

AD-A156 190


SPECIFYING AND VERIFYING CONCURRENT PROGRAMS(U) SRI
INTERNATIONAL MENLO PARK CA L LAMPORT FEB 85
ARO-20628. 5-EL DARG29-83-K-0119


1/1

UNCLASSIFIED

F/G 9/2

NL





NATIONAL BUREAU OF STANDARDS
MICROCOPY RESOLUTION TEST CHART

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

2

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER ARO 20628.5-EL		2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Specifying and Verifying Concurrent Programs		5. TYPE OF REPORT & PERIOD COVERED 15 Aug 83 - 14 Nov 84 Final Report	
AUTHOR(s) Leslie Lamport		6. PERFORMING ORG. REPORT NUMBER	
PERFORMING ORGANIZATION NAME AND ADDRESS SRI International Menlo Park, California		8. CONTRACT OR GRANT NUMBER(s) DAAG29-83-K-0119	
CONTROLLING OFFICE NAME AND ADDRESS U. S. Army Research Office Post Office Box 12211 Research Triangle Park, NC 27709		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE Feb 1985	
		13. NUMBER OF PAGES	
		15. SECURITY CLASS. (of this report) Unclassified	
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	

16. DISTRIBUTION STATEMENT (of this Report)
Approved for public release; distribution unlimited.

DTIC
SELECTED
JUN 26 1985
S
A

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES
The view, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)
Specification
Verification
Semantics
Computer Programs
Computer Programming

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)
The goal of this project was the development of formal methods for the specification and verification of concurrent programs to help avoid software errors in concurrent systems. This involved research in three areas:

We feel that previous work provides an adequate formal foundation for verification, so we have concentrated on specification and semantics.

AD-A156 190

DTIC FILE COPY

SPECIFYING AND VERIFYING CONCURRENT PROGRAMS

Final Report

February 1985

By: Leslie Lamport

Prepared for:

U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, North Carolina 27709-2211
Attention: Richard O. Ulsh, Chief
Information Processing Office

Contract No. DAAG29-83-K-0119

Approved for public release; distribution unlimited.

SRI International
333 Ravenswood Avenue
Menlo Park, California 94025-3493
Telephone: (415) 326-6200
Cable: SRI INTL MPK
TWX: 910-373-2046
Telex: 334 486



A-1

85-06 6 0 93

1 Introduction

The goal of this project was the development of formal methods for the specification and verification of concurrent programs to help avoid software errors in concurrent systems. This involved research in three areas:

- Specification.
- Verification.
- Semantics.

We feel that previous work provides an adequate formal foundation for verification [5,6,7], so we have concentrated on specification and semantics. Our work in these two areas is summarized in the following two sections.

2 Specification

A formal specification should describe precisely what it means for an implementation to be correct. This requires a precise definition of what it means for a program to satisfy a specification. Without such a definition, a formal specification is at best incomplete, and may be little better than an informal one for producing correct software.

The precise connection between a specification and its implementation is subtle. The specification is generally given in terms of large, high-level operations, which the program implements as a series of lower-level operations. Most research on specification has ignored the question of what it means for a series of low-level operations to implement a single high-level one. As a result, the specifications written with most formalisms are incomplete. The specification of a FIFO queue is one of the standard toy examples. However, in very few of the specifications can one decide if the specification is supposed to be satisfied by writing a program or building a piece of hardware. It can be argued that this is a lower-level question that these specifications are not attempting to address. However, it is disturbing that the question of how this could be specified is not addressed. A specification that does not specify whether an implementation should consist of a program or a piece of hardware cannot help the designer avoid the more subtle interface problems that plague real systems.

The reason for the incompleteness is that a specification of a module in a concurrent system must have two components: an *internal* part that describes the invisible, internal behavior, and an *interface* part that specifies precisely how the module interacts with its environment. While the internal part is independent of the implementation details, the interface specification must be given in terms of implementation-level concepts. Most specification methods consider only the internal specification and ignore the interface specification. It

is the interface specification that determines if a queue is implemented in code or silicon.

Ignoring the interface specification can lead to misleading specifications—ones that are either meaningless or incapable of being implemented. This is indicated by the surprising result that no current method seems capable of providing a formal specification of first-come-first-served priority. Thus, for example, the informal requirement in the Ada programming language that requests for interprocess communication obey such priority seems incapable of being made formal.

A detailed discussion of the problem of the formal connection between a specification and its implementation is given in [3], which is attached as Appendix A of this report. This work also explains our assertion that first-come-first-served priority cannot be adequately specified by current methods.

3 Semantics

Just as a specification method should define what it means for a program to implement a specification, we believe that a formal semantics should define what it means for the “machine-language” program generated by a compiler to be a correct implementation of a program. Thus, we regard a program formally to be a specification of a lower-level, compiled version. Our work on semantics has thus been driven by the need to establish a formal connection between the programming language's operations and the lower-level machine-language operations with which they are implemented.

We have developed a general method for specifying the semantics of concurrent programming languages. Our approach of viewing a program as a specification of its compiled version led to a semantics in which the meaning of a program is a temporal logic specification similar to the ones described in [2]. Our method can be used to define a formal semantics for all concurrent programming constructs that we have been able to think of. The method is described in [1], which is included as Appendix B of this report.

Since the semantics of a language consists not of a single specification but of an algorithm for generating a specification for every program, several logical concepts not present in [2] had to be introduced. One of these concepts was a new way of reasoning about aliasing.

An important goal of our method was *compositionality*—obtaining the meaning of a program from the meaning of its components. Combining components of a program leads to implicit aliasing; for example, when two processes are “hooked up” to a unidirectional communication channel, there is an implicit aliasing between the sender's output buffer and the receiver's input buffer. Such aliasing is of a more general nature than the identification of program variables usually understood as aliasing. To handle it, we (in collaboration with Fred Schneider) developed a method of reasoning about invariant relations—ordinary

aliasing being an invariant relation of equality between the values of two variables. This method provides an elegant new method for handling aliasing and typing relations in ordinary sequential programs that is described in [4], which is included as Appendix C of this report.

References

- [1] Leslie Lamport. *An Axiomatic Semantics of Concurrent Programming Languages*. Springer-Verlag, Berlin, 1985. To appear.
- [2] Leslie Lamport. Specifying concurrent program modules. *ACM Transactions on Programming Languages and Systems*, 5(2):190–222, April 1983.
- [3] Leslie Lamport. What it means for a concurrent program to satisfy a specification: why no one has specified priority. In *Proceedings of the Twelfth ACM Symposium on Principles of Programming Languages*, ACM SIGACT-SIGPLAN, New Orleans, January 1985.
- [4] Leslie Lamport and Fred B. Schneider. Constraints: a uniform approach to aliasing and typing. In *Proceedings of the Twelfth ACM Symposium on Principles of Programming Languages*, ACM SIGACT-SIGPLAN, New Orleans, January 1985.
- [5] Leslie Lamport and Fred B. Schneider. The “Hoare logic” of CSP, and all that. *ACM Transactions on Programming Languages and Systems*, 6(2):281–296, April 1984.
- [6] Susan Owicki and David Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, 6(4):319–340, 1976.
- [7] Susan Owicki and Leslie Lamport. Proving liveness properties of concurrent programs. *ACM Transactions on Programming Languages and Systems*, 4(3):455–495, July 1982.

APPENDIX A

**WHAT IT MEANS FOR CONCURRENT PROGRAM TO
SATISFY A SPECIFICATION:
WHY NO ONE HAS SPECIFIED PRIORITY**

What It Means for a Concurrent Program to
Satisfy a Specification:
Why No One Has Specified Priority

Leslie Lamport
Computer Science Laboratory
SRI International¹

2 July 1984
minor revision 27 Octobe. 1984

To appear in *Proceedings of the 12th Annual
ACM Symposium on the Principles of Program-
ming Languages* (January, 1985)

¹This work was supported in part by the National Science Foundation under grant number MCS-8104459, and the Army Research Office under grant number DAAG29-83-K-0119.

What It Means for a Concurrent Program to Satisfy a Specification: Why No One Has Specified Priority

Leslie Lamport
Computer Science Laboratory
SRI International*

Abstract

The formal correspondence between an implementation and its specification is examined. It is shown that existing specifications that claim to describe priority are either vacuous or else too restrictive to be implemented in some reasonable situations. This is illustrated with a precisely formulated problem of specifying a first-come-first-served mutual exclusion algorithm, which it is claimed cannot be solved by existing methods.

pri-or'i-ty (pri-ōr'ē-tī), *n.*; *pl.* -TIES(-tīz). **3.** Order of preference based on urgency, importance, or merit. [1]

1 Introduction

Specification and Implementation

A formal specification method should reduce the question of whether a program satisfies its specification to a precisely formulated mathematical problem. This reduction is what distinguishes a formal method from an informal one. Most researchers developing specification formalisms have concentrated upon the formal semantics of the specification language, apparently believing that such a semantics, together with a formal semantics for the programming language, provides the necessary reduction. However, formal semantics for the specification and programming languages

*This work was supported in part by the National Science Foundation under grant number MCS-8104459, and the Army Research Office under grant number DAAG29-83-K-0119.

are not enough; one must also define the correspondence between the two semantics.

As a trivial example, consider a program to compute the square of an integer. The specification might be given in terms of mathematical integers, while the program's semantics might be defined in terms of bit strings. To determine if the program meets its specification, we must define the correspondence between the implementation-level semantic concept of *bit string* and the specification-level concept of *integer*. Although this is easy to do, it is important that it be done; a program that expects input values to be in two's-complement representation may produce an incorrect answer when given an input value encoded in sign-magnitude representation.

For sequential programs, specified in terms of input and output values, the correspondence between implementation and specification concepts is, in principle, simple: it is just a mapping between two domains of values. However, this is not the case for concurrent programs, where the specification involves the program's behavior. The granularity of operations can be very different at the two levels; an atomic operation at the specification level may correspond to a large number of atomic program operations. The formal correspondence between the two semantic levels requires careful examination. In this paper, I consider the implications of this correspondence for the particular problem of specifying priority.

Priority

In concurrent systems, priority denotes the order of preference in which processes obtain service. It may be based upon the nature of the service being requested, the importance of the requesting process, or the order in which requests are issued. All popular methods for specifying concurrent systems allow one to write simple specifications that appear to describe priority. However, I will show that, depending upon how they are interpreted, these specifications are either too restrictive to be implementable in all situations or else they are vacuous, being satisfied by any program.

I will concentrate upon a particular example of priority: first-come-first-served (FCFS). There is nothing special

about FCFS—the same problem arises in specifying other types of priority: FCFS just provides a simple, well-studied case. It is also an important case; the Ada language requires an FCFS queuing discipline in the implementation of the rendezvous mechanism, and problems in formally specifying this requirement may be of some interest to the Ada community.

My claim that current methods cannot specify priority is a controversial one, and provokes arguments when presented to computer scientists. I have therefore formulated a challenge to those who feel that they know how to specify priority: the specification of a precisely-defined FCFS mutual exclusion algorithm. I believe that anyone claiming to have a general method for specifying concurrent programs should be able to write the required specification. By a “general” method, I mean one that permits implementations in a reasonably broad class of programming languages. Given some particular programming language having an FCFS synchronization primitive, it is easy to specify FCFS priority for programs written in that language by requiring the implementation to use the FCFS primitive. To prevent this kind of “cheating”, the challenge specifies two simple programming languages that must be handled.

The challenge is presented first, before any explanation of what makes specifying priority difficult. I urge the reader to study it and decide if it is reasonable before reading the rest of the paper. There are no tricks in the challenge; the problem that arises in trying to specify priority is a fundamental one. For a cry of “foul” to be taken seriously, it should be issued on the basis of the challenge alone, not on the ensuing discussion.

2 The Challenge

Let Blaise be a simple concurrent programming language with an atomic assignment statement, concatenation (;), if and while statements with atomic tests, a **cobegin**, and integer and boolean shared variables, but with no explicit synchronization or communication commands. The **cobegin** is assumed to be fair, meaning that a nonterminated process will eventually execute its next atomic operation, but no bound is assumed on the relative execution speeds of the different processes. All classical shared-variable multi-process algorithms can easily be written as Blaise programs.

Let the language Tony be the same as Blaise except with no shared variables, instead using CSP-style communication primitives. Moreover, assume an appropriate fairness requirement on communication so that a Blaise program can be simulated in the obvious way by a Tony program in which shared variables are replaced by extra processes.¹ Reading the value of a Blaise shared variable is simulated by a “?” operation in the Tony program, and writing its value is simulated by a “!”.

Mutual exclusion algorithms that can be written in Blaise

¹ Without some fairness constraint on communication, a Tony program cannot guarantee the fairness condition for a Blaise process that accesses a shared variable.

have been studied for years [3], and there is a general agreement that certain algorithms are FCFS and others are not. It is less clear what it means for a Tony program to be FCFS, but it is easy to write Tony programs that are obviously not FCFS—for example, an algorithm with a central scheduling process that does not always grant requests in the order it receives them.

The challenge is to specify program statements $entry_p$, and $exit_p$, for $p = 1, \dots, 17$, such that if the statement

$$entry_p; \text{ critical section}; exit_p \quad (1)$$

is embedded in the sequential process π_p , then

$$\text{cobegin } \pi_1 \parallel \dots \parallel \pi_{17} \text{ coend}$$

is an FCFS mutual exclusion algorithm in which π_p requests entry to its critical section by initiating the execution of (1). (There may also be declarations of the shared variables in a Blaise program and extra processes in a Tony program.) The specification, and the specification method, must have the following properties.

1. For any Blaise or Tony implementations of the $entry_p$ and $exit_p$ statements, the method defines a mathematical formula C and a formal system L such that the implementation satisfies the specification if and only if C is a valid formula of L .

[This is a precise statement of the requirement that a formal method reduce the question of whether a particular program satisfies the specification to a well-defined mathematical problem.]

2. (a) Any Blaise implementation that is generally regarded to be an FCFS mutual exclusion algorithm must satisfy the specification.
(b) A Tony simulation of such a Blaise program must also satisfy the specification.
3. Any Blaise or Tony program that is generally regarded not to be an FCFS mutual exclusion algorithm must not satisfy the specification.

I will attempt to answer all serious responses to this challenge. To meet the challenge, you must provide the specification and indicate how one constructs the C and L of condition 1 for any Blaise or Tony program. I will then attempt to present one or more programs that violate condition 2 or 3, in which case you must show that these conditions are not violated. The construction of C and L and the refutation of my counterexamples need not be given in full mathematical detail, but they must be rigorous enough to convince a competent computer scientist that a completely formal exposition is, in principle, possible.

3 What's So Hard About The Challenge?

Why Current Methods Don't Work

Let us now consider how one might specify the FCFS condition of the challenge. Intuitively, FCFS means that requests to enter the critical section are serviced in the order in which they are issued. To specify this more precisely, one must recognize two kinds of operations—a *request* operation and a *critical.section* operation. To each *critical.section* operation there corresponds a *request* operation, issued before it by the same process. We identify operations by subscripts, letting *request_i* and *critical.section_i* denote corresponding operations. The FCFS priority condition is usually expressed as follows:

- (*) For any distinct operations *critical.section_i* and *critical.section_j*, if *request_i* precedes *request_j*, then *critical.section_i* must precede *critical.section_j*.

The operations *request_i* and *request_j* need not be atomic actions. The condition "*request_i* precedes *request_j*," means that *request_i* finishes before *request_j* begins. If these operations are nonatomic, then they can be concurrent, meaning that neither precedes the other. In this case, condition (*) does not specify the order of the operations *critical.section_i* and *critical.section_j*. Allowing the *request* operations to be nonatomic means that the order of service does not matter (is not specified) if the requests are issued "too close together".

All the formal specification methods I know of—including [4], [5], [8], [11], [13], [14], [15], [16], [17], and [18]—specify FCFS with condition (*), although the formal expression of this condition differs with the different methods. These differences are irrelevant to the fundamental problem with condition (*).

To verify that a Blaise program satisfies (*), one must state what Blaise operations correspond to the operations *request_i* and *critical.section_i*. The *critical.section_i* operation clearly corresponds to an execution of the critical section, but what about the *request_i* operation? Let *enter_i* denote the operation by which a process initiates execution of its *entry_i* statement, so *enter_i* is an execution of the atomic action that puts the process's control at the beginning of that statement—in other words, the last atomic action before the process executes *entry_i*. What is the relation between the atomic action *enter_i* and the operation *request_i*? There are two possibilities:

1. *enter_i* equals *request_i*, making *request_i* an atomic action.
2. *enter_i* is the first action of the nonatomic operation *request_i*.

I will examine each of them in turn.

In the first case, where *request_i* is the atomic action putting the process at its *entry_i* statement, condition (*)

cannot be satisfied by any implementation. There is no way for two *entry* statements to determine in which order they were entered. Hence, no algorithm can ensure that the *critical.section* operations occur in the order required by condition (*).

It might seem unfair not to make *enter_i* part of the the *entry_i* statement, and one might define *enter_i* to be the first atomic action of *entry_i*. However, this does not solve the problem because an atomic action of a Blaise program can either read or write a shared variable, but cannot do both. Thus, if *request_i* is the atomic action *enter_i*, then the *request_i* operation cannot both announce the process's desire to enter its critical section and check for the presence of other processes waiting to enter their critical sections.² If the two operations *request_i* and *request_j* occur too close together, no algorithm can determine which one happened first, even though the semantics of Blaise specifies that they occur in some definite order. Hence, a Blaise implementation still cannot satisfy condition (*).

Now consider the second case. If *enter_i* is only the first action of the *request_i* operation, when does the operation end? This question is not answered by the specification. Since the end of the *request_i* operation happens while executing the *entry_i* statement, which is provided by implementor, he must be the one who decides where the *request_i* operation ends. In order to prove that his implementation meets condition (*), the implementor may define the end of the *request_i* operation to be anywhere he wishes. In particular, he can define *request_i* to include the entire execution of the statement *entry_i*. With this definition, any algorithm that enforces mutual exclusion of *critical.section* operations trivially satisfies condition (*). Thus, in this case, the condition is vacuous.

What Does Priority Really Mean?

What do we mean when we say that something is an FCFS algorithm? FCFS was defined in [10] by condition (*), under the interpretation in which *request_i* is a nonatomic operation, with the following additional constraint:

- (†) The *request_i* operation does not involve any waiting for other processes.

To prove that his program is FCFS, the implementor is free to define *request_i* any way he likes so long as condition (†) is satisfied. For a Blaise program, in which "busy waiting" is the only kind of waiting possible, (†) is satisfied if the execution of *request_i* takes a bounded number of program steps.³

There is no obvious way to define absence of waiting for a Tony program, which can make it difficult to decide

²This appears to be a folk theorem, having been known to a number of people but never published.

³Note that this does not mean that the execution takes a bounded length of time; Blaise does not guarantee any bound on how long a process may wait before executing one step.

whether or not a Tony program is an FCFS algorithm. This is the reason the challenge requires that a Tony simulation of an FCFS Blaise program be regarded as FCFS. The simulation can be viewed as a "compiled version" of the Blaise program, and it is reasonable to expect the compiled version of an FCFS program also to be an FCFS program.

To specify FCFS, one needs both (*) and (†), so one must be able to define what waiting means. Moreover, the definition of waiting should be independent of the programming language, since the specification of the *request_i* operation should not depend upon how *entry_i* and *exit_i* are implemented. For example, *entry_i* and *exit_i* might call subroutines that invoke special-purpose hardware to perform the necessary synchronization.

Specifying other kinds of priority poses exactly the same difficulty as specifying FCFS. Consider the classical readers/writers problem with writer priority [2]. In this problem, a process that has issued a request to write has precedence over a reader that has not yet begun to read. Letting *start_i.rd_i* denote the operation of starting to read, this is expressed as follows.

(*)' For any operations *write_i* and *read_j*, if *request_i* precedes *start_i.rd_j*, then *write_i* must precede *read_j*.

Condition (*)' has the same trouble as condition (*). If *request_i* is the operation *enter_i* that begins the request, then no Blaise program can meet this specification. On the other hand, if the implementor is allowed to define the endpoint of this operation, then condition (*)' is vacuous because *request_i* can be defined to extend until the beginning of *write_i*.

4 A Closer Look at Specification

An Informal Look

The difficulty in specifying priority should convince the reader that we need to examine more closely what it means for a program to implement a specification. To write a specification, there must be an object to be specified and a well-defined interface between the object and its environment. We can ask for a specification of a telephone exchange because we know both the object to be specified (the exchange) and its interface with the environment (the wires leading to the telephones on the exchange and to other exchanges). It is meaningless to ask for a specification of the solar system because we have no idea what the interface is between the solar system and its environment.

I will call the object being specified a "module". A complete specification of a module must contain all the information needed to determine if a particular implementation is correct, where correctness means that the module interacts properly with its environment. An examination of the specifications presented to illustrate most methods—for example, the specification of a queue (bounded buffer) in [8], [17], or [18]—reveals that they are incomplete. From these specifications, one cannot tell whether the operations

are initiated by calling a subroutine or by raising a voltage on a wire. A program and a piece of hardware cannot both interact properly with the same environment. Only in [11] is the interface specified, being defined as a simple subroutine-calling mechanism, but there was no explanation of why this implementation-level detail was introduced into a paper on specification.

A complete specification must have two parts: a specification of the module's interface and a specification of its internal behavior. The internal behavior can be specified in terms of high-level abstractions like queues and *write* operations. However, since the interface determines whether an operation is initiated by calling a Pascal subroutine named *put* or by raising the voltage on line number 7 to 4.5 ± 1 volts, it must be specified in terms of implementation-level concepts like subroutine names and voltages. We want to make the interface specification as small as possible, specifying as much as we can in terms of the internal behavior, which can be described with nice, high-level concepts; but the interface specification is necessary. Most specification methods ignore the interface and consider only the internal behavior.

The implementor should have complete freedom in implementing the objects and operations that describe the internal behavior. If the specification contains an internal operation that puts an object on a queue, then the implementor can define that operation to be the act of storing an item in an array, of adding it to a linked list, or of setting the voltage levels on the wires leading to some special device. On the other hand, the interface must be completely specified at the implementation level. The need to partition a specification into an internal part, which is implementation-independent, and an interface specification, which depends upon the implementation, was recognized in [6], and is embedded in the Larch system [7].

In the challenge, the interface is described by requiring the implementation to consist of Blaise or Tony *entry_i* and *exit_i* statements that are used in a particular way. Because I ignored the problem of how shared variables and extra processes are declared, I could pretend that the Blaise and Tony implementations had the same interface specification. In a more formal approach, the interface would have to be specified somewhat differently for the two languages.

A Formal View

As described in condition 1 of the challenge, a formal method for proving that an implementation meets its specification must convert the specification and its implementation into a mathematical formula *C* in some formal system *L* such that the implementation is correct if and only if *C* is a valid formula of *L*. I now give a very vague, high-level discussion of how this is done.

A formal specification is written in a language having a formal semantics, which means that the specification can be translated to a mathematical object *S* in some formal system *S*. Similarly, a formal semantics for the implement-

tation language describes the implementation as a semantic object I in a formal system \mathbf{I} . To be able to speak formally about the correctness of the implementation, there must be a mapping \mathcal{F} from objects in the system \mathbf{S} to objects in the system \mathbf{I} , so that $\mathcal{F}(S)$ is an object of \mathbf{I} . The object $\mathcal{F}(S)$ is the formal representation of the specification in the semantic domain of the implementation.

The formal system \mathbf{L} of the challenge is the system \mathbf{I} , and the formula \mathcal{M} that expresses the correctness of the implementation is the formula of \mathbf{I} that means " I satisfies $\mathcal{F}(S)$ ". Exactly how \mathcal{M} is constructed from $\mathcal{F}(S)$ and I depends upon the specification method. I will illustrate with two examples: a pure axiomatic approach and a pure behavioral approach.

In a pure axiomatic approach, an axiomatic semantics is given for both the specification and the implementation. An axiomatic semantics defines S to be a formula of the logical system \mathbf{S} —the conjunction of all the "axioms" comprising the specification—and I to be a formula of \mathbf{I} . The mapping \mathcal{F} is a function from the formulas of \mathbf{S} to those of \mathbf{I} . For example, suppose the specification is in terms of the value of a queue q , which is implemented with an array a . To talk about the correctness of the implementation, for every possible value a of the array a we must know the value $Q(a)$ of q that it represents.⁴ For any formula \mathcal{R} of \mathbf{S} , the formula $\mathcal{F}(\mathcal{R})$ of \mathbf{I} is obtained by substituting $Q(a)$ for q in \mathcal{R} . Thus $\mathcal{F}(\mathcal{R})$ is obtained by translating the statement \mathcal{R} , which is an assertion about the specification-level object q , into an assertion about the implementation-level object a .

In this approach, \mathcal{M} is the formula $I \supset \mathcal{F}(S)$. In other words, the implementation is correct if and only if the axioms comprising the semantics of the implementation imply the axioms of the specification, after the latter are translated by \mathcal{F} into assertions about the implementation. This is discussed at greater length in [12] for one particular axiomatic method.

