The Existence of Refinement Mappings

Martín Abadi and Leslie Lamport

August 14, 1988

©Digital Equipment Corporation 1988

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of the Systems Research Center of Digital Equipment Corporation in Palo Alto, California; an acknowledgment of the authors and individual contributors to the work; and all applicable portions of the copyright notice. Copying, reproducing, or republishing for any other purpose shall require a license with payment of fee to the Systems Research Center. All rights reserved.

Author's Abstract

Refinement mappings are used to prove that a lower-level specification correctly implements a higher-level one. We consider specifications consisting of a state machine (which may be infinite-state) that specifies safety requirements, and an arbitrary supplementary property that specifies liveness requirements. A refinement mapping from a lower-level specification $\mathbf{S_1}$ to a higher-level one $\mathbf{S_2}$ is a mapping from $\mathbf{S_1}$'s state space to $\mathbf{S_2}$'s state space. It maps steps of $\mathbf{S_1}$'s state machine to steps of $\mathbf{S_2}$'s state machine and maps behaviors allowed by $\mathbf{S_1}$ to behaviors allowed by $\mathbf{S_2}$. We show that, under reasonable assumptions about the specifications, if $\mathbf{S_1}$ implements $\mathbf{S_2}$, then by adding auxiliary variables to $\mathbf{S_1}$ we can guarantee the existence of a refinement mapping. This provides a completeness result for a practical, hierarchical specification method.

Capsule Review

This report deals with the problem of proving that *implementations* satisfy their *specifications*. Suppose, for example, that a client asks a circuit fabricator to build a box with S inside and with certain external signals (inputs and outputs). The circuit fabricator returns later with an epoxy brick that has an appropriate number of wires sticking out for the external signals, but that actually contains not S but some other ciruit I. In order to guarantee that the client cannot detect the substitution without breaking open the brick (and thereby voiding the warranty), the fabricator must be sure that for any possible behavior of I there corresponds at least one behavior of S that produces identical external signals.

In some cases, such a correspondence between behaviors can be proved using a *refinement mapping*, a function that maps states of I to states of S and that satisfies certain conditions. The refinment mapping technique reduces a problem of proving something about arbitrary behaviors to one of proving something about single state transitions. Unfortunately there are many cases in which a correct implementation I cannot be related to its specification S by a refinement mapping.

This report shows that it is possible in a very large class of cases to augment a legal implementation with some extra state components (*history* and *prophecy* variables) in a way that places no constraints on the behavior of the implementation but that makes it possible to produce an appropriate refinement mapping to the specification. This result broadens considerably the domain of applicability of the refinement mapping technique.

Jim Saxe

Contents

1 Introduction **1**
1.1 Specifications 1
1.2 Proving That One Specification Implements Another 3

2 Preliminaries **6**
2.1 Sequences 6
2.2 Properties 7
2.3 Specifications 8
2.4 Refinement Mappings 11

3 Finite Invisible Nondeterminism **12**

4 Safety Properties **16**

5 Auxiliary Variables **20**
5.1 History Variables 20
5.2 Simple Prophecy Variables 22
5.3 Prophecy Variables That Add Stuttering 25

6 Internal Continuity **28**

7 The Completeness Theorem **30**

8 Whence and Whither? **35**

References **38**

1 Introduction

1.1 Specifications

A system may be specified at many levels of abstraction, from a description of its highest-level properties to a description of its implementation in terms of microcode and circuitry. We address the problem of proving that a lower-level specification is a correct implementation of a higher-level one.

Unlike simple programs, which can be specified by input/output relations, complex systems can be adequately specified only by describing their behaviors—that is, their possible sequences of inputs and outputs. We consider specification methods in which a behavior is represented by a sequence of states and a system is specified by a set of permitted behaviors. Input and output are represented in the state—for example, by including a keyboard state describing which keys are currently depressed and a screen state describing what is currently displayed.

A specification should describe only the *externally visible* components of a system's state. However, it is often helpful to describe its behavior in terms of unobservable *internal* components. For example, a natural way to specify a queue includes a description of the sequence of elements currently in the queue, and that sequence is not externally visible. Although internal components are mentioned, the specification prescribes the behavior of only the externally visible components. The system may exhibit the externally visible behavior

$$\langle\!\langle e_0,\ e_1,\ e_2,\ \ldots \rangle\!\rangle$$

where e_i is a state of the externally visible component, if there exist states y_i of the internal component such that the *complete* behavior

$$\langle\!\langle (e_0, y_0),\ (e_1, y_1),\ (e_2, y_2),\ \ldots \rangle\!\rangle$$

is permitted by the specification. (We use $\langle\!\langle\ \rangle\!\rangle$ to denote a sequence.)

A specification may allow steps in which only the internal state component changes—for example, a sequence

$$\langle\!\langle (e_0, y_0),\ (e_1, y_1),\ (e_1, y_1'),\ (e_1, y_1''),\ (e_2, y_2),\ \ldots \rangle\!\rangle$$

Such internal steps are not externally visible, so the sequence of external states $\langle\langle e_0, e_1, e_1, e_1, e_2, \ldots\rangle\rangle$ should be equivalent to the sequence $\langle\langle e_0, e_1, e_2, \ldots\rangle\rangle$ obtained by removing the "stuttering" steps from e_1 to e_1. Let $\Gamma\langle\langle e_0, e_1, \ldots\rangle\rangle$ be the set of all sequences obtained from $\langle\langle e_0, e_1, \ldots\rangle\rangle$ by repeating states and deleting repeated states—that is, by adding and removing finite amounts of stuttering. We consider only specifications in which a sequence $\langle\langle e_0, e_1, \ldots\rangle\rangle$ is allowed only if all sequences in $\Gamma\langle\langle e_0, e_1, \ldots\rangle\rangle$ are allowed. Such specifications are said to be *invariant under stuttering*.

The behaviors permitted by a specification can be described as the set of sequences satisfying a safety and a liveness property [AS86, Lam77]. Intuitively, a safety property asserts that something bad does not happen and a liveness property asserts that something good does eventually happen. In specifying a queue, the safety property might assert that the sequence of elements removed from the queue is an initial prefix of the sequence of elements added to the queue. The liveness property might assert that an operation of putting an element into the queue is eventually completed if the queue is not full, and an operation of removing an element from the queue is eventually completed if the queue is not empty. (What operations are in progress and what elements they are adding to or have removed from the queue would be described by the externally visible state.)

We are concerned with specifications in which the safety property is described by an "abstract" nondeterministic program; a behavior satisfies the property if it can be generated by the program. Liveness properties are described either directly by writing axioms or indirectly by placing fairness constraints on the abstract program. In a specification of a queue, the program describes the sequence of actions by which an element is added to or removed from the sequence of queued elements, ensuring the safety property that the correct elements are removed from the queue. Additional fairness constraints assert that certain actions must eventually occur, ensuring the liveness property that operations that should complete eventually do complete.

Many proposed specification methods involve writing programs and fairness conditions in this way [LS84, Lam83, LT87]. (Some methods do not consider liveness at all and just specify safety properties with programs.)

To describe specifications formally, we represent a program by a state machine (whose set of states may be infinite) and we represent the fairness constraints by an arbitrary *supplementary* condition. For our results, it does not matter if the supplementary condition specifies a liveness property.

1.2 Proving That One Specification Implements Another

A specification $\mathbf{S_1}$ implements a specification $\mathbf{S_2}$ if every externally visible behavior allowed by $\mathbf{S_1}$ is also allowed by $\mathbf{S_2}$. To prove that $\mathbf{S_1}$ implements $\mathbf{S_2}$, it suffices to prove that if $\mathbf{S_1}$ allows the behavior

$$\langle\langle (e_0, z_0),\ (e_1, z_1),\ (e_2, z_2),\ \ldots \rangle\rangle$$

where the z_i are internal states, then there exist internal states y_i such that $\mathbf{S_2}$ allows

$$\langle\langle (e_0, y_0),\ (e_1, y_1),\ (e_2, y_2),\ \ldots \rangle\rangle$$

In general, each y_i can depend upon the entire sequence $\langle\langle (e_0, z_0), (e_1, z_1), (e_2, z_2), \ldots \rangle\rangle$, and proving the existence of the y_i may be quite difficult. The proof is easier if each y_i depends only upon e_i and z_i, so there exists a function f such that $(e_i, y_i) = f(e_i, z_i)$. To verify that $\langle\langle f(e_0, z_0), f(e_1, z_1), f(e_2, z_2), \ldots \rangle\rangle$ satisfies the safety property of $\mathbf{S_2}$, it suffices to show that f preserves state machine behavior—that is, it maps executions of $\mathbf{S_1}$'s state machine to executions (possibly with stuttering) of $\mathbf{S_2}$'s state machine. Proving that f preserves state machine behavior involves reasoning about states and pairs of states, not about sequences. Verifying that f preserves liveness—meaning that $\langle\langle f(e_0, z_0), f(e_1, z_1), f(e_2, z_2), \ldots \rangle\rangle$ satisfies the liveness property of $\mathbf{S_2}$—usually also requires only local reasoning, with no explicit reasoning about sequences. A mapping f that preserves state machine behavior and liveness is called a *refinement mapping*.

In the example of a queue, the internal state y_i of specification $\mathbf{S_2}$ might describe the sequence of elements currently in the queue, and the internal state z_i of specification $\mathbf{S_1}$ might describe the contents of an array that implements the queue. To prove that $\mathbf{S_1}$ implements $\mathbf{S_2}$, one would construct a refinement mapping f such that $f(e_i, z_i) = (e_i, y_i)$, where y_i describes the state of the queue that is represented by the contents of the array described by state z_i.

Several methods for proving that one specification implements another are based upon finding a refinement mapping [LS84, Lam83]. In practice, if $\mathbf{S_1}$ implements $\mathbf{S_2}$, then these methods usually can prove that the implementation is correct—usually, but not always. The methods fail if the refinement mapping does not exist. Three reasons why the mapping might not exist are:

- $\mathbf{S_2}$ may specify an internal state with "historical information" not needed by $\mathbf{S_1}$. For example, suppose $\mathbf{S_2}$ requires that the system

display up to three of the least-significant bits of a three-bit clock. This specification is implemented by a lower-level specification $\mathbf{S_1}$ that alternately displays zero and one, with no internal state. A refinement mapping does not exist because there is no way to define the internal state of a three-bit clock as a function of its low-order bit.

- $\mathbf{S_2}$ may specify that a nondeterministic choice is made before it has to be. For example, consider two specifications $\mathbf{S_1}$ and $\mathbf{S_2}$ for a system that displays ten nondeterministically chosen values in sequence. Suppose $\mathbf{S_2}$ requires that all values be chosen before any is displayed, while $\mathbf{S_1}$ requires each value to be chosen as it is displayed. Both specifications describe the same externally visible behaviors, so each implements the other. However, $\mathbf{S_2}$ requires the internal state to contain all ten values before any is displayed, while $\mathbf{S_1}$ does not specify any internal state, so no refinement mapping is possible.

- $\mathbf{S_2}$ may "run slower" than $\mathbf{S_1}$. For example, let $\mathbf{S_1}$ and $\mathbf{S_2}$ both specify clocks in which hours and minutes are externally visible and seconds are internal. Suppose that in $\mathbf{S_2}$ each step increments the clock by one second, while in $\mathbf{S_1}$ each step increments the clock by ten seconds. Both specifications allow the same externally visible behaviors. To show that $\mathbf{S_2}$ implements $\mathbf{S_1}$, we can use the refinement mapping f that rounds the time down to the nearest multiple of ten seconds. For any complete behavior $\langle\langle s_0, s_1, s_2, \ldots\rangle\rangle$ allowed by $\mathbf{S_2}$, the behavior $\langle\langle f(s_0), f(s_1), f(s_2), \ldots\rangle\rangle$ is a complete behavior allowed by $\mathbf{S_1}$ that contains nine "stuttering" steps for every step that changes the state.

  On the other hand, a complete behavior $\langle\langle s_0, s_1, s_2, \ldots\rangle\rangle$ specified by $\mathbf{S_1}$ may produce an externally visible change every six steps. For any mapping f, the sequence $\langle\langle f(s_0), f(s_1), f(s_2), \ldots\rangle\rangle$ may also produce an externally visible change every six steps. This is not allowed by $\mathbf{S_2}$, which requires fifty-nine internal steps for every externally visible one. Hence, no refinement mapping can prove that $\mathbf{S_1}$ implements $\mathbf{S_2}$.

If a refinement mapping does not exist, it can often be made to exist by adding *auxiliary variables* to the lower-level specification. An auxiliary variable is an internal state component that is added to a specification without affecting the externally visible behavior. The three situations described above in which refinement mappings cannot be found are handled as follows:

- Historical information missing from the internal state specified by $\mathbf{S_1}$

can be provided by adding a *history variable*—a well-known form of auxiliary variable that merely records past actions [Owi75].

- If $\mathbf{S_2}$ requires that a nondeterministic choice be made before it has to be, then $\mathbf{S_1}$ can be modified so the choice is made sooner by adding a *prophecy variable*. A prophecy variable is a new form of auxiliary variable that is the mirror image of a history variable—its formal definition is almost the same as the definition of a history variable with past and future interchanged, but there is an asymmetry due to behaviors having a beginning but not necessarily an end.

