

# Trace-Based Network Proof Systems: Expressiveness and Completeness

JENNIFER WIDOM  
IBM Almaden Research Center

and

DAVID GRIES and FRED. B. SCHNEIDER  
Cornell University

---

We consider incomplete trace-based network proof systems for safety properties, identifying extensions that are necessary and sufficient to achieve relative completeness. We investigate the expressiveness required of any trace logic to encode these extensions.

Categories and Subject Descriptors: D.1.3 [Programming Techniques]: Concurrent Programming; D.2.4 [Software Engineering]: Program Verification—*correctness proofs*; D.3.1 [Programming Languages]: Formal Definitions and Theory—*semantics*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—*proof theory*

General Terms: Verification

Additional Key Words and Phrases: Process networks, safety properties, temporal logics, trace logics

---

## 1. INTRODUCTION

In *trace-based network proof systems*, one specifies and reasons about traces (histories) of the values transmitted along the communication channels of the networks under consideration. Such proof systems generally allow specifications for networks to be deduced from specifications for the networks'

---

This paper is based on portions of J. Widom's Cornell University Ph.D. thesis. The work was supported in part by the National Science Foundation under grants DCR-8320274 and CCR-8701103. F. B. Schneider is also supported by the Office of Naval Research under contract N00014-86-K-0092 and by Digital Equipment Corporation. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not reflect the views of these agencies.

Authors' addresses: J. Widom, IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120, D. Gries and F. B. Schneider, Computer Science Department, Cornell University, Ithaca, NY 14853

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1992 ACM 0164-0925/92/0700-0396 \$01.50

ACM Transactions on Programming Languages and Systems, Vol. 14, No. 3, July 1992, Pages 396-416

components. Network proof systems based on first-order predicates over channel traces are given in [4], [7], and [12], but unfortunately, these logics are known to be incomplete [2, 8, 13]. Relative completeness can be achieved by permitting reasoning over the full interleaving of communication events in addition to the individual channel traces [3, 6, 9] or by using temporal logic formulas so that the interleaving is implicitly present in the semantics of the specifications [14]. Both approaches introduce more information than necessary, the modifications tend to be extensive and cumbersome, and the simplicity of the underlying trace logic is lost.

In this paper we analyze the sources of incompleteness in trace-based network proof systems and formalize the extensions needed to achieve relative completeness. A simple but logically incomplete proof system for safety properties is defined. It is representative of other incomplete trace-based network proof systems appearing in the literature. Two example networks are presented to show incompleteness of our logic. Surprisingly, both examples consist of only one process, indicating that, although compositionality is an important feature of trace-based systems, composition is not the cause of incompleteness. Our examples suggest two extensions that are necessary for relative completeness; we show that these two extensions are sufficient as well. The first source of incompleteness is the inability to state and reason about constraints on the ordering of network events. The second source is the inability to assert that the sequence of values transmitted along a communication channel is a prefix of that channel's sequence at each later point. We argue that these two properties—the temporal ordering and prefix properties—must be part of any relatively complete trace-based proof system.

The temporal ordering and prefix properties cannot be expressed in a logic based on first-order predicates over channel traces. However, they can be expressed using only a subset of temporal logic, indicating that the full power of temporal logic (or of explicit reasoning over the interleaving of communication events) is not needed. We introduce a hierarchy of subsets of temporal logic and show that a subset consisting of first-order predicates over traces with a version of the linear-time temporal *Always* operator has necessary and sufficient expressive power for relative completeness.

Section 2 defines the class of synchronous process networks we consider and introduces a formal model of computation. Section 3 defines Simple Network Logic (SNL), a trace-based network proof system that captures the essence of most such systems. In Section 4 we show that SNL is incomplete, define the temporal ordering and prefix axioms, and prove that they are necessary and sufficient for relative completeness. In Section 5 we define a hierarchy of temporal logic subsets and identify a subset of temporal logic with exactly the right expressive power for relative completeness. Finally, in Section 6 we summarize, discuss alternatives and extensions to our work, and relate our work to previous research.

## 2. PROCESS NETWORKS

Consider networks of processes that communicate and synchronize solely by message passing. Processes and communication channels are uniquely named.

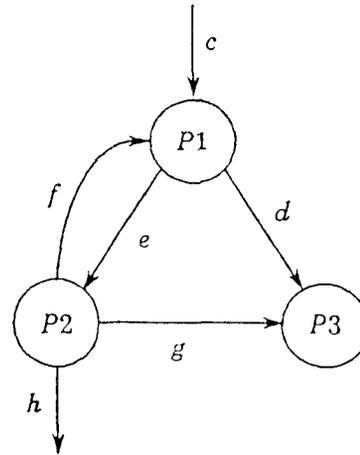


Fig. 1. Network of processes

Each channel is either *internal* or *external* with respect to a network. An internal channel connects two processes of the network; an external channel is connected to only one process, permitting communication between the network and its environment. Channels are unidirectional, and communication along them is synchronous,<sup>1</sup> so both processes incident to an internal channel must be prepared to communicate before a value is actually transmitted. Without loss of generality, we assume that

- input or output on an external channel occurs whenever the single incident process is ready;
- each message transmission occurs instantaneously;
- two message transmissions cannot occur simultaneously, thus there is a total order on the communication events of a given computation; and
- there is a fixed domain of values that can be transmitted on communication channels, and processes send and receive values in this domain only.