In a pure behavioral approach, the formal semantics of the implementation and specification are sets of behaviors: S is the set of all behaviors allowed by the specification, I is the set of all behaviors that could be produced by the implementation, and \mathbf{S} and \mathbf{I} are formal systems for reasoning about sets of behaviors. For a behavior b in the specification domain, $\mathcal{F}(b)$ is the corresponding behavior in the implementation domain. In the mutual exclusion example, the operation *critical section* _{i} is a single action in the behavior b ; it corresponds to a set of actions in $\mathcal{F}(b)$ —namely, the set of all the Blaise program steps in a single execution of the critical section.

One can define the formula \mathcal{M} to be $I \subset \mathcal{F}(S)$, where $\mathcal{F}(S) = \{\mathcal{F}(b) : b \in S\}$. In other words, the implementation is correct if and only if every possible behavior of the implementation is allowed by the specification. Some specification methods define \mathcal{M} to be $I = \mathcal{F}(S)$, requiring that the implementation be able to produce all the behaviors

⁴The mapping Q may be partial, since $Q(a)$ need only be defined for values a of a that can arise during the program's execution.

described by the specification.

There are other possibilities—for example, an axiomatic semantics for the specification and a behavioral semantics for the implementation. In any case, the definition of correctness of the implementation involves the mapping \mathcal{F} . For the specification of FCFS, the mapping \mathcal{F} is what determines which operations at the implementation level correspond to the specification's *request* _{i} operation. A complete specification must include not only S , but also the part of \mathcal{F} that determines the interface. For the queue example, it is this part of \mathcal{F} that specifies whether one puts an element in the queue by calling a subroutine or raising a voltage level. Correctness means that there exists some \mathcal{F} , part of which is determined by the specification, such that $\mathcal{F}(S)$ satisfies I . The implementor is free to define the rest of \mathcal{F} , specifying the correspondence between the implementation and the internal part of the specification any way he wishes in order to prove the correctness of his implementation. The specification places no constraint on any part of the implementation other than the interface.

5 Conclusion

Having described the difficulty in specifying priority, it would be nice if I could either explain how it can be done or else prove that it is impossible. Unfortunately, I can do neither. I believe that one cannot write a satisfactory general specification of priority—one that works for a variety of implementation domains. The difficulty in expressing priority arises from the requirement that the *request* operation should involve no waiting for other processes. Waiting is an implementation-level concept that I feel cannot be expressed in a general way. However, this conjecture, like Church's thesis, is not susceptible to formal proof. At best, one can prove only that some particular formalism cannot express priority.

If priority is not expressible by current formal specification techniques, how should we specify concurrent systems? Priority is generally regarded to be a fundamental concept that must be specified. Must we add new primitives to express it? My tentative answer is no. I believe that priority cannot be expressed precisely in those situations when it is not a fundamental property.

Remember that condition (*) does express FCFS priority if the *request* _{i} operation is interpreted to be the interface operation *enter* _{i} . The atomicity of *enter* _{i} is irrelevant; what matters is that *request* _{i} be the interface operation. Priority is a basic system requirement only when its effect is directly visible to the user, which is the case only when the *request* operation is externally visible—that is, when it is part of the interface. For example, suppose we want transactions issued by certain users to receive higher priority. The *request* operation can then be defined as the entire sequence of actions performed by the user in issuing the request, from the first keystroke to his notification that the request has been accepted by the system.

When the *request* operation is not externally visible, then

must also do the same for other language constructs besides the “;”. We will see that there is a standard prescription for doing this.

The axioms in $M[S]$ define a set of behaviors for S —namely, the set of all behaviors satisfying the temporal logic axioms starting in states that satisfy the axioms for the starting state. Although this defines a set of behaviors for every statement, it is different from a semantics in which $M[S]$ is taken to be a set of behaviors because the meaning of S is obtained from the meaning of its components by “composing” axioms, not by composing behaviors.

2.3 Is This Fair?

It can be argued that the semantics of a programming language should be defined in terms of constructive operations rather than with axioms. One should give a procedure for constructing the set of behaviors of a program rather than a set of axioms to describe it.

While a purely constructive approach would be nice, it seems to be impossible to deal with fairness constructively. Even a behavioral semantics, which looks constructive, really includes axioms for fairness. A behavioral semantics defines the meaning of a fair cobegin in terms of fair interleaving. The definition of a fair interleaving of two behaviors goes something like this:

Construct all interleavings and then throw away the ones that do not satisfy the fairness condition.

This is remarkably similar to the definition of the set of behaviors obtained from a set of actions and a set of constraints, which can be expressed as:

Construct all behaviors generated by the set of actions and then throw away those that do not satisfy the constraints.

One might argue that fair interleaving is a simple, basic concept, and I have given a particularly jaundiced expression of it. However, there are many different fairness constraints one might want to define, each of which would require a different definition of fair interleaving. For example, consider two coin-flipping processes, one with the single action *head* and the other with two actions: *tail* and *coin lost*. The first process generates only the sequence of all heads, the second process generates either a sequence of all tails or a finite string of tails followed by a sequence of *coin lost* actions. The behaviors resulting from executing the two processes concurrently are defined to consist of all possible fair sequences of heads and tails plus all

cobegin $\pi_1 \square \pi_2$ coend

is just the union of the sets of possible actions of π_1 and π_2 . Action semantics have long been favorites of theoretical computer scientists [15] because they lead to mathematically well-behaved formalisms. Unfortunately, these semantics are unsatisfactory because they cannot express fairness. Consider a coin-flipping program with two possible actions: toss a head and toss a tail. It can be viewed as the parallel composition of two processes: one that generates only heads and the other that generates only tails. An unfair coin flipper can generate any infinite sequence of heads and tails, while a fair one can generate only sequences containing infinite numbers of both heads and tails. Both the fair and the unfair coin flipper have the same set of actions (toss a head and toss a tail), so an action semantics cannot distinguish between the two.

2.2.3 Action-Axiom Semantics

The problem of fairness is solved by using an *action-axiom* semantics in which the meaning of a statement consists of a set of actions together with a set of temporal logic axioms that state conditions under which an action must eventually occur. For example, the fair coin flipper requires two axioms:

- At any time, a head must eventually occur.
- At any time, a tail must eventually occur.

The meaning of

assign processor to S

consists of the meaning of S plus the following additional axiom:

- If S is being executed—more precisely, if control is in S —then an action of S must eventually occur.

We are thus led to let $\mathcal{M}[S]$ consist of a set of actions, a set of temporal logic axioms, and a set of starting states. But how do we specify the actions? Instead of introducing some new method for specifying actions, I will specify the actions as well as the fairness properties with temporal logic axioms. Starting states will be specified by ordinary, nontemporal axioms.

To give a compositional action semantics, we must define the axioms of $\mathcal{M}[S_1; S_2]$ in terms of the axioms of $\mathcal{M}[S_1]$ and $\mathcal{M}[S_2]$; and, of course, we

cobegin $S_1 \parallel S_2$ coend

are obtained by forming interleavings of behaviors from S_1 and S_2 , and interleavings are rather awkward mathematically—especially for a fair cobegin, where only fair interleavings are allowed.

A more serious problem is raised by the language construct

assign processor to S

Intuitively, this statement causes the compiler to assign a physical (or virtual) processor to execute S . In terms of behaviors, it means that any behavior that reaches S must either subsequently reach the end of S or else include an infinite number of actions of S . In other words, a process cannot be “starved” while it is executing this statement. This is a perfectly reasonable—and compilable—statement. It can be used to construct a fair cobegin from an unfair one as follows:

**unfair cobegin assign processor to $S_1 \parallel$
 assign processor to S_2 coend**

More complicated uses of the **assign processor** statement are also possible.

Considered completely by themselves, the statements S and

assign processor to S

have the same sets of behaviors, so it is not clear how one could apply to the **assign processor** statement the same approach used above to define $\mathcal{M}[S_1; S_2]$.

2.2.2 Action Semantics

Instead of defining $\mathcal{M}[S]$ to be the set of behaviors itself, one can define it to be something that can be used to construct the set of behaviors. Since a behavior is generated by a sequence of actions starting in some state, an obvious approach is to let $\mathcal{M}[S]$ be the set of all possible actions together with the set of all possible starting states. I will call such a semantics an *action semantics*. Given an action semantics, one can define the behaviors of S to be the set of all behaviors that can be obtained from these actions starting from the specified starting states.

An action semantics is well-suited to expressing parallelism, since the set of possible actions of

the α_i are actions of that process. I will explain later exactly what states and actions are.

The semantics that I am aiming for is an axiomatic one, in which the meaning of a program is a set of axioms in a formal system. An important advantage of an axiomatic semantics is that it is very formal. A formal mathematical system is one in which reasoning can be reduced to a strict application of axioms and inference rules. Automated deduction systems can usually be applied only to a formal system. A semantics in which $M[S]$ is defined to be a set of sequences is really semi-formal, based upon the informal mathematical concepts of sets and sequences. Formalizing it requires formalizing these mathematical concepts. With an axiomatic semantics, this extra step is unnecessary; $M[S]$ is already a set of axioms in a formal system.

The problem with an axiomatic semantics is that one can understand the meaning of a formal logical system only by constructing a semantic model for it in terms of concepts that we already understand. Having constructed a semantics in which $M[S]$ is a set of axioms in some formal system is only half the job; we also have to define a semantics for the formal system in terms of well-understood mathematical concepts.

I will give a temporal logic semantics—one in which the axioms are temporal logic formulas. I will rely upon the usual semantic model of temporal logic, described later, to provide a basis for an intuitive understanding of the axioms.

2.2 Different Kinds of Semantics

2.2.1 Behavioral Semantics

An obvious method of defining a semantics for concurrent programs is to let the meaning of a statement be its set of possible behaviors, and to explicitly construct the behaviors in $M[S]$ from the behaviors of its components. For example, the set of behaviors $M[S_1; S_2]$ consists of all infinite behaviors in $M[S_1]$ together with all concatenations of finite behaviors in $M[S_1]$ with behaviors in $M[S_2]$. This can be expressed formally by:

$$M[S_1; S_2] = \{\sigma \in M[S_1] : \sigma \text{ infinite}\} \\ \cup \{\sigma\tau : \sigma \in M[S_1], \tau \in M[S_2], \text{ and } \sigma \text{ finite}\}$$

I will call such a semantics a *behavioral* semantics.

Behavioral semantics have their problems. While they work well for sequential programming constructs, they are less satisfactory for concurrent languages. The behaviors of

to reason about things like the control state that are internal to the program. However, years of experience reasoning about concurrent programs has led me to conclude that one should think about them in terms of the complete state, including externally invisible components of the state. I will not attempt to justify this conclusion here, and will simply adopt the second approach, defining the meaning of a program in terms of complete behaviors that describe the internal as well as the externally-visible effects of program operations.

Having decided that the meaning of a program is its set of possible behaviors, we must decide what a behavior is. The simplest notion of a behavior is a sequence of states. Each action of the program transforms the state. Nondeterminism, leading to sets of behaviors, appears when there are several choices of a possible next state from the same current state.

It is sometimes argued that a sequence of states cannot adequately model the execution of a concurrent program because it has no notion of concurrent activity, and that one should instead use a partially ordered set of actions. However, a partially ordered set contains exactly the same information as the set of all total orderings consistent with the partial order. Since the meaning of a statement is the *set* of behaviors, which includes all possible sequences that represent the real, partially ordered set of actions, nothing has been lost by considering sequences. The basic assumptions being made are that the execution of a program consists of discrete atomic actions, and the possible effect of an atomic action depends only upon the current state. It appears that any digital system can be accurately modeled in this way by making the atomic actions small enough and including enough information in the current state.

It turns out that to define the semantics of concurrent languages, one needs more information about a behavior than just the sequence of states; one must also know "who" performed the actions. For example, the natural definition of a fair *cobegin* states that in each infinite behavior, every nonterminating process performs infinitely many actions. Formalizing this definition requires the ability to decide which process performs each action. I will therefore define a behavior to be a sequence of the form:

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots$$

where the s_i are states and the α_i are actions. Fairness of a *cobegin* can be expressed by stating that for every process and every n : if there is no state s_i with $i > n$ in which the process has terminated, then infinitely many of

of S . For example, $M[S_1; S_2]$ should be defined in terms of $M[S_1]$ and $M[S_2]$. I will say that a semantics with this property is *compositional*.¹

When defining a formal semantics, the first thing one has to decide is what kind of object $M[S]$ should be. The execution of a statement in a sequential program is usually considered to start in some input state and produce an output state, and $M[S]$ is defined to be a mathematical object that describes the relation between the input and the output states. One way of doing this is to define $M[S]$ to be a set of ordered pairs of states.

Concurrent programs cannot be described with such a simple input/output semantics. Consider the following two program statements, where angle brackets denote indivisible atomic operations.

1. $\langle x := x + 1 \rangle$
2. **begin** $\langle x := x + y \rangle$;
 $\langle x := x - y + 1 \rangle$
end

These statements both have the same relation between input and output states—they both increment the value of x by one. However, they are not equivalent when used as part of a concurrent program. Executing the first always has the effect of adding one to x , but executing the second can have a very different effect if the value of y is changed by some other process between the two assignments to x .

A semantics for a concurrent programming language must define the meaning of a statement in terms of its behavior. There are two fundamentally different approaches to doing this. The first approach is to define the meaning of a statement S in terms of the effects it produces that are “visible” outside S . For example, in a shared-variable language, the only visible effects of executing a statement are changes to shared variables. The alternative approach to defining the semantics of concurrent languages defines the meaning of a statement in terms of complete behaviors, which include all the effects of a statement’s actions, whether externally visible or not. In most languages, these invisible effects include changes to the control state (the values of “program counters”).

An approach that mentions only visible effects is very appealing, and it has been taken by a number of researchers [3,11]. For many years, I regarded it as the proper way to think about programs, and found it unnatural

¹The terms *denotational*, *syntax-directed*, and *modular* have also been used to denote this property.

2 An Introduction to Semantics

2.1 What Are Semantics?

The syntax of a programming language defines the set of syntactically well-formed programs of that language. However, a program is more than just a string of characters; there should be a well-defined set of possible results of executing the program. The purpose of a semantics is to assign a mathematical *meaning* to each syntactically correct program that describes the effect of executing it.

I will regard a program Π to be a syntactic object, and denote by $\mathcal{M}[\Pi]$ the mathematical object denoting its meaning. To define a formal semantics, one must specify the mapping $\Pi \rightarrow \mathcal{M}[\Pi]$.

What is the purpose of a formal semantics? One purpose is to help us to understand the language. However, "understanding" is too vague to usefully characterize a formalism. I propose that a formal semantics should provide a formal basis for the following:

1. Deducing properties of a program written in the language.
2. Deciding if a compiler is correct, given a formal semantics for the target language into which the programs are compiled.

A semantics should provide a formal foundation, but not necessarily a practical method, for doing these things. A method for deducing properties of a program is called a *proof system*. A proof system is used to decide if a program works properly; a semantics is used to decide if a programming language is defined properly. One wants to reason about programs at a high level, hiding as much detail of the language as possible; a semantics should expose the language details. Although a semantics allows one, in principle, to verify properties of programs, its real purpose is to explain the language. A semantics should be used to verify the correctness of a proof system; it need not provide a practical method for reasoning about programs.

Programs can be very large and complicated. We want to reduce the problem of understanding a complex program to that of understanding its components. The meaning of a program should therefore be defined in terms of the meanings of its components. We must therefore define the meaning not just of an entire program, but of individual components—usually individual program statements. So, $\mathcal{M}[S]$ must be defined for any program statement S , and it must be defined in terms of the substatements

The first idea was developed by Susan Owicki and myself in the late 1970's, and was published in [6]. (It was developed independently, in different contexts, by other researchers [1].) The second and third ideas were developed by me shortly afterwards. The second was also used in [6], though not featured prominently there. The third idea has never appeared in print, though I have talked about it in lectures starting in 1981. The fourth idea was discovered by Fred Schneider and myself in the spring of 1984 [10]. The fifth idea has been present in all of my work on temporal logic, starting with [5]. I was originally led to it by my philosophical objections to the "next time" operator; only later did I recognize its practical significance [8].

The first two ideas were used in [14], but they are not enough to permit a compositional semantics based upon a simple temporal logic. Combined with the third idea, they do permit a compositional semantics, but a semantics that I did not find satisfying. It seemed like a large, complicated structure had to be erected solely to reason about program control, making the enterprise of dubious merit. It was the fourth idea that gelled the method into a coherent form. The apparatus for handling program control was no longer an *ad hoc* "Kludge". Rather, it was the appropriate structure to deal with aliasing. Aliasing was not considered in other approaches, but it is a problem that must be dealt with in any realistic language, if only to handle procedure calls. The fifth idea is not needed for the semantics itself; in fact the semantics would be somewhat easier to understand had I abandoned it and employed the next-time operator favored in most other temporal logic approaches. However, allowing stuttering actions enables the semantics to address the practical issue of what it means for a compiler to be correct.

In this paper, I develop these five ideas, and show how they lead to a method for defining the semantics of concurrent programming languages. A complete semantics is given only for a simple language. However, the approach is "meta-compositional" in the sense that the meaning of each language construct is defined independently of the other constructs in the language. The semantics of a richer language can be given by defining the meanings of its additional constructs, without changing the meanings of the constructs from the simple language. (This is not the case in [14] where, for example, the axioms for the assignment statement would be invalid if an unfair *cobegin* were added to the language.) I will indicate the power of my method by informally describing how the meanings of some more complicated language constructs can be defined.

1 An Introduction to this Paper

A large body of research on the logic of concurrent programs may be characterized as the "axiomatic" school. Members of this school reason about safety properties ("something bad never happens") in terms of invariance, and liveness properties ("something good eventually does happen") using temporal logic.

While they are quite successful at proving properties of a given program, axiomatic methods have not provided a satisfactory semantics for concurrent programming languages. Axiomatic methods usually reason about the entire program, while a semantics should be compositional—deriving the meaning of a program from the meanings of its components. Even the Generalized Hoare Logic described in [4] and [9], which looks compositional, actually assumes a context of a complete program. The only attempt we know of at a truly compositional axiomatic semantics for concurrent programs that handles both safety and liveness properties is given in [14]. However, while it is axiomatic in a strict logical sense, that approach is not in the spirit of the axiomatic school because it essentially defines a new temporal logic operator for every programming-language construct.

In this paper, I present a new compositional, truly axiomatic semantics for concurrent programming languages. It is based upon temporal logic, but employs five fundamental ideas beyond those found in most temporal logic methods:

1. The addition of action predicates to describe "who" performs an action.
2. Defining an assertion to be true of a statement only if it is true of every program containing that statement.
3. The introduction of renaming operations that map an assertion about a statement S into an assertion about a larger statement containing S as a substatement.
4. Defining the relations between control points, described in [4] as state predicates, to be aliasing relations among variables.
5. Allowing "stuttering" actions, so an atomic operation is represented by a finite sequence of actions, only the last one having any effect.

6.1.4	Logical Name Variables	45
6.2	Starting States	45
6.3	Behavior Axioms	46
6.3.1	Assignment	47
6.3.2	The if Statement	49
6.3.3	The while Statement	49
6.3.4	The new Statement	50
6.3.5	The cobegin Statement	51
6.3.6	Sequences of Statements	51
6.3.7	A Complete Program	51
7	Other Language Features	52
7.1	Constructs That Constrain Their Environment	52
7.1.1	The assign processor Command	52
7.1.2	Atomic Actions	52
7.1.3	Write Protection	53
7.2	Synchronization and Communication	53
7.2.1	Semaphores	53
7.2.2	CSP-Like Communication Primitives	54
7.3	Procedures	55
7.4	More General Types and Aliasing	56
7.5	Nonatomic Operations	57
8	Conclusion	59

Contents

1	An Introduction to this Paper	1
2	An Introduction to Semantics	3
2.1	What Are Semantics?	3
2.2	Different Kinds of Semantics	6
2.2.1	Behavioral Semantics	6
2.2.2	Action Semantics	7
2.2.3	Action-Axiom Semantics	8
2.3	Is This Fair?	9
2.4	Programs and Implementations	10
2.4.1	Correctness of an Implementation	10
2.4.2	The Interface	13
3	The Programming Language	16
4	States and Actions	21
4.1	Program Variables	21
4.2	Control Variables	23
4.3	Are There Other State Components?	25
4.4	Renaming	26
4.5	Actions	27
4.6	States and Actions: A Formal Summary	28
5	Temporal Logic	30
5.1	Predicates	30
5.2	The Unary Temporal Operators	31
5.3	The Binary Temporal Operators	32
5.4	Renaming	34
5.5	Temporal Logic as Semantics	34
5.6	There Won't Be a Next Time	35
5.7	Implementation Mappings	36
6	The Semantics of Language L	39
6.1	Syntactic Predicates	39
6.1.1	Aliasing	41
6.1.2	Syntactic Typing Relations	44
6.1.3	Reasoning About Syntactic Expressions	44



An Axiomatic Semantics of Concurrent Programming Languages

Leslie Lamport

19 September 1984

To appear in the lecture notes of the Advanced Seminar on Logics and Models for Verification and Specification of Concurrent Systems, held in La-Colle-Sur-Loup, France in October, 1984.

Work Supported in part by the National Science Foundation under grant number MCS-8104459 and by the Army Research Office under grant number DAAG29-83-K-0119.

APPENDIX B

**AN AXIOMATIC SEMANTICS OF CONCURRENT
PROGRAMMING LANGUAGES**

priority is a mechanism, not an end in itself. In the internal specification, we give writers priority not because correctness requires them to have precedence, but rather to ensure that they receive adequate service. For example, a common use of writer priority is to guarantee absence of starvation—a waiting writer eventually writes despite a continual stream of readers. What we need to specify is the required service, not the mechanism used to achieve it. The absence of starvation belongs to a fundamental class of properties, known as liveness properties, that are easily expressed in temporal-logic based methods like the ones of [8], [11], [13], [16], [17], and [18]. It is the basic requirement that should be specified, not the priority mechanism used to achieve it.

Acknowledgement

The difficulty in specifying priority was discovered during a discussion with Richard Schwartz and Michael Melliar-Smith, in which they kept poking holes in my attempts to specify FCFS until we all finally recognized the fundamental problem. In writing this paper, I was aided by the helpful comments of Brent Hailpern and Fred Schneider and the enthusiastic objections of the members of IFIP Working Group 2.2.

Bibliography

- [1] John Bethel, ed. *Webster's New Collegiate Dictionary*. G. & C. Meriam Co., Springfield, Mass. (1956).
- [2] P. J. Courtois, F. Heymans and D. L. Parnas. Concurrent Control with "Readers" and "Writers". *Comm. ACM* 14, 10 (October 1971), 667-668.
- [3] E. W. Dijkstra. Solution of a Problem in Concurrent Programming Control. *Comm. ACM* 17, 11 (November 1974), 643-644.
- [4] Susan Gerhart et al. An Overview of AFFIRM: A Specification and Verification System, *IFIP Congress 80*, (Oct. 1980).
- [5] Irene Greif. A Language for Formal Problem Specification. *Comm. ACM* 20, 12 (Dec. 1977), 931-935.
- [6] J. V. Guttag and J. J. Horning. Formal Specification as a Design Tool. *Proc. ACM Symposium on Principles of Programming Languages*, Las Vegas, (January 1980), 251-261.
- [7] J. V. Guttag and J. J. Horning. An Introduction to the Larch Shared Language. *Proc. IFIP Congress '83*, Paris, (1983).
- [8] Brent Hailpern. *Verifying Concurrent Processes Using Temporal Logic*, Lecture Notes in Computer Science 129, Springer-Verlag (1982).
- [9] Howard Katseff. A Solution to the Critical Section Problem with a Totally Wait-free FIFO Doorway. Technical Memorandum, Computer Science Division, University of California, Berkeley.
- [10] Leslie Lamport. A New Solution of Dijkstra's Concurrent Programming Problem, *Comm. ACM* 17, 8 (Aug. 1974), 453-455.
- [11] Leslie Lamport. Specifying Concurrent Program Modules, *ACM Trans. on Prog. Lang. and Systems* 5, 2 (Apr. 1983), 190-222.
- [12] Leslie Lamport. What Good is Temporal Logic? *Information Processing 83*, R. E. Mason (ed.), Elsevier Science Publishers (North-Holland), 1983, 657-668.
- [13] Amy Lansky and Susan S. Owicki. GEM: A Tool for Concurrency Specification and Verification. *Proc. of the Second Annual ACM Symposium on Principles of Distributed Computing* (Aug. 1983), 198-212.
- [14] Peter E. Lauer, P. Torrigiani and M. Shields. COSY: A System Specification Language Based on Paths and Processes, *Acta Informatica* 12 (1979), 109-158.
- [15] Robin Milner. *A Calculus of Communicating Systems*, Lecture Notes in Computer Science 92, Springer-Verlag (1980).
- [16] Richard L. Schwartz and P. M. Melliar-Smith. Temporal Logic Specification of Distributed Systems. *Proc. of the IEEE Conference on Distributed Systems*, (Apr. 1979).
- [17] Richard L. Schwartz, P. M. Melliar-Smith and F. H. Vogt. An Interval Logic for Higher-Level Temporal Reasoning, *Proc. of the Second Annual ACM Symposium on Principles of Distributed Computing* (Aug. 1983), 198-212.
- [18] Pierre Wolper. Specification and Synthesis of Communicating Processes Using an Extended Temporal Logic. *Proc. of the Conference on the Principles of Programming Languages*, Albuquerque (Jan. 1982).

sequences consisting of a finite number of heads and tails followed by nothing but *coin lost* actions.