- If $\mathbf{S_2}$ runs slower than $\mathbf{S_1}$, then an auxiliary variable must be added to $\mathbf{S_1}$ to slow it down. We will define prophecy variables in such a way that they can perform this slowing.

Our main result is a completeness theorem. It states that, under three hypotheses about the specifications, if $\mathbf{S_1}$ implements $\mathbf{S_2}$ then one can add auxiliary history and prophecy variables to $\mathbf{S_1}$ to form an equivalent specification $\mathbf{S_1^{hp}}$ and find a refinement mapping from $\mathbf{S_1^{hp}}$ to $\mathbf{S_2}$. The three hypotheses, and their intuitive meanings, are:

$\mathbf{S_1}$ *is machine closed.* Machine closure means that the supplementary property (the one normally used to specify liveness requirements) does not specify any safety property not already specified by the state machine. In other words, the state machine does as much of the specifying as possible.

$\mathbf{S_2}$ *has finite invisible nondeterminism.* This denotes that, given any finite number of steps of an externally visible behavior allowed by $\mathbf{S_2}$, there are only a finite number of possible choices for its internal state component.

$\mathbf{S_2}$ *is internally continuous.* A specification is internally continuous if, for any complete behavior that is not allowed, we can determine that it is not allowed by examining only its externally visible part (which may be infinite) and some finite portion of the complete behavior.

We will show by examples why these three hypotheses are needed.

We will prove that any safety property has a specification with finite invisible nondeterminism, any specification of a safety property is internally continuous, and any property has a machine-closed specification. Therefore,

our completeness theorem implies that if the specifications are written in a suitable form and $\mathbf{S_2}$ specifies only a safety property then one can ensure that a refinement mapping exists. We will also show that, even when $\mathbf{S_2}$ is not internally continuous, a refinement mapping exists to show that $\mathbf{S_1}$ satisfies the safety property specified by $\mathbf{S_2}$. Therefore, by writing suitable specifications, refinement mappings can always be used to prove the safety property of a specification if not its liveness property. We do not know if anything can be said about proving arbitrary liveness properties.

Throughout this report, proofs are written in a self-explanatory structured format. The format permits very careful proofs that can be read to any desired level of detail by ignoring lower-level statements. Writing proofs in this format helped us to eliminate many errors and greatly increased our confidence in the correctness of the results.

A glossary/index of notations and conventions appears at the end of this report, along with an index. We hope they will help the reader cope with the formalism.

2 Preliminaries

2.1 Sequences

We now define some useful notations for sequences. In these definitions, σ denotes the sequence $\langle\langle s_0, s_1, s_2, \ldots\rangle\rangle$ and τ denotes the sequence $\langle\langle t_0, t_1, t_2, \ldots\rangle\rangle$. These sequences may be finite or infinite. If σ is finite, we let $\|\sigma\|$ denote its length and $last(\sigma)$ denote its last element, so $\|\langle\langle s_0, \ldots, s_{m-1}\rangle\rangle\| = m$ and $last(\langle\langle s_0, \ldots, s_{m-1}\rangle\rangle) = s_{m-1}$. An infinite sequence is said to be *terminating* iff (if and only if) it is of the form $\langle\langle s_0, s_1, \ldots, s_n, s_n, s_n, \ldots\rangle\rangle$—in other words, if it reaches a final state in which it stutters forever.

As usual, a mapping on elements is extended to a mapping on sequences of elements by defining $g(\sigma)$ to equal $\langle\langle g(s_0), g(s_1), \ldots\rangle\rangle$, and to a mapping on sets of elements by defining $g(S)$ to equal $\{g(s) : s \in S\}$.

The sequence σ is said to be *stutter-free* if, for each i, either $s_i \neq s_{i+1}$ or the sequence is infinite and $s_i = s_j$ for all $j \geq i$. Thus, a nonterminating sequence is stutter-free iff it never stutters, and a terminating sequence is stutter-free iff it stutters only after reaching its final state. We define $\natural\sigma$ to be the stutter-free form of σ—that is, the stutter-free sequence obtained by replacing every maximal finite subsequence $s_i, s_{i+1}, \ldots, s_j$ of identical

elements with the single element s_i. For example,

$$\natural\langle\langle 0,1,1,2,2,2,3,3,3,3,4\ldots\rangle\rangle \;=\; \langle\langle 0,1,2,3,4,\ldots\rangle\rangle$$

We define $\sigma \simeq \tau$ to mean that $\natural\sigma = \natural\tau$, so $\sigma \simeq \tau$ iff σ and τ are equivalent up to stuttering, and we define $\Gamma\sigma$ to be the set $\{\tau : \tau \simeq \sigma\}$. If S is a set of sequences, $\Gamma(S)$ is the set $\{\tau : \exists\sigma \in S.\tau \in \Gamma\sigma\}$. A set of sequences S is *closed under stuttering* if $S = \Gamma(S)$. Thus, S is closed under stuttering iff for every pair of sequences σ, τ with $\sigma \simeq \tau$, if $\sigma \in S$ then $\tau \in S$.

We use "$\cdot$" to denote concatenation of sequences—that is, if $\|\sigma\| = m$, then $\sigma \cdot \tau = \langle\langle s_0,\ldots,s_{m-1},t_0,t_1,\ldots\rangle\rangle$. If $\|\sigma\| \geq m$, we let $\sigma|_m$ denote $\langle\langle s_0,s_1,\ldots,s_{m-1}\rangle\rangle$, the prefix of σ of length m.

For any set Σ, let Σ^ω denote the set of all infinite sequences of elements in Σ. An infinite sequence $\langle\langle\sigma_0,\sigma_1,\sigma_2,\ldots\rangle\rangle$ of sequences in Σ^ω is said to *converge* to the sequence σ in Σ^ω iff for all $m \geq 0$ there exists an $n \geq 0$ such that $\sigma_i|_m = \sigma|_m$ for all $i \geq n$. In this case, we define $\lim\sigma_i$ to be σ. This definition of convergence gives rise to a topology on Σ^ω. We now recall some other definitions.

Let σ be an element of Σ^ω and let S be a subset of Σ^ω. We say that σ is a *limit point* of S iff there exist elements σ_i in S such that $\lim\sigma_i = \sigma$. The set S is *closed* iff S contains all its limit points. The *closure* of S, denoted $\overline{S}$, consists of all limit points of S; it is the smallest closed superset of S.

2.2 Properties

We can only say that one specification implements another if we are given a correspondence between the externally visible states of the two specifications. For example, if $\mathbf{S_2}$ asserts that the initial value of a particular register is the integer -3 and $\mathbf{S_1}$ asserts that the register's initial value is the sequence of bits 1111100, then we can't say whether or not $\mathbf{S_1}$ implements $\mathbf{S_2}$ without knowing how to interpret a sequence of bits as an integer. In general, to decide if $\mathbf{S_1}$ implements $\mathbf{S_2}$, we must know how to interpret an externally visible state of $\mathbf{S_1}$ as an externally visible state of $\mathbf{S_2}$. Given such an interpretation, we can translate $\mathbf{S_1}$ into a specification with the same set of externally visible states as $\mathbf{S_2}$. Thus, there is no loss of generality in requiring that $\mathbf{S_1}$ and $\mathbf{S_2}$ have the same set of externally visible states.

We therefore assume that all specifications under consideration have the same fixed set Σ_E of externally visible states. A *state space* Σ is a subset of $\Sigma_E \times \Sigma_I$ for some set Σ_I of internal states. We let Π_E be the obvious

projection mapping from $\Sigma_E \times \Sigma_I$ onto Σ_E. The set Σ_E itself is considered to be a state space for which Π_E is the identity mapping.

If Σ is a state space, then a Σ-*behavior* is an element of Σ^ω. A Σ_E-behavior is called an *externally visible behavior*. A Σ-*property* P is a set of Σ-behaviors that is closed under stuttering. A Σ_E-property is called an *externally visible property*. If P is a Σ-property, then $\Pi_E(P)$ is a set of externally visible behaviors but is not necessarily an externally visible property because it need not be closed under stuttering. The externally visible property *induced* by a Σ-property P is defined to be the set $\Gamma(\Pi_E(P))$.

If Σ is clear from context or is irrelevant, we use the terms *behavior* and *property* instead of Σ-behavior and Σ-property. We sometimes add the adjective "complete", as in "complete behavior", to distinguish behaviors and properties from externally visible behaviors and properties.

A property P that is closed ($P = \overline{P}$) is called a *safety property*. Intuitively, a safety property is one asserting that something bad does not happen. To see that our formal definition of a safety property as a closed set captures this intuitive meaning, observe that if something bad happens, then it must happen within some finite period of time. Thus, P is a safety property iff, for any sequence σ not in P, one can tell that σ is not in P by looking at some finite prefix $\sigma|_i$ of σ. In other words, $\sigma \notin P$ iff there exists an i such that for all τ if $\tau|_i = \sigma|_i$ then $\tau \notin P$. Hence, $\sigma \in P$ iff for all i there exists a $\tau_i \in P$ such that $\tau_i|_i = \sigma|_i$. But $\lim \tau_i = \sigma$, which implies that $\sigma \in \overline{P}$; thus, $\sigma \in P$ iff $\sigma \in \overline{P}$. Therefore, P satisfies the intuitive definition of a safety property only if $P = \overline{P}$.

Even though we do not use the formal definition, it is interesting to note that a Σ-property L can be defined to be a liveness property iff it is dense in Σ^ω—in other words, if $\overline{L} = \Sigma^\omega$. This means that L is a liveness property iff any finite sequence of elements in Σ can be extended to a behavior in L. In a topological space, every set can be written as the intersection of a closed set and a dense set, so any property P can be written as $M \cap L$, where M is a safety property and L is a liveness property. Moreover, M can be taken to be $\overline{P}$.

2.3 Specifications

A *state machine* is a triple (Σ, F, N) where

- Σ is a state space. (Recall that this means $\Sigma \subseteq \Sigma_E \times \Sigma_I$ for some set Σ_I of internal states.)

- F, the set of *initial* states, is a subset of Σ.
- N, the *next-state relation*, is a subset of $\Sigma \times \Sigma$. (Elements of N are denoted by pairs of states enclosed in angle brackets, like $\langle s, t \rangle$.)

The *(complete) property generated by* a state machine (Σ, F, N) consists of all infinite sequences $\langle\langle s_0, s_1, \ldots \rangle\rangle$ such that $s_0 \in F$ and, for all $i \geq 0$, either $\langle s_i, s_{i+1} \rangle \in N$ or $s_i = s_{i+1}$. This set is closed under stuttering, so it is a Σ-property. The *externally visible property generated by* a state machine is the externally visible property induced by its complete property.

We now show that the complete property P generated by a state machine is a safety property. This requires proving that if $\lim \sigma_i = \sigma$ and each $\sigma_i \in P$, then $\sigma \in P$. For any behavior $\tau = \langle\langle s_0, s_1, \ldots \rangle\rangle$ and any $j \geq 0$, let τ^j be the terminating behavior $\langle\langle s_0, s_1, \ldots, s_j, s_j, s_j, \ldots \rangle\rangle$. Then τ is in P iff each τ^j is in P. Since $\lim \sigma_i = \sigma$, each σ^j equals $(\sigma_i)^j$ for some i. Since each σ_i is in P, each $(\sigma_i)^j$ is in P, which implies that σ is also in P. Hence, P is closed, so the complete property generated by a state machine is a safety property. However, we will show in Section 3 that the externally visible property generated by a state machine need not be a safety property.

A state machine (Σ, F, N) is a familiar type of nondeterministic automaton, where F is the set of starting states and N describes the possible state transitions. (However, remember that Σ may be an infinite set.) The set of sequences generated (or accepted) by such an automaton is usually defined to be the set A of all sequences starting with a state in F and progressing by making transitions allowed by N. However, we also allow stuttering transitions, so we have defined the property generated by the state machine to be $\Gamma(A)$ together with all terminating sequences obtained from finite prefixes of behaviors in $\Gamma(A)$ by infinite stuttering.

A *specification* $\mathbf{S}$ is a four-tuple (Σ, F, N, L), where (Σ, F, N) is a state machine and L is a Σ-property, called the *supplementary property* of the specification. The property M generated by the state machine (Σ, F, N) is called the *machine property* of $\mathbf{S}$. The *(complete) property defined by* $\mathbf{S}$ is defined to be $M \cap L$, and the *externally visible property defined by* $\mathbf{S}$ is defined to be $\Gamma(\Pi_E(M \cap L))$, the externally visible property induced by $M \cap L$.