A network made up of processes  $P_1, P_2, \dots, P_n$  is denoted by  $P_1 \parallel P_2 \parallel \dots \parallel P_n$  indicating parallel execution of  $P_1, P_2, \dots, P_n$ . Figure 1 illustrates a network of three processes and six communication channels. Subsequently, we use the term *network* to refer to either a single process or a network composed of several processes.<sup>2</sup>

## 2.1 Model of Computation

To reason about proof systems for networks, we introduce a formal model of network behavior. A single point, or state, during the computation of a

<sup>1</sup>Extension to asynchronous message passing is straightforward and does not affect our results [16].

<sup>2</sup>Although we assume here that networks are composed directly of primitive processes, our work easily generalizes to a hierarchical structure in which processes may be implemented by subnetworks [16].

network is represented by the histories of the network's communication channels up to that point. (The model does not include internal process state, since internal computations of processes are unimportant when reasoning about network behavior, except as they affect values sent and received.) A single computation is represented by an infinite sequence of states,<sup>3</sup> subject to constraints on the initial state and between each pair of adjacent states. The behavior of a network is represented by the set of all sequences representing its possible computations.

We now formalize the model. A *trace* of a communication channel  $c$  is a finite sequence  $\langle c.1, c.2, \dots, c.k \rangle$  of values that have been transmitted along channel  $c$  up to some point in time. An empty trace is denoted by  $\langle \rangle$ . Let  $N$  be a network with incident channels  $c_1, c_2, \dots, c_m$ . A *state* of  $N$  is a tuple containing a trace for each channel  $c_1, c_2, \dots, c_m$ . A *computation* of  $N$  is an infinite sequence of states of  $N$  such that all channel traces are empty in the initial state and each subsequent state extends at most one trace of the preceding state by at most one element (i.e., at most one message transmission occurs between each pair of states). Thus, an infinite sequence of states is a computation iff the following four computation conditions are satisfied:

- (CC1) All traces are empty in the initial state.
- (CC2) Each trace in each state of the sequence is a prefix of the corresponding trace in the subsequent state.
- (CC3) The length of each trace in each state of the sequence (except the first) is at most one more than the length of the corresponding trace in the preceding state.
- (CC4) At most one trace changes between every state and its subsequent state.

The *computation set* for  $N$  is the set of all computations representing potential execution sequences of  $N$ .<sup>4</sup>

We show that our computational model is compositional; that is, the set representing the behavior of a network can be constructed from the sets representing the behavior of the network's component processes. Let  $N$  be a network composed of  $P_1, \dots, P_n$ . Let  $CS(P_1), \dots, CS(P_n)$  denote the computation sets for  $P_1, \dots, P_n$ , respectively.  $CS(N)$ , the computation set for  $N$ , can be computed directly from  $CS(P_1), \dots, CS(P_n)$  as follows: Suppose  $\kappa$  is a computation of  $N$ . Define  $Proj(\kappa, P_i)$  to be the projection of  $\kappa$  onto those channels of  $N$  incident to  $P_i$ ; that is,  $Proj(\kappa, P_i)$  removes from the states of  $\kappa$  all traces of channels in  $N$  not incident to  $P_i$ .  $CS(N)$  is the set of all

<sup>3</sup>We choose only infinite sequences since, for our purposes, any finite sequence can be converted into an equivalent infinite sequence by indefinitely repeating the final state [10, 16].

<sup>4</sup>Our definition allows arbitrarily many (adjacent) repetitions of any state in a computation sequence, producing a very large computation set for a network. In particular, if a given computation is in a computation set, then so is that computation with any state repeated once, repeated twice, repeated three times, etc. Permitting this repetition facilitates subsequent definitions and proofs without affecting the usefulness of the model.

computations  $\kappa$  built from possible states of  $N$  such that

- $\text{Proj}(\kappa, P_i) \in \text{CS}(P_i)$ ,  $i = 1 \dots n$ , and
- $\kappa$  satisfies conditions (CC1)–(CC4) above.

Informally, then, the computations of  $N$  are all possible well-formed combinations of the computations of  $N$ 's component processes in which communications on shared channels agree.

### 3. SIMPLE NETWORK LOGIC

We now introduce a formalism, *Simple Network Logic (SNL)*, for specifying and verifying safety properties [9] of networks of processes. SNL concisely captures the essence of most trace-based proof systems.

A *specification* is a first-order predicate over channel traces; it is intended to be satisfied throughout every execution of the network it specifies [4, 7, 12].<sup>5</sup> A channel name  $c$  appearing as a free variable in a specification represents the trace of  $c$ . We use  $|c|$  to denote the length of  $c$ , and  $c1 \subseteq c2$  to denote that trace  $c1$  is a prefix of trace  $c2$ .