This is a perfectly reasonable example, which the reader may find more familiar if he replaces *coin lost* by *abort program*. A behavioral semantics for this way of combining processes would require a more complicated definition of fair interleaving, and a formal statement of this definition would look a lot like a temporal logic axiom. Fair interleaving is not a simple concept. One particular type of fair interleaving has been used so commonly that we tend to take it for granted and forget that we have never seen a constructive definition of it.

Fairness does not appear to be a constructive concept. One specifies fairness by adding axioms to exclude unfair behaviors rather than by explicitly constructing only the fair ones. Infinite objects, such as behaviors, are constructed as limits of finite approximations—a method often described as “denotational”. This does not work with fairness because there exist sequences of fair behaviors whose limits are unfair—for example, let $\sigma_1, \sigma_2, \dots$ be the sequence of coin-flipping behaviors in which all actions of σ_n are heads, except for every 2^n th action, which is a tail. Each σ_n is fair, but the limit as n goes to infinity is the behavior having only heads, which is unfair.

The topological approach of [2] solves this problem by considering only convergent sequences and defining a topology in which sequences like the above diverge. However, one might view this approach as:

Construct all sequences obtainable from the actions and throw away those that do not converge.

This looks suspiciously like the more overtly axiomatic approach. The whole distinction between constructive and axiomatic methods is probably illusory, disappearing when methods are examined closely enough.

2.4 Programs and Implementations

2.4.1 Correctness of an Implementation

One question that a semantics of a programming language should answer is: What does it mean for a compiler to be correct? Given a program Π in the high-level language, the compiler transforms it into a program π in some lower-level language. Correctness of the compiler means that π is a correct implementation of Π , but what does that mean? To speak of correctness, we must have formal semantics for both the high-level and the low-level

languages, so $\mathcal{M}[\Pi]$ and $\mathcal{M}[\pi]$ are defined. However, this is not enough to determine what it means for $\mathcal{M}[\pi]$ to represent a correct implementation of $\mathcal{M}[\Pi]$.

Consider the case of sequential programs, in which the semantics of a program is a relation on the set of program states, the pair (s, t) being in the relation $\mathcal{M}[\Pi]$ if and only if it is possible for program Π to start in state s and terminate in state t . In this case, $\mathcal{M}[\Pi]$ and $\mathcal{M}[\pi]$ are relations on two different sets of states. The states of Π specify the values of program variables like x and y ; the states of π might specify the values of machine registers like *memory location 3124* or the *program counter*. Correct implementation means that there is a correspondence between the sets of states of Π and π such that, under this correspondence, every possible execution of π is a possible execution of Π .

More formally, to establish a correspondence between the semantics of the two sequential programs, we must define a mapping F from the states of π to the states of Π . For example, suppose the variable x in Π of type *integer* is implemented in π as a two-byte integer stored in bytes 3124 and 3125 of memory. If, in a state s of π , bytes 3124 and 3125 have the values 12 and 97, then the value of x in the state $F(s)$ of Π is $12 \times 256 + 97$. In general, correctness of the implementation means that for each pair (s, t) in $\mathcal{M}[\pi]$, the pair $(F(s), F(t))$ must be in $\mathcal{M}[\Pi]$.

What about concurrent programs? As we have seen, the meaning of a concurrent program must be expressed in terms of its behavior—either directly, with a behavioral semantics, or indirectly with axioms about its behavior. Let us therefore consider first a behavioral semantics, in which $\mathcal{M}[\Pi]$ and $\mathcal{M}[\pi]$ are sets of behaviors. Intuitively, π is a correct implementation of Π if every possible behavior of π represents a possible behavior of Π . We therefore need some way of interpreting behaviors of π as possible behaviors of Π —that is a mapping F such that for any behavior σ in $\mathcal{M}[\pi]$, $F(\sigma)$ is a sequence of states and actions of Π . We can then say that π is a correct implementation of Π if, for every σ in $\mathcal{M}[\pi]$, $F(\sigma)$ is in $\mathcal{M}[\Pi]$.

In defining this mapping F , we are faced by the problem that Π and π may have different grains of atomicity. An atomic operation of Π may be implemented by a sequence of 42 atomic operations of π . For example, the atomic operation

$$(x := x + 1)$$

of Π might be implemented in π by 42 machine-language operations. Moreover, interleaved among these 42 atomic operations of π might be other

machine-language operations that belong to the implementation of an operation from a different process of Π . It would therefore seem that the mapping F must be quite complicated, taking sets of actions into single actions.

There is very simple solution to this problem—we require that to every action of π there correspond a single action of Π . The execution of a single atomic operation of Π might therefore be represented by 42 actions in a behavior in $\mathcal{M}[\Pi]$. The first 41 of these actions will be “stuttering” actions that do not change the state of Π ; the 42nd will do all the work. This makes it conceptually very easy to define the mapping F from behaviors of π to behaviors of Π . As in the sequential case, there must be a mapping F from states of π to states of Π . We also assume that F maps actions of π to actions of Π —for example, every machine-language instruction executed by π corresponds to the execution of some atomic operation of Π .² To extend F to a mapping on behaviors, if σ is the behavior

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots$$

of π , we define $F(\sigma)$ to be the behavior

$$F(s_0) \xrightarrow{F(\alpha_1)} F(s_1) \xrightarrow{F(\alpha_2)} \dots$$

The implementation is correct if, for every behavior σ of $\mathcal{M}[\pi]$, $F(\sigma)$ is a behavior in $\mathcal{M}[\Pi]$.

This seems nice in theory, but how can it be achieved in practice? The first 41 machine-language operations in the implementation of the atomic assignment must change the state in such a way that these changes are invisible when viewed at the higher level. More precisely, the 41 intermediate states of the computation must all be mapped by F into the same state as the starting state. How is this possible?

A complete answer to this question is beyond the scope of this paper. The trick lies in the definition of F , which must “unscramble” the intermediate states in the appropriate way. I will not explain here how it is done. I will only mention that, while it sounds like magic, it in fact is a simple extension of the basic idea of invariance that underlies most concurrent

²A single machine-language instruction could actually be used in the implementation of several atomic actions of Π —for example, if it were part of a subroutine called during the execution of several different statements of Π . The mapping F should therefore take state, action pairs into actions, so the action α_i of π is mapped into the action $F(s_{i-1}, \alpha_i)$ of Π . In other words, the state of π determines which atomic statement of Π is being executed by the execution of a machine-language statement.

program verification. An explanation and examples can be found in [6] and [8].

I won't consider the problem of compiler correctness. The purpose of this discussion is to point out that in order to permit a simple definition of correctness of an implementation, I cannot define a semantics in which the execution of an atomic program statement is always represented as a single atomic action. I must allow "stuttering" actions. In the action-axiom semantics, the specification of an action α must allow a finite series of null transitions $s \xrightarrow{\alpha} s$ as well as the final action $s \xrightarrow{\alpha} t$ that "does the work".

I have described correctness of an implementation in terms of a behavioral semantics, where $\mathcal{M}[\Pi]$ is a set of behaviors. In an action-axiom semantics, the meaning $\mathcal{M}[\Pi]$ of a program Π is a set of axioms that determines the set of possible behaviors. Section 5.7 explains how this concept of correctness is translated into a relation between the sets of axioms $\mathcal{M}[\Pi]$ and $\mathcal{M}[\pi]$. The only observation I will make here is that the axioms of $\mathcal{M}[\Pi]$ must permit stuttering actions. More precisely, these axioms should not be able to distinguish stuttering; if an axiom is true for a behavior σ , then it should also be true for the behavior obtained from σ by adding stuttering actions. This will be guaranteed by using a temporal logic in which no formula can distinguish stuttering—a temporal logic with no "next-time" operator.

2.4.2 The Interface

The semantics of a program is traditionally defined by describing how it affects the values of variables. However, program variables are internal to the program; all that a user sees is what he types into the program and what the program types out to him. A semantics of a program should describe its input and output, not just how it affects internal objects like variables.

Given the machine-dependence of most input and output, an explicit semantics for input and output seems like a useless exercise. Instead, observe that input and output can be represented by variables. A terminal screen can be represented as a Boolean array, each element representing the presence or absence of light at one point on the screen. Keyboard input can be simulated through a variable whose value represents the sequence of characters that have been typed but not yet processed. I will use the term *interface variables* to describe variables that represent input and output.

In general, an interface variable describes the interaction between the program and its environment. They are *global* or *free* variables, in contrast

to the *local* or *bound* variables that are declared in ordinary program declarations. For example, variables declared in a Pascal var declaration are local.

Let us again consider the mapping F , introduced above to define what it means for a lower-level program π to be correct implementation of a higher-level program Π . Recall that F describes how the variables of Π are implemented in terms of the "variables" of π —the machine registers, if π is a machine-language program. We really don't care how the local variables of Π are implemented, since they are not externally visible. The compiler is free to implement local variables any way it wishes.

The compiler does not have such freedom in its implementation of interface variables. The implementation of the interface variables must be defined *a priori* if the program is to interact with its environment in a useful way. For example, suppose that the terminal screen is represented by a Boolean array. The semantics of the program Π would provide no information about real output if the compiler could define the array elements to represent completely arbitrarily points on the screen—or to represent the values of arbitrary one-bit registers in the machine.

I have considered the implementation of the states of Π in terms of states of π , but what about the implementation of actions? Just as there are local and interface state functions, there are internal and external actions. Most actions in a program behavior are internal, being caused by program execution. However, some actions represent operations external to the program—for example, the actions that represent the entering of an input character. The semantics of Π does not distinguish this operation, which changes the value of the interface variable representing the input buffer, from program operations that change the value of variables—for example, the program operation that removes a character from the input buffer. The compiler is free to implement internal actions of Π by any internal actions of π . However, the external actions of Π must be implemented by fixed actions of π , which may be internal or external. For example, the sequence of actions of Π that add a character to the input buffer may be implemented by a sequence of actions external to π , representing external operations that put the character into an input register and actions of π that move the character from the input register into the memory registers that implement the input buffer. The compiler would be of little use if it could implement the operation of typing a character, defined in the semantics of Π simply as an operation that changes the variable representing the input buffer, as an internal operation of π that adds a randomly chosen character to the buffer.

Thus, the representation of local variables and internal actions of Π by F may be arbitrary, but the representation of interface variables and external actions must be fixed. The meaning $\mathcal{M}[\Pi]$ of Π can be defined in a completely machine-independent fashion. The machine dependency, which exists for any real compiler, is contained in the details of how interface variables and external actions are to be implemented.

Thus far, I have been talking only about implementing a complete program Π . We should consider the problem of implementing a single statement S . In this case, all the global (undeclared) variables of S must be regarded as interface variables, and their implementations must be fixed *a priori*. For example, if statements S and T were to be implemented independently, their implementations could be combined to implement $S;T$ only if a variable x common to both were implemented as the same set of machine registers.

Of course, one is seldom interested in implementing a single statement of a program. These considerations would apply to a language that allows separate compilation of components such as subroutines. I will not consider the problem of separate compilation. My purpose in discussing implementation of individual statements is to point out that the concept of global and local variables occurs at all levels of a program. Variables global to a statement S may be local to a larger statement containing S .

3 The Programming Language

The goal of this paper is to explain how the semantics of any programming language can be defined, and not to give a complete semantics for a particular language. However, to show how the formalism works, it is helpful to define rigorously the semantics of some language. I will therefore formally define the semantics of a simple language called L, and will indicate informally how the semantics of language primitives other than those in L can be defined.

The language L contains an atomic assignment statement—one whose execution is an indivisible, atomic action. L has the usual sequential control structures: concatenation (;), **if** and **while** statements, plus a **fair cobegin**. The tests in **while** and **if** statements are also taken to be atomic. L has a **new** statement that declares a local variable, so

new x : integer in S ni

declares x to be a local variable of type **integer** whose scope consists of the statement S . The **new** statement has an optional **init** clause to specify the initial value, so

new x : integer init $2 * x$ in S ni

declares that the initial value of x in S is twice the value of the variable x whose scope includes the **new** statement. The assignment of the initial value to x is assumed to be an atomic action. The **new** statement also has an optional **alias** clause that is used to declare that the new variable is the alias for something else. For example,

new x : integer alias y in S ni

declares x to be an alias for y .

It may seem strange to introduce aliasing—a concept usually ignored in simple examples—in the language L. Aliasing is an important concept because it underlies the semantics of procedure calls. If *proc* is a procedure defined with single integer-valued a call-by-name parameter *param*, then the call *proc*(*arg*) can be simulated by the statement

new *param* : integer alias *arg* in S ni

where S is the body of *proc*. (Call by value and call by reference can be simulated with call by name through the use of auxiliary variables.)

For reasons that will be clear later, the concept of aliasing is central to our semantics, and we will need to understand a more general kind of

aliasing than real programming languages usually allow. In particular, L will allow a variable to be aliased to an expression. To understand what that means, consider the declaration

new f : real alias $9 * c/5 + 32$ in S ni

In this case, we can think of f and c as representing a single temperature, where f is its value in degrees Fahrenheit and c is its value in degrees Celsius. The two assignment statements $f := 32$ and $c := 0$ have exactly the same effect; executing either one changes the value of f to 32 and the value of c to zero.

As another example, assume a type **gaussian** which represents a Gaussian integer—a number of the form $m + n\sqrt{-1}$, where m and n are integers. If x and y are variables of type **integer**, then

new z : gaussian alias $x + y * \sqrt{-1}$ in S ni

defines z to be a variable of type **gaussian** whose real part is aliased to x and whose imaginary part is aliased to y . Assigning a value to z in S also assigns values to x and y so that the relation

$$z = x + y * \sqrt{-1}$$

holds throughout the execution of S . Similarly, changing the value of x in S also changes the value of z .

In the examples of **alias** clauses given so far, assigning a value to any variable produces a well-defined result. However, this need not be the case. Inside the body S of the statement

new c : integer alias $a + b$ in S ni

assigning a value to a or b changes the value of c in the obvious way, but what is the result of assigning a value to c ? I define an assignment to c to be a nondeterministic statement that can change the values of a and b in any way such that $a + b$ equals the new value of c .

However, I will assume that the aliasing relations are such that they can always be maintained by the proper choice of values. More precisely, a program is considered illegal if its execution would force the aliasing relations to be violated. For example, the statement

new b : integer alias \sqrt{a} in S ni

is illegal if, at any time during its execution, the value of a is not a perfect square.

This approach to aliasing is similar to the one I will take for type constraints—namely, a program is illegal if its execution would force a type violation. For example, the statement

`new b : boolean init $\neg c$ in ...`

is illegal in any context in which c is not declared to be of type `boolean`.

While typing consistency is easy for a compiler to enforce in language L , the consistency of aliasing relations can be determined at compile time only if the kind of expression that can appear in an alias is restricted in some way. In fact, some restriction is obviously necessary if the compiler is to have any chance at compiling the code. Those restrictions are irrelevant to our semantics, so they are not discussed.

The basic syntax of L is given by the syntax diagrams in Figure 1. I will not bother to give a formal syntax for identifiers. The only types that I will use in L are `integer` and `boolean`. Expressions are assumed to be the usual ones constructed from variable names and the ordinary operations on integers and booleans—for example, an expression like

$(x * y + z = 17) \supset (x > y \vee \neg b)$


I will enclose `if` and `while` tests, assignment statements, and the `init` clause of a `new` statement in angle brackets to emphasize their atomicity.

In addition to the usual information, the syntax diagrams of Figure 1 also have labels attached to the nonterminal components. These labels are called *primitive selectors*. A primitive selector identifies a component of a compound statement—for example, the primitive selector *then* identifies the “then-clause” of an `if` statement. The “...” label in the specifications of the `cobegin` indicates that the primitive selectors for the clauses of a `cobegin` are integers, and likewise for a list of statements.


In more formal terms, the primitive selectors label the edges in the parse tree of a statement or program.³ A *selector* for a statement S is a sequence of primitive selectors that represents a path starting from the root in S 's parse tree. A selector identifies a component of a program or statement. For example, the selector *else, body, 2* identifies the substatement $(x := x + 4)$ in the following statement:

³Trivial nodes that have only a single son are eliminated from the parse tree, which is why there is no primitive selector associated with the first box in the syntax diagram defining a statement.

program:



statement:



simple statement:

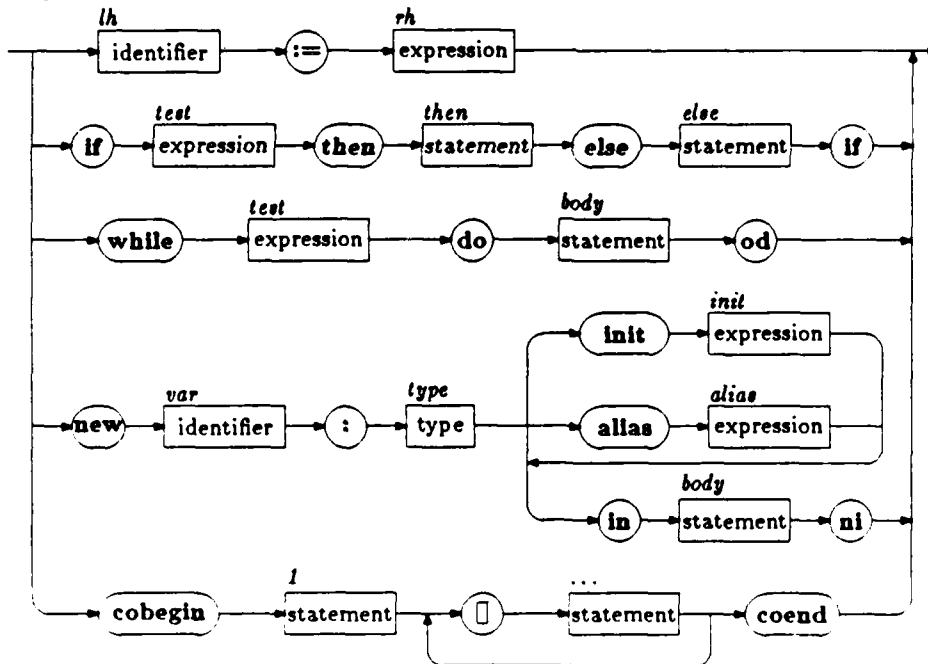


Figure 1: Basic syntax of language L.

```

if (x > 0)
  then (x := x + 1)
  else while (y > 0)
    do (y := y - 1);
      (x := x + 4)
    od;
  (y := 17)
fi

```

More formally, given a program or statement S , a substatement of S consists of a pair S, γ , where γ is a selector for S . A substatement of S is, when viewed by itself, a statement. I will often write something like: " T is the substatement S, γ of S ." This means that the substatement, when viewed alone, is the same as the statement T . However, T and S, γ are formally two different kinds of objects—one is a complete statement and the other is part of a statement.

The null selector selects the entire statement, so " S ," denotes S viewed as a substatement of itself. Since " S ," looks rather strange, I will simply write S to denote both the entire statement S and that statement viewed as a substatement of itself.

4 States and Actions

The meaning $\mathcal{M}[S]$ of a statement S will be a set of temporal logic axioms defining the behaviors of S and a set of nontemporal axioms defining its set of initial states. To give a semantics for these axioms, I must define the set $\mathcal{S}(S)$ of all possible states of S and the set $\mathcal{A}(S)$ of all possible actions of S . The initial-state axioms then define a set of states—namely, the set of all states in $\mathcal{S}(S)$ that satisfy those axioms; and the temporal axioms define a set of behaviors—namely, the set of all sequences

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots$$

with $s_i \in \mathcal{S}(S)$ and $\alpha_i \in \mathcal{A}(S)$ that satisfy those axioms.

Intuitively, the state of a statement at some time during its execution contains all the information needed to describe its possible behavior at future times. To define the set $\mathcal{S}(S)$, we must consider what information must be in the state of S .

4.1 Program Variables

The future behavior of a program certainly depends upon the current values of its variables, so a state must specify the values of all program variables. More precisely, a state in $\mathcal{S}(S)$ must include a mapping *val* from the set of variables of S to a set of values. For the simple language L , in which all variables are of type integer or boolean, the set of values consists of the set $Z \cup \{true, false\}$, where Z denotes the set of all integers.

Let S be the statement

cobegin S_1 \square S_2 **coend**

and suppose that S_1 and S_2 both contain new x statements. Each of these statements declares a different variable, but both variables have the same name x . Both of the variables named x may be defined at the same time, and may have different values. To facilitate the discussion, I will use the term *identifier* to denote the syntactic object constituting the name of a variable, and the term *variable* to denote the variable itself. Thus, S has two different variables having the same identifier x . This situation does not arise in a sequential program because, at any instant during its execution, there is at most one currently active variable for any identifier. However, it does arise in concurrent programs and must be considered.

To define *val*, we must define the value it assigns to each of the variables of *S*, which requires giving different names to different variables. Assigning unique names to variables is a nontrivial problem, since different variables may be represented in the program by the same identifier. It is solved with selectors. I let $x(S, \gamma)$ be the name of the variable with identifier x that is declared in a **new** statement whose selector in *S* is γ —in other words, where γ is the path in the parse tree of *S* leading to the **new** statement. A “global” variable with identifier x —that is, the variable denoted by an occurrence of the identifier x outside the scope of any **new** x statement—is given the name $x()$.

I consider $x()$ to be a variable of any statement *S*, even if the identifier x never appears in *S*. For example, suppose *S* is the statement:

```

< y := y + 1 >;
new z : in < z := y > ni

```

Then the variables of *S* consist of the single “bound” variable $z(S, 2)$ plus the infinite set of “free” variables $x()$, $y()$, $z()$, \dots , only one of which actually appears in *S*. Even though the variable $x()$ does not appear in *S*, it may appear in other statements in the complete program. The correctness of a program containing

```

cobegin S □ T coend

```

may depend upon the obvious fact that *S* does not change the value of $x()$. The value of $x()$ is included as part of *S*’s state so we can say formally that *S* does not change that value.

To summarize, defining the set of states $\mathcal{S}(S)$ of *S* requires defining the set of variables of *S*. The variables of *S* consist of the following:

- For any **new** x statement of *S* with selector γ , the variable named “ $x(S, \gamma)$ ”.
- For any identifier x , the variable named “ $x()$ ”.

The mapping *val* assigns a value to each of these variables.

The names of variables come up quite often when talking about programs. If *S* is a hundred page program, then the name $x(S, \gamma)$ takes up one hundred pages. Writing even the simplest statements about *S* would therefore require quite a bit of paper. Such a practical consideration is as irrelevant for the semantics of a programming language as is the cost of tape for the theory of Turing-machine computability. However, it does pose

a problem in writing examples, since a simple assertion about a five-line program might take one or two pages. The solution is, of course, to give names to statements and substatements. I will use the ordinary labeling convention to do this. For example, consider the program.

```

s: if (x > 0)
    then (x := x + 1)
    else t: new y
           in (y := x);
              (x := y + 2)
           ni;
           (z := 17)
    fi

```

The variable y declared by the `new y` statement will be called simply $y(t$ of $s)$. However, you should remember that its complete formal name is:

$$y \left(\begin{array}{l} \text{if}(x > 0) \\ \text{then} \dots \\ \text{else} \dots \\ \text{fi} \end{array} , \text{ else, } 1 \right)$$

4.2 Control Variables

There is more to a state than the values of program variables. To determine the future behavior at some point during the execution of the statement

```

u: begin s: (x := x + 1);
      t: (y := y + 1)
    end

```

we need to know whether control is at the beginning of statement s , at the beginning of statement t , or at the end of statement t . Since the state must determine the statement's possible future behavior, it must contain this control information.

I will describe this control information in terms of the boolean-valued control variables *at*, *in*, and *after*. For any substatement S, γ , there are control variables $at(S, \gamma)$, $in(S, \gamma)$, and $after(S, \gamma)$, where the values of these variables equal *true* when

$at(S, \gamma)$: control is at the beginning of substatement S, γ

Section 2.4 of the dual mapping F applies equally well to F^* , and I will not discuss further how F^* is actually constructed.

Since temporal logic formulas are constructed from predicates and temporal operators, there is an obvious extension of F^* to a mapping from $\mathcal{TL}(\Pi)$ to $\mathcal{TL}(\pi)$. For example,

$$F^*(in(\Pi, \rho) \triangleleft x(\Pi, \gamma) > 0) = F^*(in(\Pi, \rho)) \triangleleft F^*(x(\Pi, \gamma) > 0)$$

It follows from these definitions that for any behavior σ of π and any formula A in $\mathcal{TL}(\Pi)$:

$$\sigma \models F^*(A) \equiv F(\sigma) \models A$$

In terms of behaviors, π correctly implements Π if, for every possible behavior σ of π , $F(\sigma)$ is a possible behavior of Π . For simplicity, let us ignore the initial-state specification, so the meaning $\mathcal{M}[\Pi]$ of Π in an action-axiom semantics is a set of temporal logic axioms, and $F(\sigma)$ is a possible behavior of Π if and only if $F(\sigma) \models A$ is true for all $A \in \mathcal{M}[\Pi]$. But $F(\sigma) \models A$ is true if and only if $\sigma \models F^*(A)$ is, so π correctly implements Π if and only if $\sigma \models F^*(A)$ is true for all $A \in \mathcal{M}[\Pi]$ and all behaviors σ of π . The behaviors of π consist of the sequences satisfying all the axioms of $\mathcal{M}[\pi]$. It follows from this that π correctly implements Π if, for every axiom A in $\mathcal{M}[\Pi]$, $F^*(A)$ is implied by the axioms in $\mathcal{M}[\pi]$. Thus, proving correctness of the implementation involves deducing, from the axioms for π , the truth of $F^*(A)$ —the translation of A into an assertion about π —for every axiom A in $\mathcal{M}[\Pi]$.

