned}$$

First, we decompose the sequential composition,

apply(rule *AnnSeq*)

1. $\vdash \text{AnnCond}_1 \{x. x = 0\} \{x. x = 0\} (\text{AnnBasic} \{x. x = 0\} (\lambda x. x + 1))$
 $(\text{AnnBasic} \{x. x = 0\} (\lambda x. 0)) \text{ pre } (\text{AnnBasic} \{x. x = 1\} (\lambda x. x + 2))$
2. $\vdash \text{AnnBasic} \{x. x = 1\} (\lambda x. x + 2) \{x. x = \text{Suc } 2\}$

If we start by the first command of the sequential composition and apply the rule *AnnCond*₁, the second part of the program is moved 3 subgoals down:

apply(rule *AnnCond*₁)

1. $\{x. x = 0\} \cap \{x. x = 0\} \subseteq \text{pre } (\text{AnnBasic} \{x. x = 0\} (\lambda x. x + 1))$
2. $\vdash \text{AnnBasic} \{x. x = 0\} (\lambda x. x + 1) \text{ pre } (\text{AnnBasic} \{x. x = 1\} (\lambda x. x + 2))$
3. $\{x. x = 0\} \cap \neg \{x. x = 0\} \subseteq \text{pre } (\text{AnnBasic} \{x. x = 0\} (\lambda x. 0))$
4. $\vdash \text{AnnBasic} \{x. x = 0\} (\lambda x. 0) \text{ pre } (\text{AnnBasic} \{x. x = 1\} (\lambda x. x + 2))$
5. $\vdash \text{AnnBasic} \{x. x = 1\} (\lambda x. x + 2) \{x. x = \text{Suc } 2\}$

In general it is difficult to keep track of the number of subgoals that are generated. However, if we start processing the subgoals “bottom-up”, the subgoals with lower numbers remain where they are, no matter how many new subgoals are generated.

A.3 VCG for Parallel Programs

For the constructors shared between atomic and parallel programs the *vcg*-tactic is exactly like *HoareRuleTac*. We only have to add to this tactic the case where a command is a parallel composition.

First, we derive a rule where the pre- and postcondition of the *Parallel* rule are generalized:

lemma *ParallelConseqRule*:

$$\begin{aligned} & \llbracket p \subseteq (\bigcap i \in \{i. i < \text{length } Ts\}. \text{pre } (\text{the } (\text{com } (Ts!i))))); \\ & \quad \Vdash (\bigcap i \in \{i. i < \text{length } Ts\}. \text{pre } (\text{the } (\text{com } (Ts!i)))) \\ & \quad \quad (\text{Parallel } Ts) \\ & \quad (\bigcap i \in \{i. i < \text{length } Ts\}. \text{post } (Ts!i)); \\ & \quad (\bigcap i \in \{i. i < \text{length } Ts\}. \text{post } (Ts!i)) \subseteq q \rrbracket \Longrightarrow \Vdash p \text{ (Parallel } Ts) \text{ } q \end{aligned}$$

Then, the *Parallel* rule should be applied backwards on the subgoal that results from the second premise of this rule. As a result, two new subgoals, namely, the derivability of the component programs and their interference freedom, are generated. However, the first subgoal, which results from the first premise of the rule for parallel composition *Parallel*, is quite complicated because of the universal quantifier. We derive a variant of the rule *Parallel* that is more suitable for backwards application.

The predicate $[\vdash] Ts$ indicates if all elements in the list of specifications of component programs Ts are derivable:

$$\begin{aligned} \mathbf{constdefs} \text{ map-ann-hoare} &:: (\alpha \text{ ann-com-op} \times \alpha \text{ assn}) \text{ list} \Rightarrow \text{bool} \quad ([\vdash] -) \\ [\vdash] Ts &\equiv \forall i < \text{length } Ts. \exists c \text{ } q. Ts!i = (\text{Some } c, q) \wedge \vdash c \text{ } q \end{aligned}$$

This definition corresponds to the first premise of the *Parallel* rule. We substitute it in original proof rule:

$$\begin{aligned} \mathbf{lemma} \text{ ParallelRule} &: \llbracket [\vdash] Ts; \text{interfree } Ts \rrbracket \\ &\Longrightarrow \Vdash (\bigcap i \in \{i. i < \text{length } Ts\}. \text{pre } (\text{the } (\text{com } (Ts!i)))) \\ & \quad \quad \text{Parallel } Ts \\ & \quad (\bigcap i \in \{i. i < \text{length } Ts\}. \text{post } (Ts!i)) \end{aligned}$$

To automate the proof of $[\vdash] Ts$ we derive two rules that distinguish whether the list is empty or not, and whether it is a parameterized list of component programs:

lemma *MapAnnEmpty*: $[\vdash] []$

lemma *MapAnnList*: $\llbracket \vdash c \text{ } q; [\vdash] xs \rrbracket \Longrightarrow [\vdash] (\text{Some } c, q)\#xs$

lemma *MapAnnMap*:

$$\forall k. a \leq k \wedge k < b \longrightarrow \vdash (c \text{ } k) (Q \text{ } k) \Longrightarrow [\vdash] \text{map } (\lambda k. (\text{Some } (c \text{ } k), Q \text{ } k)) [a..b[]]$$

By using these rules we avoid dealing with the quantifier of the original *Parallel* rule. Observe that for the case of parameterized programs, if we

apply *MapAnnMap* backwards, eliminate the quantifier in the premise and “move” $a \leq k \wedge k < b$ to the assumptions, we obtain the subgoal $a \leq k \wedge k < b \implies \vdash c k Q k$ for an arbitrary but fixed value of k . Consequently, proving the derivability of a parameterized list of programs is reduced to proving derivability of an arbitrary but fixed component without need for induction.

Proving the predicate *interfree* *Ts* can also be optimized by new derived rules. First, we define a function such that given a component program x and a list of component programs xs , it checks if the assertions in x are preserved by all atomic actions in xs and if all assertions in xs are preserved by the atomic actions in x :

```
constdefs interfree-swap :: ( $\alpha$  ann-triple-op  $\times$   $\alpha$  ann-triple-op list)  $\Rightarrow$  bool
  interfree-swap  $\equiv$   $\lambda(x, xs). \forall y \in \text{set } xs. \text{interfree-aux } (\text{com } x, \text{post } x, \text{com } y)$ 
     $\wedge$  interfree-aux (com  $y$ , post  $y$ , com  $x$ )
```

With this function we derive proof rules for *interfree*:

lemma *interfree-Empty*: *interfree* []

lemma *interfree-List*:

[[*interfree-swap* (x, xs); *interfree* xs]] \implies *interfree* ($x\#xs$)

lemma *interfree-Map*:

$\forall i j. a \leq i \wedge i < b \wedge a \leq j \wedge j < b \wedge i \neq j \longrightarrow \text{interfree-aux } (c\ i, Q\ i, c\ j)$
 $\implies \text{interfree } (\text{map } (\lambda k. (c\ k, Q\ k)) [a..b()])$