State machines are easier to work with than arbitrary sets of sequences, so one would like to specify a property purely in terms of state machines. However, the complete property generated by a state machine is a safety property. The supplementary property of a specification is needed to introduce liveness requirements. However, if we were to place no additional

requirement on our specifications, we could use the supplementary property to do all the specifying. To see why this leads to trouble, let $\mathbf{S_2}$ be a specification consisting of any arbitrary state machine that generates an externally visible safety property O together with the trivial supplementary property that contains all behaviors. Define $\mathbf{S_1}$ to be the specification with state space Σ_E whose state machine is the trivial one that generates all Σ_E-behaviors and whose supplementary property is O. Obviously $\mathbf{S_1}$ implements $\mathbf{S_2}$. The existence of a refinement mapping from $\mathbf{S_1}$ to $\mathbf{S_2}$ implies that $\mathbf{S_1}$'s state machine implements $\mathbf{S_2}$'s state machine. However, $\mathbf{S_1}$ has the trivial state machine and no internal state. As we will see, auxiliary variables are added to a specification's state machine without affecting or being affected by the supplementary property. (This is what makes the addition of auxiliary variables practical.) No sound method of adding auxiliary variables can transform the trivial machine into one that implements an arbitrary state machine. Therefore, we need some constraint on the supplementary property.

In practice, we specify a desired property P by writing P as the intersection $M \cap L$ of a safety property M and a liveness property L. We try to construct L so that it does not specify any safety property, meaning that it does not rule out any finite behavior. More precisely, we try to choose L to be a liveness property such that any finite sequence of states generated by the state machine is the prefix of a behavior in P. For our results, it is not necessary that L be a liveness property; we need only require that L does not specify any safety property not implied by M. To express this requirement formally, we say that a specification $\mathbf{S}$ having machine property M and supplementary property L is *machine closed* iff $M = \overline{M \cap L}$.

The following lemma implies that, for a machine-closed specification, we can ignore the supplementary property and consider only the state machine when we are interested in finite portions of behaviors.

Lemma 1 *If $M = \overline{P}$, then every prefix of a behavior in M is the prefix of a behavior in P.*

Proof of Lemma 1
Given: A1. $M = \overline{P}$.
A2. $\sigma \in M$.
A3. $m \geq 0$.
Prove: C1. There exists $\tau \in P$ such that $\tau|_m = \sigma|_m$.
Pf: 1. Choose $\sigma_i \in P$ such that $\lim \sigma_i = \sigma$.

Pf: A1, A2, and the definition of $\overline{P}$.

2. Choose $n \geq 0$ such that, for all $i \geq n$, $\sigma_i|_m = \sigma|_m$.
   Pf: Definition of convergence.
3. C1 holds.
   Pf: Let τ be σ_n.

End Proof of Lemma 1

The converse of this lemma is also true when M is generated by a state machine, but we will not need it.

2.4 Refinement Mappings

A specification $\mathbf{S_1}$ *implements* a specification $\mathbf{S_2}$ iff the externally visible property induced by $\mathbf{S_1}$ is a subset of the externally visible property induced by $\mathbf{S_2}$. In other words, $\mathbf{S_1}$ implements $\mathbf{S_2}$ iff every externally visible behavior allowed by $\mathbf{S_1}$ is also allowed by $\mathbf{S_2}$.

A *refinement mapping* from a specification $\mathbf{S_1} = (\Sigma_1, F_1, N_1, L_1)$ to a specification $\mathbf{S_2} = (\Sigma_2, F_2, N_2, L_2)$ is a mapping $f : \Sigma_1 \to \Sigma_2$ such that

R1. For all $s \in \Sigma_1$: $\Pi_E(f(s)) = \Pi_E(s)$. (f preserves the externally visible state component.)

R2. $f(F_1) \subseteq F_2$. (f takes initial states into initial states.)

R3. If $\langle s, t\rangle \in N_1$ then $\langle f(s), f(t)\rangle \in N_2$ or $f(s) = f(t)$. (A state transition allowed by N_1 is mapped by f into a [possibly stuttering] transition allowed by N_2.)

R4. $f(P_1) \subseteq L_2$, where P_1 is the property defined by $\mathbf{S_1}$. (f maps behaviors allowed by $\mathbf{S_1}$ into behaviors that satisfy $\mathbf{S_2}$'s supplementary property.)

Conditions R1–R3 are local, meaning that they can be checked by reasoning about states or pairs of states rather than about behaviors. Condition R4 is not local, but checking it is simplified by the fact that f is not an arbitrary mapping on sequences, but is obtained from a mapping on states. Thus, one can apply local methods like well-founded induction to prove R4.

Proposition 1 *If there exists a refinement mapping from* $\mathbf{S_1}$ *to* $\mathbf{S_2}$*, then* $\mathbf{S_1}$ *implements* $\mathbf{S_2}$.

Proof of Proposition 1.

Given: A1. $\mathbf{S_i} = (\Sigma_i, F_i, N_i, L_i)$, for $i = 1, 2$.
A2. M_i is the machine property of $\mathbf{S_i}$, for $i = 1, 2$.
A3. f is a refinement mapping from $\mathbf{S_1}$ to $\mathbf{S_2}$.
A4. $\eta \in \Gamma(\Pi_E(M_1 \cap L_1))$

Prove: C1. $\eta \in \Gamma(\Pi_E(M_2 \cap L_2))$.

Pf: 1. $f(M_1) \subseteq M_2$.
Given: A1.1. $\sigma = \langle\langle s_0, s_1, \ldots \rangle\rangle \in M_1$.
Prove: C1.1. $f(\sigma) \in M_2$.
Pf: 1.1. $f(s_0) \in F_2$.
Pf: By A1.1, A2, the definition of machine property (which permits stuttering), A3, and property R2 in the definition of refinement mapping.
1.2. For all $i \geq 0$: $\langle f(s_i), f(s_{i+1}) \rangle \in N_2$ or $f(s_i) = f(s_{i+1})$.
Pf: By A1.1, A2, the definition of machine property, A3, and property R3.
1.3. C1.1 holds.
Pf: By 1.1, 1.2, the definition of $f(\sigma)$ (it equals $\langle\langle f(s_0), f(s_1), \ldots \rangle\rangle$), A2, and the definition of machine property.
2. $f(M_1 \cap L_1) \subseteq M_2 \cap L_2$.
Pf: By 1, A3, and R4, since $g(S \cap T) \subseteq g(S) \cap g(T)$ for any sets S and T and any mapping g.
3. Choose $\sigma = \langle\langle s_0, s_1, \ldots \rangle\rangle \in M_1 \cap L_1$ such that $\Pi_E(\sigma) \simeq \eta$.
Pf: Such a σ exists by A4 and the definition of Γ.
4. $\Pi_E(f(\sigma)) = \Pi_E(\sigma)$.
Pf: By A3 and R1.
5. $\Pi_E(f(\sigma)) \simeq \eta$.
Pf: By 3 and 4.
6. $\Pi_E(f(\sigma)) \in \Pi_E(M_2 \cap L_2)$.
Pf: By 3 and 2.
7. C1 holds.
Pf: By 5, 6, and the definition of Γ.

End Proof of Proposition 1.

3 Finite Invisible Nondeterminism

The machine property M of a specification is a safety property. However, the property that is really being specified by the specification's state machine

is the externally visible property $\Gamma(\Pi_E(M))$ induced by M. The following example shows that this externally visible property is not necessarily a safety property.

Let Σ_E be the set $\mathbf{N}$ of natural numbers, and define the state machine (Σ, F, N) by:

- Σ equals $\Sigma_E \times \mathbf{N}$.
- F equals $\{(0,0)\}$.
- N is the union of the following two sets:
  - $\{\langle (0,0), (1,n)\rangle : n \in \mathbf{N}\}$,
  - $\{\langle (m, n+1), (m+1, n)\rangle : m, n \in \mathbf{N}\}$.

A stutter-free behavior of this machine starts in state $(0,0)$, goes to state $(1,n)$ for some arbitrary $n \geq 0$, then goes through the sequence of states $(2, n-1), (3, n-2), \ldots, (n-i+1, i)$ for some $i \geq 0$, and terminates (stutters forever) in the state $(n-i+1, i)$.

The set of externally visible behaviors induced by this state machine consists of all sequences obtainable by stuttering from a sequence σ_n of the form $\langle\langle 0, 1, 2, \ldots, n, n, n, \ldots\rangle\rangle$. This set is not closed, because $\lim \sigma_n = \langle\langle 0, 1, 2, 3, \ldots\rangle\rangle$, and $\langle\langle 0, 1, 2, 3, \ldots\rangle\rangle$ is not in the set. The externally visible property specified by this state machine is the conjunction of two properties:

1. The set of all behaviors that start in state 0 and change state only by adding 1 to the previous state.
2. The set of terminating behaviors.

The first property is a safety property, but the second is a liveness property; their intersection is neither a safety nor a liveness property.

The purpose of a specification is to specify an externally visible property. We feel that the externally visible property specified by a state machine should be a safety property, so we want to restrict the class of allowed state machines.

The reason the externally visible property defined by the state machine in our example is not a safety property can be traced to the existence of infinitely many state transitions $\langle (0,0), (1,n)\rangle$ that correspond to the same externally visible transition $\langle 0, 1\rangle$. It is this type of infinite invisible nondeterminism that allows the introduction of liveness into the externally visible

property of a state machine. To ensure that a state machine specifies only safety properties, we must restrict it to having finite invisible nondeterminism.

Instead of defining the concept of finite invisible nondeterminism for a state machine, it is more general to define it for a property. A state machine is defined to have finite invisible nondeterminism iff the property it generates does.

Definition 1 *Let P be a property and O its induced externally visible property $\Gamma(\Pi_E(P))$. We say that P is* fin *(for* finitely invisibly nondeterministic*) iff for all $\eta \in O$ and all $n \geq 0$, the set*

$$\{\sharp(\sigma|_m) : (m > 0) \wedge (\sigma \in P) \wedge (\Pi_E(\sigma|_m) \simeq \eta|_n)\}$$

is finite. We say that a specification is fin iff the complete property of the specification is fin.

In other words, property P is fin iff every finite prefix $\eta|_n$ of any externally visible behavior η is the projection of only finitely many inequivalent (under $\simeq$) finite prefixes $\sigma|_m$ of complete behaviors σ in P.

If a property M is fin then every stronger property P is also fin. (Property P is stronger than property M iff $P \subseteq M$.) In our main theorem, instead of requiring that the state machine of $\mathbf{S_2}$ is fin, we make the weaker assumption that $\mathbf{S_2}$ is fin. This is strictly weaker only if $\mathbf{S_2}$ is not machine closed, since a machine-closed specification is fin iff its state machine is fin.

The following proposition asserts that the externally visible property of a fin state machine is a safety property. It is a simple corollary of the subsequent lemma, which will be used later as well.

Proposition 2 *If a safety property P is fin, then the externally visible property $\Gamma(\Pi_E(P))$ that it induces is also a safety property.*

Lemma 2 (Closure and nondeterminism) *Let property P be fin and let O be the externally visible property that it induces. If δ is a limit point of O then there is a limit point ρ of P such that $\Pi_E(\rho) \simeq \delta$.*

Proof of Lemma 2

Given: A1. P is fin.

A2. $O = \Gamma(\Pi_E(P))$.

A3. δ is a limit point of O.

Prove: C1. There exists ρ such that:

C1a. ρ is a limit point of P.

C1b. $\Pi_E(\rho) \simeq \delta$.

Pf: 1. Let Θ_n equal $\{\natural(\sigma|_m) : (m > 0) \wedge (\sigma \in P) \wedge (\Pi_E(\sigma|_m) \simeq \delta|_n)\}$. For all n, the set Θ_n is finite. (Θ_n is the set of stutter-free prefixes of behaviors in P that are externally equivalent to $\delta|_n$.)

Pf: By A3, we can choose $\eta \in O$ such that $\eta|_n = \delta|_n$. Statement 1 then follows from A1 and Definition 1.

2. For all n, the set Θ_n is nonempty.

Pf: 2.1. Choose $\eta \in O$ such that $\eta|_n = \delta|_n$.

Pf: A3 implies the existence of η.

2.2. Choose $\sigma \in P$ such that $\Pi_E(\sigma) \simeq \eta$.

Pf: A2 and definition of Γ imply the existence of σ.

2.3. There exists m such that $\Pi_E(\sigma|_m) \simeq \eta|_n$.

Pf: 2.2 and the definition of $\simeq$.

2.4. $\natural(\sigma|_m) \in \Theta_n$, so Θ_n is nonempty.

Pf: $\sigma \in P$ (by 2.2), and $\Pi_E(\sigma|_m) \simeq \delta|_n$ (by 2.3 and 2.1), so 2.4 follows from 1 (the definition of Θ_n).

3. For finite sequences σ and τ, let $\sigma \preceq \tau$ iff there is a (possibly empty) sequence χ such that $\tau = \sigma \cdot \chi$. For all n and all $\theta \in \Theta_{n+1}$ there exists $\theta' \in \Theta_n$ such that $\theta' \preceq \theta$.

Pf: By 1 (the definition of Θ_n), since if $\tau|_m \simeq \delta|_{n+1}$, then there exists $m' \leq m$ such that $\tau|_{m'} \simeq \delta|_n$.

4. There is an infinite sequence $\rho_1 \preceq \rho_2 \preceq \rho_3 \preceq \ldots$ with each $\rho_i \in \Theta_i$.

Pf: By 1, 2, 3 and König's Lemma [Knu73, pages 381–383].