As explained in Section 2.4, a compiler is free to implement local variables and internal actions in any fashion, but interface (global) variables and external actions have a fixed implementation. The mapping F^* is defined on state predicates by defining $F^*(v)$ in terms of the variables of π , for every variable v of Π . The definition of $F^*(v)$ is arbitrary for a local variable v , but is fixed for an interface variable. To prove the correctness of an implementation, we are allowed to define $F^*(v)$ any way we like if v is a local variable, but must use the predetermined definition if v is an interface variable. Similar comments apply to actions.

I find it helpful to think of the semantics $\mathcal{M}[\Pi]$ of Π as the specification of a lower-level implementation. When viewed this way, there is an implicit existential quantification over the names of all local variables and internal actions. More precisely, the specification consists of the conjunction of all the axioms in $\mathcal{M}[\Pi]$, with existential quantification over these variable

$$s \xrightarrow{\alpha} s \xrightarrow{\alpha} \dots \xrightarrow{\alpha} s \xrightarrow{\alpha} t$$

It is the inability to distinguish stuttering that makes it easy to talk about a lower-level program implementing a higher-level one.

5.7 Implementation Mappings

I can now continue the discussion, begun in Section 2.4, of what it means for a lower-level program to correctly implement a higher-level one. Let Π be the higher-level program and π be its lower-level implementation. From the point of view of behaviors, we saw that there should be mappings F from the states and actions of π to the states and actions of Π so that if σ is the behavior

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots$$

of π , then $F(\sigma)$, which is defined to be

$$F(s_0) \xrightarrow{F(\alpha_1)} F(s_1) \xrightarrow{F(\alpha_2)} \dots$$

is a behavior of Π .

How are the mappings F defined? In action-axiom semantics, one never mentions states, just state predicates—mappings from the state into a set of Booleans. A state is determined by the values of all state predicates. To define a mapping $F : S(\pi) \rightarrow S(\Pi)$, one defines a mapping F^* that maps state predicates of Π into state predicates of π . Intuitively, F^* defines the state predicates of Π in terms of the state predicates of π . For example, $F^*(x(\Pi, \gamma) > 0)$ is the state predicate of π that “implements” the state predicate $x(\Pi, \gamma) > 0$ of Π ; in other words, it is the translation of the high-level statement that the value of the variable $x(\Pi, \gamma)$ is positive into a lower-level statement involving the values of memory registers, program counters, etc. Defining the mapping F^* requires describing how the variables (both program and control variables) of Π are implemented by the “variables” (machine registers) of π . The mappings F and F^* are related by

$$s \models F^*(P) \equiv F(s) \models P$$

for any state s in $S(\pi)$ and state predicate P in $\mathcal{PR}(\Pi)$.

In a similar way, F^* is defined to map action predicates of Π into action predicates of π , so $F^* : \mathcal{PR}(\Pi) \rightarrow \mathcal{PR}(\pi)$. Finding the mapping F^* is the heart of the proof that π correctly implements Π . The discussion in

$\sigma \models A = \text{true}$ for all $\sigma \in \Sigma$. The valid formulas for a program S are those that are valid for the set of all behaviors of S .

Note that $\models_{\Sigma} \text{false}$ is true if and only if Σ is the empty set. The semantics I give can produce contradictory sets of axioms for a program—axioms from which one can deduce the formula *false*. This is not an inconsistency in the system; rather it is an indication that there are no legal behaviors of the program, so the program is illegal. This will be the case, for example, if a program assigns a boolean value to a variable of type integer.

I consider the notion \models_{Σ} of semantic validity only for sets Σ having the property that for any $\sigma \in \Sigma$ and any $n \geq 0$: $\sigma^{+n} \in \Sigma$. Intuitively, this means that the temporal logic does not assume any preferred starting state. Formally, this means that the truth of $\models A$ implies the truth of $\models \Box A$ —a rule of inference known to logicians as the *Necessitation Rule*. This rule implies that whenever we give a predicate P as a temporal-logic axiom, we are really asserting that $\Box P$ is true.

The validity of the Necessitation Rule means that it is impossible to write a temporal logic formula which asserts that the program is executed only in certain starting states. Thus, one should define the semantics $\mathcal{M}[S]$ of S to consist of both a set of temporal logic axioms that constrain the allowed behaviors of S and a set of nontemporal axioms—that is, predicates—that constrain the starting state. The semantic meaning $\mathcal{M}[S]$ defines the set of behaviors of S to be the set of all behaviors σ such that:

- $\sigma \models A$ for every temporal axiom $A \in \mathcal{M}[S]$, and
- $s_0 \models A$ for every nontemporal axiom $A \in \mathcal{M}[S]$, where s_0 is the starting state of σ .

However, as I will show, it is not necessary to specify any initial states for a substatement S of a program. The only initial-state specification that must be added is that a complete program starts at its entry point.

5.6 There Won't Be a Next Time

An important “feature” of the temporal logic I am using is that there is no “next time” operator. There is no way in this logic to express the concept of the next state in the behavior. In fact, no formula in the logic can distinguish between two behaviors that differ only in the addition of “stuttering” actions—that is, where an action $s \xrightarrow{\alpha} t$ in the behavior is replaced by the finite sequence of actions

5.4 Renaming

Having already extended the renaming mappings to predicates, it is easy to extend them to temporal logic formulas constructed from predicates. For example, for any variable names v and w , we have

$$\rho_{S,\gamma}(\Box(v \vee \Diamond w)) = \Box(\rho_{S,\gamma}(v) \vee \Diamond \rho_{S,\gamma}(w))$$

Thus, if T is the substatement S, γ of S ,

$$\rho_{S,\gamma} : \mathcal{TL}(T) \rightarrow \mathcal{TL}(S)$$

The renaming mappings do *not* induce any mappings on behaviors. This is because they map states and atomic-action names in opposite directions:

$$\begin{aligned} \rho_{S,\gamma}^s : S(S) &\rightarrow S(T) \\ \rho_{S,\gamma}^a : A(T) &\rightarrow A(S) \end{aligned}$$

Since a behavior consists of an alternating sequence of states and action names, the renaming mappings do not work on behaviors. This may be the source of the difficulties encountered in trying to give a behavioral semantics—one in which $\mathcal{M}[[S]]$ is a set of behaviors—to concurrent programming languages

5.5 Temporal Logic as Semantics

For each statement S of the programming language, I have defined a set $\mathcal{TL}(S)$ of temporal logic formulas, and a notion of semantic validity \models for these formulas. In an action-axiom semantics, the meaning of S includes a set of temporal logic formulas that must be satisfied by the behaviors of S . This set of formulas is specified by giving axioms and inference rules, which means that we have a logical system and a notion of a provable formula. I will not discuss provability, and will restrict myself to validity.

I have defined $\sigma \models A$ for a behavior σ and a temporal logic formula A , but I have not defined the concept $\models A$ —validity of a formula. For any formula $A \in \mathcal{TL}(S)$, one usually defines $\models A$ to equal *true* if and only if $\sigma \models A$ equals *true* for all behaviors σ in $\mathcal{B}(S)$.

The formulas A for which $\models A$ is true are those that are true for all sequences of states and actions from S , so their truth rests only on the properties of S 's sets of states and actions, not on properties of S 's dynamic behavior—for example, the formula $\Box(x() \in Z \supset (x^2 \geq x))$. A formula A is said to valid for all behaviors in some subset Σ of $\mathcal{B}(S)$, written $\models_{\Sigma} A$, if

Intuitively, $A \sqsubseteq B$ means that B holds for at least as long as A does—that is, A holds for a length of time \leq the length of time that B holds, so it represents a “temporal \leq ”.

I will extend the definition of $\mathcal{TL}(S)$ to include temporal formulas constructed with the operator \sqsubseteq as well as \square and \diamond . The unary operators can be defined in terms of \sqsubseteq ; for example, $\square A \equiv true \sqsubseteq A$. Thus, the single operator \sqsubseteq is all we need.

The operator \triangleleft is defined in terms of \sqsubseteq by

$$A \triangleleft B \stackrel{\text{def}}{=} (A \vee \neg B) \sqsubseteq B$$

A little thought shows that

$$\sigma \models (A \triangleleft B) \stackrel{\text{def}}{=} \forall n : (\forall m < n : \sigma^{+m} \models A) \supset \sigma^{+n} \models B$$

so $A \triangleleft B$ means that A holds for a length of time $<$ the length of time that B holds. The operators \triangleleft and \sqsubseteq obey the same transitivity relations that $<$ and \leq do. For example,

$$(A \sqsubseteq B) \wedge (B \triangleleft C) \supset (A \triangleleft C)$$

I will use \sqsubseteq to define a new type of temporal formula that is useful for specifying actions. For any action name $\alpha \in \mathcal{A}(S)$ and any predicates P and Q , I define $\{P\}\langle\alpha\rangle\{Q\}$ to be the temporal logic formula that means that executing α starting in a state in which P is true can produce a state in which Q is true. (It is just the ordinary Hoare triple for the atomic action α , viewed as a temporal formula.) However, we must allow stuttering actions of α which do nothing, and hence leave P true. The formal definition is

$$\{P\}\langle\alpha\rangle\{Q\} \stackrel{\text{def}}{=} P \supset (Act(\alpha) \triangleleft P \vee Q)$$

Intuitively, $\sigma \models \{P\}\langle\alpha\rangle\{Q\}$ asserts that if, P is true in the initial state of σ and the first one or more actions of σ are α actions, then $P \vee Q$ remains true through the first state before the first action that is not an α action. If $Q \supset \neg Act(\alpha)$ holds, which is the only case in which this formula will be used, then Q will be true only after the last of these initial α actions. Hence, it asserts that α can perform a series of stuttering actions leaving P true, and can also “finish” by making Q true.

The formula $\{P\}\langle\alpha\rangle\{Q\}$ is used to describe how an action can change the state. It is also necessary to state that an action does not change something. I therefore introduce the formula $e \overset{\alpha}{\not\neq}$, which asserts that the action α does not change the value of the expression e . It is defined by

$$e \overset{\alpha}{\not\neq} \stackrel{\text{def}}{=} \forall \eta : (e = \eta) \supset Act(\alpha) \triangleleft (e = \eta)$$

$$s_n \xrightarrow{\alpha_{n+1}} s_{n+1} \xrightarrow{\alpha_{n+2}} \dots$$

unless σ is finite and n is less than the length m of σ , in which case σ^{+n} is defined to be the sequence consisting of the single state s_m . We define $\sigma \models A$ inductively as follows.

- If A is a state predicate, then $\sigma \models A \stackrel{\text{def}}{=} s_0 \models A$. (The value of a state predicate is its value in the starting state.)
- If A is an action predicate, then $\sigma \models A \stackrel{\text{def}}{=} \alpha_1 \models A$. (The value of an action predicate is its value for the first action.) However, if σ consists of the single state s_0 with no actions, then $\sigma \models A \stackrel{\text{def}}{=} \text{false}$.
- The logical connectives "distribute" in the obvious way. For example,

$$\sigma \models (A \vee B) \stackrel{\text{def}}{=} (\sigma \models A) \vee (\sigma \models B)$$

- The temporal operators are defined by

$$\begin{aligned} \sigma \models \Box A &\stackrel{\text{def}}{=} \forall n : \sigma^{+n} \models A \\ \sigma \models \Diamond A &\stackrel{\text{def}}{=} \exists n : \sigma^{+n} \models A \end{aligned}$$

Note that \Diamond is the dual of \Box —that is, $\Box A \equiv \neg \Diamond \neg A$ for any A . The operator \rightsquigarrow is defined by

$$A \rightsquigarrow B \stackrel{\text{def}}{=} \Box (A \supset \Diamond B)$$

Note also that $\Box \Diamond A$ has the intuitive meaning that A is true infinitely often.

5.3 The Binary Temporal Operators

While the unary temporal operator \Box , and the operators derivable from it, are quite natural and easy to understand, they are not sufficiently expressive. We need an additional binary temporal operator. There are many binary operators that are equivalent in the sense that one can be represented by another. My favorite one, introduced in [6], is the operator \trianglelefteq , whose semantics is defined as follows.

$$\sigma \models (A \trianglelefteq B) \stackrel{\text{def}}{=} \forall n : (\forall m \leq n : \sigma^{+m} \models A) \supset \sigma^{+n} \models B$$

is a predicate containing the free logical value variable χ , the bound logical value variable η , and the two program variables $x()$ and $y(S, \gamma)$ in $\mathcal{V}(S)$. For any predicate P and state s of $S(S)$, $s \models P$ is a formula involving values and value variables.

An action predicate of S is an expression of the form $Act(S, \gamma)$, where S, γ is a substatement of S . The action predicate $Act(S, \gamma)$ defines a boolean-valued function on the set $A(S)$ that has the value *true* on an action-name α if and only if α is the name of an atomic action of the substatement S, γ . I write $\alpha \models Act(S, \gamma)$ to denote the value of $Act(S, \gamma)$ on α . Remembering that atomic-action names are just components S, μ of S , we see that $S, \mu \models Act(S, \gamma)$ equals true if and only if $\mu = \gamma, \nu$ for some ν .

Let $\mathcal{PR}(S)$ denote the set of all state and action predicates of S . Since state predicates are built out of variable names and action predicates are all of the form $Act(S, \gamma)$, the renaming mappings induce mappings on predicates in the obvious way. If T is the substatement S, γ of S , then

$$\rho_{S, \gamma} : \mathcal{PR}(T) \rightarrow \mathcal{PR}(S)$$

These renaming mappings satisfy the expected relation (2). Moreover, if P is a tautology of $\mathcal{PR}(T)$, then $\rho_{S, \gamma}$ is a tautology of $\mathcal{PR}(S)$.

5.2 The Unary Temporal Operators

I will begin with the simpler form of temporal logic, using only unary temporal operators. The formulas of this logic are constructed from predicates, the usual logical operations, and the two unary temporal operators \diamond and \square . More precisely, for any statement S , the set $\mathcal{TL}(S)$ of temporal logic formulas of S consists of all formulas constructed from $\mathcal{PR}(S)$ with the logical operators and the unary operators \diamond and \square .

Just as predicates are true or false for states, temporal logic formulas are true or false for behaviors. Let $\mathcal{B}(S)$ denote the set of all finite and infinite sequences of the form

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots \quad (5)$$

where the s_i are states in $S(S)$ and the α_i are atomic-action names in $A(S)$. We give a semantics for temporal logic formulas by defining $\sigma \models A$ for any behavior σ in $\mathcal{B}(S)$ and any temporal logic formula A in $\mathcal{TL}(S)$.

If σ is the sequence (5), for any nonnegative integer n let σ^{+n} be the sequence

5 Temporal Logic

In the action-axiom semantics, I use temporal logic to express the constraints describing when an action must eventually occur. Temporal logic, introduced into the study of concurrent programs by Pnueli [13], is now quite familiar. I will therefore only sketch the logic that I will need, and refer the reader to [5] and the appendix of [6] for more details.

5.1 Predicates

The building-blocks of our temporal logic are *predicates*. For any program statement S , I define a set of predicates. There are two kinds of predicates: state predicates and action predicates.

A state predicate of S is just an expression constructed from variable names in $\mathcal{V}(S)$, including control variable names. For example,

$$at(S, \gamma) \vee \neg b() \supset x(S, \mu) = y() + 1$$

I will also include as predicates such expressions as $v \in Z$, where v is a variable name and Z denotes the set of integers.

Since a state in $S(S)$ assigns a value to all variable names in $\mathcal{V}(S)$, it assigns a value to a predicate. For any state s of $S(S)$, I denote by $s \models P$ the value assigned to the state predicate P by the state s .

A predicate is normally a boolean-valued expression, but I have not restricted predicates in this way; $y() + 17$ is just as much a predicate as $\neg b()$. The reason is that there is no way of knowing whether an expression has a boolean value without knowing the types of all its variables, and the types of undeclared variables are not known. We must have rules for computing the value of $y() + 17$ even when the value of $y()$ is *true*. I will handle this problem by adding an additional *undefined* value, and define *true* + 17 to equal *undefined*.

The presence of an *undefined* value means that we must be careful when manipulating expressions, since the usual rules of arithmetic and logic don't hold. For example, $x + 1 > x$ does not equal *true* if x is a boolean. However, $x \in Z \supset (x + 1 > x)$ should always have the value *true*.

It is necessary to allow predicates to have *logical value variables* (not to be confused with program and control variables) and quantifiers. Thus,

$$\forall \eta : x() + \chi > y(S, \gamma) + \eta$$

These renaming mappings satisfy (2) and (4).

Since action names are just the names of substatements, the renaming mappings can be applied to them in the usual way. Thus, if T is the substatement S, γ of S , then

$$\rho_{S, \gamma} : \mathcal{A}(T) \rightarrow \mathcal{A}(S)$$

is defined in the obvious way—namely, $\rho_{S, \gamma}(T, \mu) = S, \gamma, \mu$.

4.6 States and Actions: A Formal Summary

For every statement S in the language L , I have defined the following:

- A set $\mathcal{V}(S)$ of variable names, consisting of:
 - all program-variable names of the form $x()$ for every identifier x and of the form $x(S, \gamma)$, where γ is the selector in S of a new x statement.
 - all control-variable names of the form $at(S, \gamma)$, $in(S, \gamma)$, and $after(S, \gamma)$, for all substatements S, γ of S .
- The set $\mathcal{S}(S)$ of states of S , which is defined to be the set of all mappings

$$val : \mathcal{V}(S) \rightarrow Z \cup \{true, false\}$$

- The set $\mathcal{A}(S)$ of atomic-action names of S , defined to be the set of all components of the form S, γ where S, γ is a **while** or **if test**, an atomic assignment, or an **init** clause of a **new** statement.
- If T is the substatement S, γ of S , the renaming mappings

$$\rho_{S, \gamma} : \mathcal{V}(T) \rightarrow \mathcal{V}(S)$$

$$\rho_{S, \gamma}^* : \mathcal{S}(S) \rightarrow \mathcal{S}(T) \quad \rho_{S, \gamma} : \mathcal{A}(T) \rightarrow \mathcal{A}(S)$$

$$x := x^2$$

of program Π , and consider the 42 steps in the machine-language implementation π of Π that execute statement S . As mentioned earlier, 41 of them will be stuttering actions that leave the value of x unchanged. Before the first 41 steps have been executed, π 's state may no longer have the information needed to deduce the initial value of x . For example, after 20 steps, the state of π may show that x will wind up with the value 4, but may not show whether it started equal to 2 or -2 . This means that the single nonstuttering action must be among the first 20 steps, and the remaining steps must be stuttering actions of the statement following S . If S is the last statement of Π , then these remaining steps are λ actions.

$$\rho_{S,\gamma,\delta} = \rho_{S,\gamma} \circ \rho_{T,\delta} \quad (2)$$

By taking this equality to be a definition, we can formally define $\rho_{S,\gamma}$ for any selector γ by defining it for all primitive selectors. This formal definition should be obvious and is omitted.

Let T be the substatement S, γ of S . A state val of $S(S)$ is a mapping from $\mathcal{V}(S)$ to values, and $\rho_{S,\gamma}$ is a mapping from $\mathcal{V}(T)$ to $\mathcal{V}(S)$. The composition $val \circ \rho_{S,\gamma}$ is therefore a mapping from $\mathcal{V}(T)$ to values, which is a state in $S(T)$. Thus, the mapping $\rho_{S,\gamma}$ induces a mapping

$$\rho_{S,\gamma}^* : S(S) \rightarrow S(T)$$

from states of S to states of T , defined by

$$\rho_{S,\gamma}^*(val) \stackrel{\text{def}}{=} val \circ \rho_{S,\gamma} \quad (3)$$

for any $val \in S(S)$.⁴

It follows easily from (2) that the mappings $\rho_{S,\gamma}^*$ satisfy the following "adjoint" form of (2).

$$\rho_{S,\gamma,\delta}^* = \rho_{T,\delta}^* \circ \rho_{S,\gamma}^* \quad (4)$$

4.5 Actions

The states of S are defined using the set $\mathcal{V}(S)$ of variable names. The actions of S will be defined in terms of a set $\mathcal{A}(S)$ of atomic-action names of S .

The atomic actions of S are the components written in angle brackets. In the language L , there are just four kinds of atomic actions: assignment statements, if tests, while tests, and init clauses (of new statements). (I assume that the initial-value assignment of the new statement is performed as a single atomic action.) The set $\mathcal{A}(S)$ of atomic-action names of S consists of the set of all components S, γ of S such that γ is a selector for one of the following: an assignment statement, the *test* component of an if statement, the *test* component of a while statement, or the *init* component of a new statement. For reasons having to do with defining compiler correctness that are irrelevant to the remainder of this paper, if Π is a complete program, then $\mathcal{A}(\Pi)$ is defined to contain one additional action name: the name λ , which is the name of a null action.⁵

⁴The renaming mapping $\rho_{S,\gamma}^*$ has no connection with the mapping F discussed in Section 2.4.1 between the states of an implementation and the states of a higher-level program.

⁵The following example shows why the λ action is needed. Let S be the statement

It would therefore seem that we should add types and aliasing information to the state. In fact, we needn't. The reason is that, in language L, types and aliasing relations are *static* properties; they do not change during execution of the program. Executing an action of L does not change the type of a variable or any aliasing relations. (We sometimes think of executing a new statement by first executing its declarations, but that makes no sense because declarations are not actions.)

In a more complex language, types and aliasing relations can be dynamic. For example, in Pascal, if x is a variable of type pointer, then the aliasing relation " x is aliased to y " is dynamic, since its truth is changed by assigning a new value to x . In these cases, it may be necessary to add types and aliasing information to the state. However, in most languages, aliasing relations among control variables will be static, and can be handled the same way as in language L.

4.4 Renaming

For any statement S , let $\mathcal{V}(S)$ denote the set of names of variables of S . A state *val* of $S(S)$ is a mapping that assigns a value to each variable name in $\mathcal{V}(S)$. For a compositional semantics, we must be able to derive information about the states of S from information about the states of its component substatements. This requires the fundamental concept of a *renaming* mapping.

Let statement T be the substatement S, γ of S . Every variable of T is a variable of S , except that it may be known by a different name. I will define $\rho_{S, \gamma}$ to be the mapping on names such that if v is the name of a variable in T , then $\rho_{S, \gamma}(v)$ is the name of the corresponding variable in S . Hence,

$$\rho_{S, \gamma} : \mathcal{V}(T) \rightarrow \mathcal{V}(S)$$

The variable $x(T, \mu)$, which is the variable of T with identifier x that is declared in the new statement T, μ , is called by the name $x(S, \gamma, \mu)$ when it is regarded as a variable of S . Thus,

$$\rho_{S, \gamma}(x(T, \mu)) = x(S, \gamma, \mu)$$

A variable that has the name $x()$ as a variable in T is undeclared in T . If it is undeclared in S , then it has the same name as a variable of S , so $\rho_{S, \gamma}(x()) = x()$. However, if it is declared in the new x statement S, ν , so ν is a prefix of γ , then $\rho_{S, \gamma}(x()) = x(S, \nu)$.

The renaming mappings compose in the natural way. If T is the substatement S, γ of S , then for any substatement $T \delta$ of T , we have

I therefore prefer to use the term *implicit* variables for variables other than ordinary program variables. Some languages employ other implicit variables besides control variables. For example, a language that provides a buffered message-passing primitive will contain implicit variables whose values describe the set of messages in the queues.

Ordinary program variables may be free (undeclared), like $x()$, or bound (declared), like $x(S, \gamma)$. I have written all control variables as bound variables, but are they really bound? Remember that the free variables are interface variables and bound variables are internal ones. In order to use a compiled version of a statement S , one must know where its starting and ending control points are, but need know nothing of its internal control points. This suggests that $at(S)$, $in(S)$, and $after(S)$ are interface variables for statement S , while, for any non-null selector γ , $at(S, \gamma)$, $in(S, \gamma)$, and $after(S, \gamma)$ are internal variables. The control variables $at(S)$, $in(S)$, and $after(S)$ are best viewed as undeclared and might better be written as $at()$, $in()$, and $after()$. (They are not written that way both for historical reasons and because it would tend to be confusing.) These variables are implicitly declared, and aliased to other control variables, when S is written as part of a larger statement.

4.3 Are There Other State Components?

Does a mapping *val* from variable names of S values tell us everything we need to know about the current state of S in order to determine its future behavior? At first glance, it might seem that it doesn't. For example, what is the effect of executing

$s: x := y + 1$

when the value of y is 17? The answer depends upon the type of x . If x is of type **integer**, then the execution sets x to 18. However, if x is of type **boolean**, then executing s produces an error.

Moreover, suppose x is of type **integer** and $y = 17$, so executing s changes the value of x to 18. What does this execution do to the value of y ? If y is not aliased to x , then its value is left unchanged. However, if s appears inside the statement

new y : integer alias x in ... ni

then the value of y is also changed to 18.