The definitions of *interfree-swap* and *interfree-aux* can also be expressed as inference rules suitable for automation. For *interfree-swap* we derive the following three rules:

lemma *interfree-swap-Empty*: *interfree-swap* ($x, []$)

lemma *interfree-swap-List*:

[[*interfree-aux* (*com* x , *post* x , *com* y); *interfree-aux* (*com* y , *post* y , *com* x);
interfree-swap (x, xs)]] \implies *interfree-swap* ($x, y\#xs$)

lemma *interfree-swap-Map*:

$\forall k. a \leq k \wedge k < b \longrightarrow \text{interfree-aux } (\text{com } x, \text{post } x, c\ k)$
 $\wedge \text{interfree-aux } (c\ k, Q\ k, \text{com } x)$
 $\implies \text{interfree-swap } (x, \text{map } (\lambda k. (c\ k, Q\ k)) [a..b()])$

And for *interfree-aux* the next three ones:

lemma *interfree-aux-rule1*: *interfree-aux* (*co*, *q*, *None*)

lemma *interfree-aux-rule2*:

$\forall (R, r) \in (\text{atomics } a). \Vdash (q \cap R) r q \implies \text{interfree-aux } (None, q, \text{Some } a)$

lemma *interfree-aux-rule3*:

$\forall (R, r) \in (\text{atomics } a). \Vdash (q \cap R) r q \wedge (\forall p \in (\text{assertions } c). \Vdash (p \cap R) r p)$
 $\implies \text{interfree-aux } (\text{Some } c, q, \text{Some } a)$

The premises of the last two rules are, due to the universal quantifiers, not yet suitable for automation. The solution lies in proving a separate rule for all possible instances of the commands c and a in the conclusion. For each case, the assertions and atomic commands involved are known. The idea is to first extract the assertions that have to be checked for interference and then the atomic commands that should preserve those assertions. In the end, we obtain a subgoal for each of the Hoare triples involved in the interference freedom test.

The rules for isolating the assertions are:

lemma *AnnBasic-assertions*:

$\llbracket \text{interfree-aux } (None, r, \text{Some } a); \text{interfree-aux } (None, q, \text{Some } a) \rrbracket \implies$
 $\text{interfree-aux } (\text{Some } (\text{AnnBasic } r f), q, \text{Some } a)$

lemma *AnnSeq-assertions*:

$\llbracket \text{interfree-aux } (\text{Some } c_1, q, \text{Some } a); \text{interfree-aux } (\text{Some } c_2, q, \text{Some } a) \rrbracket \implies$
 $\text{interfree-aux } (\text{Some } (\text{AnnSeq } c_1 c_2), q, \text{Some } a)$

lemma *AnnCond₁-assertions*:

$\llbracket \text{interfree-aux } (None, r, \text{Some } a); \text{interfree-aux } (\text{Some } c_1, q, \text{Some } a);$
 $\text{interfree-aux } (\text{Some } c_2, q, \text{Some } a) \rrbracket \implies$
 $\text{interfree-aux } (\text{Some } (\text{AnnCond}_1 r b c_1 c_2), q, \text{Some } a)$

lemma *AnnCond₂-assertions*:

$\llbracket \text{interfree-aux } (None, r, \text{Some } a); \text{interfree-aux } (\text{Some } c, q, \text{Some } a) \rrbracket \implies$
 $\text{interfree-aux } (\text{Some } (\text{AnnCond}_2 r b c), q, \text{Some } a)$

lemma *AnnWhile-assertions*:

$\llbracket \text{interfree-aux } (None, r, \text{Some } a); \text{interfree-aux } (None, i, \text{Some } a);$
 $\text{interfree-aux } (\text{Some } c, q, \text{Some } a) \rrbracket \implies$
 $\text{interfree-aux } (\text{Some } (\text{AnnWhile } r b i c), q, \text{Some } a)$

lemma *AnnAwait-assertions*:

$\llbracket \text{interfree-aux } (None, r, \text{Some } a); \text{interfree-aux } (None, q, \text{Some } a) \rrbracket \implies$
 $\text{interfree-aux } (\text{Some } (\text{AnnAwait } r b c), q, \text{Some } a)$

By repeatedly applying these rules backwards only subgoals of the form *interfree-aux* ($None, r, \text{Some } a$) are left. That is, we eliminate quantification over assertions. Finally, the assertion r has to be checked for invariance under all atomic actions of a . This quantifier is “unfolded” by using rules that distinguish on the command a :

lemma *AnnBasic-atomics*:

$\Vdash (q \cap r) (Basic\ f)\ q \implies interfree\ aux\ (None, q, Some\ (AnnBasic\ r\ f))$

lemma *AnnSeq-atomics*:

$\llbracket interfree\ aux\ (c, q, Some\ a1); interfree\ aux\ (c, q, Some\ a2) \rrbracket \implies interfree\ aux\ (c, q, Some\ (AnnSeq\ a1\ a2))$

lemma *AnnCond₁-atomics*:

$\llbracket interfree\ aux\ (c, q, Some\ a1); interfree\ aux\ (c, q, Some\ a2) \rrbracket \implies interfree\ aux\ (c, q, Some\ (AnnCond_1\ r\ b\ a1\ a2))$

lemma *AnnCond₂-atomics*:

$interfree\ aux\ (c, q, Some\ a) \implies interfree\ aux\ (c, q, Some\ (AnnCond_2\ r\ b\ a))$

lemma *AnnWhile-atomics*:

$interfree\ aux\ (c, q, Some\ a) \implies interfree\ aux\ (c, q, Some\ (AnnWhile\ r\ b\ i\ a))$

lemma *Annatom-atomics*:

$\Vdash (q \cap r)\ a\ q \implies interfree\ aux\ (None, q, Some\ (AnnAwait\ r\ \{x.\ True\}\ a))$

lemma *AnnAwait-atomics*:

$\Vdash (q \cap (r \cap b))\ a\ q \implies interfree\ aux\ (None, q, Some\ (AnnAwait\ r\ b\ a))$

By repeatedly applying these rules backwards, only subgoals that state the derivability of a Hoare-triple remain. Then, the verification conditions are automatically generated with the `vcg-tactic` for atomic programs *HoareRuleTac* since the interference freedom tests involve only atomic actions.

The full tactic for the generation of the verification conditions of a parallel program is a combination of the tactics and rules above. Because of the interdependencies between them they must be defined together connected by the keyword *and*²:

ML `{*`

fun `WlpTac` *i* = *rtac* (*thm* `SeqRule`) *i* *THEN* `HoareRuleTac` `false` (*i*+1)

and `HoareRuleTac` *precond* *i* =

`WlpTac` *i* *THEN* `HoareRuleTac` *precond* *i*

ORELSE

`FIRST`[*rtac* (*thm* `SkipRule`) *i*,

rtac (*thm* `BasicRule`) *i*,

`EVERY`[*rtac* (*thm* `ParallelConseqRule`) *i*,

`ParallelTac` (*i*+1)],

`EVERY`[*rtac* (*thm* `CondRule`) *i*,

`HoareRuleTac` `false` (*i*+2),

²This tactic is a slight simplification of the real one which can be found in the source theories.