5. For all i, choose ρ'_i such that:

5a. $\rho'_i \simeq \rho_i$.

5b. $\|\rho'_i\| \geq i$.

5c. $\rho'_1 \preceq \rho'_2 \preceq \rho'_3 \preceq \ldots$.

Pf: The existence of the ρ'_i is proved by induction using 4, where the length of ρ'_i is increased by stuttering the last element when necessary.

6. Let $\widehat{\rho}_i$ be an element of P such that ρ'_i is a prefix of $\widehat{\rho}_i$.

Pf: Since $\rho_i \in \Theta_i$ (by 4), the definition of Θ_i (1) implies that there exists a stutter-free sequence $\psi_i \in P$ such that ρ_i is a prefix of ψ_i. By 5a and the assumption that P (like all properties) is invariant under stuttering, $\widehat{\rho}_i$ can be obtained by adding stuttering to ψ_i.

7. Let ρ equal $\lim \widehat{\rho}_i$.

Pf: ρ exists by 6 (ρ'_i a prefix of $\widehat{\rho}_i$), 5b, and 5c.

8. C1a holds.
   Pf: Follows immediately from 7, 6 ($\widehat{\rho}_i \in P$), and the definition of limit point.
9. For every i there exists an $m \geq i$ such that $\Pi_E(\widehat{\rho}_i|_m) \simeq \delta|_i$.
   Pf: 5a, 5b, 6, 4 ($\rho_i \in \Theta_i$), and 1 (the definition of Θ_i).
10. $\lim \Pi_E(\widehat{\rho}_i) \simeq \delta$.
    Pf: Follows immediately from 9.
11. C1b holds.
    Pf: By 6 ($\widehat{\rho}_i \in P$), 7, and 10, since $\lim \xi_i = \xi$ implies $\lim \Pi_E(\xi_i) = \Pi_E(\xi)$.

End Proof of Lemma 2

For a state machine to be fin, it may not make an infinite nondeterministic choice unless all but a finite part of that choice is immediately revealed in the externally visible state. We can weaken our definition by requiring only that the choice eventually be revealed. Formally, this means defining a property P with induced externally visible property O to be fin iff for every η in O and $n \geq 0$ there exists an $n' \geq n$ such that the set

$$\{\natural(\sigma|_m) \ : (m > 0) \wedge (\sigma \in P) \wedge (\Pi_E(\sigma|_m) \simeq \eta|_n) \\ \wedge \exists m' : (\Pi_E(\sigma|_{m'}) \simeq \eta|_{n'})\}$$

is finite. However, using this weaker definition of finite invisible nondeterminism would require somewhat more powerful prophecy variables and would complicate our proofs, so we will stick with our original definition.

4 Safety Properties

Alpern and Schneider [AS87] and others have observed in the finite-state case that there is a correspondence between state machines and externally visible safety properties. We extend their results to the infinite-state case for state machines with finite invisible nondeterminism. We also prove a result that allows us to apply our completeness theorem to safety properties even when the internal continuity hypothesis defined later is not satisfied.

Proposition 2 implies that the externally visible property generated by a fin state machine is a safety property. We now prove the converse.

Proposition 3 *Every externally visible safety property can be generated by a state machine with finite invisible nondeterminism.*

Proof of Proposition 3

Given: A1. O is a Σ_E-property.

A2. $O = \overline{O}$.

Prove: C1. There exists a state machine (Σ, F, N) generating a (complete) property M such that

C1a. M is fin.

C1b. $O \subseteq \Gamma(\Pi_E(M))$.

C1c. $\Gamma(\Pi_E(M)) \subseteq O$.

Pf: 1. Define the state machine (Σ, F, N) as follows:

- $\Sigma = \{(last(\theta|_n), \theta|_n) : n \geq 1 \wedge \theta \in O\}$. ($\Sigma$ consists of all pairs $(e_i, \langle\!\langle e_0, e_1, \ldots, e_i \rangle\!\rangle)$ such that $\langle\!\langle e_0, e_1, \ldots, e_i \rangle\!\rangle$ is a prefix of a sequence in O.)
- $F = \{(e, \langle\!\langle e \rangle\!\rangle) \in \Sigma\}$. (The starting states are ones whose internal components have length one.)
- $N = \{\langle (e, h), (e', h \cdot \langle\!\langle e' \rangle\!\rangle) \rangle \in \Sigma \times \Sigma\}$ (The machine can go from state $(e_i, \langle\!\langle e_0, \ldots, e_i \rangle\!\rangle)$ only to state $(e_{i+1}, \langle\!\langle e_0, \ldots, e_i, e_{i+1} \rangle\!\rangle)$ for some e_{i+1}.)

2. A stutter-free sequence $\langle\!\langle (e_0, h_0), (e_1, h_1), \ldots \rangle\!\rangle$ is in M iff, for all $i \geq 0$, $h_i = \langle\!\langle e_0, e_1, \ldots, e_i \rangle\!\rangle$ and there exists $\eta_i \in O$ such that $h_i = \eta_i|_{i+1}$.

   Pf: Follows easily by induction from the definition of the state machine (Σ, F, N) and of the property that it generates.

3. C1a holds.

   Pf: By Definition 1, we must show that for any $\eta \in O$ and all $n \geq 0$ the set

   $$\{\natural(\sigma|_m) : (m > 0) \wedge (\sigma \in M) \wedge (\Pi_E(\sigma|_m) \simeq \eta|_n)\}$$

   is finite. However, it follows from 2 that if $\eta = \langle\!\langle e_0, e_1, \ldots \rangle\!\rangle$ then this set contains only the single element

   $$\langle\!\langle (e_0, \langle\!\langle e_0 \rangle\!\rangle), (e_1, \langle\!\langle e_0, e_1 \rangle\!\rangle), \ldots, (e_{n-1}, \langle\!\langle e_0, \ldots, e_{n-1} \rangle\!\rangle) \rangle\!\rangle$$

4. C1b holds.

   Pf: For any $\eta = \langle\!\langle e_0, e_1, \ldots \rangle\!\rangle$ in O, statement 2 implies that $\sigma = \langle\!\langle \ldots, (e_i, \eta|_{i+1}), \ldots \rangle\!\rangle$ is in M, and obviously $\Pi_E(\sigma) = \eta$.

5. $\Pi_E(M) \subseteq O$.

   Given: A5.1. $\langle\!\langle (e_0, h_0), (e_1, h_1), \ldots \rangle\!\rangle \in M$.

   Prove: C5.1. $\langle\!\langle e_0, e_1, \ldots \rangle\!\rangle \in O$.

   Pf: 5.1. For all $i \geq 0$ choose $\eta_i \in O$ such that $\eta_i|_{i+1} = \langle\!\langle e_0, \ldots, e_i \rangle\!\rangle$.

   Pf: By 2, the η_i exist.

   5.2. $\lim \eta_i = \langle\!\langle e_0, e_1, \ldots \rangle\!\rangle$.

   Pf: Follows immediately from 5.1.

5.3. C5.1 holds.

Pf: By 5.1 (which asserts that $\eta_i \in O$), 5.2, A2, and the definition of $\overline{O}$.

6. C1c holds.

Pf: By 5 and the assumption that O is a property (A1), so $\Gamma(O) = O$.

End Proof of Proposition 3

If specification $\mathbf{S_2}$ is not internally continuous, it is possible for it to be implemented by a specification $\mathbf{S_1}$ without there being a refinement mapping from $\mathbf{S_1}$ to $\mathbf{S_2}$. (Internal continuity was mentioned in the introduction and will be defined formally in Section 6.) However, since safety properties are internally continuous, we would expect to be able to prove that, whenever $\mathbf{S_1}$ implements $\mathbf{S_2}$, the externally visible machine property of $\mathbf{S_1}$ implements the externally visible machine property of $\mathbf{S_2}$. Combined with our main theorem, the following result shows that this is always possible if $\mathbf{S_1}$ is machine closed and the machine property of $\mathbf{S_2}$ is fin.

Theorem 1 (Separate safety proofs) *Let $P_1 = M_1 \cap L_1$ and $P_2 = M_2 \cap L_2$, where the L_i are arbitrary properties and the M_i are safety properties; and let O_i and O_i^M be the externally visible properties induced by P_i and M_i, respectively. If $M_1 = \overline{P_1}$, M_2 is fin, and $O_1 \subseteq O_2$, then $O_1^M \subseteq O_2^M$.*

Proof of Theorem 1

Given: A1. For $i = 1, 2$:

A1a. $P_i = M_i \cap L_i$.

A1b. M_i closed.

A1c. $O_i = \Gamma(\Pi_E(P_i))$.

A1d. $O_i^M = \Gamma(\Pi_E(M_i))$.

A2. $M_1 = \overline{P_1}$.

A3. M_2 is fin.

A4. $O_1 \subseteq O_2$.

Prove: C1. $O_1^M \subseteq O_2^M$.

Pf: 1. For any set Q of behaviors $\Gamma(\overline{Q}) \subseteq \overline{\Gamma(Q)}$.

Given: A1.1. $\sigma \in \Gamma(\overline{Q})$.

Prove: C1.1. $\sigma \in \overline{\Gamma(Q)}$.

Pf: 1.1. There exists $\sigma' \in \overline{Q}$ such that $\sigma' \simeq \sigma$.

Pf: A1.1 and the definition of Γ.

1.2. There exists a function r such that, for all $i \geq 0$, $\sigma|_i \simeq \sigma'|_{r(i)}$.

Pf: 1.1 and the definition of $\simeq$.

1.3. For all $i \geq 0$ there exists $\tau'_i \in Q$ such that $\tau'_i|_{r(i)} = \sigma'|_{r(i)}$.
Pf: Definition of $\overline{Q}$ and 1.1.

1.4. $\sigma|_i \simeq \tau'_i|_{r(i)}$.
Pf: 1.2 and 1.3.

1.5. For each i, let $\tau'_i = \langle\!\langle t_{i,0}, t_{i,1}, \ldots \rangle\!\rangle$ and define τ_i to equal $\sigma|_i \cdot \langle\!\langle t_{i,r(i)}, t_{i,r(i)+1}, \ldots \rangle\!\rangle$. Then $\tau_i \simeq \tau'_i$.
Pf: 1.4.

1.6. $\tau_i \in \Gamma(Q)$.
Pf: $\tau_i \simeq \tau'_i$ (by 1.5), $\tau'_i \in Q$ (by 1.3), and the definition of Γ.

1.7. $\lim \tau_i = \sigma$.
Pf: By 1.5 and the definition of convergence.

1.8. C1.1 holds.
Pf: 1.6, 1.7, and the definition of closure.

2. For any set Q of behaviors $\Pi_E(\overline{Q}) \subseteq \overline{\Pi_E(Q)}$.
Given: A2.1. $\eta \in \Pi_E(\overline{Q})$.
Prove: C2.1. $\eta \in \overline{\Pi_E(Q)}$.
Pf: 2.1. There exists $\sigma \in \overline{Q}$ such that $\eta = \Pi_E(\sigma)$.
Pf: A2.1.

2.2. For all $i \geq 0$ choose τ_i in Q such that $\tau_i|_i = \sigma|_i$.
Pf: 2.1 and the definition of $\overline{Q}$.

2.3. For all $i \geq 0$, $\Pi_E(\tau_i)|_i = \eta|_i$.
Pf: 2.1 and 2.2, since $\Pi_E(\psi|_i) = (\Pi_E(\psi))|_i$ for any sequence ψ.

2.4. $\Pi_E(\tau_i) \in \Pi_E(Q)$.
Pf: By 2.2 ($\tau_i \in Q$).

2.5. C2.1 holds.
Pf: By 2.3, which implies $\lim \Pi_E(\tau_i) = \eta$, and 2.4.

3. $O_1^M \subseteq \overline{O_1}$.
Pf: 3.1. $O_1^M = \Gamma(\Pi_E(\overline{P_1}))$.
Pf: A2 and A1d.

3.2. $\overline{O_1} = \overline{\Gamma(\Pi_E(P_1))}$.
Pf: A1c.

3.3. $\Pi_E(\overline{P_1}) \subseteq \overline{\Pi_E(P_1)}$.
Pf: 2.

3.4. $O_1^M \subseteq \Gamma(\overline{\Pi_E(P_1)})$
Pf: 3.1, 3.3, and monotonicity of Γ.

3.5. $\Gamma(\overline{\Pi_E P_1}) \subseteq \overline{\Gamma(\Pi_E(P_1))}$.
Pf: 1.

3.6. 3 holds.

Pf: 3.4, 3.5, and 3.2.

4. $\overline{O_1} \subseteq \overline{O_2}$.
   Pf: A4 and monotonicity of the closure operation.
5. $O_2 \subseteq O_2^M$.
   Pf: A1a, A1c, A1d, and the monotonicity of Π_E and Γ.
6. $\overline{O_2} \subseteq \overline{O_2^M}$.
   Pf: 5 and monotonicity of closure.
7. $O_2^M = \overline{O_2^M}$.
   Pf: A1b, A1d, A3, and Proposition 2.
8. C1 holds.
   Pf: 3, 4, 6, and 7.