A specification for a network  $N$  is a predicate  $\phi$  over the traces of  $N$ 's communication channels. (The only free variables permitted in specifications for  $N$  are those corresponding to channels of  $N$ .) We say that  $N$  *satisfies*  $\phi$ , written  $N \text{ sat } \phi$ , if, at every point during any computation of  $N$ , the traces of values transmitted on  $N$ 's channels satisfy  $\phi$ . For example, consider a process (or network)  $N$  that repeatedly reads an integer from channel  $c1$  and sends its successor on channel  $c2$ . This behavior can be specified in SNL as

$$N \text{ sat } (|c1| - 1 \leq |c2| \leq |c1|) \wedge (\forall i: 1 \leq i \leq |c2|: c2.i = c1.i + 1).$$

Some network proof systems provide facilities for constructing specifications for sequential primitive processes [4, 7, 19]. In others, it is assumed that existing logics for sequential programs are used for this, or that specifications for primitive processes are given [8, 12, 14]. Without loss of generality, we adopt the latter approach. Thus, we assume that some programming language and proof system are used initially to construct primitive processes (following our communication model) and to verify trace logic specifications for the behavior of each. We are then interested only in deducing specifications for a network from specifications for its component processes. Hence, let the axioms of SNL be formulas  $P \text{ sat } \phi$ , where  $P$  is a primitive process and  $\phi$  is a valid specification for  $P$ ; that is,  $\phi$  is satisfied by every execution of  $P$ . We assume that included is at least one such axiom (but not necessarily all) for each process  $P$  of interest.

<sup>5</sup>Actually, in [12], specifications consist of predicate pairs, but for our purposes it is adequate to consider only single predicates that remain invariant throughout a computation [16].

Specifications for networks can be derived from specifications for their component processes using the following inference rule of SNL:

*Definition 3.1. Network Composition Rule*

$$\frac{P_1 \mathbf{sat} \phi_1, P_2 \mathbf{sat} \phi_2, \dots, P_n \mathbf{sat} \phi_n}{P_1 \parallel P_2 \parallel \dots \parallel P_n \mathbf{sat} \bigwedge_i \phi_i}$$

Conjoining process specifications using this rule results in “linking” any shared channels in network  $N$ , because in  $\bigwedge_i \phi_i$  all  $c$ 's, say, denote the same channel trace.

We also need a rule for deducing valid specifications for a network from other valid specifications, since several valid specifications may exist for a given network. For this, we rely on the validity of formulas in our underlying trace logic. If  $\phi$  is a specification over channel traces and  $S$  is a state (a mapping from channel variables to trace values), let  $S \models \phi$  denote that formula  $\phi$  is satisfied by state  $S$ . If  $\phi$  is satisfied by every state, we write  $\models \phi$ . The SNL Consequence Rule is then as follows:

*Definition 3.2. Consequence Rule*

$$\frac{N \mathbf{sat} \phi_1, \models \phi_1 \Rightarrow \phi_2}{N \mathbf{sat} \phi_2}$$

The use of  $\models \phi_1 \Rightarrow \phi_2$  in this rule, expressing that every state that satisfies  $\phi_1$  also satisfies  $\phi_2$ , is further discussed below.

These two rules, or variants thereof, form the basis of most trace-based systems we know of, including those in [4], [7], [12], [14], and [18].

### 3.1 Soundness and Completeness

We can use the computational model introduced in Section 2.1 to prove that SNL is *sound* [1]: If  $N \mathbf{sat} \phi$  is a theorem of SNL, then  $\phi$  is indeed valid for network  $N$ . To do so, we must first define validity of specifications with respect to computations in the formal model. Recall that for a network  $N$ ,  $CS(N)$  denotes the set of all computations corresponding to  $N$ 's possible behaviors. A specification  $\phi$  is valid for  $N$  iff  $\phi$  is true in every state of every computation in  $N$ 's computation set.

*Definition 3.3. Validity of Specifications.* Specification  $\phi$  is valid for network  $N$  iff  $\kappa.i \models \phi$  for all  $\kappa \in CS(N)$  and  $i \geq 0$ , where  $\kappa = \langle \kappa.0, \kappa.1, \kappa.2, \dots \rangle$ .<sup>6</sup>

Using this definition we establish the soundness of SNL.

**THEOREM 3.4. Soundness of SNL.** *Let  $N$  be a network and  $\phi$  a specification such that  $N \mathbf{sat} \phi$  is a theorem of SNL. Then  $\phi$  is valid for  $N$ .*

<sup>6</sup>Note that indexing of computations begins with 0, while indexing of channel traces begins with 1. This notation facilitates subsequent definitions and proofs.

PROOF. See the Appendix.  $\square$

We would also like SNL to be *complete*: If any specification  $\phi$  is valid for a network  $N$ , then  $N \text{ sat } \phi$  is provable in SNL. However, an additional assumption must be made. A specification for a network is derived from specifications for its component processes using the Network Composition Rule. If the given process specifications are valid but too weak, then a valid network specification might not be provable. What we really want to know is whether  $N \text{ sat } \phi$  can be proved when the given specifications for  $N$ 's component processes are as “strong” as possible [14].