HoareRuleTac false (i+1)],
EVERY[rtac (thm WhileRule) i,
HoareRuleTac true (i+1)]]
THEN (if precondition then (K all-tac i) else rtac (thm subset-refl) i)

and AnnWlpTac i = rtac (thm AnnSeq) i
THEN AnnHoareRuleTac (i+1)

and AnnHoareRuleTac i =
AnnWlpTac i THEN AnnHoareRuleTac i
ORELSE
FIRST[rtac (thm AnnSkipRule) i,
EVERY[rtac (thm AnnAtomRule) i,
HoareRuleTac true (i+1)],
rtac (thm AnnWaitRule) i,
rtac (thm AnnBasic) i,
EVERY[rtac (thm AnnCond₁) i,
AnnHoareRuleTac (i+3),
AnnHoareRuleTac (i+1)],
EVERY[rtac (thm AnnCond₂) i,
AnnHoareRuleTac (i+1)],
EVERY[rtac (thm AnnWhile) i,
AnnHoareRuleTac (i+2)],
EVERY[rtac (thm AnnAwait) i,
HoareRuleTac true (i+1)]]

and ParallelTac i = EVERY[rtac (thm ParallelRule) i,
interfree-Tac (i+1),
MapAnn-Tac i]

and MapAnn-Tac i =
FIRST[rtac (thm MapAnnEmpty) i,
EVERY[rtac (thm MapAnnList) i,
MapAnn-Tac (i+1),
AnnHoareRuleTac i],
EVERY[rtac (thm MapAnnMap) i,
rtac (thm allI) i, rtac (thm impI) i,
AnnHoareRuleTac i]]

and interfree-swap-Tac i =
FIRST[rtac (thm interfree-swap-Empty) i,

$EVERY[rtac (thm interfree-swap-List) i,$
 $interfree-swap-Tac (i+2),$
 $interfree-aux-Tac (i+1),$
 $interfree-aux-Tac i],$
 $EVERY[rtac (thm interfree-swap-Map) i,$
 $rtac (thm allI) i, rtac (thm impI) i, rtac (thm conjI) i,$
 $interfree-aux-Tac (i+1),$
 $interfree-aux-Tac i]]$

and $interfree-Tac i =$

$FIRST[rtac (thm interfree-Empty) i,$
 $EVERY[rtac (thm interfree-List) i,$
 $interfree-Tac (i+1),$
 $interfree-swap-Tac i],$
 $EVERY[rtac (thm interfree-Map) i,$
 $rtac (thm allI) i, rtac (thm allI) i, rtac (thm impI) i,$
 $interfree-aux-Tac i]]$

and $interfree-aux-Tac i =$

$FIRST[rtac (thm interfree-aux-rule1) i,$
 $dest-assertions-Tac i]$

and $dest-assertions-Tac i =$

$FIRST[EVERY[rtac (thm AnnBasic-assertions) i,$
 $dest-atomics-Tac (i+1),$
 $dest-atomics-Tac i],$
 $EVERY[rtac (thm AnnSeq-assertions) i,$
 $dest-assertions-Tac (i+1),$
 $dest-assertions-Tac i],$
 $EVERY[rtac (thm AnnCond₁-assertions) i,$
 $dest-assertions-Tac (i+2),$
 $dest-assertions-Tac (i+1),$
 $dest-atomics-Tac i],$
 $EVERY[rtac (thm AnnCond₂-assertions) i,$
 $dest-assertions-Tac (i+1),$
 $dest-atomics-Tac i],$
 $EVERY[rtac (thm AnnWhile-assertions) i,$
 $dest-assertions-Tac (i+2),$
 $dest-atomics-Tac (i+1),$
 $dest-atomics-Tac i],$


```

EVERY[rtac (thm AnnAwait-assertions) i,
      dest-atomics-Tac (i+1),
      dest-atomics-Tac i],
dest-atomics-Tac i]

```

and *dest-atomics-Tac i* =

```

FIRST[EVERY[rtac (thm AnnBasic-atomics) i,
             HoareRuleTac true i],
      EVERY[rtac (thm AnnSeq-atomics) i,
            dest-atomics-Tac (i+1),
            dest-atomics-Tac i],
      EVERY[rtac (thm AnnCond1-atomics) i,
            dest-atomics-Tac (i+1),
            dest-atomics-Tac i],
      EVERY[rtac (thm AnnCond2-atomics) i,
            dest-atomics-Tac i],
      EVERY[rtac (thm AnnWhile-atomics) i,
            dest-atomics-Tac i],
      EVERY[rtac (thm Annatom-atomics) i,
            HoareRuleTac true i],
      EVERY[rtac (thm AnnAwait-atomics) i,
            HoareRuleTac true i]]

```

*}

Note that subgoals are always treated counting downwards, to avoid problems when subgoals are added or deleted. The final tactic is given the name *oghoare*:

```

ML {*
fun oghoare-tac i thm = SUBGOAL (fn (term, -) =>
  (HoareRuleTac true i)) i thm
*}

```

Notice that the tactic for parallel programs *oghoare-tac* is initially invoked with the value *true* for the parameter *precond*.

Parts of the tactic can be also individually used to generate the verification conditions for annotated sequential programs and to generate verification conditions out of interference freedom tests like in §3.3.3:

```

ML {*
fun annhoare-tac i thm = SUBGOAL (fn (term, -) =>

```

```

(AnnHoareRuleTac i) i thm

fun interfree-aux-tac i thm = SUBGOAL (fn (term, -) =>
  (interfree-aux-Tac i) i thm
*)

```

The so defined ML tactics are then “exported” to be used in Isabelle proofs.

```

method-setup oghoare = {*
  Method.no-args (Method.SIMPLE-METHOD' HEADGOAL (oghoare-tac)) *}
  verification condition generator for the oghoare logic

```

```

method-setup annhoare = {*
  Method.no-args
  (Method.SIMPLE-METHOD' HEADGOAL (annhoare-tac)) *}
  verification condition generator for the ann-hoare logic

```

```

method-setup interfree-aux = {*
  Method.no-args
  (Method.SIMPLE-METHOD' HEADGOAL (interfree-aux-tac)) *}
  verification condition generator for interference freedom tests

```

The three tactics for generating verification conditions, *oghoare* for parallel programs, *annhoare* for annotated sequential programs and *interfree-aux* for parts of interference freedom tests, are the only tactics used to verify the examples presented throughout this work. They are invoked simply with **apply**, i.e. **apply** *oghoare*, etc.

Appendix B

Formal Declaration of Concrete Syntax

This section presents the specification of syntax and translations needed to obtain a user-friendly external representation for programs and assertions. As explained in §2.7.1 we use the quote/antiquote technique.

syntax

-quote :: $\beta \Rightarrow (\alpha \Rightarrow \beta)$ ($\ll\text{-}\gg$)
-antiquote :: $(\alpha \Rightarrow \beta) \Rightarrow \beta$ (')

The syntax and translations declared here correspond to the programming language used for the Owicki-Gries formalization. The counterpart for the language of the rely-guarantee formalization is very similar and thus not shown here.