$in(S, \gamma)$: control is at the beginning of or inside S, γ , but not at its exit point. Note that $at(S, \gamma) \supset in(S, \gamma)$ is always true.

$after(S, \gamma)$: control is at the exit point of S, γ —that is, at the point just after its execution is completed.

In addition to complete substatements, the at , in , and $after$ variables are also defined for certain parts of statements that denote atomic operations—namely, the $test$ of an if or $while$ statement and the $init$ clause of a new statement (if it has one). Also, the control variables $at(\Pi)$, $in(\Pi)$, and $after(\Pi)$ are defined for a complete program Π .

The statement u above thus has eight control variables: $at(u)$, $in(u)$, $after(u)$, $at(s \text{ of } u)$, $in(s \text{ of } u)$, $after(s \text{ of } u)$, $at(t \text{ of } u)$, $in(t \text{ of } u)$, and $after(t \text{ of } u)$. They are not all independent, however, since we have

$$\begin{aligned} at(u) &= at(s \text{ of } u) \\ after(u) &= after(t \text{ of } u) \\ in(u) &= in(s \text{ of } u) \wedge in(t \text{ of } u) \\ after(s \text{ of } u) &= at(t \text{ of } u) \end{aligned} \tag{1}$$

These equations represent aliasing relations between the control variables. The study of these aliasing relations is deferred until later.

The mapping val that assigns values to variables must assign values to the control variables as well as the ordinary program variables. Of course, we must assume that at , in , and $after$ are not identifiers, so they cannot be used for ordinary program variables.

Variables like at , in , and $after$ are sometimes called “dummy” or “ghost” variables. This seems to imply that they are not as real as ordinary program variables. Indeed, I have found that many computer scientists regard their use as somewhat distasteful—perhaps even immoral. Control variables are every bit as real as ordinary program variables. They differ from program variables only in that the programmer does not explicitly write them. Every programmer knows that he can often simplify a program’s control structure—that is, eliminate control variables—by adding program variables; and, conversely, he can eliminate program variables by using a more complex control structure—that is, by adding control variables. A compiler handles both kinds of variables in very much the same way; in the compiled version of a program, the values of program variables and control variables are both encoded in terms of the contents of memory registers and program-location counters.

and action names. The names of interface variables and external actions represent fixed, externally defined objects.

6 The Semantics of Language L

With these preliminaries out of the way, I can now give the semantics of language L. This is done by defining the meaning $\mathcal{M}[S]$ of S , where S is any statement or complete program. I define $\mathcal{M}[S]$ to consist of a set of temporal logic axioms that specify the set of behaviors of S . As discussed below, for a complete program Π , I will also need one nontemporal axiom—that is, a predicate—to specify the starting state.

The basic idea behind achieving a compositional semantics is the requirement that any axiom asserted about a statement T must be valid for any statement containing T as a substatement. Of course, an axiom about T must be renamed to become an axiom about a statement containing T . The formal statement of this idea is:

Composition Principle: If T is substatement S, γ of S , then for any formula A : if $A \in \mathcal{M}[T]$ then $\rho_{S, \gamma} A \in \mathcal{M}[S, \gamma]$.

6.1 Syntactic Predicates

I observed in Section 4.3 that there is information we need in order to define $\mathcal{M}[S]$ that is not in the state of S —namely, type and aliasing information. This information is not in the state because it is determined syntactically and does not change during execution of S . Unfortunately, it may be determined not by the syntax of S , but by the syntax of the complete program containing S . For example, the aliasing relations defined by a new statement are not known when defining $\mathcal{M}[S]$ for a statement S in its body.

For our simple language L, typing information can be handled by ordinary axioms; the fact that a variable v is of type integer is expressed by the requirement that the value of v always be an integer. Aliasing relations can also be expressed by similar requirements—for example the aliasing relation defined by

new z : gaussian alias $x + y\sqrt{-1}$ in ...

is expressed by requiring that the value of $z(S)$ always equals the value of $x() + y()\sqrt{-1}$. However, the fact that $z(S)$ is *not* aliased to the variable $a()$ cannot be expressed in this way.

The absence of aliasing relations is expressed with a new relation \perp , where $v \perp w$ means intuitively that assigning a value to the variable named v does not change the value of the variable named w , and vice-versa. An

ordinary state predicate such as $v = w$, which asserts that the values of v and w are equal, is true or false for a particular state. However, the truth of the expression $v \perp w$ depends only upon the syntactic structure of the program; it is true for one state of S if and only if it is true for all states of S .

An expression like $v \perp w$, whose value is a boolean that depends only on the program syntax, is called a *syntactic predicate*. Unfortunately, if v and w are undeclared variables of S , then the value of $v \perp w$ depends upon the syntax of the program that contains S , and its value is not determined when we are defining $\mathcal{M}[S]$. Thus, a syntactic predicate either has a definite boolean value, or else has an undetermined value.

I will allow syntactic predicates to appear in a temporal logic formula of $\mathcal{TL}(S)$ anywhere that an ordinary state predicate can. However, there is no reason to write $\Diamond(v \perp w)$, since if $v \perp w$ is ever true, then it is always true for every state of S . Formally, a syntactic predicate in a temporal logic formula of $\mathcal{TL}(S)$ is viewed as a boolean constant if its value is determined by S , and as a logical variable if its value is undetermined.

Formally, a syntactic predicate appearing in an axiom of $\mathcal{M}[S]$ is a constant if its value is determined by S , and it is a logical variable if its value is not determined. Thus, writing the syntactic predicate $x(S, \gamma) \perp y(S, \mu)$ is simply an "abbreviation" for either *true* or *false*, since the aliasing relations of variables declared in S are determined. On the other hand, a syntactic predicate such as $x() \perp y()$ represents a logical variable, since aliasing relations between undeclared variables are undetermined. Because there is an implicit universal quantification over all free logical variables in an axiom, an axiom containing a syntactic predicate is asserted to be true whatever value is assigned to it.

We can apply renaming mappings to syntactic predicates in the obvious way. Thus, if P is a syntactic predicate for T , and T is the substatement S, γ of S , then $\rho_{S, \gamma}(P)$ is a syntactic predicate for S . When the predicate P occurs in an axiom A of $\mathcal{M}[T]$, the expression $\rho_{S, \gamma}(P)$ occurs in $\rho_{S, \gamma}(A)$, which, by the Composition Principle, is an axiom of $\mathcal{M}[S]$. A little thought reveals that, to ensure the validity of the Composition Principle, we want the following property to hold:

Syntactic Composition Property: For any syntactic predicate P :
if the value of P is defined for T , then the value of $\rho_{S, \gamma}(P)$ is also defined for S and equals the value of P .

The use of syntactic predicates is not really necessary. I could include

the information that they express in the state. Had I done so, a syntactic predicate having an undetermined value would become a component of the state, and $S(S)$ would include states having all possible values of that predicate. A syntactic predicate whose value is determined in a statement S could be represented either as a state component constrained to have only one possible value, or as a constant.

6.1.1 Aliasing

The absence of aliasing will be expressed by the relation \perp between variable names in $\mathcal{V}(S)$. This will be done axiomatically by defining a logical system for deriving \perp relations. To do this, I must first introduce a relation \prec , where $v \prec \{w_1, \dots, w_n\}$ means that the variable name v is not directly aliased to any variable names other than w_1, \dots, w_n . It is convenient to extend this relation to a relation between sets of variable names, where $\{v_1, \dots, v_m\} \prec \{w_1, \dots, w_n\}$ means that each of the variable names v_i is not directly aliased to any variable names other than the w_j . We then have the obvious inference rule:

For any sets $V, W \in \mathcal{V}(S)$: if $V' \subseteq V$, $W \subseteq W'$, and $V \prec W$,
then $V' \prec W'$.

The relation \perp on $\mathcal{V}(S)$ is defined so that $v \perp w$ means that neither v nor w is aliased, directly or indirectly, to the other. In other words, it means that $v \neq w$ and there do not exist both a chain of \prec relations from v and a chain of \prec relations from w that lead to a common variable name. This leads to the following rules for deriving \perp relations.

- If $v \prec \emptyset$, $w \prec \emptyset$, and $v \neq w$, then $v \perp w$.
- If $v \perp w$ then $w \perp v$.
- If $v \prec \{w_1, \dots, w_n\}$, $w_1 \perp w$, \dots , $w_n \perp w$, and $v \neq w$, then $v \perp w$.

I extend \perp to a relation on finite sets by letting $\{v_1, \dots, v_n\} \perp \{w_1, \dots, w_n\}$ denote $\forall i, j : v_i \perp w_j$.

Having given general rules for reasoning about \prec , I must define the relation for variable names in $S(S)$ for an arbitrary statement S . The value of a syntactic predicate $V \prec W$ or $V \perp W$ will be undetermined if V and W both contain the names of undeclared variables of S . To define the values of the ones that are determined, I will take the Syntactic Composition Property as an axiom, and give a recursive definition based upon the structure of S .

The first observation is that a program variable cannot be aliased to a control variable, and vice-versa. I therefore require that $v \perp w$ equal *true* whenever v is a program variable and w is a control variable.

Since the only dependency relations on program variables are introduced by the *alias* clauses of *new* statements, all dependency relations among program variables are obtained from the Syntactic Composition Property and the following axiom:

If S is the statement *new* $x \dots$ *alias* $exp \dots$ and y_1, \dots, y_n are all the variable names in exp , then $x(S) \prec \{y_1(), \dots, y_n()\}$.

I must now define the dependency relations on control variable names. I will do this by assuming the Syntactic Composition Property and defining the relations introduced by each language construct. There are a number of aliasing relations that are similar to the ones introduced by an *alias* clause, except that the aliasing relations for the control variables are implicit in the program structure rather stated explicitly in a *new* statement. To define the \prec relations, I will write down these aliasing equations, where the control variable comprising the left-hand side of an equation is considered to depend upon each of the variables on the right-hand side. There is one set of equations for each programming language construct.

Besides these aliasing equalities, some other aliasing relations are given as boolean expressions—that is, asserting that the boolean expressions are true. No dependency relations are implied by these expressions, but they are listed here for future reference.

There is only one axiom that explicitly defines \perp relations; it is given for the *cobegin* statement.

assignment $in(S) = at(S)$
 $\neg(at(S) \wedge after(S))$

if $in(S, test) = at(S, test)$
 $after(S, test) = at(S, then) \vee at(S, else)$
 $at(S) = at(S, test)$
 $in(S) = in(S, test) \vee in(S, then) \vee in(S, else)$
 $after(S) = after(S, then) \vee after(S, else)$
 $\neg(at(S, test) \wedge (in(S, then) \vee in(S, else) \vee after(S)))$
 $\neg(in(S, then) \wedge in(S, else))$

while $at(S) = at(S, test)$
 $in(S, test) = at(S, test)$
 $after(S, test) = at(S, body) \vee after(S)$
 $after(S, body) = at(S, test)$
 $in(S) = at(S, test) \vee in(S, body)$
 $\neg(at(S, test) \wedge (after(S) \vee in(S, body)))$

new There are two cases. If there is no **init** clause, then:

$$\begin{aligned}
 at(S) &= at(S, body) \\
 after(S) &= after(S, body) \\
 in(S) &= in(S, body)
 \end{aligned}$$

If there is an **init** clause, then:

$$\begin{aligned}
 at(S) &= at(S, init) \\
 in(S, init) &= at(S, init) \\
 after(S, init) &= at(S, body) \\
 in(S) &= in(S, init) \vee in(S, body) \\
 after(S) &= after(S, body) \\
 &\neg(at(S, init) \wedge in(S, body))
 \end{aligned}$$

cobegin If there are n clauses in the **cobegin**, then

$$\begin{aligned}
 at(S) &= at(S, 1) \wedge \dots \wedge at(S, n) \\
 after(S) &= after(S, 1) \wedge \dots \wedge after(S, n) \\
 in(S) &= in(S, 1) \wedge \dots \wedge in(S, n) \\
 \{in(S, i), after(S, i)\} &\perp \{in(S, j), after(S, j)\} \text{ for } i \neq j
 \end{aligned}$$

sequence If S is $S_1; \dots; S_n$, then for all $i = 1, \dots, n$:

$$\begin{aligned}
 after(S, i-1) &= at(S, i) \text{ if } i > 0 \\
 in(S) &= in(S, i) \\
 &\neg(in(S, i) \wedge in(S, j)) \text{ for } i \neq j
 \end{aligned}$$

program If S is the complete program, then

$$\begin{aligned}
 at(S) &= at(S, body) \\
 in(S) &= in(S, body) \\
 after(S) &= after(S, body)
 \end{aligned}$$

A close study of these aliasing relations reveals that we can prove a relation such as $in(S, \gamma) \perp at(S, \mu)$ if and only if the substatements S, γ and S, μ lie in different clauses of a **cobegin**.

The \perp relations among program variables and the aliasing and \perp relations among control variables are regarded as axioms in a separate system for reasoning about syntactic expressions. However, they play the same function as axioms of $\mathcal{M}[S]$. For example, if S, γ is an assignment statement, then the aliasing relation $\neg(at(S, \gamma) \wedge after(S, \gamma))$ allows us to deduce $\Diamond(\neg at(S, \gamma))$ from $\Diamond after(S, \gamma)$.

6.1.2 Syntactic Typing Relations

Because the type structure of our language L is so simple, no explicit reference to types need appear in its semantics. However, this is not the case for a language in which the action of an assignment statement is affected by the types of its left- and right-hand sides—for example, if coercion was performed. We would also have to introduce explicit reference to types if a type mismatch in an assignment statement produced a run-time error or an indeterminate result, or if it halted the process executing the assignment.

Explicit reference to types is done by introducing predicates such as $type(x) = integer$. If the types of variables are determined syntactically by the program text, then these predicates would be syntactic predicates. Otherwise, they would have to be ordinary state predicates, and the state would have to include components that determine their values.

6.1.3 Reasoning About Syntactic Expressions

Although a syntactic predicate like $v \perp w$ resembles an ordinary state predicate like $v = 7$, it is logically quite different. The variable name “ v ” denotes the value of the variable in the expression $v = 7$, while it denotes the name itself in $v \perp w$. For example, from the expressions $v = 7$ and $w = v$ we can deduce $w = 7$. However, from the syntactic expression $u \perp v$ and the ordinary expression $w = v$ we cannot in general deduce $u \perp w$; just because the *values* of two variables happen to be equal in some state does not imply that the variables have the same aliasing relations. We can only make that conclusion if $w = v$ is a syntactic equality of names, rather than an expression denoting equality of values.

By introducing syntactic predicates as a class of entities separate from ordinary state predicates, with their own logical system for reasoning about

them, I have circumvented the need to distinguish between the use of a variable name as a name and as a value. In a syntactic predicate, a variable name represents itself. In a state predicate, it represents the value of the variable. Using two different logical systems avoids confusion. One cannot make invalid deductions, like deducing a \perp relation from the equality of the values of v and w , because inferences about \perp can be made only in the logic for reasoning about syntactic predicates, whereas equality of values can be expressed only with state predicates, and one reasons about them with a separate logic.

For languages in which types and aliasing relations are dynamic properties, so they must be reflected in the state, we cannot use this trick for separating the two different uses of variable names. We must then write $value(v)$ rather than the variable name v to denote the value of v . Equality of values is denoted by the predicate $value(w) = value(v)$, and $w = v$ denotes equality of names.

6.1.4 Logical Name Variables

Just as I introduced logical value variables in state predicates, I will also introduce *logical name variables* for syntactic predicates. A logical value variable is a logical variable with an implicit range in the set of values that a variable may have. Similarly, a logical name variable is a logical variable with an implicit range in the set of names that a variable may have. I will use the letter ν to denote a logical name variable.

The use of logical name variables has an important implication with respect to renaming. Consider an axiom of the form $\forall \nu : A(\nu)$. Viewed as a formula in $\mathcal{TL}(S)$, it is equivalent to an infinite conjunction of the form $A(v_1) \wedge A(v_2) \wedge \dots$, where the v_i are all the names in $\mathcal{V}(S)$. However, the two formulas behave differently under a renaming mapping ρ . In particular, $\rho(\forall \nu : A(\nu))$ equals $\forall \nu : \rho(A(\nu))$, so the renamed formula includes a quantification over variable names not present in $\rho(A(v_1) \wedge \dots)$.

6.2 Starting States

One might expect that the meaning $M[S]$ of a statement S should include a set of axioms that determine the set of starting states. However, consider what the initial value of a program variable should be. The user has no way of specifying it, since an `init` clause of a new statement is interpreted as an executable action that replaces the initial value with the specified one. One

might want to specify that the initial value of a variable v of type **integer** should be an integer. However, $M[S]$ will contain an axiom asserting that this is true for every state during the execution of S , so it is therefore true of the initial state. Similarly, the axioms in $M[S]$ will assert that the aliasing relations specified by **new** statements are true throughout the execution, so they are also constrained to hold in the initial state.

What about the initial values of control variables? Surely we should require that a statement S should start in a state in which $at(S)$ is true. However, this would be a mistake because it would violate the Composition Principle, since $\rho_{T;S,2}(at(S))$ should not be true of the starting state of the sequence of statements $T;S$, and our whole approach is based upon the Composition Principle.

Remember that the only reason for specifying the starting state is to be able to obtain from our semantics a set of behaviors. However, we are really interested only in the set of behaviors of a complete program, not of its substatements. There is no reason to constrain the starting states of substatements; we need only constrain the starting state of a complete program, which we do by simply assuming that $at(\Pi)$ is true of the initial state of a complete program Π . We can do this without violating the Composition Principle because a complete program cannot be part of any larger statement.

6.3 Behavior Axioms

I now define the set $M[S]$ of behavioral axioms for any statement and complete program S . This will, of course, be done compositionally, giving a set axioms for each language construct. Remember that in addition to the axioms given explicitly below, $M[S]$ also contains all the axioms implied by the Composition Principle.

I will include in $M[S]$ axioms to assert that the appropriate aliasing relations hold throughout the execution of S . For control variables, those aliasing relations were already described in Section 6.1.1. Rather than write them over again, I will simply assume that the aliasing relations described there appear as axioms in $M[S]$ for the appropriate construct describing S . For example, the list of axioms for the assignment given below are assumed implicitly to include the axioms $in(S) = at(S)$ and $\neg at(S) \wedge after(S)$ from Section 6.1.1. (However, the \perp relations given for the **cobegin**, being syntactic predicates, are not axioms in $M[S]$.)

In addition to the aliasing relations for control variables, we should also

assert their types. Therefore, we implicitly add the axiom $v \in \{true, false\}$ to $\mathcal{M}[S]$ for every control variable v in $\mathcal{V}(S)$.

There are also axioms relating the action predicate $Act(S)$ to the action predicates of its components. For example, the axioms for a **while** statement S would include the following:

- $Act(S) \equiv Act(S, test) \vee Act(S, body)$
- $\neg(Act(S, test) \wedge Act(S, body))$

The first asserts that the only actions of S are the *test* action and the actions of its body; the second asserts that the *test* action is not an action of the body. These and similar axioms are assumed for all the constructs and are not included. Note that these axioms are given only for compound statements; there is no such axiom for the assignment statement.

In the following description of the axioms, formal axioms are followed by their informal explanations. For any programming-language expression exp , I let $exp()$ denote the expression obtained by replacing every identifier y in exp by the variable name $y()$.

6.3.1 Assignment

If S is the statement $\langle x := exp \rangle$, then $\mathcal{M}[S]$ contains the following axioms:

1. $Act(S) \supset at(S)$
The atomic statement S can be executed only when control is at S .
2. $\forall \eta : \{at(S) \wedge exp() = \eta\} \langle S \rangle \{after(S) \wedge x() = \eta\}$
Executing S sets the value of x to exp and changes control from $at(S)$ to $after(S)$.
3. $\forall \nu : \{x(), at(S), after(S)\} \perp \nu \supset \nu \not\stackrel{S}{\neq}$
(Note that ν is a logical name variable.) The statement S does not modify any variable not aliased to x , $at(S)$, or $after(S)$.
4. $\Box \Diamond Act(S) \supset \Diamond \neg at(S)$
There cannot be infinitely many actions of S while control remains forever at S . (The reader may find this easier to understand if he replaces the implication by a disjunction.) In other words, there can be only finitely many stuttering actions of S before the assignment is executed.

An understanding of these axioms for assignment is crucial to an appreciation of how action-axiom semantics works, so some further discussion of them is in order. The four axioms are indeed action axioms, since they describe the behavior of the assignment action S . The four axioms assert the following:

1. When the action *may* occur.
2. What changes to the state components executing the action *may* perform.
3. What state components the action may not change.
4. When the action *must* change the state.

Every atomic program action is described by four similar axioms.

Note that axioms 1-3 assert safety properties, while axiom 4 states a liveness property. From the axioms for the other statements, it will follow that in language L , if $at(S)$ ever becomes true, it can be made false only by executing action S . Axiom 2 asserts that $at(S)$ can then become false only when $after(S)$ becomes true and the assignment of exp to x occurs. In a richer language, executing another statement might make $at(S)$ become false—for example, by aborting the process containing statement S . However, Axioms 1-4 would still be valid.

Observe that Axiom 2 determines the value of x immediately after execution of S . However, it asserts nothing about x 's value after the execution of any other action.

For language L , Axiom 4 implies that if $at(S)$ is true then eventually it will become false (thereby making $after(S)$ true). However, this depends upon the fact that L does not have any form of unfair cobegin. The axiom is valid for more general languages that do have these features.

It is instructive to consider what these axioms imply in case statement S appears inside declarations that produce a type mismatch—say in which x is of type **integer** and exp of type **boolean**. The axioms for those declarations will imply that the value of x is always an integer and the value of exp is always a boolean. It then follows from Axiom 2 that executing an S action can never make $at(S)$ false, since doing so would require setting the value of x to a boolean, contradicting the axioms for the declarations. However, I have already observed that, for language L , $at(S)$ must eventually become false. Thus, the set of axioms for the incorrect program—the one producing a type mismatch in statement S —are contradictory, implying that only

the empty set of behaviors satisfy them. However, in a richer language, if S were contained inside an unfair cobegin, then the axioms might not be contradictory, and might be satisfied by a behavior in which a process remained stalled forever with $at(S)$ true. In this case, the type mismatch would force that process to “die”, allowing other processes to proceed.

6.3.2 The if Statement

If S is the statement **if** (exp) **then** ... , then the following axioms are in $\mathcal{M}[S]$. They are the standard four action axioms—in this case, for the *test* action. Note their similarity to the corresponding axioms for the assignment statement.

1. $Act(S, test) \supset at(S, test)$
The test can be executed only when control resides at it.
2. $\{at(S, test)\}\langle S, test \rangle\{\{at(S, then) \wedge exp()\} \vee \{at(S, else) \wedge \neg exp()\}\}$
Control remains at the beginning of the test until it either reaches the entry point of the **then** clause with exp true, or else it reaches the entry point of the **else** clause with exp false.
3. $\forall \nu : \{at(S, test), after(S, test)\} \perp \nu \supset \nu \not\stackrel{S, test}{\neq}$
The test does not modify any variable it shouldn't. (Again, ν is a logical name variable.)
4. $\Box \Diamond Act(S, test) \supset \Diamond \neg at(S, test)$
There can be only finitely many stuttering actions of the test before it is really executed. This is the only liveness axiom for the if statement.

6.3.3 The while Statement

The axioms for the statement **while** (exp) **do** ... are analogous to the ones for the if statement, and are given without comment.

1. $Act(S, test) \supset at(S, test)$
2. $\{at(S, test)\}\langle S, test \rangle\{\{at(S, body) \wedge exp()\} \vee \{after(S) \wedge \neg exp()\}\}$
3. $\forall \nu : \{at(S, test), after(S, test)\} \perp \nu \supset \nu \not\stackrel{S, test}{\neq}$
4. $\Box \Diamond Act(S, test) \supset \Diamond \neg at(S, test)$

6.3.4 The new Statement

The **new** statement is a declaration. If it has no **init** clause, then it performs no new action. The axioms describing this statement therefore do not follow the pattern for action axioms followed by the preceding statements. Instead, they assert relations that hold throughout the execution.

If S is the statement

new $x : type$ **in** ...

then the following axiom is in $\mathcal{M}[S]$, where we identify **integer** with the set Z and **boolean** with the set $\{true, false\}$.

1. $x(S) \in type$

The value of x is always consistent with the type declaration.

If S is the statement

new $x : type$ **alias** exp **in** ...

then $\mathcal{M}[S]$ contains the above axiom plus the following:

2. $x(S) = exp()$

The aliasing relation always holds.

If S is the statement

new $x : type$ **init** exp **in** ...

then the following axioms hold. The first is, of course, the same as for the other versions of the **new** statement. The last four are the action axioms for the initial-assignment action, following the standard pattern. They are almost identical to the corresponding axioms for the assignment statement, the only difference (in axiom 3 below) indicating that the **init** clause performs an assignment to the variable $x(S)$ declared in the **new** statement rather than to the undeclared variable $x()$.

1. $x(S) \in type$

2. $Act(S, init) \supset at(S, init)$

3. $\forall \eta : \{at(S, init) \wedge exp() = \eta\} \langle S \rangle \{after(S, init) \wedge x(S) = \eta\}$

4. $\forall \nu : \{x(S), at(S), after(S)\} \perp \nu \supset \nu \neq \overset{S,init}{\nu}$

5. $\square \diamond Act(S, init) \supset \diamond \neg at(S, init)$

6.3.5 The cobegin Statement

If S is the statement

. cobegin $S_1 \square \dots \square S_n$ coend

then the following axiom is in $\mathcal{M}[S]$.