*Definition 3.5. Precise Specifications.* A specification  $\phi$  is *precise* for a network  $N$  iff

- (1)  $\phi$  is valid for  $N$ ; and
- (2) if  $\kappa$  is any computation<sup>7</sup> containing traces for the channels in  $N$  and  $\kappa.j \models \phi$  for all  $j \geq 0$ , then  $\kappa \in CS(N)$ .

Thus, a specification is precise if it is valid and if every computation satisfying the specification is a possible computation of the network being specified. For completeness, we are interested in proving  $N \text{ sat } \phi$  whenever  $\phi$  is valid and whenever the specifications for the processes in  $N$  are precise.

The SNL Consequence Rule relies on the validity of formulas in the specification language of SNL. SNL specifications can involve elements of the data domain from which messages are drawn, sequences of such elements (the channel traces), and lengths of sequences. Since arithmetic itself is incomplete [15], when designing a program proof system one aims for *relative completeness* (as in [5]): Assuming that one can prove any valid formula of the underlying logic—which in this case is a trace logic that includes predicate logic, arithmetic, and the data domain of the network being considered—is the proof system complete?<sup>8</sup> SNL is not relatively complete, as we now show.

#### 4. INCOMPLETENESS OF SIMPLE NETWORK LOGIC

Two examples are given to show the incompleteness of SNL. Each example illustrates an inherent property of network behavior that cannot be expressed in SNL, but that is necessary for relative completeness in a trace-based proof system.

##### 4.1 Temporal Ordering Property

Consider the single-process network of Figure 2. As an informal description of process  $P$ , we are given four facts: (1)  $P$  reads at most one value from

<sup>7</sup>Recall from Section 2.1 that a computation is an infinite sequence of states such that the initial state contains only empty channel traces and each subsequent state extends at most one trace of the preceding state by at most one element.

<sup>8</sup>Most proof systems make assumptions about both the provability of statements in the underlying logic and the expressiveness of the specification language involved. This is sometimes referred to as *Cook completeness* [1, 5]. We, too, have made an expressiveness assumption in our supposition that precise specifications for primitive processes can be written in SNL. See Section 6 for further discussion.

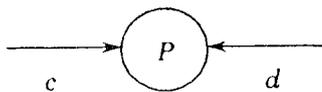


Fig. 2. Simple network.

channel  $c$ , (2)  $P$  reads at most one value from channel  $d$ , (3)  $P$  reads a value from  $c$  before reading from  $d$ , and (4)  $P$  reads a value from  $d$  before reading from  $c$ . By direct translation, a formal specification is

$$P \text{ sat } \phi_1: |c| \leq 1 \wedge |d| \leq 1 \wedge |d| \leq |c| \wedge |c| \leq |d|. \quad (1)$$

Let the data domain for this network be  $\{a\}$ . The following specification is valid for  $P$  and is equivalent to (1):

$$P \text{ sat } \phi_2: (c = \langle \rangle \wedge d = \langle \rangle) \wedge (c = \langle a \rangle \wedge d = \langle a \rangle). \quad (2)$$

$P$  is always in one of two states: Either no values have been read from  $c$  or  $d$ , or a value  $a$  has been read from each. However, since two values cannot be transmitted simultaneously (condition (CC4) from Section 2.1),  $P$  can reach a state in which  $(c = \langle a \rangle \wedge d = \langle a \rangle)$  only by being in a state in which either  $(c = \langle a \rangle \wedge d = \langle \rangle)$  or  $(c = \langle \rangle \wedge d = \langle a \rangle)$ , neither of which is permitted by specification (2) (or specification (1)). Thus,  $P$  will never read a value from either  $c$  or  $d$ , so a third valid and equivalent specification for  $P$  is

$$P \text{ sat } \phi_3: c = \langle \rangle \wedge d = \langle \rangle. \quad (3)$$

All three specifications are valid and, in fact, precise. Any computation satisfying  $\phi_1$ ,  $\phi_2$ , or  $\phi_3$  is a computation of  $P$ : no values are ever read on  $c$  or  $d$ . However, consider an attempt at proving (3) given, say, precise specification  $\phi_2$  of (2). Since there is only a single process, the Network Composition Rule is irrelevant; the only SNL inference rule applicable is the Consequence Rule. But  $\models \phi_2 \Rightarrow \phi_3$  does not hold. Hence, (3) is unprovable from (2), even though it is valid.

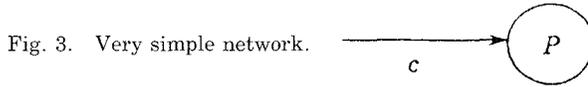
We need a way to formalize the reasoning about event ordering used to obtain specification (3). It must assert the following property:

*Definition 4.1. Temporal Ordering Property.* Suppose  $c1$  and  $c2$  are channels of a network  $N$ ,  $c1.x$  and  $c2.y$  are transmitted as a result of distinct communication events, and in any computation of  $N$

- (1)  $c1.x$  is transmitted before  $c2.y$ , and
- (2)  $c2.y$  is transmitted before  $c1.x$ .

Then  $(|c1| < x \wedge |c2| < y)$  holds throughout any computation of  $N$ , i.e., neither message is ever transmitted.