End Proof of Theorem 1

5 Auxiliary Variables

Although in practice refinement mappings usually exist, they do not always exist. To construct a refinement mapping, it may be necessary to add auxiliary variables. We now formally define two types of auxiliary variables: the well-known *history variable* and the new *prophecy variable*. These auxiliary variables are added to a specification's state machine; the supplementary property is essentially left unchanged.

5.1 History Variables

Adding a history variable means augmenting the state space with an additional component Σ_H and modifying the state machine in such a way that this additional component records past information but does not affect the behavior of the original state components. Formally, a specification $\mathbf{S^h} = (\Sigma^h, F^h, N^h, L^h)$ is said to be *obtained from* the specification $\mathbf{S} = (\Sigma, F, N, L)$ *by adding a history variable* iff the following five conditions are satisfied. In these conditions, we identify $(\Sigma_E \times \Sigma_I) \times \Sigma_H$ with $\Sigma_E \times (\Sigma_I \times \Sigma_H)$ (so H1 implies that Σ^h is a state space), and we let $\Pi_{[H]}$ be the obvious projection mapping from $\Sigma \times \Sigma_H$ onto Σ. (In the intuitive explanation, we say that a Σ^h-behavior σ *simulates* the Σ-behavior $\Pi_{[H]}(\sigma)$.)

H1. $\Sigma^h \subseteq \Sigma \times \Sigma_H$ for some set Σ_H.

H2. $\Pi_{[H]}(F^h) = F$. (A state in Σ is an initial state of $\mathbf{S}$ iff it is the first component of an initial state of $\mathbf{S^h}$.)

H3. If $\langle (s,h),(s',h')\rangle \in N^h$ then $\langle s,s'\rangle \in N$ or $s = s'$. (Every step of $\mathbf{S^h}$'s state machine simulates a [possibly stuttering] step of $\mathbf{S}$'s state machine.)

H4. If $\langle s,s'\rangle \in N$ and $(s,h) \in \Sigma^h$ then there exists $h' \in \Sigma_H$ such that $\langle (s,h),(s',h')\rangle \in N^h$. (From any state, $\mathbf{S^h}$'s state machine can simulate any possible step of $\mathbf{S}$'s state machine.)

H5. $L^h = \Pi^{-1}_{[H]}(L)$. (A Σ^h-behavior is in L^h iff the Σ-behavior that it simulates is in L.)

The following result shows that adding a history variable leaves an implementation essentially unchanged.

Proposition 4 (Soundness of history variables) *If* $\mathbf{S}^h$ *is obtained from* $\mathbf{S}$ *by adding a history variable, then the two specifications define the same externally visible property.*

Proof of Proposition 4

Given: A1. $\mathbf{S} = (\Sigma, F, N, L)$, $\mathbf{S^h} = (\Sigma^h, F^h, N^h, L^h)$, and H1–H5 hold.
A2. M and M^h are the machine properties of $\mathbf{S}$ and $\mathbf{S}^h$, respectively.
A3. $P = M \cap L$ and $P^h = M^h \cap L^h$.
A4. $O = \Gamma(\Pi_E(P))$ and $O^h = \Gamma(\Pi_E(P^h))$.

Prove: C1. $O^h \subseteq O$.
C2. $O \subseteq O^h$.

Pf: 1. $\Pi_{[H]}(M^h) \subseteq M$.
Pf: Follows from A2, A1 (conditions H2 and H3), and the definition of the machine property of a specification.

2. $\Pi_{[H]}(P^h) \subseteq P$.
Pf: From A3, 1, and H5, since $g(S \cap T) \subseteq g(S) \cap g(T)$ for any function g and sets S and T.

3. C1 holds.
Pf: From 2, A4, and the fact that $\Pi_E(\Pi_{[H]}(s)) = \Pi_E(s)$ for any $s \in \Sigma^h$.

4. $P \subseteq \Pi_{[H]}(P^h)$.
Given: A4.1. $\sigma = \langle\!\langle s_0, s_1, \ldots \rangle\!\rangle$ in P.
Prove: C4.1. There exists $\tau \in P^h$ such that $\Pi_{[H]}(\tau) = \sigma$.
Pf: 4.1. $s_0 \in F$ and, for all $i \geq 0$, $\langle s_i, s_{i+1}\rangle \in N$.
Pf: A3 and the definition of machine property.

4.2. For all $i \geq 0$ choose h_i inductively such that $(s_0, h_0) \in F^h$ and $\langle (s_i, h_i), (s_{i+1}, h_{i+1}) \rangle \in N^h$.
Pf: The existence of h_0 follows from 4.1 ($s_0 \in F$) and H2; for $i \geq 0$, the existence of h_{i+1} follows from 4.1 ($\langle s_i, s_{i+1} \rangle \in N$) and H4.

4.3. Let $\tau = \langle\!\langle (s_0, h_0), (s_1, h_1), \ldots \rangle\!\rangle$. Then $\tau \in M^h$.
Pf: 4.2, A2, and the definition of machine property.

4.4. $\Pi_{[H]}(\tau) = \sigma$.
Pf: By definition of τ (4.3).

4.5. $\tau \in L^h$.
Pf: 4.4, H5, and A4.1.

4.6. C4.1 holds.
Pf: 4.5, 4.3, and A3, which imply that $\tau \in P^h$, and 4.4.

5. C2 holds.
Pf: From 4, A4, the monotonicity of Γ and Π_E, and the fact that $\Pi_E(\Pi_{[H]}(s)) = \Pi_E(s)$ for any $s \in \Sigma^h$.

End Proof of Proposition 4

5.2 Simple Prophecy Variables

A prophecy variable is the dual of a history variable; its definition is almost that of a history variable with time running backwards. Intuitively, whereas a history variable records past behavior, a prophecy variable guesses future behavior. Using notation similar to that used in defining history variables, we define a specification $\mathbf{S^P} = (\Sigma^p, F^p, N^p, L^p)$ to be *obtained from* $\mathbf{S} = (\Sigma, F, N, L)$ *by adding a prophecy variable* iff the following conditions are satisfied. (Conditions P2′ and P4′ will be replaced in Section 5.3.)

P1. $\Sigma^p \subseteq \Sigma \times \Sigma_P$ for some set Σ_P.

P2′. $F^p = \Pi^{-1}_{[P]}(F)$. (This is the expected correspondence between the initial states of the two specifications.)

P3. If $\langle (s, p), (s', p') \rangle \in N^p$ then $\langle s, s' \rangle \in N$ or $s = s'$. (Every step of $\mathbf{S^P}$'s state machine simulates a [possibly stuttering] step of $\mathbf{S}$'s state machine.)

P4′. If $\langle s, s' \rangle \in N$ and $(s', p') \in \Sigma^p$ then there exists $p \in \Sigma_P$ such that $\langle (s, p), (s', p') \rangle \in N^p$. (From every state in Σ^p, the state machine of $\mathbf{S^P}$ can take a backwards step that simulates any possible backwards

step of $\mathbf{S}$'s state machine. This is the time-reversed version of condition H4.)

P5. $L^p = \Pi_{[P]}^{-1}(L)$. (The supplementary property of $\mathbf{S^P}$ is the set of behaviors that simulate behaviors in the supplementary property of $\mathbf{S}$.)

P6. For all $s \in \Sigma$, the set $\Pi_{[P]}^{-1}(s)$ is finite and nonempty. (To every state of $\mathbf{S}$ there corresponds some nonzero finite number of states of $\mathbf{S^P}$.)

Condition P6 is the only one not corresponding to any condition for history variables. It is needed because time reversal is asymmetric—all behaviors have initial states but only terminating behaviors have final states. The second example below indicates why it is needed.

We now give two examples to illustrate the definition of prophecy variables. We mention only the state machines; the supplementary property can be taken to be the trivial one containing all behaviors.

For our first example, we take a state machine that nondeterministically generates an integer between 0 and 9. To do this, the machine counts up by one until it either decides to stop or else reaches 9, at which point it stutters forever. The set Σ_E of externally visible states is the set $\mathbf{N}$ of natural numbers, and the internal state component is a Boolean that becomes true when the final value is reached. (The Boolean values are written $\mathbf{t}$ and $\mathbf{f}$.)

- $\Sigma = \mathbf{N} \times \{\mathbf{t}, \mathbf{f}\}$.
- $F = \{(0, \mathbf{f})\}$.
- N is the union of the following two sets:
  - $\{\langle (i-1, \mathbf{f}), (i, \mathbf{f})\rangle : 0 < i < 10\}$,
  - $\{\langle (i, \mathbf{f}), (i, \mathbf{t})\rangle : i \in \mathbf{N}\}$.

The set of stutter-free behaviors generated by this state machine consists of all sequences of the forms

$$\langle\!\langle (0, \mathbf{f}), (1, \mathbf{f}), \ldots, (n, \mathbf{f}), (n, \mathbf{t}), (n, \mathbf{t}), (n, \mathbf{t}), \ldots \rangle\!\rangle$$

and

$$\langle\!\langle (0, \mathbf{f}), (1, \mathbf{f}), \ldots, (n, \mathbf{f}), (n, \mathbf{f}), (n, \mathbf{f}), \ldots \rangle\!\rangle$$

with $0 \leq n < 10$.

We now add a prophecy variable whose value is a natural number. This variable "predicts" the maximum number of nonstuttering steps that the state machine will take. The precise definition of the new state machine is:

- Σ^p is the union of the following two sets:
  - $\{(i, \mathbf{f}, j) : 0 \leq i,\ 0 \leq j, \text{ and } i + j < 10\}$,
  - $\{(i, \mathbf{t}, 0) : 0 \leq i < 10\}$.
- $F^p = \{(0, \mathbf{f}, j) \in \Sigma^p\}$.
- N^p is the union of the following two sets:
  - $\{\langle (i-1, \mathbf{f}, j+1), (i, \mathbf{f}, j)\rangle \in \Sigma^p \times \Sigma^p\}$,
  - $\{\langle (i, \mathbf{f}, 0), (i, \mathbf{t}, 0)\rangle \in \Sigma^p \times \Sigma^p\}$.

The reader can check that the conditions P1–P4′ and P6 given above are satisfied. (Condition P5 is satisfied if L and L^p are the trivial properties that contain all behaviors.) Observe that although condition P4′ is satisfied, condition H4 is not. The state machine can take a backwards step from the state $(6, \mathbf{f}, 0)$ but not a forward step.

The only stutter-free behaviors of (Σ^p, F^p, N^p) starting from the state $(0, \mathbf{f}, n)$ are of the forms

$$\langle\!\langle (0, \mathbf{f}, n), (1, \mathbf{f}, n-1), \ldots, (n, \mathbf{f}, 0), (n, \mathbf{t}, 0), (n, \mathbf{t}, 0), \ldots\rangle\!\rangle$$

and

$$\langle\!\langle (0, \mathbf{f}, n), (1, \mathbf{f}, n-1), \ldots, (i, \mathbf{f}, n-i), (i, \mathbf{f}, n-i), \ldots\rangle\!\rangle$$

with $0 \leq i \leq n$. The set of externally visible behaviors generated by the two state machines is the same; the stutter-free behaviors have the form $\langle\!\langle 0, 1, \ldots, n, n, n, \ldots\rangle\!\rangle$ for some n less than 10. State machine (Σ, F, N) decides nondeterministically when it is going to stop counting, while in state machine (Σ^p, F^p, N^p) this choice is made by the initial value of the prophecy variable.

As our second example, replace "10" by "∞" in the definitions of the two state machines. Conditions P1–P4′ still hold, but P6 does not; for each state $(i, \mathbf{f})$ of Σ there are an infinite number of states $(i, \mathbf{f}, j)$ in Σ^p. The externally visible stutter-free behaviors of (Σ^p, F^p, N^p) consist of sequences of the form $\langle\!\langle 0, 1, \ldots, n, n, n, \ldots\rangle\!\rangle$ for any natural number n. The state machine (Σ, F, N) generates all these behaviors plus the additional behavior $\langle\!\langle 0, 1, 2, 3, \ldots\rangle\!\rangle$ that never terminates. Because the finiteness condition P6 is not satisfied, adding the auxiliary variable changed the specification by ruling out this nonterminating behavior—effectively adding a liveness condition.

We can use our last example to indicate why we need the hypothesis of finite invisible nondeterminism for our completeness theorem. Let $\mathbf{S_2}$ be the specification consisting of the state machine (Σ^p, F^p, N^p) we just constructed (the one with "10" replaced by "∞") and the trivial supplementary property containing all Σ^p-behaviors. Let $\mathbf{S_1}$ be the specification with state machine (Σ, F, N) and supplementary property L consisting of all terminating behaviors. Both specifications define the same set of externally visible behaviors—all behaviors obtainable by stuttering from ones of the form $\langle\langle 0, 1, \ldots, n, n, n \rangle\rangle$. To construct a refinement mapping, we would have to add to $\mathbf{S_1}$ a prophecy variable that "guesses" the value of the last component of a state of Σ^p. However, no such prophecy variable can be constructed that satisfies P6, since for any starting state of $\mathbf{S_1}$ there are an infinite number of corresponding starting states of $\mathbf{S_2}$.