1. $\forall i$ s.t. $1 \leq i \leq n : (\square \diamond Act(S)) \supset (\square \diamond Act(S, i))$
If S performs infinitely many actions, then each process of S performs infinitely many actions. In other words, if S is never starved, then no subprocess of S is starved. This is the fairness axiom.

6.3.6 Sequences of Statements

No new axioms are needed for the sequence of statements $S_1; \dots; S_n$. All necessary properties are obtained from the aliasing relations among its control variables, the relations among its action predicates, and the Composition Principle.

6.3.7 A Complete Program

If Π is a complete program, then the only additional axiom in $\mathcal{M}[\Pi]$ is:

1. $in(\Pi) \supset Act(\Pi, body)$
The complete program never stops executing until it reaches the end, whereupon $in(\Pi)$ becomes false.

This axiom asserts the absence of any external actions while control is in program Π , reflecting the absence of any explicit input or output in language L .

Constraints:
A Uniform Approach to Aliasing and Typing

Leslie Lamport¹
SRI International

Fred B. Schneider²
Cornell University

24 July 1984
revised 4 September 1984
minor revision 26 October 1984

To appear in *Proceedings of the 12th Annual ACM
Symposium on the Principles of Programming Lan-
guages* (January, 1985)

¹Work supported in part by the National Science Foundation under grant number MCS-8104459 and by the Army Research Office under grant number DAAG29-83-K-0119. Current address: SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025.

²Research supported in part by NSF grant DCR-8320274 and by a Faculty Development Award from IBM Corp. Current address: Department of Computer Science, Cornell University, Ithaca, NY 14853.

APPENDIX C

**CONSTRAINTS:
A UNIFORM APPROACH TO ALIASING AND TYPING**

- [12] S. S. Owicki and L. Lamport. Proving Liveness Properties of Concurrent Programs. *ACM Trans. on Prog. Lang. and Systems* 4, 3 (1982), 455-495.
- [13] A. Pnueli. The Temporal Logic of Programs. *Proceedings of the 18th IEEE Symposium on the Foundations of Computer Science* (1977), Providence Rhode Island.
- [14] H. Barringer, R. Kuiper. and A. Pnueli. Now You May Compose Temporal Logic Specifications. *Sixteenth ACM Symposium on the Theory of Computing*, (May, 1984).
- [15] J. Sifakis. A Unified Approach for Studying the Properties of Transition Systems. *Theoretical Computer Science* 18 (1982), 227-258.

References

- [1] Karl M. Abrahamson. *Decidability and Expressiveness of Logics of Processes*. Ph. D. Thesis, issued as Technical Report No. 80-08-01, Department of Computer Science, University of Washington. (August 1980).
- [2] Denotational Semantics of Concurrency. J. W. de Bakker and J. C. Zucker. *Fourteenth ACM Symposium on the Theory of Computing*, San Francisco, (May, 1982), 153-158.
- [3] Z. C. Chen and C. A. R. Hoare. Partial Correctness of Communicating Sequential Processes. *Proceedings of the Second IEEE International Conference on Distributed Computer Systems*, (1981) 1-12.
- [4] L. Lamport. The "Hoare Logic" of Concurrent Programs. *Acta Informatica* 14 (1980), 21-37.
- [5] L. Lamport. "Sometime" Is Sometimes "Not Never". *Proceedings of the Seventh Annual ACM Conference on the Principles of Programming Languages*, (January 1980) 174-185.
- [6] L. Lamport. Specifying Concurrent Program Modules. *ACM Transactions on Prog. Logic and Sys.* 5, 2 (April 1983) 190-222.
- [7] L. Lamport. Reasoning About Nonatomic Operations. *Proceedings of the Tenth Annual ACM Conference on the Principles of Programming Languages*, (January 1983) 28-37.
- [8] L. Lamport. What Good Is Temporal Logic? *Information Processing 83*, R. E. Mason (ed.). Elsevier Science Publishers, North Holland (1983). 657-668.
- [9] L. Lamport and F. B. Schneider. The Hoare Logic of CSP and All That. *ACM Transactions on Prog. Logic and Sys.* 6, 2 (April 1984) 281-296.
- [10] L. Lamport and F. B. Schneider. Constraints: A Uniform Approach to Aliasing and Typing. To appear in *Proceedings of the Twelfth Annual ACM Conference on the Principles of Programming Languages*, (January 1985).
- [11] R. Milner. *A Calculus of Communicating Systems*. Lecture Notes in Computer Science, Number 92. Springer-Verlag, Berlin (1980).

Acknowledgements

I wish to thank Fred Schneider and Willem-Paul de Roever for their detailed comments on an earlier draft of this paper.

from the axioms for the corresponding construct given here in Section 6.3, together with the Composition Principle.

As described in [9], other logical systems for proving safety properties of concurrent programs can be described in terms of GHL, so the soundness of GHL can be used to prove the soundness of the other systems. GHL is manifestly not a complete system for reasoning about concurrent programs, since it does not address questions of liveness. It is not clear how to use our semantics to prove completeness of GHL for the class of properties it can express.

A method for proving liveness properties of programs is given in [12]. It considers a simple language that is essentially the same as language L except without the new statement. The method explicitly assumes a complete program Π , and is based upon temporal logic plus the following single axiom:

Atomic Action Axiom: For any atomic action Π, γ of Π :

$$at(\Pi, \gamma) \supset \Diamond after(\Pi, \gamma)$$

To prove the soundness of this axiom, we must show that

$$\Box (in(\Pi, \gamma) \supset \Diamond Act(\Pi, \gamma))$$

holds for every substatement and atomic action Π, γ of Π . This is intuitively clear, since the language contains only fair cobegin statements, and is derivable from our axioms by induction on the size of Π . The above Atomic Action Axiom then follows easily from our liveness axiom for complete programs, the liveness axioms for the individual statements, plus the Composition Principle. The additional axioms given in [12] for weakly and strongly fair semaphore operations can similarly be derived from the ones I gave earlier.

8 Conclusion

I have given an axiomatic semantics for a simple concurrent programming language L , and have indicated how the same method can be applied to more complicated language constructs. Most of this paper has been devoted to developing the fundamental ideas upon which the method is based. The axioms themselves are reasonably simple—simple enough so I feel that they do provide an understanding of the language constructs. For example, the difference between a weakly fair and a strongly fair semaphore is described quite concisely and precisely by their respective axioms.

A programming language semantics provides a logical basis for a proof system for reasoning about programs in the language. One can talk about the soundness and completeness of the proof system in terms of the semantics. Note that it makes no sense to talk about soundness and completeness of the semantics. Indeed, the semantics $\mathcal{M}[S]$ of a program can include contradictory axioms; this simply means that there are no valid behaviors for S , so there is something wrong with the program, not with the semantics.

The obvious task now is to investigate existing proof systems in terms of this semantics. Unfortunately, such an undertaking is beyond the scope of this paper. However, some brief remarks are in order. The Generalized Hoare Logic (GHL) presented in [4] and [9] introduced *at*, *in*, and *after* as predicates rather than variables. The relation \parallel used in [4] is just the relation \perp .

The semantics of GHL formulas was not stated with sufficient precision in [4], since the relation between the statement S and its name, denoted ' S ', was never made clear. A close examination of GHL reveals that there is an implicit complete program Π , and that if S is the substatement Π, γ of Π , then a formula written in terms of S should really be written in terms of Π, γ .

To verify the soundness of GHL, one must express the GHL formula $\{P\} S \{Q\}$ as a temporal logic formula. As explained in [9], it suffices to consider the case $P = Q$, for which the definition is simply:

$$\{P\} S \{P\} \stackrel{\text{def}}{=} \{P\} \langle \Pi, \gamma \rangle \{P\}$$

where S is the substatement Π, γ of the implied complete program Π . The soundness of the general rules for reasoning about GHL formulas follows easily from their interpretation as temporal logic formulas. The soundness of the axioms and rules given in [4] for each language construct can be deduced

$$\forall \eta : \{in(S) \wedge [S](x() = \eta)\} \langle S \rangle \{after(S) \wedge x() = \eta\}$$

Note that the rules for reasoning about these generalized dynamic logic predicates imply that

$$at(S) \supset ([S](x() = \eta) \equiv (exp = \eta))$$

The liveness axiom for a nonatomic assignment is simply

$$\Box \Diamond Act(S) \supset \Diamond \neg in(S)$$

liveness requirements one could make in this case, since the value of the expression exp could change. One reasonable possibility is the following:

$$\Box \Diamond (Act(S) \wedge type_valid(x(), exp())) \supset \Diamond \neg at(S)$$

Allowing a more general form of aliasing, such as the one defined in [10], presents a similar problem if one requires that an assignment which would violate an aliasing constraint cause the process to hang up. One approach to this is to put the aliasing constraints in the state, just as I did with type constraints. The new state components would correspond to the "location" values often used to handle aliasing.

7.5 Nonatomic Operations

Every construct that I have mentioned specifies the atomic actions. For example, I have defined the semantics only of an atomic assignment statement. It is easy to give the semantics of an assignment statement with smaller atomic operations. For example, an assignment

$$\langle x \rangle := \langle exp \rangle$$

in which the evaluation of exp and the changing of x are distinct atomic operations can be represented by

$$\langle t := exp \rangle; \langle x := t \rangle$$

where t is an implicit variable. A similar translation is possible when the evaluation of the right-hand side is broken into smaller atomic operations; it is described in [4].

The situation changes when no atomicity is specified. For example, consider an assignment statement $x := y + 1$ that has the expected effect only if x and y are not modified by any other operation during the course of its execution. If any such modification does take place, then x may be set to any value consistent with its type. We can think of this assignment as a compound statement for which we know nothing about its internal structure except its partial correctness property (when executed alone) and the fact that it always terminates (unless the process executing it is starved).

Handling such nonatomic operations requires a new class of state predicate—the "generalized dynamic logic" predicates $[S]P$ introduced in [7]. The second assignment axiom for an atomic assignment is replaced by the following one for a nonatomic assignment $x := exp$:

While this method of handling procedure calls works only for nonrecursive procedures, the basic idea applies to recursive ones as well. Replacing a procedure call by the body of the procedure produces an infinite program text for recursive procedures; but nowhere have I made use of the assumption that the program text is finite. Of course, the compositional method of recursively defining $M[S]$ no longer terminates with a finite set of axioms. However, the definition can be viewed as an algorithm for enumerating an infinite collection of axioms.

Thus, adding recursion means that $M[S]$ consists of an infinite set of axioms. It is in this case that the distinction between a semantics and a proof system becomes evident. An infinite set of axioms is unsatisfactory as a proof system, because ordinary logic provides no way of deducing a conclusion whose correctness is based upon an infinite set of assumptions. Such deductions are required to prove nontrivial properties of recursive programs. Thus, I have not provided a proof system for programs with recursive procedures.

On the other hand, a semantics is concerned with validity, not proof. The meaning of a program Π is the set of behaviors that satisfy the axioms in $M[\Pi]$, and this is well-defined even for an infinite set of axioms. The problem of proof systems is discussed in the conclusion.

7.4 More General Types and Aliasing

Let us now consider a language in which a type mismatch does not produce an illegal program, but generates "incorrect" behavior. As mentioned earlier, this requires adding predicates of the form $type(x) = \dots$, which are syntactic predicates if types can be determined syntactically and state predicates if types are dynamic.

First, suppose that a type mismatch in the assignment $x := exp$ causes x to be set nondeterministically to any value in its range. This is easily represented by changing axiom 2 of the assignment to the following, where $type_valid(x, \eta)$ is true if and only if the type of x permits it to be assigned the value η .

$$\forall \eta : \{at(S) \wedge exp() = \eta\} \langle S \rangle \\ \{after(S) \wedge (x() = \eta \vee \neg type_valid(x(), \eta))\}$$

Next, suppose that a type mismatch causes the assignment to "hang up", effectively deadlocking the process. This requires that axiom 4 be changed so it does not demand termination in this case. There are several different

Note that $\rho_{S,i,\gamma}(S, i, \gamma, left())$ is the *exp* of $exp! \xi$, with all component variables appropriately renamed, and similarly for $\rho_{S,j,\mu}(S, j, \mu, left())$.

It is straightforward to extend this approach to guarded communication commands such as $\langle exp \rightarrow x? \xi \rangle$, which means that the communication action may be carried out only if *exp* has the value *true*. The new safety axiom is obtained from the above in much the same way that the safety axiom for the $P(s)$ semaphore operation is obtained from the corresponding axiom for the assignment statement—the guards here playing the part of the enabling condition $s > 0$ for the $P(s)$ operation.

There are several different choices of liveness properties that one can require of these channels. They are all basically simple to express with temporal logic formulas. However, their formal statement requires some careful manipulation of syntactic predicates, which I won't bother doing.

The safety properties of CSP-like communication primitives are expressed more easily with a formal semantics based only upon externally observable actions, such as [11]. When shared variables are not allowed, such a semantics can define the meaning of a process as the set of possible communications it can engage in. However, this kind of semantics does not seem capable of handling liveness properties easily.

7.3 Procedures

Although language *L* does not have procedures, its new statement contains the basic mechanism needed for procedure calls. A call of a nonrecursive procedure can be simulated by replacing the procedure call by new statements plus the body of the procedure. For example, let *proc* be a procedure with a declaration

```
procedure proc(a : integer, var b : boolean ) body
```

in which its first argument is call by value and its second is call by name. The call $proc(x + y, z)$ can be translated to

```
new a init x + y in new b alias z  
  in body ni ni
```

To handle call by reference parameters, one needs to introduce pointer variables into the language. Of course, aliasing and procedure calls become more interesting when pointers and arrays are introduced, but a discussion of the problems raised by pointers and arrays is beyond the scope of this paper.

occur when s has the value zero. Several liveness axioms have been proposed for the semaphore. Probably the most common are weak liveness, expressed by

$$(\Box(s > 0) \wedge \Box \Diamond Act(P(s))) \supset \Box \neg at(P(s))$$

and strong liveness, expressed by

$$(\Box \Diamond(s > 0) \wedge \Box \Diamond Act(P(s))) \supset \Box \neg at(P(s))$$

(They are discussed in [12].) In both these cases, the $V(s)$ operation is just an ordinary atomic assignment.

More complicated versions of the semaphore impose a specific queueing discipline, like first-come-first-served, on the execution of competing $P(s)$ operations. They may require adding a queue of waiting processes to the state, plus predicates to describe the state of the queue.

7.2.2 CSP-Like Communication Primitives

The easiest way to model the CSP “!” and “?” operations is in terms of channels. We include the operations $\langle x?\xi \rangle$ and $\langle exp!\xi \rangle$ for any variable x and expression exp . They denote CSP-like synchronous communication over a channel named ξ . We modify the **cobegin** statement by adding a clause of the form **channels** ξ_1, \dots, ξ_m , which declares the channel names ξ_i .

As explained in [9], we consider communication actions to be actions of the channel, so $Act(S)$ is identically *false* if S is a ! or ? operation. A channel ξ has a separate atomic action for every pair of statements $\langle x?\xi \rangle$, $\langle exp!\xi \rangle$ contained in different clauses of the **cobegin** in which ξ is declared. This atomic action is axiomatized much like the assignment statement $\langle x := exp \rangle$, except that its execution changes the values of the four control variables $at(x?\xi)$, $after(x?\xi)$, $at(exp!\xi)$, and $after(exp!\xi)$.

To do this formally, we must extend our variable-naming convention in the obvious way to channel variables and add new syntactic predicates $S, \gamma \in !v$ and $S, \gamma \in ?v$ to assert that the substatement S, γ is a ! or ? operation of the channel named v . The safety axiom for the declaration of channel ξ will be something like:

$$\begin{aligned} \forall \gamma, \mu, i, j \text{ s.t. } i \neq j : S, i, \gamma \in !\xi(S) \wedge S, j, \mu \in ?\xi(S) \supset \\ \forall \eta : \{ at(S, i, \gamma) \wedge at(S, j, \mu) \wedge \rho_{S, i, \gamma}(S, i, \gamma, left()) = \eta \} \\ \{ \xi(S) \} \{ after(S, i, \gamma) \wedge after(S, j, \mu) \wedge \rho_{S, j, \mu}(S, j, \mu, left()) = \eta \} \end{aligned}$$

7.1.3 Write Protection

Imagine a situation in which one wants the variable x to be modified only in a particular statement, but to be accessible elsewhere. This might be expressed by the following statement S :

encapsulate x in S'

The semantics of this statement are described formally by:

$$\forall \eta : x() = \eta \supset (\neg Act(S)) \triangleleft (x() = \eta)$$

which asserts that the value of x remains unchanged while any action not in S is executed.

7.2 Synchronization and Communication

The bread and butter of concurrent programming language constructs are the synchronization and interprocess communication mechanisms. I will discuss only two.

7.2.1 Semaphores

The usual semaphore P and V operations are variants of the atomic assignment statement: $P(s)$ looking much like the assignment $\langle s := s - 1 \rangle$ and $V(s)$ looking like $\langle s := s + 1 \rangle$. There are two basic differences. First of all, the $P(s)$ operation may be performed only when s is positive. One way of expressing this is to change the first axiom of the assignment statement to:

$$Act(P(s)) \supset (at(P(s)) \wedge s > 0)$$

However, this would require changing other axioms, since deadlock is represented by the absence of any possible actions, and the axiom given above for the complete program asserts that this is impossible.

The other way of handling this is to allow only stuttering actions of $P(s)$ to occur when $s \leq 0$. This is achieved by replacing the second axiom of the assignment statement with the following:

$$\forall \eta : \{at(P(s)) \wedge (x() = \eta)\} \langle P(s) \rangle \{after(P(s)) \wedge (x() = \eta - 1 \geq 0)\}$$

The second change that must be made to the assignment axioms is in the liveness condition. We can no longer require that an infinite number of actions of $P(s)$ cause the operation to be completed, since they might all

7 Other Language Features

While I have given a formal semantics only for the simple language L, action-axiom semantics can be used to describe a wider variety of concurrent programming language constructs than any other method I know of. In this section, I will consider a few interesting constructs. In doing so, I will not bother to give the usual axioms that describe the relations among control variables and among action predicates.

7.1 Constructs That Constrain Their Environment

Most language constructs constrain the behavior of their components. For example, an if statement determines when its then and else clauses can be executed. The following three language constructs constrain the behavior of a larger program containing them. They are therefore impossible to specify in a compositional, purely behavioral semantics. It is the Composition Principle that makes them expressible with action-axiom semantics.

7.1.1 The assign processor Command

As described above, the statement

assign processor to ...

directs the compiler to guarantee that the body of the statement gets its share of computing cycles, so it is not starved. This is expressed by the axiom:

$$in(S) \supset \Diamond Act(S)$$

7.1.2 Atomic Actions

One might want to introduce "angle brackets" as a language construct, so $\langle S \rangle$ denotes that S is to be executed as an indivisible atomic action. This is done by requiring that no other actions are interleaved with the executions of S , expressed formally by:

$$Act(\langle S \rangle) \supset in(\langle S \rangle) \trianglelefteq Act(\langle S \rangle)$$

Constraints: A Uniform Approach to Aliasing and Typing

Leslie Lamport*
SRI International

Fred B. Schneider†
Cornell University

Abstract

A constraint is a relation among program variables that is maintained throughout execution. Type declarations and a very general form of aliasing can be expressed as constraints. A proof system based upon the interpretation of Hoare triples as temporal logic formulas is given for reasoning about programs with constraints. The proof system is shown to be sound and relatively complete, and example program proofs are given.

1 Introduction

Type declarations and aliasing relations have traditionally been thought of as unrelated concepts. However, both can be viewed as specifying properties that do not change during program execution. This view leads to a uniform method for reasoning about types and aliasing in which the usual Hoare logic triples are regarded as temporal logic formulas.

Aliasing two variables x and y means they always have the same value. This is usually implemented by allocating the same memory location to x and y , thereby ensuring that both variables are changed whenever either is assigned a new value. However, they could be allocated separate memory locations and both updated on an assignment to either. Viewing aliasing as defining certain relationships

*Work supported in part by the National Science Foundation under grant number MCS-8104459 and by the Army Research Office under grant number DAAG29-83-K-0119. Current address: SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025.

†Research supported in part by NSF grant DCR-8320274 and by a Faculty Development Award from IBM Corp. Current address: Department of Computer Science, Cornell University, Ithaca, NY 14853.

between the values of variables, with no implication about storage allocation, allows more general kinds of aliasing and leads to a simple method for reasoning about aliasing.

To express a more general form of aliasing, we introduce the `var` statement. To illustrate its use, suppose a program computes a temperature, and that some times it is convenient to refer to that temperature in degrees Fahrenheit and other times in degrees Celsius. We will write the statement

```
var  $f, c$  : real constraints  $f = 9 * c / 5 + 32$  in  $S$ 
```

which declares variables f and c within statement S to be of type *real* and to be *aliased*, so that if the value of f is a temperature in degrees Fahrenheit, then the value of c is that temperature in degrees Celsius. Changing f causes a corresponding change to c , and vice-versa. Notice that this more general form of aliasing cannot be implemented simply by allocating overlapping memory locations to f and c .

The `constraints` clause of a `var` statement is a directive that a specified predicate—in our example, the aliasing relation $f = 9 * c / 5 + 32$ —be maintained as an invariant, which means that execution is aborted if the predicate becomes false.

A type declaration can also be viewed as an invariant, so it can be specified in a `constraints` clause. If we take the view that the type of a variable defines the set of values that variable can have, then declaring a variable f to be of type *real* is the same as requiring that the predicate $f \in \mathcal{R}$ be true throughout execution, where \mathcal{R} is the set of real numbers.¹ Thus, we could eliminate the “: *real*” from the above `var` statement and add the constraint $x, y \in \mathcal{R}$. Since doing so would make the statement less readable, we will retain the customary syntax for type declarations.

Aliasing and typing can be viewed in terms of constraints because they are *static* properties. While *dynamic* properties, such as the values of variables, can be changed by execution of a program statement, static properties cannot. (In most languages, like the one considered here, a

¹For simplicity, we assume \mathcal{R} is the infinite set that mathematicians call the real numbers, thereby avoiding the problems that round-off errors would introduce for reasoning about equality of expressions.

declaration is not a complete statement but rather part of a statement.) The methods we develop for reasoning about aliasing and types can be used to reason about any static property.

Returning to aliasing, consider a more complicated example in which a program refers a point in terms of both its Cartesian coordinates x, y and its polar coordinates r, θ . Variables x, y, r , and θ are declared as follows.

```
var x, y, r: real,  $\theta$ : [0, 2 *  $\pi$ )
  constraints  $x = r * \cos(\theta)$  and  $y = r * \sin(\theta)$ 
  in S
```

(The type declaration for θ states that it is a real in the range $0 \leq \theta < 2\pi$.) We would like this declaration to mean that when x is changed, r and θ are changed according to the constraints, but y is not. However, the fact that y should not change is based upon the knowledge that x and y are independent coordinates, which is not something discernible in the above statement. An additional constraint is needed to specify that assigning to x should not change the value of y and vice-versa; we write this constraint as $x \perp y$. Similarly, r and θ should be independent, so assigning a value to either r or θ does not change the other. Hence, the additional constraint $r \perp \theta$ is needed. The following declaration of x, y, r , and θ gives the desired aliasing relations.

```
var x, y, r: real,  $\theta$ : [0, 2 *  $\pi$ )
  constraints  $x = r * \cos(\theta)$  and  $y = r * \sin(\theta)$ 
             and  $x \perp y$  and  $r \perp \theta$ 
  in S
```

Finally, observe that the `var` statement can express forms of aliasing traditionally implemented by overlapping storage. The statement

```
var full, right_4: natural
  constraints right_4 = full mod 16
  in S
```

aliases variable `right_4` to the right-most four bits of `full`, where `natural` denotes the nonnegative integers. Moreover, the declaration ensures the desired semantics even on a computer where integers are not stored in binary.

It is probably impossible for a compiler to handle our form of aliasing in all its generality. While the Fahrenheit/Celsius and `full/right_4` examples do not pose difficult compiling problems, consider what happens if the following statements appear in the body of the above `var x, y, r, θ` statement:

```
read(x, y); write( $\theta$ )
```

Input values a, b with $a \neq 0$ produce the output value $\arctan(b/a)$ —something no present-day compiler is likely to figure out.

We are interested in our general form of aliasing in order to reason about *implicit variables*—variables representing portions of the program state that are not directly visible to the programmer. For example, in a concurrent

programming language with primitives to perform buffered message-passing, messages sent but not yet delivered are part of the state that must be described by implicit variables. (The ρ and σ multisets of [18] are such variables.) Implicit variables often involve complex aliasing relations. For some message-passing schemes, a channel is modelled by having an implicit variable in a sender aliased to an implicit variable in the receiver. Even more complex aliasing occurs when a channel emanating from a network is aliased to the union of the channels emanating from its components. The CSP language [10] supports such a hierarchical channel-naming scheme.