Formalizing this property would allow  $\phi_3$  to be deduced from  $\phi_2$ , making (3) provable from (2). Unfortunately, the Temporal Ordering Property cannot be expressed in the trace logic underlying SNL. This is discussed further in Section 4.4.1 and is proved in Section 5.



#### 4.2 Prefix Property

Consider the network of Figure 3. Suppose  $\{a, b\}$  is the data domain, and let a precise specification for process  $P$  be

$$P \text{ sat } \phi_4: c = \langle \rangle \vee c = \langle a \rangle \vee c = \langle b, a \rangle. \quad (4)$$

Since  $P$  sends one value at a time on channel  $c$  (condition (CC3) from Section 2.1), disjunct  $c = \langle b, a \rangle$  can never hold. (It would hold only if  $c = \langle b \rangle$  held first, and this is prohibited by  $\phi_4$ .) Therefore, (4) can be simplified to

$$P \text{ sat } \phi_5: c \subseteq \langle a \rangle. \quad (5)$$

However,  $\phi_4$  does not imply  $\phi_5$ , so (5) cannot be proved from precise specification (4). Here we need the following property:

*Definition 4.2. Prefix Property.* For any channel  $c$  and integers  $0 \leq x \leq y$ , the trace of  $c$  after  $x$  values have been transmitted is always a prefix of the trace of  $c$  after  $y$  values have been transmitted.

By using this property in conjunction with  $\phi_4$ , we could eliminate disjunct  $c = \langle b, a \rangle$  and obtain (5). Like the Temporal Ordering Property, however, the Prefix Property cannot be expressed in the trace logic underlying SNL.

#### 4.3 Incorporating the Properties in SNL

Consider any SNL proof that establishes  $N \text{ sat } \phi$  for a network  $N = P_1 \parallel \dots \parallel P_n$ . We assume that as axioms we are given  $P_1 \text{ sat } \phi_1, \dots, P_n \text{ sat } \phi_n$ , where  $\phi_1, \dots, \phi_n$  are precise. The first rule to be applied in any such proof is the Network Composition Rule, so we immediately infer  $N \text{ sat } \bigwedge_i \phi_i$ . All remaining steps in the proof have to be applications of the Consequence Rule. By transitivity of implication, any string of Consequence Rule applications can be collapsed into one, so  $N \text{ sat } \phi$  is provable iff  $\models \bigwedge_i \phi_i \Rightarrow \phi$ .

By the following theorem, we know that specification  $\bigwedge_i \phi_i$  is precise for  $N$ : the conjunction of precise process specifications results in a network specification that is also precise. (A similar theorem for temporal logic is proved in [14].)

**THEOREM 4.3. Preciseness Preservation.** *Let  $\phi_i$  be a precise specification for  $P_i$ ,  $1 \leq i \leq n$ , and let  $N = P_1 \parallel \dots \parallel P_n$ . Then  $\bigwedge_i \phi_i$  is a precise specification for  $N$ .*

**PROOF.** See the Appendix.  $\square$

Thus, our proof system would be relatively complete if  $\models \phi_1 \Rightarrow \phi_2$  whenever  $\phi_1$  is a precise specification for a network  $N$  and  $\phi_2$  is a valid specification for  $N$ . However, the examples given in Sections 4.1 and 4.2 illustrate that this implication does not always hold.

To obtain a relatively complete system, the implication in the hypothesis of the Consequence Rule must be modified so that all valid specifications can be deduced from precise specifications. We do this by strengthening the antecedent of the implication, adding a set of axioms such that if, say,  $A$  is the conjunction of axioms in the set, then  $\models (\phi_1 \wedge A) \Rightarrow \phi_2$  whenever  $\phi_1$  and  $\phi_2$  are precise and valid, respectively, for a given network. The Temporal Ordering and Prefix Properties are the basis for such a set of axioms.

#### 4.4 Temporal Ordering and Prefix Axioms

We now prove that axiomatizations of the Temporal Ordering and Prefix Properties are necessary and sufficient for deducing  $\phi_2$  from  $\phi_1$  whenever  $\phi_1$  and  $\phi_2$  are precise and valid, respectively, for a given network. There is a fundamental difference, however, between any axiomatization of the Temporal Ordering (or Prefix) Property and specifications  $\phi_1$  and  $\phi_2$ : Event ordering is considered with respect to an entire computation, whereas  $\phi_1$  and  $\phi_2$  are considered with respect to the individual states of a computation. Since  $(\phi_1 \wedge A) \Rightarrow \phi_2$  must be considered with respect to computations, we resolve the discrepancy by introducing an operator  $\square$ , which converts single-state specifications to operate over sequences of states:  $\square \phi$  is valid for a sequence of states  $\kappa$ , denoted by  $\kappa \models \square \phi$ , iff  $\phi$  is true in every state of the sequence.<sup>9</sup> That is,

$$\kappa \models \square \phi \quad \text{iff} \quad \kappa.i \models \phi \text{ for all } i \geq 0.$$