The complete property P_2 defined by this specification $\mathbf{S_2}$ is a safety property, and we will see that this implies that $\mathbf{S_2}$ is internally continuous. Moreover, specification $\mathbf{S_1}$ is machine closed. Nevertheless, adding auxiliary variables to $\mathbf{S_1}$ will not allow us to construct a refinement mapping to prove that it implements $\mathbf{S_2}$. Our completeness theorem does not apply because P_2 is not fin.

In this example, the prophecy variable we wanted to add would not satisfy P6. However, the supplementary property happened to ensure that adding the prophecy variable did not change the externally visible behavior. If we were to replace P6 by the weaker requirement that $\mathbf{S^P}$ have the same externally visible property as $\mathbf{S}$, then we could find a refinement mapping. However, this requirement is precisely what we had to prove in the first place—namely, that $\mathbf{S_1}$ implements $\mathbf{S_2}$.

5.3 Prophecy Variables That Add Stuttering

We now generalize our definition of a prophecy variable to allow it to introduce stuttering. Condition P2′ asserts that a state $(s, p) \in \Sigma^p$ is an initial state of $\mathbf{S^P}$'s state machine iff s is an initial state of $\mathbf{S}$'s state machine. We relax this condition by requiring only that such a state (s, p) be reachable from an initial state by steps that simulate stuttering steps. Formally, we replace P2′ by:

P2. (a) $\Pi_{[P]}(F^p) \subseteq F$.

(b) For all $(s, p) \in \Pi^{-1}_{[P]}(F)$ there exist $p_0, p_1, \ldots, p_n = p$ such that $(s, p_0) \in F^p$ and, for $0 \leq i < n$, $\langle (s, p_i), (s, p_{i+1}) \rangle \in N^p$.

Similarly, we relax condition P4′ by allowing $\mathbf{S^P}$'s state machine to simulate the step in $\mathbf{S}$'s state machine from state s to state s' by a sequence of $n+1$ steps, the last n of which simulate stuttering steps. The precise condition that replaces P4′ is:

P4. If $\langle s, s' \rangle \in N$ and $(s', p') \in \Sigma^p$ then there exist $p, p'_0, \ldots, p'_{n-1}, p'_n = p'$ such that $\langle (s, p), (s', p'_0) \rangle \in N^p$ and, for $0 \leq i < n$, $\langle (s', p'_i), (s', p'_{i+1}) \rangle \in N^p$.

As with history variables, the addition of prophecy variables leaves an implementation essentially unchanged.

Proposition 5 (Soundness of prophecy variables) *If $\mathbf{S}^p$ is obtained from $\mathbf{S}$ by adding a prophecy variable, then the two specifications define the same externally visible property.*

Proof of Proposition 5

Given: A1. $\mathbf{S} = (\Sigma, F, N, L)$, $\mathbf{S^P} = (\Sigma^p, F^p, N^p, L^p)$, and P1–P6 hold.
A2. M and M^p are the machine properties of $\mathbf{S}$ and $\mathbf{S^P}$, respectively.
A3. $P = M \cap L$ and $P^p = M^p \cap L^p$.
A4. $O = \Gamma(\Pi_E(P))$ and $O^p = \Gamma(\Pi_E(P^p))$.

Prove: C1. $O^p \subseteq O$.
C2. $O \subseteq O^p$.

Pf: 1. C1 holds.

Pf: The proof is identical to the proof of the corresponding condition for history variables in Proposition 4.

2. $P \subseteq \Pi_{[P]}(P^p)$.

Given: A2.1. $\sigma = \langle\langle s_0, s_1, \ldots \rangle\rangle \in P$.

Prove: C2.1. There exists $\tau \in P^p$ such that $\Pi_{[P]}(\tau) \simeq \sigma$.

Pf: 2.1. Let $\mathcal{G}$ be the directed graph with

Nodes: the set $\Sigma^p \times \mathbf{N}$.

Edges: there is an edge between $((s_i, p), i)$ and $((s_j, p'), j)$ iff $j = i + 1$ and either $(s_i, p) = (s_{i+1}, p')$ or there exist $p_0, p_1, \ldots, p_n = p'$ in Σ_P such that $\langle (s_i, p), (s_{i+1}, p_0) \rangle \in N^p$ and, for all $0 \leq k < n$, $\langle (s_{i+1}, p_k), (s_{i+1}, p_{k+1}) \rangle \in N^p$.

Let $\mathcal{G}'$ be the subgraph of $\mathcal{G}$ reachable from nodes of the form $((s_0, p), 0)$. Then $\mathcal{G}'$ is acyclic, with finite branching and a finite set of sources.

Pf: It is obviously acyclic, since there is an edge from $((s, p), i)$ to $((s', p'), i')$ only if $i' = i + 1$. Its sources are all the nodes

of the form $((s_0, p), 0)$. For each j, P6 implies that there are only a finite set of p such that $(s_j, p) \in \Sigma^p$, so $\mathcal{G}'$ has a finite set of sources and is finitely branching.

2.2. For all $n \geq 0$ and all $(s_n, p_n) \in \Sigma^p$ there exist elements $p_0, \ldots, p_{n-1}$ in Σ_P such that $\langle\!\langle((s_0, p_0), 0), \ldots, ((s_n, p_n), n)\rangle\!\rangle$ is a path in $\mathcal{G}'$.

Pf: The proof is by induction on n. The case $n = 0$ is trivial. For $n > 0$, condition P4 implies the existence of the required p_{n-1}, and the induction hypothesis provides $p_0, \ldots, p_{n-2}$.

2.3. Choose elements $p_i \in \Sigma_P$ such that $\langle\!\langle((s_0, p_0), 0), ((s_1, p_1), 1), \ldots\rangle\!\rangle$ is an infinite path in $\mathcal{G}'$.

Pf: The existence of this path follows from 2.1, 2.2, and König's Lemma.

2.4. Let $\rho = \langle\!\langle(s_0, p_0), \ldots, (s_i, p_i), \ldots\rangle\!\rangle$. Choose a sequence $\rho' = \langle\!\langle(s'_0, p'_0), \ldots, (s'_i, p'_i), \ldots\rangle\!\rangle$ such that:

2.4a. $\Pi_{[P]}(\rho') \simeq \sigma$.

2.4b. For all $i \geq 0$: $\langle(s'_i, p'_i), (s'_{i+1}, p'_{i+1})\rangle \in N^p$ or $(s'_i, p'_i) = (s'_{i+1}, p'_{i+1})$.

2.4c. $(s'_0, p'_0) = (s_0, p_0)$.

Pf: Let ρ' be the supersequence of ρ obtained by inserting between (s_i, p_i) and (s_{i+1}, p_{i+1}) the sequence $\langle\!\langle(s_{i+1}, p'_0), (s_{i+1}, p'_1), \ldots, (s_{i+1}, p'_{k-1})\rangle\!\rangle$ of elements in Σ^p whose existence is guaranteed by 2.3 and the definition of edges in $\mathcal{G}'$ (2.1). (Recall that $\sigma = \langle\!\langle s_0, s_1, \ldots\rangle\!\rangle$.)

2.5. Choose $\tau = \langle\!\langle(t_0, q_0), (t_1, q_1), \ldots\rangle\!\rangle$ such that:

2.5a. $\Pi_{[P]}(\tau) \simeq \sigma$.

2.5b. For all $i \geq 0$: $\langle(t_i, q_i), (t_{i+1}, q_{i+1})\rangle \in N^p$ or $(t_i, q_i) = (t_{i+1}, q_{i+1})$.

2.5c. $(t_0, q_0) \in F^p$.

Pf: By A2.1, we have $s_0 \in F$. By P2, there exists a finite sequence $\langle\!\langle(s_0, p''_0), \ldots, (s_0, p''_n)\rangle\!\rangle$ of elements in Σ^p such that $(s_0, p''_0) \in F^p$, each $\langle(s_0, p''_i), (s_0, p''_{i+1})\rangle \in N^p$, and $p''_n = p_0$. Let $\tau = \langle\!\langle(s_0, p''_0), \ldots, (s_0, p''_{n-1})\rangle\!\rangle \cdot \rho'$.

2.6. $\tau \in M^p$.

Pf: By A2, 2.5b, and 2.5c.

2.7. $\tau \in P^p$.

Pf: By A3, 2.6, and P5.

2.8. C2.1 holds.

Pf: 2.7, 2.5a.

3. C2 holds.

   Pf: From 2, A1–A4, and the fact that $\Pi_E(\Pi_{[P]}(t)) = \Pi_E(t)$ for any $t \in \Sigma^p$.

End Proof of Proposition 5

6 Internal Continuity

We now define internal continuity, which appears in the third hypothesis of our main theorem. But first, we give an example that indicates why the hypothesis is needed for our completeness theorem.

Let $\Sigma_E = \mathbf{N}$, let η_i be the terminating sequence $\langle\langle 0, 1, \ldots, i, i, i, \ldots\rangle\rangle$, and let η be the nonterminating sequence $\langle\langle 0, 1, 2, \ldots\rangle\rangle$. Let $\langle\langle e_0, e_1, \ldots\rangle\rangle \times x$ denote the sequence $\langle\langle (e_0, x), (e_1, x), \ldots\rangle\rangle$. We construct a specification $\mathbf{S_2}$ that defines the property whose stutter-free sequences consist of all sequences $\eta_i \times \mathbf{t}$ together with the sequence $\eta \times \mathbf{f}$. Formally, $\mathbf{S_2} = (\Sigma_2, F_2, N_2, L_2)$, where

- $\Sigma_2 = \mathbf{N} \times \{\mathbf{t}, \mathbf{f}\}$. (The internal component is a Boolean.)
- $F_2 = \{(0, \mathbf{t}), (0, \mathbf{f})\}$. (Behaviors start with their visible components equal to 0.)
- $N_2 = \{\langle (i, b), (i+1, b)\rangle\}$. (The external component is incremented by 1 and the internal component remains constant.)
- L_2 consists of all behaviors *except* ones of the form $\sigma \times \mathbf{f}$ with σ terminating, and $\sigma \times \mathbf{t}$ with σ nonterminating.

The externally visible property O_2 defined by $\mathbf{S_2}$ consists of the behaviors η_i, the behavior η, and all behaviors obtained from them by stuttering. Specification $\mathbf{S_2}$ is fin and machine closed.

The externally visible property O_2 is also defined by the simpler specification $\mathbf{S_1} = (\Sigma_1, F_1, N_1, L_1)$, where

- $\Sigma_1 = \Sigma_E = \mathbf{N}$. (There is no internal component.)
- $F_1 = \{0\}$. (All behaviors start at 0.)
- $N_1 = \{\langle i, i+1\rangle\}$. (The state is incremented by 1.)
- $L_1 = \Sigma_1^\omega$ (the trivial property that allows all behaviors).

Obviously, $\mathbf{S_1}$ implements $\mathbf{S_2}$. Let $\mathbf{S_1^p} = (\Sigma_1^p, F_1^p, N_1^p, L_1^p)$ be any specification obtained from $\mathbf{S_1}$ by adding a prophecy variable. We now show that there does not exist a refinement mapping from $\mathbf{S_1^p}$ to $\mathbf{S_2}$; in fact there does not exist any mapping from Σ_1^p to Σ_2 that proves that $\mathbf{S_1^p}$ implements $\mathbf{S_2}$.

Let P_1^p be the property defined by $\mathbf{S_1^p}$. We show by contradiction that there does not exist any mapping $f : \Sigma_1^p \rightarrow \Sigma_2$ such that (i) $\Pi_E(f(i,p)) = i$ and (ii) $f(P_1^p) \subseteq P_2$. For each i let $\eta_i' \in P_1^p$ be a behavior with $\Pi_{[P]}(\eta_i') \simeq \eta_i$. Moreover, P5 implies that we can choose η_i' to have no repeated nonfinal states, meaning that for $j < i$ and $k > 1$, there is no segment $\langle\langle (j,p_1), (j,p_2), \ldots, (j,p_k) \rangle\rangle$ of η_i' with $p_1 = p_k$. By (i), we then have that for every i and m with $i < m$ there is an l such that $\Pi_E(\eta_m'|_l) \simeq \eta_i|_{i+1}$. Moreover, P6 and the absence of repeated nonfinal states imply that for each i there is an integer $\pi(i) > i$ such that $l \leq \pi(i)$ for all such m. We can choose π so that $\pi(i+1) \geq \pi(i)$ for all i.

For any n, the set $\{\eta_j'|_{\pi(n)}\}$ is finite (by P6). Therefore, we can inductively construct the sequence θ_n of length $\pi(n)$ such that θ_n is a prefix of infinitely many of the η_j' and is also a prefix of θ_{n+1}. Let $\eta' = \lim \theta_n$; then $\Pi_E(\eta') \simeq \eta$. Since each θ_n is a prefix of some η_j', clearly η' is in the machine property of $\mathbf{S_1^p}$. Property P5 then implies that $\eta' \in P_1^p$. By definition of η_i', assumption (ii) implies that $f(\eta_i') \simeq \eta_i \times \mathbf{t}$, which implies that $f(\eta') \simeq \eta \times \mathbf{t}$. We then have $\eta' \in P_1^p$ and $f(\eta') \notin P_2$, which contradicts assumption (ii).