We start with the syntax for component programs. Mixfix notation for commands can be defined by declaring new syntax constants.

syntax

-Assign :: $idt \Rightarrow \beta \Rightarrow \alpha \text{ com}$ (' - := -)
-Seq :: $\alpha \text{ com} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ com}$ (-,;/ -)
-Cond :: $\alpha \text{ bexp} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ com}$ (**if - then - else - fi**)
-While-inv :: $\alpha \text{ bexp} \Rightarrow \alpha \text{ assn} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ com}$ (**while - inv - do - od**)

-AnnAssign :: $\alpha \text{ assn} \Rightarrow idt \Rightarrow \beta \Rightarrow \alpha \text{ com}$ (- ' - := -)
-AnnSeq :: $\alpha \text{ ann-com} \Rightarrow \alpha \text{ ann-com} \Rightarrow \alpha \text{ ann-com}$ (-;;/ -)
-AnnCond₁ :: $\alpha \text{ assn} \Rightarrow \alpha \text{ bexp} \Rightarrow \alpha \text{ ann-com} \Rightarrow \alpha \text{ ann-com}$

$$\Rightarrow \alpha \text{ ann-com } (- \text{ if } - \text{ then } - \text{ else } - \text{ fi})$$

-AnnCond₂ :: $\alpha \text{ assn} \Rightarrow \alpha \text{ bexp} \Rightarrow \alpha \text{ ann-com} \Rightarrow \alpha \text{ ann-com} (- \text{ if } - \text{ then } - \text{ fi})$

-AnnWhile :: $\alpha \text{ assn} \Rightarrow \alpha \text{ bexp} \Rightarrow \alpha \text{ assn} \Rightarrow \alpha \text{ ann-com} \Rightarrow \alpha \text{ ann-com}$
(- while - inv - do - od)

-AnnAwait :: $\alpha \text{ assn} \Rightarrow \alpha \text{ bexp} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ ann-com}$
(- await - then - end)

We also introduce external syntax for commands which are simply abbreviations of the existing ones.

syntax

-Skip :: $\alpha \text{ com}$ (**skip**)

-Cond₂ :: $\alpha \text{ bexp} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ com}$ (**if - then - fi**)

-AnnSkip :: $\alpha \text{ assn} \Rightarrow \alpha \text{ ann-com}$ (- **skip**)

-AnnAtom :: $\alpha \text{ assn} \Rightarrow \alpha \text{ com} \Rightarrow \alpha \text{ ann-com}$ (- ⟨-⟩)

-AnnWait :: $\alpha \text{ assn} \Rightarrow \alpha \text{ bexp} \Rightarrow \alpha \text{ ann-com}$ (- **wait - end**)

The external syntax for assertions is:

syntax

-Assert :: $\alpha \Rightarrow \alpha \text{ set}$ (⟨-⟩)

Part of the corresponding translations of these concrete syntax constants into internal (abstract) syntax can be done simply by directed rewriting equations:

translations

$\langle\{b\}\rangle \mapsto \text{Collect } \langle\langle b \rangle\rangle$

$\langle x := a \rangle \mapsto \text{Basic } \langle\langle (-\text{update-name } x \ a) \rangle\rangle$

$c_1;; c_2 \mapsto \text{Seq } c_1 \ c_2$

if b **then** c_1 **else** c_2 **fi** $\mapsto \text{Cond } \langle\{b\}\rangle \ c_1 \ c_2$

while b **inv** i **do** c **od** $\mapsto \text{While } \langle\{b\}\rangle \ i \ c$

$r \langle x := a \rangle \mapsto \text{AnnBasic } r \langle\langle (-\text{update-name } x \ a) \rangle\rangle$

$c_1;; c_2 \mapsto \text{AnnSeq } c_1 \ c_2$

r **if** b **then** c_1 **else** c_2 **fi** $\mapsto \text{AnnCond}_1 \ r \ \langle\{b\}\rangle \ c_1 \ c_2$

r **if** b **then** c **fi** $\mapsto \text{AnnCond}_2 \ r \ \langle\{b\}\rangle \ c$

r **while** b **inv** i **do** c **od** $\mapsto \text{AnnWhile } r \ \langle\{b\}\rangle \ i \ c$

r **await** b **then** c **end** $\mapsto \text{AnnAwait } r \ \langle\{b\}\rangle \ c$

skip \rightleftharpoons *Basic id*
if b then c fi \rightleftharpoons **if b then c else skip fi**

r skip \rightleftharpoons *AnnBasic r id*
r <c> \rightleftharpoons *r await True then c end*
r wait b end \rightleftharpoons *r await b then skip end*

These equations translate in the direction of the arrow. Expressions on the left-hand side are the input syntax, and those on the right-hand side are the internal representation. The one-to-one translations, denoted by \rightleftharpoons are automatically performed by Isabelle's parser and printer. This is the case for the sequential composition and the command's abbreviations. For the rest of the commands the translation from internal to external representation is a little more complicated because boolean conditions and program variables have to be translated according to the quote/antiquote technique. When the translations are too complicated as to be expressed by an equation, Isabelle allows so-called parse and print translation functions to be programmed in ML.

For parallel programs there is also some concrete syntax of the form

cobegin $c_1 \{q_1\} \parallel \dots \parallel c_n \{q_n\}$ **coend**

defined formally as follows:

nonterminals

prgs

syntax

-*PAR* :: *prgs* \Rightarrow α (**cobegin** - **coend**)

-*prg* :: $[\alpha, \alpha] \Rightarrow$ *prgs* (- -)

-*prgs* :: $[\alpha, \alpha, \textit{prgs}] \Rightarrow$ *prgs* (- - || -)

The following one-direction translations translate the external representation into a *Parallel* command with argument the corresponding list of component programs:

translations

-*prg a c* \mapsto [(*Some a*, *c*)]

-*prgs a c ps* \mapsto (*Some a*, *c*) # *ps*

-*PAR ps* \mapsto *Parallel ps*

The definition of syntax and the corresponding translation for program schemas are:

syntax

-prg-scheme :: [α , α , α , α , α] \Rightarrow *prgs* (**scheme** [- \leq - < -] - -)

translations

-prg-scheme *j i k c q* \Rightarrow (*map* (λi . (*Some* *c*, *q*)) [*j..k*()])

Notice that the internal representation is just a list of parallel programs. Thus, it has to be enclosed in a **cobegin**–**coend** environment, where it can be further composed with other parameterized or concrete parallel programs.

With these equations, the translation from external into internal representation is almost finished, only one constant is still not understandable for Isabelle, namely *-quote*. To translate quotations into internal syntax we use a *parse-translation* function which uses the basic quote/antiquote translations predefined in the theory Isabelle/Pure (see `Syntax.quote_tr` and `Syntax.quote_tr'`).

parse-translation {*

```

let
  fun quote-tr [t] = Syntax.quote-tr -antiquote t
    | quote-tr ts = raise TERM (quote-tr, ts);
in [(-quote, quote-tr)] end
*}

```

This last step completes the translation from external into internal syntax.