(Note that by Definition 3.3 of validity of specifications, specification  $\phi$  is thus valid for a network  $N$  iff  $\kappa \models \square \phi$  for every  $\kappa$  in the computation set of  $N$ .) Using  $\square$  and  $\models$ , we modify the Consequence Rule as follows:

*Definition 4.4. Modified Consequence Rule*

$$\frac{N \text{ sat } \phi_1, \models (\square \phi_1 \wedge A) \Rightarrow \square \phi_2}{N \text{ sat } \phi_2}$$

Notice that this rule requires that the underlying logic is extended to include the operator  $\square$  and that formula  $(\square \phi_1 \wedge A) \Rightarrow \square \phi_2$  is interpreted over sequences of states rather than over single states. That is,  $\models (\square \phi_1 \wedge A) \Rightarrow \square \phi_2$  holds iff every sequence of states satisfying  $(\square \phi_1 \wedge A)$  also satisfies  $\square \phi_2$ . Next, we consider the axioms comprising  $A$ .

**4.4.1 Temporal Ordering Axiom.** The Temporal Ordering Axiom will formalize the Temporal Ordering Property. Suppose some communication  $c1.x$  must happen before some  $c2.y$ . Then  $\square(|c1| < x \Rightarrow |c2| < y)$ . This assertion captures ordering of communication events for any channels  $c1$  and  $c2$  and any indexes  $x$  and  $y$ , even if  $x = y$  or  $c1$  and  $c2$  are the same channel. We are interested only in the ordering of distinct events, so the case in which

<sup>9</sup>This is a weakened version of the “always,” or “henceforth,” operator (also  $\square$ ) of temporal logic [10, 11], since here  $\phi$  cannot contain other temporal operators. Temporal operators of varying strengths are discussed in Sections 4.4.2 and 5.

$c1.x$  and  $c2.y$  are produced by the same event (i.e.,  $x = y$  and  $c1$  and  $c2$  are the same channel) is not of interest. Now, if  $\Box(|c2| < y \Rightarrow < y \Rightarrow |c1| < x)$  holds as well (with  $x \neq y$  or  $c1$  and  $c2$  distinct), then neither  $c1.x$  nor  $c2.y$  can ever occur; equivalently,  $\Box(|c1| < x \wedge |c2| < y)$ . Hence, we state the Temporal Ordering Axiom as follows:

*Definition 4.5. ORDERING.* If  $c1$  and  $c2$  are channels,  $x \geq 1$  and  $y \geq 0$  are integers, and either  $x \neq y$  or  $c1$  and  $c2$  are distinct; then<sup>10</sup>

$$\Box(|c1| < x \Leftrightarrow |c2| < y) \Rightarrow \Box(|c1| < x \wedge |c2| < y).$$

Allowing  $y = 0$  permits the assertion that an empty channel trace temporally precedes all communications events on that channel.<sup>11</sup> We disallow  $x = y = 0$ , however, since this results in a pathological situation in which the antecedent of the implication is trivially true but the consequent is trivially false.

We can now prove soundness of *ORDERING* with respect to our computational model.

**THEOREM 4.6. Soundness of ORDERING.** *If  $\kappa$  is a computation (recall Section 2.1) then  $\kappa \models \text{ORDERING}$ .*

**PROOF.** See the Appendix.  $\square$

**4.4.2 Prefix Axiom.** To formulate an axiom for the Prefix Property, we introduce a more powerful version of  $\Box$  in which  $\Box$  may be applied to formulas that themselves contain  $\Box$ 's. (This is the usual linear-time temporal logic interpretation for  $\Box$  [11].) Now,  $\Box\phi$  is valid for a sequence of states iff  $\phi$  is valid for every suffix of that sequence:

$$\kappa \models \Box\phi \quad \text{iff} \quad \langle \kappa.i, \kappa.(i+1), \kappa.(i+2), \dots \rangle \models \phi \text{ for all } i \geq 0.$$

When  $\phi$  contains no  $\Box$  operators,  $\langle \kappa.i, \kappa.(i+1), \kappa.(i+2), \dots \rangle \models \phi$  is usually interpreted to be true iff  $\phi$  is true in the first state, that is, iff  $\kappa.i \models \phi$ . (For more detailed and rigorous discussions of the semantics of temporal operators, see, e.g., [11] and [16].)

The Prefix Axiom can be stated using  $\Box$  as follows:

*Definition 4.7. PREFIX.* If  $c$  is a channel,  $x \geq 1$  is an integer, and  $v$  is a value in the data domain of messages, then

$$\Box(c.x = v \Rightarrow \Box(c.x = v)).$$

<sup>10</sup>Technically, this is an axiom scheme rather than an axiom, since substitution for metasymbols  $c1$ ,  $c2$ ,  $x$ , and  $y$  is permitted.

<sup>11</sup>Suppose, for the sake of contradiction, that  $c$  is nonempty in the initial state of some computation satisfying *ORDERING*. Then  $\Box(|c| < x \Leftrightarrow |c| < 0)$  for some  $x \geq 1$ . However,  $\Box(|c| \geq 0)$ , so the conclusion of *ORDERING* does not hold.