This proof can be extended to the case where $\mathbf{S_1}$ is replaced by any specification $\mathbf{S_1^h}$ obtained from it by adding a history variable. We just replace η with any behavior allowed by $\mathbf{S_1^h}$ that simulates it, and replace η_i with an initial prefix of this new η. Thus, first adding a history variable still does not allow one to construct the refinement mapping.

The problem with specification $\mathbf{S_2}$ is that $\eta \times \mathbf{t}$ is not in P_2 even though $\Pi_E(\eta \times \mathbf{t})$ is in O_2 and any finite portion of $\eta \times \mathbf{t}$ is the same as the corresponding portion of some behavior $\eta_i \times \mathbf{t}$ in P_2. The sequence $\eta \times \mathbf{t}$ is not in P_2 even though we cannot tell that it isn't by looking either at its externally visible component or at any finite part of the complete behavior. To rule out this possibility, we must add to our completeness theorem the hypothesis that P_2 is *internally continuous*.

Definition 2 *A Σ-property P with induced externally visible property O is* internally continuous *iff, for any Σ-behavior σ, if $\Pi_E(\sigma) \in O$ and $\sigma \in \overline{P}$, then $\sigma \in P$. A specification is internally continuous iff the (complete) property it defines is internally continuous.*

Suppose $P = M \cap L$ and $M = \overline{P}$. Then $\lim \sigma_i = \sigma$ for $\sigma_i \in P$ iff $\sigma \in M$. It follows from this that, for a machine-closed specification, internal continuity is equivalent to the condition that a complete behavior is allowed iff it is generated by the state machine and its externally visible component is allowed. In particular, safety properties are internally continuous.

Since the machine property M is closed, if $\lim \sigma_i = \sigma$ for $\sigma_i \in M \cap L$, then $\sigma \in L$ iff $\sigma \in M \cap L$. This implies that if L is internally continuous, then $M \cap L$ is internally continuous. Hence, for any specification, if the supplementary property is internally continuous, then the specification is internally continuous. The converse is not true, since if M is the empty property, then $M \cap L$ is internally continuous for any L.

Any specification can be made internally continuous by adding to L all sequences σ in M such that $\Pi_E(\sigma) \in O$. Expanding L in this way obviously adds no new externally visible behaviors, so the resulting specification is equivalent to the original one. The expansion could introduce infinite internal nondeterminism, but not if M is fin.

7 The Completeness Theorem

We can now prove our main result.

Theorem 2 (Completeness) *If the machine-closed specification* $\mathbf{S_1}$ *implements the internally continuous, fin specification* $\mathbf{S_2}$*, then there is a specification* $\mathbf{S_1^h}$ *obtained from* $\mathbf{S_1}$ *by adding a history variable and a specification* $\mathbf{S_1^{hp}}$ *obtained from* $\mathbf{S_1^h}$ *by adding a prophecy variable such that there exists a refinement mapping from* $\mathbf{S_1^{hp}}$ *to* $\mathbf{S_2}$*.*

Proof of Theorem 2

Given: A1. For $i = 1, 2$: $\mathbf{S_i} = (\Sigma_i, F_i, N_i, L_i)$, M_i is the machine property of $\mathbf{S_i}$, $P_i = M_i \cap L_i$, and $O_i = \Gamma(\Pi_E(P_i))$.

A2. $O_1 \subseteq O_2$.

A3. $\mathbf{S_1}$ is machine closed.

A4. $\mathbf{S_2}$ is fin.

A5. $\mathbf{S_2}$ is internally continuous.

Prove: C1. There exist specifications $\mathbf{S_1^h}$ and $\mathbf{S_1^{hp}}$ such that:

C1a. $\mathbf{S_1^h}$ is obtained from $\mathbf{S_1}$ by adding a history variable.

C1b. $\mathbf{S_1^{hp}}$ is obtained from $\mathbf{S_1^h}$ by adding a prophecy variable.

C1c. There exists a refinement mapping f from $\mathbf{S_1^{hp}}$ to $\mathbf{S_2}$.

Pf: 1. Let $\mathbf{S_1^h}$ equal $(\Sigma_1^h, F_1^h, N_1^h, L_1^h)$, where

- $\Sigma_1^h = \{(last(\sigma|_n), \sigma|_n) : n > 0 \text{ and } \sigma \in P_1\}$. (The history component h of any state (s, h) is a finite prefix of a behavior in P_1 that ends in state s.)
- $F_1^h = \{(s, h) \in \Sigma_1^h : \|h\| = 1\}$.
- $N_1^h = \{\langle (s, h), (s', h') \rangle \in \Sigma_1^h \times \Sigma_1^h : h' = h \cdot \langle\!\langle s' \rangle\!\rangle\}$. (A step of $\mathbf{S_1^h}$'s state machine simulates a step of $\mathbf{S_1}$'s state machine and adds the new state to the history component.)
- $L_1^h = \Pi_{[H]}^{-1}(L_1)$. (As required by H5.)

Then C1a holds.

Pf: 1.1. H1, H3, and H5 hold.

Pf: Follows immediately from the definition of $\mathbf{S_1^h}$.

1.2. $\Pi_{[H]}(F_1^h) \subseteq F_1$.

Pf: Immediate from the definition of F_1^h.

1.3. $F_1 \subseteq \Pi_{[H]}(F_1^h)$

Pf: For any $s \in F_1$, the sequence $\langle\!\langle s, s, s, \ldots \rangle\!\rangle \in M_1$. Therefore, A3 and Lemma 1 imply that $\langle\!\langle s \rangle\!\rangle$ is a prefix of a behavior in P_1, so $(s, \langle\!\langle s \rangle\!\rangle) \in F_1^h$ and $s = \Pi_{[H]}((s, \langle\!\langle s \rangle\!\rangle))$.

1.4. H2 holds.

Pf: 1.2 and 1.3.

1.5. H4 holds.

Pf: For any $\langle s, s' \rangle \in N_1$ and $(s, h) \in \Sigma_1^h$, let $h' = h \cdot \langle\!\langle s' \rangle\!\rangle$. Then A3 and Lemma 1 imply that h' is the prefix of a behavior in P_1, so $(s', h') \in \Sigma_1^h$ by definition of Σ_1^h, and $\langle (s, h), (s', h') \rangle \in N_1^h$ by definition of N_1^h.

2. Let $\mathbf{S_1^{hp}}$ equal $(\Sigma_1^{hp}, F_1^{hp}, N_1^{hp}, L_1^{hp})$, where

- Σ_1^{hp} equals the set of triples $(s, h, \natural(\sigma|_m))$ with $(s, h) \in \Sigma_1^h$, $\sigma \in P_2$, $m > 0$, and $\Pi_E(\sigma|_m) \simeq \Pi_E(h)$, where we write (s, h, p) instead of $((s, h), p)$. (The prophecy component p of (s, h, p) is an initial stutter-free prefix of a behavior in P_2 such that p and h are externally equivalent.)
- $F_1^{hp} = \{(s, h, p) \in \Sigma^{hp} : (s, h) \in F_1^h \text{ and } \|p\| = 1\}$. (Note that this implies $s \in F_1$ and $p = \langle\!\langle t \rangle\!\rangle$ with $t \in F_2$.)
- N_1^{hp} is the set of pairs $\langle (s, h, p), (s', h', p') \rangle$ in $\Sigma_1^{hp} \times \Sigma_1^{hp}$ such that either

  (a) $p' = p \cdot \langle\!\langle last(p') \rangle\!\rangle$ and either $\langle (s, h), (s', h') \rangle \in N_1^h$ or $(s, h) = (s', h')$, or

(b) $p' = p$ and $\langle (s,h),(s',h')\rangle \in N_1^h$.

(A step of $\mathbf{S_1^{hp}}$'s state machine either increases the length of the prophecy component by one and simulates a [possibly stuttering] step of $\mathbf{S_1^h}$'s state machine, or else leaves the prophecy component unchanged and simulates a nonstuttering step of $\mathbf{S_1^h}$'s state machine.)

- $L_1^h = \Pi_{[H]}^{-1}(L_1)$. (As required by P5.)

Then C1b holds.

Pf: 2.1. P1, P3, and P5 hold.

Pf: Immediate from the definition of $\mathbf{S_1^{hp}}$.

2.2. $\Pi_{[P]}(F_1^{hp}) \subseteq F_1^h$.

Pf: Immediate from the definitions of F_1^{hp} and F_1^h.

2.3. For all $(s,h,p) \in \Pi_{[P]}^{-1}(F_1^h)$ there exist $p_0, p_1, \ldots, p_n = p$ such that $(s,h,p_0) \in F_1^{hp}$ and, for $0 \le i < n$, $\langle (s,h,p_i), (s,h,p_{i+1})\rangle \in N_1^{hp}$.

Pf: 2.3.1. Let $(s,h,p) \in \Pi_{[P]}^{-1}(F_1^h)$, and let $p = \langle\!\langle t_0, t_1, \ldots, t_n\rangle\!\rangle$. Then $h = \langle\!\langle s\rangle\!\rangle$ and $\Pi_E(p) \simeq \Pi_E(\langle\!\langle s\rangle\!\rangle)$.

Pf: By definitions of F_1^h and Σ_1^{hp}.

2.3.2. Let $p_i = \langle\!\langle t_0, \ldots, t_i\rangle\!\rangle$. Then $\Pi_E(p_i) \simeq \Pi_E(h)$.

Pf: By 2.3.1.

2.3.3. $(s,h,p_0) \in F_1^{hp}$ and $\langle (s,h,p_i), (s,h,p_{i+1})\rangle \in N_1^{hp}$ for $0 \le i < n$.

Pf: By 2.3.2 and the definitions of F_1^{hp} and N_1^{hp}.

2.4. P2 holds.

Pf: By 2.2 and 2.3.

2.5. P4 holds.

Given: A2.5.1. $\langle (s,h), (s',h')\rangle \in N_1^h$ and $(s',h',p') \in \Sigma_1^{hp}$.

Prove: C2.5.1. There exist p, p'_0, ..., p'_{n-1}, $p'_n = p'$ in Σ_P such that $\langle (s,h,p), (s',h',p'_0)\rangle \in N_1^{hp}$ and, for $0 \le i < n$, $\langle (s',h',p'_i), (s',h',p'_{i+1})\rangle \in N_1^{hp}$.

Pf: 2.5.1. $p' = \natural(\sigma|_m)$ for some $\sigma \in P_2$, and $\Pi_E(p') \simeq \Pi_E(h')$.

Pf: By A2.5.1 and the definition of Σ_1^{hp}.

2.5.2. $h' = h \cdot \langle\!\langle s'\rangle\!\rangle$.

Pf: By A2.5.1 ($\langle (s,h),(s',h')\rangle \in N_1^h$) and the definition of N_1^h.

2.5.3. Let p be the longest prefix of p' such that $\Pi_E(p) \simeq \Pi_E(h)$.

Pf: The existence of p follows from 2.5.1 and 2.5.2.

2.5.4. $p' = p \cdot \langle\!\langle t_0, \ldots, t_n\rangle\!\rangle$ where $\Pi_E(t_i) \simeq \Pi_E(s')$ for $0 \le i \le n$.

Pf: By 2.5.3, 2.5.1, and 2.5.2.

2.5.5. Let $p'_i = p \cdot \langle\!\langle t_0, \ldots, t_i \rangle\!\rangle$. Then $(s', h', p'_i) \in \Sigma_1^{hp}$ for $0 \leq i \leq n$.

Pf: By 2.5.4 and 2.5.1, we have $\Pi_E(p'_i) \simeq \Pi_E(h')$. The result then follows from the definition of Σ_1^{hp}, since A2.5.1 implies $(s', h') \in \Sigma_1^h$.

2.5.6. C2.5.1 holds.

Pf: Follows easily from 2.5.5, 2.5.1, and the definition of N_1^{hp}.

2.6. P6 holds.

Given: A2.6.1. $(s, h) \in \Sigma_1^h$.

Prove: C2.6.1. $\{p : (s, h, p) \in \Sigma_1^{hp}\}$ is finite.

C2.6.2. There exists $p \in \Sigma_P$ such that $(s, h, p) \in \Sigma_1^{hp}$.

Pf: 2.6.1. Choose $\psi \in P_1$ such that $h = \psi|_n$, and let $\eta = \Pi_E(\psi)$.

Pf: ψ exists by A2.6.1 and the definition of Σ_1^h.

2.6.2. C2.6.1 holds.

Pf: By definition of Σ_1^{hp} and η (in 2.6.1), A4, and Definition 1.

2.6.3. Choose $\sigma \in P_2$ such that $\Pi_E(\sigma) \simeq \eta$.

Pf: Such a σ exists since $\eta \in O_1$ (by 2.6.1) and $O_1 \subseteq O_2$ (by A2).