For the recovery of the external syntax from the internal representation, a similar ML program called *print-translation* is defined. As usual in Isabelle syntax translations, the part for printing is more complicated. It is only recommended for Isabelle experts and thus it is not shown here. However, we mention that the full ML *print-translation* function consists only of 50 lines of code.

As a comparative remark, it is worth mentioning that the ML program implementing the parse and print translations needed for a method to represent the state via abstraction over tuples of program variables due to [von Wright *et al.*, 1993] (see 2.7.1) consisted of 550 lines of code.

Bibliography

- [Abdulla and Jonsson, 1998] P. A. Abdulla and B. Jonsson. Verifying networks of timed processes. In B. Steffen, editor, *7th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS '98)*, volume 1384 of *Lect. Notes in Comp. Sci.*, pages 298–312, 1998.
- [Alur and Dill, 1994] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–236, 1994.
- [Andersen *et al.*, 1994] F. Andersen, K. Petersen, and J. Pettersson. Program verification using HOL-UNITY. In J. Joyce and C. Seger, editors, *Higher Order Logic Theorem Proving and Its Applications*, volume 780 of *Lect. Notes in Comp. Sci.*, pages 1–15. Springer-Verlag, 1994.
- [Apt and Kozen, 1986] K. R. Apt and D. Kozen. Limits for automatic verification of finite-state concurrent systems. *Information Processing Letters*, 22(6):307–309, 1986.
- [Apt and Olderog, 1991] K. R. Apt and E.-R. Olderog. *Verification of Sequential and Concurrent Programs*. Springer-Verlag, 1991.
- [Apt *et al.*, 1980] K. Apt, N. Francez, and W. de Roever. A proof system for communicating sequential processes. *ACM Transactions on Programming Languages and Systems*, 2(3):359–385, 1980.
- [Apt, 1981a] K. R. Apt. Recursive assertions and parallel programs. *Acta Informatica*, 15:219–232, 1981.
- [Apt, 1981b] K. R. Apt. Ten years of hoare logic: A survey – part I. *ACM Trans. on Prog. Languages and Systems*, 3:431–483, 1981.
- [Basten and Hooman, 1999] T. Basten and J. Hooman. Process algebra in PVS. In *Proceedings TACAS '99*, volume 1579 of *Lect. Notes in Comp. Sci.*, pages 270–284. Springer-Verlag, 1999.

- [Ben-Ari, 1984] M. Ben-Ari. Algorithms for on-the-fly garbage collection. *ACM Trans. Programming Languages and Systems*, 6:333–344, 1984.
- [Best, 1996] E. Best. *Semantics of Sequential and Parallel Programs*. International Series in Computer Science. Prentice Hall, 1996.
- [Bruns, 1997] G. Bruns. *Distributed Systems Analysis with CCS*. Prentice-Hall, 1997.
- [Camillieri, 1990] A. Camillieri. Mechanizing CSP trace theory in higher order logic. *IEEE Transactions on Software Engineering*, 16:993–1004, 1990.
- [Chandy and Misra, 1984] K. Chandy and J. Misra. The drinking-philosophers problem. In *ACM Trans. Programming Languages and Systems*, volume 6, pages 632–646, 1984.
- [Chetali and Heyd, 1997] B. Chetali and B. Heyd. Formal verification of concurrent programs in LP and COQ: A comparative analysis. In E. Gunter and A. Felty, editors, *Theorem Proving in Higher Order Logics*, volume 1275 of *Lect. Notes in Comp. Sci.*, pages 69–85. Springer-Verlag, 1997.
- [Clarke *et al.*, 1995] E. Clarke, O. Grumberg, and S. Jha. Verifying parameterized networks using abstraction and regular languages. In Lee and Smolka, editors, *6th Int. Conf. on Concurrency Theory (CONCUR '95)*, volume 962 of *Lect. Notes in Comp. Sci.*, pages 395–407, 1995.
- [Das *et al.*, 1999] S. Das, D. L. Dill, and S. Park. Experience with predicate abstraction. In *CAV*, volume 1633 of *Lect. Notes in Comp. Sci.*, pages 160–171. Springer-Verlag, 1999.
- [de Boer *et al.*, 1997] F. de Boer, U. Hannemann, and W.-P. de Roever. Hoare-style compositional proof systems for reactive shared variable concurrency. In *Foundations of Software Technology and Theoretical Computer Science*, volume 1346 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1997.
- [de Roever *et al.*, 1998] W.-P. de Roever, H. Langmaack, and A. Pnueli, editors. *Compositionality: The Significant Difference, Proceedings of the International Symposium COMPOS '97*, volume 1536 of *Lect. Notes in Comp. Sci.*, Malente, Germany, 7–12 September 1997 1998. Springer-Verlag.

- [de Roever *et al.*, 2000] W.-P. de Roever, F. S. de Boer, U. Hannemann, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *State-Based Proof Theory of Concurrency: from Noncompositional to Compositional Methods*. Tracts in Theoretical Computer Science. Cambridge University Press, 2000. To appear.
- [Dijkstra *et al.*, 1978] E. W. Dijkstra, L. Lamport, A. J. Martin, C. S. Scholten, and E. F. M. Steffens. On-the-fly garbage collection: An exercise in cooperation. *Communications of the ACM*, 21(11):966–975, 1978.
- [Dijkstra, 1968] E. W. Dijkstra. Cooperating sequential processes. In F. Genuys, editor, *Programming Languages: NATO Advanced Study Institute*, pages 43–112, London, 1968. Academic Press.
- [Dijkstra, 1976] E. W. Dijkstra. *A discipline of programming*. Prentice Hall, 1976.
- [Dutertre and Schneider, 1997] B. Dutertre and S. Schneider. Using a PVS embedding of CSP to verify authentication protocols. In E. Gunter and A. Felty, editors, *Theorem Proving in Higher Order Logics*, volume 1275 of *Lect. Notes in Comp. Sci.*, pages 121–136. Springer-Verlag, 1997.
- [Engberg *et al.*, 1993] U. Engberg, P. Grønning, and L. Lamport. Mechanical verification of concurrent systems with TLA. In G. v. Bochmann and D. Probst, editors, *Computer-Aided Verification (CAV '92)*, volume 663 of *Lect. Notes in Comp. Sci.*, pages 44–55. Springer-Verlag, 1993.
- [Esparza, 1995] J. Esparza. Petri nets, commutative context-free grammars, and basic parallel processes. In *Fundamentals of Computation Theory*, volume 965 of *Lect. Notes in Comp. Sci.*, pages 221–232, 1995.
- [Feijen and van Gasteren, 1999] W. Feijen and A. van Gasteren. On a method of multiprogramming. In *Monographs in Computer Science*. Springer-Verlag, 1999.
- [Filliatre, 1999] J.-C. Filliatre. *Preuve de programmes impératifs en théorie des types*. PhD thesis, Université Paris-Sud, 1999.
- [Francez and Rodeh, 1980] N. Francez and M. Rodeh. Achieving distributed termination without freezing. Technical report (TR 72), IBM Israel Scientific Center, 1980.