In the Generalized Hoare Logic (GHL) [12,14], a logic for concurrent programs, one must reason about state components that describe the control state. In the original presentation of GHL, the control state was modelled by *at*, *in*, and *after* predicates, where *at*(S) is true when control is at the entry point of statement S , *after*(S) is true when control is at the exit point of statement S , and *in*(S) is true when *at*(S) is true or control is at a component of S . Axioms were given to describe the relations among these predicates. Thus, if S is the statement $S_1; S_2$, the axioms of GHL state:

$$\begin{aligned} at(S) &\equiv at(S_1) \\ after(S) &\equiv after(S_2) \\ in(S) &\equiv at(S) \vee in(S_1) \vee in(S_2) \\ after(S_1) &\equiv at(S_2) \end{aligned}$$

GHL included *ad hoc* rules for reasoning about these control predicates. However, by viewing the control predicates as implicit variables, and considering the above relations not as equality of predicates but as aliasing relations among variables, we can reason about the control state with exactly the same rules used to reason about the values of ordinary program variables. This is described in detail in [15].

2 Primitives for Constrained Execution

A `var` statement, like the one for the Cartesian/polar coordinate example, specifies three things:

- The names of new variables— x, y, r , and θ in the example.
- Constraints the new variables must satisfy, including those given explicitly by the `constraints` clause and those implicit in the type declarations. In the example, the constraints are:

$$\begin{aligned} x \in \mathcal{R} & & x &= r * \cos(\theta) \\ y \in \mathcal{R} & & y &= r * \sin(\theta) \\ r \in \mathcal{R} & & x &\perp y \\ \theta \in \mathcal{R} \wedge 0 \leq \theta < 2\pi & & r &\perp \theta \end{aligned}$$

- Other independence constraints involving the new variables. In the example, there is the implicit assumption

that x , y , r , and θ are not aliased to any other variables, except perhaps variables declared in the body of the **var** statement. Thus, there are implicit constraints $x \perp q$ for all variables q declared outside the **var** statement, and similarly for y , r , and θ .

Instead of reasoning directly about the **var** statement, three primitive statements are introduced, each of which performs one of the above functions. These statements can be used to model the **var** statement and the aliasing of implicit variables described above.

The **new** statement is used to define new variable names, where

$$\text{new } x_1, x_2, \dots, x_n \text{ in } S$$

defines x_1, x_2, \dots, x_n to be new variable names for use within S . These variable names, plus any defined in a new statement containing this one, can be referenced from within S . The usual scoping rules apply, so that a variable x_i defined by this **new** statement is different from any other variables with the same name defined by a different **new** statement.

The **declare** statement is used to specify constraints. The statement

$$\text{declare } C \text{ in } S$$

where C is a predicate, indicates that C is to be maintained during execution of S and that abortion is to occur if this becomes impossible. If S contains a nondeterministic step, such as a nondeterministic assignment statement, then the choice must be made (if possible) so that the truth of C is maintained.

Finally, the **may alias** statement is used to specify independence relations implicit in a **var** statement. The statement

$$x \text{ may alias } x_1, x_2, \dots, x_n \text{ in } S$$

specifies that, during execution of S , the value of x is independent of all variables, other than x_1, x_2, \dots, x_n , declared in the context of this statement. Thus, this statement specifies constraints $x \perp q$ for all variables q not in the list x_1, x_2, \dots, x_n . Although **may alias** specifies constraints, it cannot be modeled with a **declare** statement because that would require explicitly writing relations $x \perp p$ for every variable name p different from the x_i , and there could be an infinite number of such names.

The **new**, **declare**, and **may alias** statements can be used to model a **var** statement. The **var** x, y, r, θ example could be represented as:

$$\begin{aligned} &\text{new } x, y, r, \theta \text{ in} \\ &\quad x \text{ may alias } x, y, r, \theta \text{ in} \\ &\quad y \text{ may alias } x, y, r, \theta \text{ in} \\ &\quad r \text{ may alias } x, y, r, \theta \text{ in} \\ &\quad \theta \text{ may alias } x, y, r, \theta \text{ in} \\ &\quad \text{declare } x \in \mathcal{R} \text{ and } y \in \mathcal{R} \text{ and} \\ &\quad \quad r \in \mathcal{R} \text{ and } \theta \in [0, 2\pi) \text{ and} \end{aligned}$$

$$\begin{aligned} &x = r * \cos(\theta) \text{ and } y = r * \sin(\theta) \\ &\text{and } x \perp y \text{ and } r \perp \theta \\ &\text{in } S \end{aligned}$$

In general, let $x_1, \dots, x_n, y_1, \dots, y_m$ be $n + m$ distinct variable names, and let C be a predicate constructed from the m variables y_i ; plus zero or more of the x_i . The statement

$$\text{var } x_1 : T_1, \dots, x_n : T_n \text{ constraints } C \text{ in } S$$

is modeled by

$$\begin{aligned} &\text{new } x_1, x_2, \dots, x_n \text{ in} \\ &\quad x_1 \text{ may alias } x_1, \dots, x_n, y_1, \dots, y_m \text{ in} \\ &\quad \dots x_n \text{ may alias } x_1, \dots, x_n, y_1, \dots, y_m \text{ in} \\ &\quad \text{declare } C \text{ and } x_1 \in T_1 \text{ and} \\ &\quad \dots \text{ and } x_n \in T_n \\ &\quad \text{in } S \end{aligned}$$

3 Reasoning About Constraints

Our goal is to prove partial correctness formulas of the form $\{P\} S \{Q\}$, as first proposed by Hoare [8]. To reason about constrained execution, we interpret such a formula as a temporal assertion about the executions of S —namely, $\{P\} S \{Q\}$ is equivalent to the assertion that any terminating execution of S beginning with a state in which P is true ends in a state in which Q is true. Thus, we are viewing $\{P\} S \{Q\}$ as a temporal formula, which is not how it is usually viewed in Hoare's logic.

One reason about such temporal assertions with temporal logic.² The only knowledge of temporal logic needed to understand this paper is that the temporal formula $\Box A$ is true of a statement S if and only if A remains true throughout every possible execution of S . The only formal rules for reasoning about temporal logic formulas that we need are the following, which are immediate consequences of the definition of \Box .

Strong Necessitation Rule:

$$\frac{P \Rightarrow Q}{\Box P \Rightarrow \Box Q}$$

Multiplicative Axiom:

$$\Box(A \wedge B) = \Box A \wedge \Box B$$

To apply temporal logic to program executions, we need to know what actions are atomic. For example, the formula $\Box A$ asserts that A is true before and after each atomic action. In general, an atomic action represents the execution of a primitive statement that can change the value of a variable. Execution of an assignment statement is the only kind of atomic action needed to describe the class of languages considered in this paper.

²See [16] for an elementary discussion of temporal logic and the appendix of [13] for the more advanced temporal logic needed to formalize our method.

A program execution is a sequence of atomic actions. The possible executions of the statement

declare C in S

are all those executions of S for which C is true throughout the execution—that is, those executions for which $\Box C$ is true. This leads immediately to an inference rule for **declare**:

declare Rule:

$$\frac{\Box C \Rightarrow \{P\} S \{Q\}}{\{P\} \text{ declare } C \text{ in } S \{Q\}}$$

In this rule, the hypothesis states that the predicate $\{P\} S \{Q\}$ is true or that C does not hold throughout all executions of S .

Note that in our temporal logic interpretation of partial correctness formulas, the pre- and postconditions are assertions about the program state, so they must be Boolean-valued functions on the state. In particular, they cannot contain temporal operators like \Box .

All the usual inference rules for partial correctness formulas (see [8]) still hold under this new interpretation. For example,

Rule of Consequence:

$$\frac{P' \vdash P, \{P\} S \{Q\}, Q \vdash Q'}{\{P'\} S \{Q'\}}$$

This rule allows the precondition of a partial correctness formula to be weakened and the postcondition to be strengthened.

The only new general rule needed for reasoning about constraints is the

Constraint Strengthening Rule:

$$\frac{\{P \wedge C\} S \{Q \vee \neg C\}}{\Box C \Rightarrow \{P\} S \{Q\}}$$

To show the validity of this rule, observe that the hypothesis asserts that every terminating execution of S that begins with $P \wedge C$ true terminates with $Q \vee \neg C$ true. Another way of saying this is:

For any terminating execution of S : C true of the initial state implies that if P is true of the initial state then Q is true of the final state or $\neg C$ is true of the final state.

The conclusion asserts the following:

For any execution of S : C true throughout the execution implies that if P is true of the initial state and the execution terminates, then Q is true of the final state.

It should now be clear why the hypothesis implies the conclusion.

The expression $x \perp y$, introduced above for stating independence, can be given a precise meaning as a temporal formula. The formal definition is given in the Appendix. Intuitively, $x \perp y$ is the temporal formula asserting that if the next atomic action of the program is an assignment to x , then the assignment does not change the value of y , and if the next atomic action is an assignment to y , then it does not change the value of x .

The most general kind of temporal formula we write is of the form

$$\Box C \Rightarrow \{P\} S \{Q\}.$$

Because temporal formulas cannot appear in pre- and postconditions, \perp relations can appear only in C , which means they appear only in the form $\Box(x \perp y)$. The temporal formula $\Box(x \perp y)$ asserts that no assignment to x during execution causes the value of y to change, and vice-versa.

Rules for reasoning about \perp could be deduced from its formal definition. Instead, we state as axioms two properties that seem to be sufficient for reasoning about programs. An obvious axiom is:

Commutativity Axiom:

$$x \perp y = y \perp x$$

Another obvious axiom states that if x is always equal to y , then $\Box(y \perp z)$ implies $\Box(x \perp z)$. While this rule is sufficient for our examples, the following generalization is sometimes needed.

Substitution Axiom: For any (single-valued) function f of n arguments and for any variables x, y_1, \dots, y_n :

$$\Box(x = f(y_1, \dots, y_n) \wedge y_1 \perp z \wedge \dots \wedge y_n \perp z) \Rightarrow \Box(x \perp z)$$

4 Axioms For a Toy Language

In the preceding section, certain general rules for reasoning about temporal logic formulas were given. We now give language-specific rules and axioms for a toy language. The language contains the usual **skip**, **assignment** ($:=$), **concatenation** ($;$), and **while** constructs, in addition to the aforementioned **new**, **declare**, and **may alias** statements. As usual, there is one rule or axiom for each language construct. The rule for **declare** was already given.

4.1 Assignment

Consider the standard way of reasoning about assignment statements in Hoare's logic. Assuming there is no aliasing, one deduces

$$\{true\} x := y + 1 \{x = y + 1\}.$$

This formula is not valid when aliasing is allowed, because the assignment might appear in the body of a `var` statement that aliases x to equal y , in which case the postcondition $x = y + 1$ would be false. However, this postcondition would be satisfied if x and y were not aliased, which means that it would be satisfied if we constrained execution of $x := y + 1$ by requiring that $x \perp y$ hold. Hence, the rule to use when aliasing is possible is

$$\square x \perp y \Rightarrow \{true\} x := y + 1 \{x = y + 1\}$$

(Remember that $x \perp y$ cannot appear in the precondition because temporal formulas may not appear in pre- or post-conditions.)

More generally, the Assignment Axiom in Hoare's logic is

$$\{Q(exp, y_1, \dots, y_n)\} x := exp \{Q(x, y_1, \dots, y_n)\}$$

where $Q(x, y_1, \dots, y_n)$ is any predicate involving only the program variables x, y_1, \dots, y_n , and exp is an expression. Again, this axiom is valid only if x , the target of the assignment, is not aliased to any of the y_i 's. Therefore, when aliasing is allowed, the correct formula is

Assignment Axiom:

$$\square(x \perp y_1 \wedge \dots \wedge x \perp y_n) \Rightarrow \{Q(exp, y_1, \dots, y_n)\} x := exp \{Q(x, y_1, \dots, y_n)\}$$

4.2 May alias

Recall that the statement

$$x \text{ may alias } x_1, x_2, \dots, x_n \text{ in } S$$

is really equivalent to

declare A **in** S

where A is a conjunction of terms of the form $x \perp y$, for an infinite number of variables y —namely, every y that is not among the x_i . Therefore, the **declare** rule gives the following:

$$\frac{\square A \Rightarrow \{P\} S \{Q\}}{\{P\} x \text{ may alias } x_1, \dots, x_n \text{ in } S \{Q\}} \quad (1)$$

Now, let y_1, y_2, \dots, y_m be a finite number of variables, all different from any of the x_i . Then,

$$A \Rightarrow (x \perp y_1 \wedge \dots \wedge x \perp y_m). \quad (2)$$

Given

$$\square(x \perp y_1 \wedge \dots \wedge x \perp y_m) \Rightarrow \{P\} S \{Q\}$$

we use (2) to deduce

$$\square A \Rightarrow \{P\} S \{Q\}$$

which is the hypothesis of (1). Therefore, we get the following inference rule.

may alias Rule:

$$\frac{\square(x \perp y_1 \wedge \dots \wedge x \perp y_m) \Rightarrow \{P\} S \{Q\}, \quad \forall i, j : y_i \stackrel{\text{syn}}{\neq} x_j}{\{P\} x \text{ may alias } x_1, \dots, x_n \text{ in } S \{Q\}}$$

where $a \stackrel{\text{syn}}{\neq} b$ means that a and b are syntactically different variable names.

4.3 New

The rule for the **new** statement is essentially the same as the one for ordinary Hoare triples—see Rule 16 of [1]. The statement

new x_1, \dots, x_n **in** S

declares that the variables x_i are different from all variables declared outside the statement. It is equivalent to substituting for all free (undeclared) occurrences of x_i in S another variable y_i that is not used anywhere in the entire program. Of course, when reasoning about S in isolation, we do not know what the entire program is. However, since, we are concerned only with a particular pre- and postcondition, it suffices to choose the y_i so that they don't appear in that pre- or postcondition or in S . This leads to the following rule, where $S[y_1/x_1, \dots, y_n/x_n]$ is the statement obtained by substituting y_i for every free occurrence of x_i in S , for $i = 1, \dots, n$.

new Rule: For any distinct variable names y_1, \dots, y_n not occurring free in P, S , or Q :

$$\frac{\{P\} S[y_1/x_1, \dots, y_n/x_n] \{Q\}}{\{P\} \text{new } x_1, \dots, x_n \text{ in } S \{Q\}}$$

Note that, unlike [1], initial values for the variables x_i in S are not assumed; executions containing arbitrary initial values are permitted.

4.4 Remaining Statements

The axioms for the remaining statements are just the ordinary Hoare logic partial correctness rules. For example, Hoare's rule for statement concatenation is:

Statement Concatenation:

$$\frac{\{P\} S \{Q'\}, \{Q'\} S' \{Q\}}{\{P\} S; S' \{Q\}}$$

5 Examples

5.1 No Aliasing

We first consider an example in which there is no aliasing—that is, there are no constraints. Let S be the statement

```
var x : real in y := y + 1;
      x := y + 3
```

We will prove the obvious relation

$$\{y = 1\} S \{y = 2\} \quad (3)$$

From the Assignment Axiom, we get

$$\begin{aligned} \Box \text{true} &\Rightarrow \{y = 1\} y := y + 1 \{y = 2\} \\ \Box x' \perp y &\Rightarrow \{y = 2\} x' := y + 3 \{y = 2\} \end{aligned}$$

Using ordinary propositional logic and the Strong Necessitation Rule, the antecedents of these implications can be strengthened to get

$$\begin{aligned} \Box(x' \perp y \wedge x' \in \mathcal{R}) &\Rightarrow \{y = 1\} y := y + 1 \{y = 2\} \quad (4) \\ \Box(x' \perp y \wedge x' \in \mathcal{R}) &\Rightarrow \{y = 2\} x' := y + 3 \{y = 2\} \quad (5) \end{aligned}$$

Combining (4) and (5) gives

$$\begin{aligned} \Box(x' \perp y \wedge x' \in \mathcal{R}) &\Rightarrow \\ \{y = 1\} y := y + 1 \{y = 2\} \wedge &\{y = 2\} x' := y + 3 \{y = 2\} \end{aligned}$$

Application of the Statement Concatenation Rule to the consequent of this yields

$$\begin{aligned} \Box(x' \perp y \wedge x' \in \mathcal{R}) &\Rightarrow \\ \{y = 1\} y := y + 1; x' := y + 3 &\{y = 2\} \end{aligned}$$

The declare Rule now allows us to conclude

$$\begin{aligned} \Box(x' \perp y) &\Rightarrow \{y = 1\} \\ &\text{declare } x' \in \mathcal{R} \text{ in} \\ &y := y + 1; x' := y + 3 \{y = 2\} \end{aligned}$$

Applying the may alias rule yields

$$\begin{aligned} \{y = 1\} x' \text{ may alias } x', y \text{ in} \\ \text{declare } x' \in \mathcal{R} \text{ in} \\ y := y + 1; x' := y + 3 \{y = 2\} \end{aligned}$$

Finally, the new Rule allows us to deduce

$$\begin{aligned} \{y = 1\} \text{ new } \quad \text{in} \\ x \text{ may alias } x, y \text{ in} \\ \text{declare } x \in \mathcal{R} \\ \text{in } y := y + 1; x := y + 3 \{y = 2\} \end{aligned}$$

The statement in this formula is equivalent to our original statement S , according to the method of modeling var statements described in Section 2, so the desired result is proved.

5.2 Simple Aliasing

Next, let S be the same as above except with x and y aliased:

$$\begin{aligned} \text{var } x : \text{real constraint } x = y \text{ in } y := y + 1; \\ x := y + 3 \end{aligned}$$

Formula (3) is no longer valid for this program. Instead, we have

$$\{y = 1\} S \{y = 5\} \quad (6)$$

The proof of this is as follows. From the Assignment Axiom, we have

$$\begin{aligned} \Box \text{true} &\Rightarrow \{y = 1\} y := y + 1 \{y = 2\} \\ \Box \text{true} &\Rightarrow \{y = 2\} x' := y + 3 \{x' = 5\} \end{aligned}$$

Combining these, using the Statement Concatenation Rule, yields

$$\Box \text{true} \Rightarrow \{y = 1\} S' \{x' = 5\}$$

where S' is the statement

$$y := y + 1; x' := y + 3$$

Applying the Rule of Consequence to this, using the tautologies

$$\begin{aligned} (x' = 5) &\Rightarrow (y = 5 \vee x' \neq y) \\ (y = 1 \wedge x' = y) &\Rightarrow (y = 1) \end{aligned}$$

we conclude

$$\Box \text{true} \Rightarrow \{y = 1 \wedge x' = y\} S' \{y = 5 \vee x' \neq y\}$$

The Constraint Strengthening Rule now allows us to deduce

$$\Box x' = y \Rightarrow \{y = 1\} S' \{y = 5\}$$

Using propositional logic and the Strong Necessitation Rule, we can strengthen the antecedent of the implication to obtain

$$\Box(x' = y \wedge x' \in \mathcal{R}) \Rightarrow \{y = 1\} S' \{y = 5\}$$

The declare Rule now yields

$$\{y = 1\} \text{ declare } x' = y \text{ and } x' \in \mathcal{R} \text{ in } S' \{y = 5\}$$

from which the may alias Rule allows us to deduce

$$\begin{aligned} \{y = 1\} x' \text{ may alias } x', y \text{ in} \\ \text{declare } x' = y \text{ and } x' \in \mathcal{R} \text{ in} \\ S' \{y = 5\} \end{aligned}$$

Finally, the new Rule allows the conclusion

$$\begin{aligned} \{y = 1\} \text{ new } x \text{ in} \\ x \text{ may alias } x, y \text{ in} \\ \text{declare } x = y \text{ and } x \in \mathcal{R} \text{ in} \\ y := y + 1 \\ x := y + 3 \{y = 5\} \end{aligned}$$

This is just what is required to prove (6).

5.3 Cartesian/Polar Coordinates

As a final example, let S be the following statement.

$$\begin{aligned} \text{var } x, y : \text{real} \\ \text{constraints } x = r * \cos(\theta) \text{ and } y = r * \sin(\theta) \\ \text{and } x \perp y \text{ in } x := 2 * x; \\ y := 2 * y \end{aligned}$$

A little trigonometry shows that if r is initially positive, then executing S should double the value of r and leave θ unchanged. Thus, the following should hold:

$$\{r = r_0 \wedge \theta = \theta_0\} S \{r = 2r_0 \wedge \theta = \theta_0\}$$

However, further reflection indicates that this is not quite valid because executing S can add any even multiple of π to θ or can negate r and add any odd multiple of π to θ . Thus, we stipulate that $r \geq 0$ and $0 \leq \theta < 2\pi$ remain true throughout execution, and prove

$$\begin{aligned} \square (r \geq 0 \wedge 0 \leq \theta < 2\pi) \Rightarrow \\ \{r = r_0 \wedge \theta = \theta_0\} S \{r = 2r_0 \wedge \theta = \theta_0\} \quad (7) \end{aligned}$$

Let S' be the statement

$$x' := 2 * x'; y' := 2 * y'.$$

From the Statement Concatenation Rule and two applications of the Assignment Axiom, we deduce

$$\begin{aligned} \square x' \perp y' \Rightarrow \{x' = r_0 \cos \theta_0 \wedge y' = r_0 \sin \theta_0\} S' \\ \{x' = 2r_0 \cos \theta_0 \wedge y' = 2r_0 \sin \theta_0\} \end{aligned}$$

Now, note that the following are tautologies:

$$\begin{aligned} (x' = r_0 \cos \theta_0 \wedge r \cos \theta = r_0 \cos \theta_0 \wedge x' = r \cos \theta) \\ \Rightarrow (x' = r_0 \cos \theta_0) \end{aligned}$$

$$(x' = 2r_0 \cos \theta_0) \Rightarrow (r \cos \theta = 2r_0 \cos \theta_0 \vee x' \neq r \cos \theta)$$

Similar tautologies apply to y' . Therefore, by the Rule of Consequence we conclude

$$\begin{aligned} \square x' \perp y' \Rightarrow \\ \left\{ \begin{array}{l} x' = r_0 \cos \theta_0 \wedge r \cos \theta = r_0 \cos \theta_0 \wedge x' = r \cos \theta \wedge \\ y' = r_0 \sin \theta_0 \wedge r \sin \theta = r_0 \sin \theta_0 \wedge y' = r \sin \theta \end{array} \right\} \\ S' \\ \left\{ \begin{array}{l} (r \cos \theta = 2r_0 \cos \theta_0 \vee x' \neq r \cos \theta) \wedge \\ (r \sin \theta = 2r_0 \sin \theta_0 \vee y' \neq r \sin \theta) \end{array} \right\} \end{aligned}$$

From this, the Constraint Strengthening Rule and Rule of Consequence allow us to derive

$$\begin{aligned} \square (x' \perp y' \wedge x' = r \cos \theta \wedge y' = r \sin \theta) \Rightarrow \\ \left\{ \begin{array}{l} r \cos \theta = r_0 \cos \theta_0 \wedge \\ r \sin \theta = r_0 \sin \theta_0 \end{array} \right\} S' \left\{ \begin{array}{l} r \cos \theta = 2r_0 \cos \theta_0 \wedge \\ r \sin \theta = 2r_0 \sin \theta_0 \end{array} \right\} \end{aligned}$$

By the Rule of Consequence, using some trigonometry and the previous theorem we can now deduce

$$\begin{aligned} \square (x' \perp y' \wedge x' = r \cos \theta \wedge y' = r \sin \theta) \Rightarrow \\ \left\{ \begin{array}{l} r = r_0 \wedge \theta = \theta_0 \wedge \\ r \geq 0 \wedge 0 \leq \theta < 2\pi \end{array} \right\} S' \left\{ \begin{array}{l} (r = 2r_0 \wedge \theta = \theta_0) \vee \\ \neg(r \geq 0 \wedge 0 \leq \theta < 2\pi) \end{array} \right\} \end{aligned}$$

We now use the Constraint Strengthening Rule to derive

$$\begin{aligned} \square (x' \perp y' \wedge x' = r \cos \theta \wedge y' = r \sin \theta \wedge \\ r \geq 0 \wedge 0 \leq \theta < 2\pi) \Rightarrow \\ \{r = r_0 \wedge \theta = \theta_0\} S' \{r = 2r_0 \wedge \theta = \theta_0\} \end{aligned}$$

It is now a simple matter to use the **declare** Rule, the **may alias** Rule, and finally the **new** Rule to obtain (7).

6 Discussion

We have introduced the idea of describing types and aliasing in terms of constraints and given general rules for reasoning about constrained execution. Our approach involves embedding the usual Hoare partial correctness formalism in temporal logic. One reasons about static properties with constraints and about dynamic properties with pre- and postconditions.

Having applied our method to a simple language, we now consider some of the problems in extending it to more complex languages. We also discuss the relation of our approach to previous work.

6.1 Types

In our toy language, we were able to handle a type declaration simply by translating it to a constraint about the values that the variable can assume. This does not work for languages that make more extensive use of type information—for example, by performing coercions in the event of a type mismatch; nor does it work for languages in which a type mismatch in an assignment generates an indeterminate result rather than abortion. (Our semantics causes abortion if executing an assignment would violate a type constraint.)

Reasoning about these more complex languages requires adding state predicates that characterize the type of a variable and modifying our Axiom of Assignment. However, care must be employed when reasoning about predicates like $\text{type}(x) = \text{'integer'}$ because $x = y$, which means that the values of x and y are equal, does not imply $\text{type}(x) = \text{type}(y)$.