This axiom (scheme) asserts that, once a value has been transmitted as  $c.x$ ,  $c.x$  remains unchanged. This is equivalent to the Prefix Property as stated in Section 4.2.<sup>12</sup>

**THEOREM 4.8.** Soundness of *PREFIX*. *If  $\kappa$  is a computation, then  $\kappa \models \text{PREFIX}$ .*

**PROOF.** Let  $\kappa$  be any computation. Then  $\kappa$  satisfies condition (CC2) (Section 2.1), and *PREFIX* follows directly.  $\square$

**4.4.3 Necessity and Sufficiency of ORDERING and PREFIX.** By letting  $A$  be *ORDERING*  $\wedge$  *PREFIX*, we can prove that, if  $\phi_1$  is a precise specification for a network  $N$  and  $\phi_2$  is a valid specification for  $N$ , then  $\models (\Box \phi_1 \wedge A) \Rightarrow \phi_2$  (from the hypothesis of Modified Consequence Rule 4.4) holds. Thus, *ORDERING* and *PREFIX* are sufficient for achieving relative completeness. In addition, we will argue that *ORDERING* and *PREFIX* are necessary: If either axiom is removed from  $A$ , then there is a network  $N$  with precise and valid specifications  $\phi_1$  and  $\phi_2$ , respectively, such that  $\Box \phi_1$  and  $A$  do not imply  $\Box \phi_2$ .

We begin with a key lemma.

**LEMMA 4.9.** *Let  $\kappa$  be any infinite sequence of states.  $\kappa$  represents a computation iff  $\kappa \models \text{ORDERING} \wedge \text{PREFIX}$ .*

**PROOF.** See the Appendix.  $\square$

With this lemma in hand, we can easily prove that our two axioms are sufficient for relative completeness.

**THEOREM 4.10.** Sufficiency of the Axioms. *If  $\phi_1$  is a precise specification for network  $N$  and  $\phi_2$  is a valid specification for  $N$ , then  $\models (\Box \phi_1 \wedge \text{ORDERING} \wedge \text{PREFIX}) \Rightarrow \Box \phi_2$ .*

**PROOF.** We show that any infinite sequence of states  $\kappa$  satisfying  $\Box \phi_1$ , *ORDERING*, and *PREFIX* also satisfies  $\Box \phi_2$ . Since  $\kappa \models \text{ORDERING} \wedge \text{PREFIX}$ , by Lemma 4.9 we know that  $\kappa$  is a computation. By Definition 3.5 of preciseness, since  $\kappa \models \Box \phi_1$  and  $\phi_1$  is precise,  $\kappa \in CS(N)$ . By validity of  $\phi_2$ , every  $\kappa \in CS(N)$  satisfies  $\Box \phi_2$ . Hence,  $\kappa$  satisfies  $\Box \phi_2$ .  $\square$

Thus, with *ORDERING* and *PREFIX*, we ensure that any valid network specification is implied by a precise specification for the network; by Precision-Preservation Theorem 4.3, a precise network specification is obtainable from precise specifications for the network's component process. That

<sup>12</sup>A different axiomatization of the Prefix Property can be given using the “next” operator of temporal logic in addition to  $\Box$  [17]. The definition given here, however, shows that the Prefix Property can be encoded using only  $\Box$ . This is of importance in Section 5, where we consider the minimal expressiveness required of any trace logic used as the basis of a relatively complete proof system.

*ORDERING* and *PREFIX* are necessary (as well as sufficient) for the implication to always hold is shown in the following theorem:

**THEOREM 4.11. Necessity of the Axioms.** *There exist networks  $N1$ ,  $N2$ , and  $N3$ , with precise specifications  $\phi_1^P$ ,  $\phi_2^P$ , and  $\phi_3^P$ , respectively, and valid specifications  $\phi_1^V$ ,  $\phi_2^V$ , and  $\phi_3^V$ , respectively, such that*

- (1)  $\not\models (\Box \phi_1^P \wedge \text{ORDERING}) \Rightarrow \Box \phi_1^V$
- (2)  $\not\models (\Box \phi_2^P \wedge \text{PREFIX}) \Rightarrow \Box \phi_2^V$ , and
- (3)  $\not\models (\Box \phi_3^P \Rightarrow \Box \phi_3^V)$ .

**PROOF.** For (1), let  $N1$  be the example network of Section 4.2. For (2), let  $N2$  be the example network of Section 4.1. Fact (3) follows directly from (1) and (2).  $\square$

## 5. STRENGTHENING THE PROOF SYSTEM

We have demonstrated that an axiomatization of the Temporal Ordering and Prefix Properties is necessary and sufficient for relative completeness of SNL. However,  $\text{ORDERING} \wedge \text{PREFIX}$  is not the only way to formalize the properties of computation that must be encoded in the logic. From the proof of Theorem 4.10, we see that the function of *ORDERING* and *PREFIX* is to characterize legal network computations. Thus, we are interested in the expressiveness required of a logic in order to encode the notion of a legal network computation.