- [Galm, 1995] N. Galm. Verifikation von IMP-Programmen mit Hilfe der Hoare-Logik mit dem Theorembeweiser Isabelle. Ausarbeitung zum Fortgeschrittenen-Praktikum, 1995.
- [German and Sistla, 1992] S. German and A. Sistla. Reasoning about systems with many processes. *Journal of the ACM*, 39(3):675–735, 1992.
- [Goldschlag, 1990] D. Goldschlag. Mechanically verifying concurrent programs with the Boyer-Moore Prover. *IEEE Transactions on Software Engineering*, 16:1005–1022, 1990.
- [Gordon and Melham, 1993] M. Gordon and T. Melham, editors. *Introduction to HOL: A Theorem-Proving Environment for Higher Order Logic*. Cambridge University Press, 1993.
- [Gordon, 1979] M. Gordon. *The Denotational Description of Programming Languages, An Introduction*. Springer-Verlag, New York, 1979.
- [Gordon, 1989] M. Gordon. Mechanizing programming logics in higher order logic. In G. Birtwistle and P. Subrahmanyam, editors, *Current Trends in Hardware Verification and Automated Theorem Proving*. Springer-Verlag, 1989.
- [Gries, 1997] D. Gries. An exercise in proving parallel programs correct. *Communications of the ACM*, 20(12):921–930, 1997.
- [Harrison, 1998] J. Harrison. Formalizing Dijkstra. In *Theorem Proving in Higher Order Logics (TPHOLs '98)*, volume 1497 of *Lect. Notes in Comp. Sci.*, pages 171–188. Springer-Verlag, 1998.
- [Havelund and Shankar, 1997] K. Havelund and N. Shankar. A mechanized refinement proof for a garbage collector. *Formal Aspects of Computing*, 3:1–28, 1997.
- [Havelund, 1996] K. Havelund. Mechanical verification of a garbage collector. In *FMPPTA*, 1996. Available at <http://ic-www.arc.nasa.gov/ic/projects/amphion/people/havelund/>.
- [Hennessy and Plotkin, 1979] M. Hennessy and G. Plotkin. Full abstraction for a simple programming language. In *Mathematical Foundations of Computer Science*, volume 74 of *Lect. Notes in Comp. Sci.*, pages 108–120. Springer-Verlag, 1979.

- [Henzinger, 1995] T. Henzinger. Hybrid automata with finite bisimulations. In *ICALP '95*, 1995.
- [Heyd and Crégut, 1996] B. Heyd and P. Crégut. A modular coding of Unity in Coq. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics*, volume 1125 of *Lect. Notes in Comp. Sci.*, pages 251–266. Springer-Verlag, 1996.
- [Hoare, 1969] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12:576–580, 1969.
- [Homeier and Martin, 1996] P. V. Homeier and D. F. Martin. Mechanical verification of mutually recursive procedures. In M. McRobbie and J. Slaney, editors, *Automated Deduction — CADE-13*, volume 1104 of *Lect. Notes in Comp. Sci.*, pages 201–215. Springer-Verlag, 1996.
- [Hooman *et al.*, 2000] J. Hooman, W.-P. de Roever, P. Pandya, H. Schepers, Q. Xu, and P. Zhou. *Compositional Theory of Concurrency*. Cambridge University Press, 2000. To appear.
- [Hooman, 1995] J. Hooman. Verifying part of the ACCESS.bus protocol using PVS. In P. S. Thiagarajan, editor, *15th Conference on the Foundations of Software Technology and Theoretical Computer Science*, volume 1026 of *Lect. Notes in Comp. Sci.*, pages 96–110, Bangalore, India, December 1995. Springer-Verlag.
- [Hooman, 1998] J. Hooman. Developing proof rules for distributed real-time systems with PVS. In *Proc. Workshop on Tool Support for System Development and Verification*, volume 1 of *BISS Monographs*, pages 120–139. Shaker Verlag, 1998.
- [Jackson, 1998] P. Jackson. Verifying a garbage collection algorithm. In J. Grundy and M. Newey, editors, *11th International Conference on Theorem Proving in Higher Order Logics (TPHOLs '98)*, *Lect. Notes in Comp. Sci.* Springer-Verlag, 1998. Available at www.dcs.ed.ac.uk/home/pbj.
- [Jones, 1981] C. B. Jones. *Development methods for computer programs including a notion of interference*. PhD thesis, Oxford University Computing Laboratory, 1981.
- [Jones, 1983] C. B. Jones. Tentative steps towards a development method for interfering programs. *ACM Transactions on Programming Languages and Systems*, 5(4):596–619, 1983.

- [Jonker, 1992] J. E. Jonker. On-the-fly garbage collection for several mutator. *Distributed Computing*, 5:187–199, 1992.
- [Jonsson and Parrow, 1993] B. Jonsson and J. Parrow. Deciding bisimulation equivalences for a class of non-finite state programs. *Information and Computation*, 107(2):272–302, 1993.
- [Kalvala, 1995] S. Kalvala. A formulation of TLA in Isabelle. In E. Schubert, P. Windley, and J. Alves-Foss, editors, *Higher Order Logic Theorem Proving and its Applications*, volume 971 of *Lect. Notes in Comp. Sci.*, pages 214–228. Springer-Verlag, 1995.
- [Kleymann, 1998] T. Kleymann. *Hoare logic and VDM: Machine-Checked Soundness and Completeness Proofs*. PhD thesis, Department of Computer Science, Univ. of Edinburgh, 1998.
- [Lamport, 1977] L. Lamport. Proving the correctness of multiprocess programs. *IEEE Transactions on Software Engineering*, 3(2):125–143, 1977.
- [Långbacka and von Wright, 1997] T. Långbacka and J. von Wright. Refining reactive systems in HOL using action systems. In E. Gunter and A. Felty, editors, *Theorem Proving in Higher Order Logics*, volume 1275 of *Lect. Notes in Comp. Sci.*, pages 183–197. Springer-Verlag, 1997.
- [Levin and Gries, 1981] G. Levin and D. Gries. A proof technique for communicating sequential processes. *Acta Informatica*, 15:281–302, 1981.
- [Misra and Chandy, 1981] J. Misra and M. Chandy. Proofs of networks of processes. *IEEE Transactions on Software Engineering*, 7(7):417–426, 1981.
- [Müller and Nipkow, 1997] O. Müller and T. Nipkow. Traces of I/O automata in Isabelle/HOLCF. In M. Bidoit and M. Dauchet, editors, *TAPSOFT '97: Theory and Practice of Software Development*, volume 1214 of *Lect. Notes in Comp. Sci.*, pages 580–594. Springer-Verlag, 1997.
- [Müller, 1998] O. Müller. *A Verification Environment for I/O Automata Based on Formalized Meta-Theory*. PhD thesis, TU München, September 1998.
- [Naraschewski and Wenzel, 1998] W. Naraschewski and M. Wenzel. Object-oriented verification based on record subtyping in higher-order logic. In J. Grundy and M. Newey, editors, *Theorem Proving in Higher Order Logics*, volume 1479 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1998.