If a type mismatch can abort execution, reasoning about type correctness requires proving total correctness properties. While we have not yet considered this problem, we feel that our approach should be ideal for proving termination properties because it is based upon temporal logic, and temporal logic is effective for proving liveness properties like termination [16].

6.2 Generalized Assignments: Expressions as Targets

In the Cartesian/polar coordinate example, x is aliased to $r \cos \theta$. Thus, assigning to x is the same as assigning to $r \cos \theta$. The obvious next step is to try writing $r \cos \theta$ on the left-hand side of an assignment statement, even if there is no variable aliased to this expression. For arbitrary expressions exp_1 and exp_2 , the statement

$$exp_1 := exp_2$$

causes the value of exp_1 after execution to be the same as the value of exp_2 before execution. To reason about this form of generalized assignment, we extend \perp to be a relation on expressions rather than just on variable names. The temporal formula $exp_1 \perp exp_2$ now means that assigning to

exp_1 does not change the value of exp_2 , and vice-versa. The axiom for generalized assignment is the same as the Assignment Axiom given above, except that x and the y_i can be arbitrary expressions. The commutativity and substitution axioms given above are also valid for these more general \perp relations. However, additional axioms are needed for deriving \perp relations between expressions from \perp relations between their components—axioms such as

$$(exp_1 \perp exp_2) \wedge (exp_1 \perp exp_3) \Rightarrow exp_1 \perp (exp_2 + exp_3)$$

We do not give a formal semantics for this here.

6.3 Arrays and Pointers

Our approach can handle arrays by regarding assignment to an element of an array as an assignment to the entire array, as described in [11]. Array assignment cannot be handled using our generalized assignment statement, where an expression like $A[i]$ appears on the left-hand side, because this does not give the usual semantics for

$$A[exp] := exp'$$

Letting the subscripts *old* and *new* denote values before and after the assignment, the semantics of generalized assignment defines the above statement to mean that $A_{new}[exp_{new}] = exp'_{old}$, while the usual meaning of array assignment is $A_{new}[exp_{old}] = exp'_{old}$. This difference helps explain why the ordinary assignment axiom is not easily extended to arrays. This also indicates why our more general assignment statement is not easily compiled, since it requires computing the value of an expression in a (new) state without knowing what that state is.

By regarding an array as a single variable, our formalism can handle aliasing relations between entire arrays. However, our current formalism does not handle simple aliasing of array elements. For example, if x is not aliased to any element of the array A , then we can easily prove that assigning a value to $A[1]$ does not change x 's value by deducing

$$\Box x \perp A \Rightarrow \{x = 7\} A[1] := 1 \{x = 7\}$$

However, we cannot do this just knowing that x is not aliased to $A[1]$, because our rules are not strong enough to prove

$$\Box x \perp A[1] \Rightarrow \{x = 7\} A[1] := 1 \{x = 7\}$$

The required generalization must replace \perp with a non-commutative relation, since the assertion that assigning to $A[i]$ does not change i is not equivalent to the assertion that assigning to i does not change $A[i]$. Moreover, aliasing relations are no longer static, since whether $A[i]$ and x are aliased may depend upon the value of i . Being dynamic, aliasing relations like \perp have to appear in pre- and postconditions, which is prohibited by the current formalism.

Similar extensions are needed for reasoning about aliasing in programs that use pointers. Moreover, in a language

with pointers, type relations might also be dynamic—for example, if a pointer can point to variables of type *real* and of type $[0, 2 * \pi)$. In this case, type relations would have to appear in pre- and postconditions.

These extensions will be described in a future paper.

6.4 Procedures

The most general form of parameter passing is call by name, since it can be used to simulate call by reference and call by value-result. With call by name, a formal parameter is essentially aliased to the corresponding argument. Thus, our approach can be used for reasoning about procedures.

Traditionally, programming languages with procedures do not allow an arbitrary expression as the argument corresponding to a formal parameter that appears on the left-hand side of an assignment, since assignment to an expression is not defined by these languages. We have defined what it means to assign a value to a variable that is aliased to an expression, so there is no semantic reason for this prohibition. However, some restriction is needed to ensure that the language can be compiled.

6.5 Related Work

Previous work on aliasing, [1,3,4,5,6,7,9], has been motivated by shared storage among arguments of a procedure call. We are aware of no work that can handle the rich aliasing structures that concern us. However, programming languages where computations are partially or completely specified in terms of constraints have been investigated [2,19,20,21,22].

In most previous work on aliasing, the program state consists of a mapping from variable names to a space of (abstract) locations plus a map from locations to values. We feel that if the language itself has no pointers, then the semantics should not be given in terms of pointers, even if the values of these pointers are abstract locations instead of real memory addresses. Moreover, the existence of semantic pointers unnecessarily complicates reasoning about programs. Also, approaches based on locations are rarely fully abstract [3]. Finally, and most importantly, the use of locations does not support reasoning about the more general form of aliasing that is not based on shared storage.

Our work resembles Reynolds' [17] handling of call by name in many ways. Reynolds defines a formal system, called Specification Logic, that contains a relation $\#$ very similar to \perp and an assignment rule much like ours. The effects of aliasing are described in terms of *interference*, which can be seen as the dual of our viewing aliasing in terms of *invariants*. (See [14] for a discussion of the relation between interference and invariance.)

The meaning of a formula in Specification Logic is based on an *environment* in addition to a state. The environment is a mapping from variable names to a space of locations and brings with it the difficulties mentioned above. In addition, in Specification Logic, assertions about the environment are made using a completely new logic, distinct from

be one used to reason about the state. In our approach, all reasoning is done in a single logical system—temporal logic. Assertions about the state are expressed by temporally trivial formulas—formulas containing no temporal operators. We do not need the concept of an environment, using instead temporal assertions about the state. Of course, Specification Logic handles forms of aliasing not considered in this paper. We are currently extending our temporal logic approach to cover the full range of language features handled by Specification Logic.

Our approach can be viewed as a generalization of one proposed by Brooks [3], and we were somewhat influenced by his work. We handle a much more general form of aliasing, but, if we restricted ourselves to aliasing relations that are simple equalities between variables, then proofs in our system and in Brooks' would be quite similar. To extend Brooks' work to handle our more general form of aliasing, it appears that a new deductive method would have to be added; we avoided this by embedding our proof system in temporal logic.

Acknowledgments

We would like to thank A. Demers, C. Cartwright, W. P. DeLoever, D. Gries, J. Hook, A. Meyer, and the members of FIP Working Groups 2.2 and 2.3 for helpful discussions.

Bibliography

- [1] K. R. Apt. Ten Years of Hoare's Logic: A Survey. *Trans. on Programming Languages and Systems* 3, 4 (October 1981), 431-483.
- [2] Alan Borning. The Programming Language Aspects of ThingLab, a Constraint-Oriented Simulation Laboratory. *Trans. on Programming Languages and Systems* 3, 4 (Oct. 1981), 353-387.
- [3] Stephen D. Brooks. A Fully Abstract Semantics And A Proof System For An ALGOL-like Language With Sharing. Preliminary Draft, Carnegie-Mellon University, Feb. 1984.
- [4] R. Cartwright and D. Oppen. The Logic of Aliasing. *Acta Informatica* 15, (1981) 365-384.
- [5] J. W. de Bakker. *Mathematical Theory of Program Correctness*, Prentice-Hall, New Jersey, 1980.
- [6] J. Halpern, A. Meyer, and B. Trakhtenbrot. From Denotational to Operational and Axiomatic Semantics for ALGOL-like Languages: An Overview. *Proc. 1983 Workshop on Logics of Programs*, Lecture Notes in Computer Science, Volume 164, Springer-Verlag, 1984.
- [7] D. Gries and G. Levin. Assignment and Procedure Call Proof Rules. *Trans. on Programming Languages and Systems* 2, 4 (Oct. 1980), 564-579.
- [8] C. A. R. Hoare. An Axiomatic Basis for Computer Programming. *Communications of the ACM* 12, 10 (October 1969), 576-580.
- [9] C. A. R. Hoare. Procedures and Parameters: An Axiomatic Approach. *Symposium on Semantics of Algorithmic Languages*, Lecture Notes in Computer Science, Volume 88, Springer-Verlag, New York, 1971.
- [10] C. A. R. Hoare. Communicating Sequential Processes. *Communications of the ACM* 21, 8 (August 1978), 666-677.
- [11] C. A. R. Hoare and N. Wirth. An Axiomatic definition of the programming language Pascal. *Acta Informatica* 2 (1973), 335-355.
- [12] L. Lamport. The "Hoare Logic" of Concurrent Programs. *Acta Informatica* 14 (1980), 21-37.
- [13] L. Lamport. Specifying Concurrent Program Modules. *Trans. on Programming Languages and Systems* 5, 2 (April 1983), 190-222.
- [14] L. Lamport and F. B. Schneider. The "Hoare Logic" of CSP and All That. *Trans. on Programming Languages and Systems* 6, 2 (April 1984), 281-296.
- [15] L. Lamport and F. B. Schneider. Fully Compositional Generalized Hoare Logic. In preparation.
- [16] S. S. Owicki and L. Lamport. Proving Liveness Properties of Temporal Logic. *Trans. on Programming Languages and Systems* 4, 3 (July 1982), 455-495.
- [17] J. Reynolds. *The Craft of Programming*. Prentice Hall International, London, 1981.
- [18] R. D. Schlichting and F. B. Schneider. Using Message Passing for Distributed Programming: Proof Rules and Disciplines. *Trans. on Programming Languages and Systems* 6, 3 (July 1984), 402-431.
- [19] R. M. Stallman and G. J. Sussman. Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis. *Artificial Intelligence* 9, 1977, 135-196.
- [20] G. L. Steele Jr., and G. J. Sussman. Constraints. *Proceedings APL'79, ACM SIGPLAN STAPL APL Quote Quad* 9, 4 (June 1979), 208-225.
- [21] G. J. Sussman, and R. M. Stallman. Heuristic Techniques in Computer-Aided Circuit Analysis. *IEEE Transactions on Circuits and Systems* CAS-22, 11 (November 1975).
- [22] I. E. Sutherland. SKETCHPAD: A Man-Machine Graphical Communication System. M.I.T. Lincoln Laboratory Technical Report 296, (January 1963).

Appendix

The Formal Semantics

A.1 The Language

We now give a formal semantics for our toy language containing **skip**, **assignment**, **concatenation**, and **while**, plus the three statements **new**, **declare**, and **may alias** introduced to model the **var** statement. The class of expressions and variable types is not specified. We assume only that expressions are built from some set *Var* of variable names, that variables assume values in some set *Val*, and that expressions are built from operators on those values. However, in the statement

declare *C* in *S*

S can involve \perp in addition to Boolean expressions.

Finally, we require the value of an expression to be defined for any values of its component variables. Thus, the expression $x + 10$ must be assigned a value, even when $x = \text{true}$. This can be done by including a special value *undefined* in *Val*; the precise details for doing this are irrelevant.

A.2 Temporal Logic

A *state* is defined to be a mapping from *Var* to *Val*, so a state *s* assigns a value $s(x)$ to every variable name $x \in \text{Var}$. Let *S* denote the set of states. A state *s* is extended to a mapping from ordinary (nontemporal) expressions to values in the obvious way—for example, $s(x + y)$ is defined to equal $s(x) + s(y)$. Let $s \models \text{exp}$ denote the assertion that $(\text{exp}) = \text{true}$. ($s \models x > 10$ is false if $s(x)$ has a nonnumeric value.)

An *action* is defined to be an element of $\text{Var} \cup \{\tau\}$, where τ is a symbol not in *Var*. For $x \in \text{Var}$, action *x* represents an assignment to *x*. Action τ represents an assignment to a variable declared in some **new** statement inside the current statement; thus τ models a variable that is “invisible” in the current context.

A *behavior* is defined to be a sequence σ of the form

$$s_0 \xrightarrow{x_1} \dots \xrightarrow{x_n} s_n \quad (8)$$

where the s_i are states and the x_i are actions. This behavior denotes an execution starting in state s_0 and terminating in state s_n , where the i^{th} action changes the state from s_{i-1} to s_i . Since partial correctness does not distinguish between aborting and infinite looping, we consider only finite (terminating) behaviors, although our definitions are easily extended to include infinite (nonterminating) ones. We allow the case $n = 0$, where s_0 is the behavior starting in state s_0 that performs no actions.

In the temporal logic of [13], a formula is composed of state predicates, action predicates, and temporal operators. The set of Boolean expressions is taken to be our state predicates. Action predicates are those of the form $\alpha(x)$,

together with the action predicate *halt*. A Boolean expression is true of a behavior if it is true of the first state in the behavior. An action predicate $\alpha(x)$ is true if *x* is the first action in the behavior; action predicate *halt* is true only if there is no next action. The semantics of this temporal logic assigns a truth value to $\sigma \models F$ for every behavior σ and every formula *F*. We write $\models F$ to denote that $\sigma \models F$ is true for all behaviors σ .

We will define a behavioral semantics for our language where $\mathcal{M}[S]$ is a set of behaviors representing all possible terminating executions of *S*. Semantic validity of a temporal logic formula *F* for a program *S* is defined by

$$\models_S F \stackrel{\text{def}}{=} \forall \sigma \in \mathcal{M}[S] : \sigma \models F$$

To define the temporal operator \perp , we first define the operator \lrcorner by letting $\sigma \models x \lrcorner y$ mean that if the first action of σ is an assignment to *x*, then that assignment does not change the value of *y*. In other words, letting σ be the sequence of (8), we have

$$\sigma \models x \lrcorner y \stackrel{\text{def}}{=} (x \stackrel{\text{def}}{=} x_1) \Rightarrow (s_0(y) = s_1(y))$$

In the temporal logic of [13],

$$x \lrcorner y \stackrel{\text{def}}{=} \forall \eta : (y = \eta) \Rightarrow (\alpha(x) \triangleleft y = \eta)$$

We define $x \perp y$ to be $(x \lrcorner y) \wedge (y \lrcorner x)$. Note that $\sigma \models x \perp y$ is true if σ is a sequence with no actions.

The formulas deduced by our method for reasoning about programs are of the form $\Box C \Rightarrow \{P\} S \{Q\}$, where *C* is a temporal logic formula. However, $\{P\} S \{Q\}$ is not a temporal logic formula because it refers to the statement *S*—a concept with no counterpart in the temporal logic. To make semantic sense out of this formula, first define $\{P\} \rightarrow \{Q\}$ to be true for the behavior (8) if and only if $s_0 \models P \Rightarrow s_n \models Q$. This is defined in terms of temporal logic operators by

$$\{P\} \rightarrow \{Q\} \stackrel{\text{def}}{=} P \Rightarrow \Box(\text{halt} \Rightarrow Q)$$

A program *S* satisfies $\Box C \Rightarrow \{P\} S \{Q\}$ if and only if $\models_S \Box C \Rightarrow \{P\} \rightarrow \{Q\}$.

A.3 The Behavioral Semantics

For any statement *S* in our language, we define $\mathcal{M}[S]$ to be a set of behaviors. The definition is by induction on the structure of *S*.

skip $\mathcal{M}[\text{skip}] \stackrel{\text{def}}{=} \{s_0 : s_0 \in S\}$

The **skip** statement generates no actions.

assignment $\mathcal{M}[x := \text{exp}] \stackrel{\text{def}}{=} \{s \xrightarrow{x} t : t(x) = s(\text{exp})\}$

An assignment generates a single action that sets the value of the left-hand side to the original value of the right-hand side. Note that $\mathcal{M}[S]$ contains behaviors that make arbitrary changes to other variables, since any such change could be caused by an appropriate aliasing.

tenation $M[S_1; S_2]$ is defined to equal

$$\{s_0 \xrightarrow{a_1} \dots \xrightarrow{a_{n+m}} s_{n+m} : \\ s_0 \xrightarrow{a_1} \dots \xrightarrow{a_n} s_n \in M[S_1] \\ \text{and } s_n \xrightarrow{a_{n+1}} \dots \xrightarrow{a_{n+m}} s_{n+m} \in M[S_2]\}$$

Note that we are including only finite (terminating) behaviors.

We define $M[\text{while } B \text{ do } S]$ inductively by

$$\begin{aligned} M[\text{while}_0 B \text{ do } S] &\stackrel{\text{def}}{=} \{s_0 \in S : s_0(B) = \text{false}\} \\ M[\text{while}_{i+1} B \text{ do } S] &\stackrel{\text{def}}{=} \\ &\{s_0 \xrightarrow{a_1} \dots \xrightarrow{a_n} s_n \in M[S; \text{while}_i B \text{ do } S] : \\ & s_0(B) = \text{true}\} \\ M[\text{while } B \text{ do } S] &\stackrel{\text{def}}{=} \bigcup_{i=0}^{\infty} M[\text{while}_i B \text{ do } S] \end{aligned}$$

Intuitively, $M[\text{while}_i \dots]$ contains the behaviors in which the body of the while statement is executed exactly i times.

re $M[\text{declare } C \text{ in } S] \stackrel{\text{def}}{=} \{\sigma \in M[S] : \sigma \models \square C\}$

The declare acts as a "filter" to eliminate any behaviors of S in which C does not always hold.

If σ is the sequence (8), let $\sigma[x_1/y_1, \dots, x_n/y_n]$ denote the sequence

$$s'_0 \xrightarrow{a'_1} \dots \xrightarrow{a'_n} s'_n$$

where

$$s'_i(v) = \begin{cases} s_i(v) & \text{if } \forall j : v \neq x_j \text{ and } v \neq y_j \\ s_i(x_j) & \text{if } v \stackrel{\text{syn}}{=} x_j \\ s_i(y_j) & \text{if } v \stackrel{\text{syn}}{=} y_j \end{cases} \quad (9)$$

$$x'_i = \begin{cases} x_i & \text{if } \forall j : x'_i \neq y_j \\ \tau & \text{otherwise} \end{cases} \quad (10)$$

Then $\sigma \in M[\text{new } x_1, \dots, x_n \text{ in } S]$ if and only if there exist variable names y_1, \dots, y_n not free in S and $\sigma' \in M[S[y_1/x_1, \dots, y_n/x_n]]$ such that $\sigma = \sigma'[x_1/y_1, \dots, x_n/y_n]$.

This formal definition captures the intuitive notion that to execute the new statement, one first executes its body with new variables y_i substituted for the x_i . The resulting execution is then modified by hiding all references to the variables y_i —replacing assignments to the y_i by τ actions and letting y_i refer once more to its external declaration—and requiring that the externally declared values of the x_i remain unchanged.

alias $M[x \text{ may alias } x_1, \dots, x_n \text{ in } S]$ is defined to equal

$$\{\sigma \in M[S] : \forall y : (\forall i : y \neq x_i) \Rightarrow \sigma \models \square(x \perp y)\}$$

This is the formal definition of our intuitive idea that the may alias is equivalent to a declaration of an an infinite number of \perp relations.

A.4 Soundness and Completeness

In the main body of this paper, we gave a set of rules for deriving formulas of the form $\square C \Rightarrow \{P\} S \{Q\}$. Having defined the set of behaviors $M[S]$, we have given a semantic meaning to these formulas, namely:

$$M[\square C \Rightarrow \{P\} S \{Q\}] \stackrel{\text{def}}{=} \models_s \square C \Rightarrow \{P\} \rightarrow \{Q\}$$

We can now discuss the soundness and completeness of our system.

Soundness means that for any formula F derived by our system, $M[F]$ equals *true*. The proof of this involves checking the validity of all our axioms and proof rules. This involves a straightforward formalization of the informal arguments given in section 4.

Completeness means that every semantically correct formula is derivable using our rules. Since completeness is impossible, one usually proves relative completeness in the sense of Cook, which, as explained in [1], means that the system is complete if we assume that:

- C1. All valid state predicates are given.
- C2. The set of state predicates is sufficiently expressive, meaning that for any state predicate P and statement S , $\text{post}_S(P)$, the strongest postcondition of P with respect to S , is a state predicate.

These assumptions are not enough to guarantee completeness in our system. In ordinary Hoare logic, one assumes the ability to reason about state predicates. Since formulas in our logic contain temporal operators, like \square and \perp , we need to reason about temporal logic formulas. We therefore assume that³

- C1'. All valid temporal logic formulas constructed from the state predicates are given.

where a *valid temporal logic formula* is one that is true for any behavior. Thus, just as assumption C1 for ordinary Hoare Logic contains only information about the state space—not information about the program—so C1' gives information about state functions and temporal operators—not about the program. Since state predicates are temporal formulas, C1' subsumes C1.

In addition to strengthening C1, we must also strengthen C2. To see why, suppose our set of state predicates did not contain the predicate *true*, but required that we use the semantically equivalent predicate $x = x$. Our Assignment Axiom does not allow us to deduce $\{x = x\} y := 1 \{x = x\}$; we can only deduce it under the irrelevant hypothesis $\square(y \perp x)$. In general, we need to assume that if a state predicate does not depend upon the value of a variable x , then we can write that predicate as an expression that does not involve x . We therefore require the following additional expressiveness condition:

³Actually, this assumption is stronger than necessary, since we are concerned only with temporal logic formulas of the form $\square C \Rightarrow \{P\} S \{Q\}$.

any expression exp , mathematical function f , and variables y_1, \dots, y_n : if the relation $exp = f(y_1, \dots, y_n)$ is valid, then $f(y_1, \dots, y_n)$ is an expression.

Having made these extra assumptions, we now complete the proof of completeness. Completeness for ordinary Hoare logic states that every valid formula is provable. In our system, the valid formulas are of the form $\Box C \Rightarrow \{P\} S \{Q\}$, where C is of the form $C' \wedge (x \perp y_1) \wedge \dots \wedge (x \perp y_n)$ and C' is an ordinary (nontemporal) expression. Completeness therefore states that every valid formula of this form is provable.

The proof is by induction on the structure of S . If S is a statement, then $\Box C \Rightarrow \{P\} S \{Q\}$ is semantically equivalent to $(C' \wedge P) \Rightarrow Q$. By assumption C1', this is provable, and we can use it, the ordinary axiom for `skip` (`{P} skip {P}`), the Rule of Consequence, and the Constraint Strengthening Rule to deduce $\Box C \Rightarrow \{P\} S \{Q\}$. (Details are left to the reader.)

Let S be the assignment statement $x := exp$. A consequence of C1' is a complete system for deriving \perp relations. Thus, we can assume that C includes all relations $x_i \perp y_i$ that are derivable from it. It is easy to see that $\Box C \Rightarrow \{P\} S \{Q\}$ is semantically equivalent to, and, by the Constraint Strengthening Rule, derivable from

$$\perp y_1 \wedge \dots \wedge x_n \perp y_n \Rightarrow \{P \wedge C'\} S \{Q \vee \neg C'\}$$

We therefore restrict consideration to the case in which C consists only of the conjunction of the constraints $x_i \perp y_i$. The formula $\{P\} S \{Q\}$ is semantically equivalent to $post_S(P) \Rightarrow Q$. By the Rule of Consequence and C2, we can let $Q = P$.

Let $Q = Q(x, z_1, \dots, z_m)$, where the z_i are different from the y_j . By C2a, we can assume that the value of Q depends only on each of the z_j . Since $Q = post_S(P)$, this means that there is a behavior $s \xrightarrow{x} t$ in $\mathcal{M}[S]$ such that $s \models P$, but $t \not\models Q$ for some state t' that differs from t only in the values of the z_j . However, since the semantics of S does not change the ending values of any variable other than x , the state t' is in $\mathcal{M}[S]$. Therefore, if $x \perp z_j$ were not one of the constraints in C , then $\Box C \Rightarrow \{P\} S \{Q\}$ would be false. Hence, the constraint C contains all the \perp relations that we need to apply our Assignment Axiom, and completeness follows by the same argument as in the ordinary Hoare logic [1].

The proof for the `declare` is immediate, since if S is the statement

declare C' in S'

$\Box C \Rightarrow \{P\} S \{Q\}$ is semantically equivalent to $\Box C \wedge C' \Rightarrow \{P\} S' \{Q\}$. Similarly, if S is the statement

x may alias x_1, \dots, x_n in S'

the result follows from the fact that $\Box C \Rightarrow \{P\} S \{Q\}$ is semantically equivalent to $\Box C \wedge C' \Rightarrow \{P\} S' \{Q\}$, where C' is $(x \perp y_1) \wedge \dots \wedge (x \perp y_n)$ and the y_i are all the variable names other than the x_j that appear in P, C , and

To finish the completeness proof, we must show that for every compound statement S , we can prove every valid formula $\Box C \Rightarrow \{P\} S \{Q\}$ under the assumption that we can prove every such formula for the components of S . This involves a separate proof for every type of compound statement. The proofs for concatenation, `while`, and `new` are similar to the ones for the ordinary Hoare logic given in [1]. The only difference in the proofs arises because of the "assertion" $\Box C \Rightarrow$. The proofs in [1] rely on the fact that $\{P\} S \{Q\}$ is valid iff $post_S(P) \Rightarrow Q$. The formula $\Box C \Rightarrow \{P\} S \{Q\}$ is valid iff $\{P\} \text{declare } C \text{ in } S \{Q\}$ is valid, by the semantic equivalence mentioned above. This, in turn, is valid iff $post_{\text{declare } C \text{ in } S}(P)$ is expressible, and implies Q . However, expressibility follows from our assumptions C2 and C2a. With this observation, the completeness proofs of [1] are now easily extended to our more general class of assertion.

END

FILMED

8-85

DTIC