2.6.4. C2.6.2 holds.

Pf: By 2.6.3 and the definition of η (in 2.6.1), we can choose m such that $\Pi_E(\sigma|_m) \simeq \Pi_E(h)$. Let $p = \natural(\sigma|_m)$. The definition of Σ_1^{hp} implies that $(s, h, p) \in \Sigma_1^{hp}$.

3. Define $f : \Sigma_1^{hp} \to \Sigma_2$ by $f((s, h, p)) = \mathit{last}(p)$. Then f is a refinement mapping.

Pf: 3.1. f satisfies R1.

Pf: By definition of Σ_1^{hp}, if $(s, h, p) \in \Sigma_1^{hp}$ then $(s, h) \in \Sigma_1^h$ and $\Pi_E(p) \simeq \Pi_E(h)$. But $(s, h) \in \Sigma_1^h$ implies $s = \mathit{last}(h)$ (by definition of Σ_1^h), so $\Pi_E(p) \simeq \Pi_E(h)$ implies $\Pi_E(s) = \Pi_E(\mathit{last}(p))$.

3.2. f satisfies R2.

Pf: By definition of F_1^{hp}, its elements are of the form $(s, \langle\!\langle s \rangle\!\rangle, \langle\!\langle t \rangle\!\rangle)$ where $t \in F_2$, so $f((s, \langle\!\langle s \rangle\!\rangle, \langle\!\langle t \rangle\!\rangle)) = t \in F_2$.

3.3. f satisfies R3.

Given: A3.3.1. $\langle (s, h, p), (s', h', p') \rangle \in N_1^{hp}$.

Prove: C3.3.1. $\langle last(p), last(p')\rangle \in N_2$ or $last(p) = last(p')$.

Pf: By definition of N_1^{hp}, A3.3.1 implies $p' = \langle\!\langle t_0, \ldots, t_n \rangle\!\rangle$ for some infinite sequence $\langle\!\langle t_0, t_1, \ldots \rangle\!\rangle \in P_2$, and either $p = p'$, in which case C3.3.1 is immediate, or $p = \langle\!\langle t_0, \ldots, t_{n-1} \rangle\!\rangle$. In the latter case, we must prove $\langle t_{n-1}, t_n\rangle \in N_2$. However, this follows immediately from the fact that $\langle\!\langle t_0, t_1, \ldots \rangle\!\rangle \in P_2 \subseteq M_2$ and the definition of the machine property of a specification.

3.4. f satisfies R4.

Given: A3.4.1. $\tau = \langle\!\langle (s_0, h_0, p_0), (s_1, h_1, p_1), \ldots \rangle\!\rangle \in P_1^{hp}$.

Prove: C3.4.1. $f(\tau) = \langle\!\langle last(p_0), last(p_1), \ldots \rangle\!\rangle \in L_2$.

Pf: 3.4.1. Let $\sigma = \langle\!\langle s_0, s_1, \ldots \rangle\!\rangle$. Then $\Pi_E(\sigma) = \Pi_E(\tau)$.

Pf: Follows immediately from R1 (by 3.1).

3.4.2. $\Pi_E(\sigma) \in O_1$.

Pf: C1a (proved in 1), C1b (proved in 2), and Propositions 4 and 5 imply that $\Pi_E(\tau) \in O_1$, so 3.4.2 follows from 3.4.1.

3.4.3. For all $n \geq 0$, $f(\tau)|_n \simeq p_n$.

Pf: By A3.4.1, $\langle (s_i, h_i, p_i), (s_{i+1}, h_{i+1}, p_{i+1})\rangle \in N_1^{hp}$ or $(s_i, h_i, p_i) = (s_{i+1}, h_{i+1}, p_{i+1})$ for all $i \geq 0$. By definition of N_1^{hp}, this implies $p_{i+1} = p_i$ or $p_{i+1} = p_i \cdot \langle\!\langle last(p_i) \rangle\!\rangle$ for all i. A simple induction proof then shows that $p_n \simeq \langle\!\langle last(p_0), \ldots, last(p_n) \rangle\!\rangle$.

3.4.4. For all $n \geq 0$ there exists $\psi_n \in P_2$ such that $\psi_n|_n = f(\tau)|_n$.

Pf: By definition of Σ_1^{hp}, there exists a sequence ϕ_n such that $p_n \cdot \phi_n \in P_2$. Let $\psi_n = f(\tau)|_n \cdot \phi_n$. By 3.4.3, $\psi_n \simeq p_n \cdot \phi_n$, so ψ_n is in P_2.

3.4.5. C3.4.1 holds.

Pf: 3.4.4 implies that $\lim \psi_n = f(\tau)$ and $\psi_n \in P_2$. By 3.4.1, 3.4.2, R1 (proved in 3.1), and A2, we have $\Pi_E(f(\tau)) \in O_2$. Since $\mathbf{S_2}$ is internally continuous (by A5) and the ψ_n are in P_2 (by 3.4.4), Definition 2 implies that $f(\tau) \in P_2$. This proves C3.4.1, since $P_2 \subseteq L_2$ (by A1).

End Proof of Theorem 2

The converse of this completeness theorem is not true. For instance, no matter how pathological a specification is, we can use the identity refinement mapping to prove that it implements itself.

The hypotheses of the internal continuity and finite invisible nondeterminism of $\mathbf{S_2}$ can be removed from our completeness theorem by generalizing the definition of a prophecy variable—namely, by replacing condition P6 with the explicit requirement that the externally visible behaviors of $\mathbf{S^p}$ be the same as those of $\mathbf{S}$. This result is proved by defining $\mathbf{S_1^h}$ as in the proof of Theorem 2, and defining $\mathbf{S_1^{hp}}$ such that

- Σ_1^{hp} is the set of 4-tuples (s, h, n, τ) with $(s, h) \in \Sigma_1^h$, $\tau \in P_2$, and $\Pi_E(h) \simeq \Pi_E(\tau|_n)$.
- F_1^{hp} is the set of all states of the form $(s, h, 1, \tau)$.
- N_1^{hp} is the set of pairs $\langle (s, h, n, \tau), (s', h', n+1, \tau) \rangle$ with either $\langle (s, h), (s', h') \rangle \in N_1^h$ or $(s, h) = (s', h')$.
- The refinement mapping is defined by letting $f((s, h, n, \tau))$ be the n^{th} element of τ.

However, the condition that replaces P6 asserts that specification $\mathbf{S^p}$ implements $\mathbf{S}$, which is precisely the type of condition we are trying to prove in the first place. This generalization of Theorem 2 is therefore of little practical value, so we will not bother to state it and prove it formally.

There is one simple way to strengthen the completeness theorem that is of some interest. The specification $\mathbf{S_2}$ is fin and internally continuous iff the property P_2 that it defines is fin and internally continuous. We can weaken the hypothesis by requiring only that there exist a fin and internally continuous property P_2' contained in P_2 that induces the same externally-visible property as P_2. Let $\mathbf{S_2'}$ be the specification obtained from $\mathbf{S_2}$ by replacing L_2 with $L_2 \cap P_2'$. The correctness of this result follows easily from Theorem 2 by replacing $\mathbf{S_2}$ with $\mathbf{S_2'}$.

8 Whence and Whither?

Refinement mappings are not new. They form the basis of the methods advocated by Lam and Shankar [LS84] and by us [Lam83], and they are used by Lynch and Tuttle [LT87] to prove that one automaton implements another. However, none of this work addresses the issue of completeness. Jonsson [Jon87] and Stark [Sta88] did prove completeness results similar to ours, but for smaller classes of specifications.

Complete methods for checking that a program implements a specification, without constructing refinement mappings, have been developed. Some of the most general are those of Alpern and Schneider [AS87], Manna and Pnueli [MP87], and Vardi [Var87]. Their methods differ from our approach in at least two important ways:

- They do not consider behaviors with different amounts of "stuttering" to be equivalent, so their definition of what constitutes a correct implementation is weaker than ours.

- They require constructing the negation of specifications. In practice, the negation of a specification may be hard to find and hard to understand.

Because of these differences, the methods may not offer practical alternatives to the use of refinement mappings for proving correctness.

Our exposition has been purely semantic. We have considered specifications, but not the languages in which they are expressed. We proved the existence of refinement mappings, but said nothing about whether they are expressible in any language. We do not know what languages can describe the necessary auxiliary variables and resulting refinement mappings.

Our results also raise the question of what properties can be described by specifications that are fin and internally continuous. If the specification language is expressive enough, then all properties can be defined by specifications without internal state, which are trivially fin and internally continuous. At the other extreme, one can easily invent artificially impoverished languages that do not allow any fin or internally continuous specifications. The question becomes interesting only for interesting specification languages, such as various forms of temporal logic. In addition, recall that the hypotheses of our completeness theorem can be weakened by requiring only that $\mathbf{S_2}$'s complete property be equivalent to a fin and internally continuous subproperty. This raises the more general question of what expressible properties have equivalent fin and continuous subproperties.

Acknowledgements

The first example we saw that demonstrated the inadequacy of history variables is due to Herlihy and Wing [HW87]. The introduction of prophecy variables was based on a suggestion of Jim Saxe, who also provided many useful suggestions for improving the exposition. We wish to thank Pierre

Wolper for making clear whence our ideas came and Gordon Plotkin for making clear that whither they will lead is not an easy question, since he couldn't answer it on the spot.

References

[AS86] Bowen Alpern and Fred B. Schneider. *Recognizing Safety and Liveness*. Technical Report TR86-727, Department of Computer Science, Cornell University, January 1986.

[AS87] Bowen Alpern and Fred Schneider. Proving boolean combinations of deterministic properties. In *Proceedings of the Second Symposium on Logic in Computer Science*, pages 131–137, IEEE, June 1987.

[HW87] M.P. Herlihy and J.M. Wing. Axioms for concurrent objects. In *Proceedings of the Fourteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 13–26, ACM, Munich, January 1987.

[Jon87] Bengt Jonsson. *Compositional Verification of Distributed Systems*. PhD thesis, Uppsala University, 1987.

[Knu73] Donald E. Knuth. *Fundamental Algorithms*. Volume 1 of *The Art of Computer Programming*, Addison-Wesley, Reading, Massachusetts, second edition, 1973.

[Lam77] Leslie Lamport. Proving the correctness of multiprocess programs. *IEEE Transactions on Software Engineering*, SE-3(2):125–143, March 1977.

[Lam83] Leslie Lamport. Specifying concurrent program modules. *ACM Transactions on Programming Languages and Systems*, 5(2):190–222, April 1983.

[LS84] Simon S. Lam and A. Udaya Shankar. Protocol verification via projections. *IEEE Transactions on Software Engineering*, SE-10(4):325–342, July 1984.

[LT87] Nancy Lynch and Mark Tuttle. Hierarchical correctness proofs for distributed algorithms. In *Proceedings of the Sixth Symposium on the Principles of Distributed Computing*, pages 137–151, ACM, August 1987.

[MP87] Zohar Manna and Amir Pnueli. Specification and verification of concurrent programs by ∀-automata. In *Proceedings of the Four-*

teenth Symposium on the Principles of Programming Languages, pages 1–12, ACM, January 1987.

[Owi75] S. Owicki. *Axiomatic Proof Techniques for Parallel Programs*. PhD thesis, Cornell University, August 1975.

[Sta88] Eugene W. Stark. Proving entailment between conceptual state specifications. *Theoretical Computer Science*, 56(1):135–154, January 1988.

[Var87] Moshe Vardi. Verification of concurrent programs: the automata-theoretic framework. In *Proceedings of the Second Symposium on Logic in Computer Science*, pages 167–176, IEEE, June 1987.

Glossary/Index of Notations and Conventions

e The externally visible component of a state, 7

f A refinement mapping, 11

h A history variable, 20

p A prophecy variable, 22

s, t States, 7

y, z Internal components of states, 7

F The set of initial states of a state machine, 9

L A supplementary property, 9

M A property, usually generated by a state machine, 9

N The next-state relation of a state machine, 9

O An externally visible property, 8

P A complete property, 8

S A set—typically a set of sequences, 7

S A specification, 9

δ, η, θ Externally visible behaviors, 8

ρ, σ, τ, ψ Sequences, usually representing complete behaviors, 8

$\Gamma\sigma$ The set of all behaviors equivalent to σ up to stuttering, 7

ΓS The set of all behaviors equivalent to behaviors in S up to stuttering, 7

Π_E The projection from states onto their external components, 7

$\Pi_{[X]}$ The projection from states that eliminates the X component, 20

Σ A set of states, 7

Σ_E A set of externally visible states, 7

Σ_I A set of internal states, 7

Σ^ω The set of all infinite sequences of elements of Σ, 7

$\langle\!\langle s_1, s_2, \ldots \rangle\!\rangle$ The sequence whose first element is s_1, whose second element is s_2, etc., 1

$\natural\sigma$ The stutter-free form of σ, 6

$\simeq$ Equivalence of sequences up to stuttering, 7

$\sigma \cdot \tau$ The concatenation of the sequences σ and τ, 7

$\sigma|_m$ The prefix of sequence σ of length m, 7

$\langle s, t \rangle$ A pair of states that is an element of the next-state relation of a state machine, 9

$\overline{S}$ The closure of the set S, 7