- [Nesi, 1994] M. Nesi. Value-passing CCS in HOL. In J. Joyce and C. Seger, editors, *Higher Order Logic Theorem Proving and Its Applications*, volume 780 of *Lect. Notes in Comp. Sci.*, pages 352–365. Springer-Verlag, 1994.
- [Nipkow and Paulson, 2001] T. Nipkow and L. C. Paulson. Isabelle/HOL – The Tutorial. Available at <http://isabelle.in.tum.de/doc/tutorial.pdf>, 2001.
- [Nipkow and Prensa Nieto, 1999] T. Nipkow and L. Prensa Nieto. Owicki/Gries in Isabelle/HOL. In J.-P. Finance, editor, *Fundamental Approaches to Software Engineering (FASE '99)*, volume 1577 of *Lect. Notes in Comp. Sci.*, pages 188–203. Springer-Verlag, 1999.
- [Nipkow, 1996] T. Nipkow. Winskel is (almost) right: Towards a mechanized semantics textbook. In V. Chandru and V. Vinay, editors, *Foundations of Software Technology and Theoretical Computer Science*, volume 1180 of *Lect. Notes in Comp. Sci.*, pages 180–192. Springer-Verlag, 1996.
- [Nipkow, 1998] T. Nipkow. Winskel is (almost) right: Towards a mechanized semantics textbook. *Formal Aspects of Computing*, 10:171–186, 1998.
- [Owicki and Gries, 1976a] S. Owicki and D. Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, 6:319–340, 1976.
- [Owicki and Gries, 1976b] S. Owicki and D. Gries. Verifying properties of parallel programs: an axiomatic approach. *Communications of the ACM*, 19:279–285, 1976.
- [Owicki, 1975] S. Owicki. *Axiomatic proof techniques for parallel programs*. PhD thesis, Computer Science Dept., Cornell University, 1975.
- [Owre *et al.*, 1995] S. Owre, J. Rushby, N. Shankar, and F. von Henke. Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, 21(2):107–125, February 1995.
- [Owre *et al.*, 1996] S. Owre, S. Rajan, J. Rushby, N. Shankar, and M. Srivas. PVS: Combining specification, proof checking, and model checking. In R. Alur and T. A. Henzinger, editors, *Computer-Aided Verification, CAV '96*, volume 1102 of *Lect. Notes in Comp. Sci.*, pages 411–414, New Brunswick, NJ, July/August 1996. Springer-Verlag.

- [Paulson, 1994] L. C. Paulson. *Isabelle: A Generic Theorem Prover*, volume 828 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1994. Isabelle home page: <http://isabelle.in.tum.de/>.
- [Paulson, 2000] L. C. Paulson. Mechanizing UNITY in Isabelle. *ACM Transactions on Computational Logic*, 1(1):3–32, 2000.
- [Paulson, 2001] L. C. Paulson. Mechanizing a theory of program composition for UNITY. *ACM Transactions on Programming Languages and Systems*, 2001. in Press.
- [Peterson, 1981] G. L. Peterson. Myths about the mutual exclusion problem. *Information Processing Letters*, 12(3):223–252, 1981.
- [Pixley, 1988] C. Pixley. An incremental garbage collection algorithm for multimutator systems. *Distributed Computing*, 3:41–50, 1988.
- [Plotkin, 1981] G. D. Plotkin. A structural approach to operational semantics. Technical report DAIMI-FN 19, Department of Computer Science, Aarhus University, 1981.
- [Prensa Nieto and Esparza, 2000] L. Prensa Nieto and J. Esparza. Verifying single and multi-mutator garbage collectors with Owicki/Gries in Isabelle/HOL. In M. Nielsen and B. Rovan, editors, *Mathematical Foundations of Computer Science (MFCS '00)*, volume 1893 of *Lect. Notes in Comp. Sci.*, pages 619–628. Springer-Verlag, 2000.
- [Prensa Nieto, 2001] L. Prensa Nieto. Completeness of the Owicki-Gries system for parameterized parallel programs. In *Formal Methods for Parallel Programming: Theory and Applications (FMPPTA '01)*, 2001.
- [Rogers, 1987] H. Rogers, Jr. *Theory of Recursive Functions and Effective Computability*. MIT Press, 1987.
- [Rushby, 2000] J. Rushby. Theorem proving for verification. In F. Cassez, editor, *Modelling and Verification of Parallel Processes: MoVEP 2k*, Nantes, France, June 2000. Tutorial presented at MoVEP: revised version to be published by Springer LNCS.
- [Russinoff, 1994] D. M. Russinoff. A mechanically verified garbage collector. *Formal Aspects of Computing*, 6:359–390, 1994.
- [Schoenfield, 1967] J. R. Schoenfield. *Mathematical Logic*. Addison-Wesley, 1967.

- [Scott and Strachey, 1971] D. S. Scott and C. Strachey. Towards a mathematical semantics for computer languages. Technical report PRG-6, Programming Research Group, University of Oxford, 1971.
- [Shankar, 1996] N. Shankar. PVS: Combining specification, proof checking, and model checking. In M. Srivas and A. Camilleri, editors, *Formal Methods in Computer-Aided Design (FMCAD '96)*, volume 1166 of *Lect. Notes in Comp. Sci.*, pages 257–264, Palo Alto, CA, November 1996. Springer-Verlag.
- [Shankar, 1998] N. Shankar. Machine-assisted verification using theorem proving and model checking. In M. Broy, editor, *Mathematical Methods in Program Development*. Springer-Verlag, 1998.
- [Søgaard-Andersen *et al.*, 1993] J. Søgaard-Andersen, S. Garland, J. Guttag, N. Lynch, and A. Pogoyants. Computer-assisted simulation proofs. In 4th *Conference on Computer-Aided Verification*, volume 697 of *Lect. Notes in Comp. Sci.*, pages 305–319. Springer-Verlag, 1993.
- [Soundararajan, 1984] N. Soundararajan. A proof technique for parallel programs. *Theoretical Computer Science*, 31:13–29, 1984.
- [Stirling, 1988] C. Stirling. A generalization of Owicki-Gries’s Hoare logic for a concurrent while language. *Theoretical Computer Science*, 58:347–359, 1988.
- [Stølen, 1990] K. Stølen. *Development of Parallel Programs on Shared Data-Structures*. PhD thesis, Computer Science Department, Manchester University, 1990.
- [Stølen, 1991] K. Stølen. A method for the development of totally correct shared-state parallel programs. In J. C. M. Baeten and J. F. Groote, editors, *CONCUR '91*, volume 527 of *Lect. Notes in Comp. Sci.*, pages 510–525. Springer-Verlag, 1991.
- [Tej and Wolff, 1997] H. Tej and B. Wolff. A corrected failure-divergence model for CSP in Isabelle/HOL. In J. Fitzgerald, C. Jones, and P. Lucas, editors, *FME '97: Industrial Applications and Strengthened Foundations of Formal Methods*, volume 1313 of *Lect. Notes in Comp. Sci.*, pages 318–337. Springer-Verlag, 1997.
- [van de Snepscheut, 1987] J. L. A. van de Snepscheut. “Algorithms for on-the-fly garbage collection” revisited. *Information Processing Letters*, 24:211–216, 1987.

- [von Oheimb, 2001] D. von Oheimb. *Analyzing Java in Isabelle/HOL – Formalization, Type Safety and Hoare Logic*. PhD thesis, TU München, 2001. <http://www4.in.tum.de/~oheimb/diss/>.
- [von Wright and Långbacka, 1993] J. von Wright and T. Långbacka. Using a theorem prover for reasoning about concurrent algorithms. In G. v. Bochmann and D. Probst, editors, *Computer-Aided Verification (CAV '92)*, volume 663 of *Lect. Notes in Comp. Sci.*, pages 56–68. Springer-Verlag, 1993.
- [von Wright *et al.*, 1993] J. von Wright, J. Hekanaho, P. Luostarinen, and T. Långbacka. Mechanizing some advanced refinement concepts. *Formal Methods in System Design*, 3:49–81, 1993.
- [Wenzel, 2001a] M. Wenzel. *Isabelle/Isar – a versatile environment for human-readable formal proof documents*. PhD thesis, TU München, 2001. <http://www4.in.tum.de/~wenzelm/diss/>.
- [Wenzel, 2001b] M. Wenzel. Miscellaneous Isabelle/Isar examples for higher order logic. Isabelle/Isar proof document, TU München, February 2001.
- [Wilson, 1992] P. R. Wilson. Uniprocessor garbage collection techniques. In *International Workshop on Memory Management*, *Lect. Notes in Comp. Sci.*, 1992. Available at www.cs.ukc.ac.uk/people/staff/rej/gc.html.
- [Winskel, 1993] G. Winskel. *The Formal Semantics of Programming Languages*. MIT Press, 1993.
- [Xu *et al.*, 1995] Q. Xu, W.-P. de Roever, and J. He. Rely-guarantee method for verifying shared variable concurrent programs. Technical report, Christian-Albrechts-Universität Kiel, 1995.
- [Xu *et al.*, 1997] Q. Xu, W.-P. de Roever, and J. He. The rely-guarantee method for verifying shared variable concurrent programs. *Formal Aspects of Computing*, 9(2):149–174, 1997.

Index

- \Rightarrow , 9
- $[\dots]$, 12
- $\#$, 11
- \equiv , 9
- ' , 12
- \longrightarrow , 12
- \Longrightarrow , 12
- $!$, 11
- \rightarrow , 10
- \Rightarrow , 11
- \times , 11
- $- -1 \rightarrow -$, 24
- $- -n \rightarrow -$, 24
- $- -* \rightarrow -$, 24
- $- -P1 \rightarrow -$, 24
- $- -Pn \rightarrow -$, 24
- $- -P* \rightarrow -$, 24
- \vdash , 33
- \vDash , 33
- \vdash_{st} , 34
- \models , 32
- \vDash , 32
- $\vDash - sat [-, -, -, -]$, 112
- $\vdash - sat [-, -, -, -]$, 114
- $\vDash - SAT [-, -, -, -]$, 113
- $\vdash - SAT [-, -, -, -]$, 116
- $- -c \rightarrow -$, 104
- $- -c* \rightarrow -$, 104
- $- -e \rightarrow -$, 104
- $- -pc \rightarrow -$, 105
- $- -pe \rightarrow -$, 105
- ' , 53
- $\ll\text{-}\gg$, 53
- $\text{-};\text{-}$, 57
- $\text{-},\text{-}$, 57
- Ω , 28
- ∞ , 127
- accessible, 72
- All-None*, 23
- and**, 18
- AnnAwait*, 19
- AnnBasic*, 18
- ann-com*, 18
- ann-com-op*, 23
- AnnCond₁*, 18
- AnnCond₂*, 19
- annhoare* (vcg), 180
- ann-hoare* (constant), 33
- ann-SEM*, 27
- ann-sem*, 27
- AnnSeq*, 18
- ann-transition*, 22
- ann-triple-op*, 23
- AnnWhile*, 19
- antiquotation, 53
- apply**, 10
- assertions*, 37
- assn*, 17
- assum*, 112
- atom-com*, 21
- atomic commands, 16
- atomic region, 57
- atomics*, 37

auto, 10
 auxiliary variables, 40

Basic, 18, 102
bexp, 17, 102
Black, 74
Blacks, 74
BtoW, 74

 closed systems, 3, 102
Collector, 80
Color-Target, 77
Com, 116
com(constant), 23
com(type), 18, 102
comm, 112
 complement set, 12
 completeness, 32, 143
 component commands, 16
 component transition, 103
 composed variable, 145
 computation, 103, 106
Cond, 18, 102
conf, 104
 configuration, 22
confs, 107
conjoin, 127
consts, 9
contsdefs, 9
cp, 107
cptn, 107
cptn-mod, 108

datatype, 9
 deadlock, 19
 deep embedding, 8
defs, 9
 denotational semantics, 22

edge, 74
edges, 74

 elementary arithmetic, 150
 environment transition, 103
 expressiveness, 150

 free list, 72
fst, 11
fwhile, 28

 garbage, 72

hd, 11
 history variable, 40
 Hoare triple, 31

 incremental, 72
inductive, 9
 inductive cases, 9
 interference freedom, 37
interfree, 38
interfree-aux, 38
interfree-aux (vcg), 180
 interleaving semantics, 26
intros, 9
 Isar, 13

last, 11
lemma, 9
length, 11
lift, 108
 location variables, 40

map, 11
 memory, 72
 module, 77
Mul-Prop, 93
Mutator, 77
Mut-init, 77
 mutually recursive, 18

node, 74
nodes, 74
None, 11

oghoare (constant), 33
oghoare (vcg), 50, 179
open systems, 3, 102
operational semantics, 22
option, 11

Parallel, 21, 103
parallel commands, 16
parameterized program, 144
par-assum, 113
par-com, 102
par-comm, 113
par-conf, 105
par-confs, 107
par-cp, 107
par-cptn, 107
par-rgformula, 116
parse_translation, 11, 184
partial correctness semantics, 27
Post, 116
post, 23
postcondition, 31
Pre, 116
pre, 20
precondition, 31
primrec, 9
print_translation, 11
print_translation, 184
program transition, 103
proof obligations, 11
proof outline, 19
Proper, 80

Queue, 93
quotation, 53

R (constant), 76
Reach, 74
reachable, 72
record, 53
recursively enumerable, 154

Redirect-Edge, 77
relative completeness, 150
Rely, 116
resolution, 10
rewriting, 10
rgformula, 111
Roots, 74
roots, 72
rule induction, 9

Safe, 76
schematic variables, 168
sem, 27
SEM, 27
semantic completeness, 150
Seq, 18, 102
set, 11
shallow embedding, 8
simplifier, 10
snd, 11
Some, 11
soundness, 32
stable, 113
standard proof outline, 19
strong soundness, 42
syntax, 10

T (constant), 76
tactic, 10
tactical, 10
the, 12
theorem, 9
tl, 11
transition, 22
translations, 10
types, 9

validity, 32
vcg (verification condition generator), 11, 49, 167
verification conditions, 11

wait-statement, 57
weakest liberal precondition (wlp),
167
While, 18, 102
White, 74