We formalize this expressiveness requirement by analyzing the relationship between valid and precise specifications. By Definition 3.3 of validity, a specification is valid for a network  $N$  iff it is satisfied by every state reachable by  $N$ . By Definition 3.5 of preciseness, a state is reachable by  $N$  iff it is reachable by a computation that always satisfies a precise specification for  $N$ . Hence, we observe the following:

*Observation 5.1.* A specification is valid for a network iff it is satisfied by every state reachable by a computation that always satisfies a precise specification for that network.

Formalizing this observation results in a relatively complete proof system. One such formalization is implication  $(\Box \phi_1 \wedge A) \Rightarrow \Box \phi_2$  of Modified Consequence Rule 4.4. We establish bounds on the expressive power required of any trace logic that formalizes Observation 5.1.

Let  $\phi$  range over trace logic formulas (i.e., over formulas in the specification language of SNL). Suppose  $K(\phi)$  is a formula in some logic  $L$  such that, for any  $\phi$ , a state satisfies  $K(\phi)$  iff the state is reachable by a computation that always satisfies  $\phi$ . Consider the following Generalized Consequence Rule, where  $K(\phi_1) \Rightarrow \phi_2$  is a formula of logic  $L$ :

*Definition 5.2. Generalized Consequence Rule*

$$\frac{N \text{ sat } \phi_1, \models K(\phi_1) \Rightarrow \phi_2}{N \text{ sat } \phi_2}$$

(Note that  $K(\phi_1) \Rightarrow \phi_2$  is being interpreted over states, not over state se-

quences.) By Observation 5.1 and the definition of  $K(\phi)$ , if  $\phi_1$  and  $\phi_2$  are precise and valid specifications, respectively, for  $N$ , then  $\models K(\phi_1) \Rightarrow \phi_2$ . Therefore, incorporating the Generalized Consequence Rule yields a proof system that is complete relative to logic  $L$ . Our goal is to isolate the power required to express  $K(\phi)$ , a formula satisfied by exactly those states reachable by a computation that always satisfies  $\phi$ . Here we provide an outline of our expressiveness results; the reader is referred to [16] for technical details.

First, we observe that formula  $K(\phi)$  can be expressed in an extended trace logic in which some explicit reasoning over computations is permitted. Suppose the communication channels of all networks under consideration are  $c_1, c_2, \dots, c_m$ , and let the data domain of transmittable values be a set  $V$ . A state can be represented by a tuple  $\bar{t} = [t_1, \dots, t_m]$  of channel traces, where each  $t_i$  is the trace of  $c_i$ . A computation up to some point in time is represented by a finite sequence  $\langle \bar{t}^0, \bar{t}^1, \dots, \bar{t}^n \rangle$  of such tuples. (Since we are interested only in states reachable by computations, we need not consider the infinite sequences representing full computations.)  $K(\phi)$  can be expressed in a logic that allows quantification over such sequences. Let

–  $\phi[\bar{c}/\bar{t}^i]$  denote SNL specification  $\phi$  with channel trace variables  $c_1, \dots, c_m$  replaced by traces  $t_1^i, \dots, t_m^i$ ; and

–  $t^i[t_j^i/t_j^i \cdot \langle v \rangle]$  denote tuple  $t^i$  with trace  $t_j^i$  extended by value  $v \in V$ .

We then define  $K(\phi)$  as follows:

*Definitions 5.3. Formula  $K(\phi)$  in Extended Trace Logic*

$$\begin{aligned}
 K_{\text{ETL}}(\phi) \equiv & \\
 & (\exists \langle \bar{t}^0, \bar{t}^1, \dots, \bar{t}^n \rangle: && \text{There exists a sequence of states} \\
 & \bar{t}^0 = [\langle \rangle, \dots, \langle \rangle] \wedge && \text{such that: in the first state} \\
 & \bar{t}^n = [c_1, \dots, c_k] \wedge && \text{all traces are empty, } \phi \text{ is} \\
 & (\forall i: 0 \leq i \leq n: \phi[\bar{c}/\bar{t}^i]) \wedge && \text{satisfied in every state, and in} \\
 & (\forall i: 0 \leq i < n: && \text{every pair of adjacent states:} \\
 & \bar{t}^{i+1} = \bar{t}^i \vee && \text{either the states are identical or the} \\
 & (\exists j, v: 1 \leq j \leq m, v \in V: && \text{second state extends exactly one trace} \\
 & \bar{t}^{i+1} = \bar{t}^i[t_j^i/t_j^i \cdot \langle v \rangle])) && \text{of the first state by exactly one element.}
 \end{aligned}$$

The free variables of  $K_{\text{ETL}}(\phi)$  are channel trace variables  $c_1, \dots, c_m$ . As illustrated by the annotation,  $K_{\text{ETL}}(\phi)$  is satisfied by exactly those states reachable by a computation that always satisfies  $\phi$ . Therefore,  $K_{\text{ETL}}(\phi)$  could be used in the Generalized Consequence Rule to obtain relative completeness.

Linear-Time Temporal Logic (TL) can be used for similar but implicit reasoning over sequences of states [11] and is generally considered more convenient for such reasoning. As indicated by the definitions of *ORDERING* and *PREFIX* given in Sections 4.4.1 and 4.4.2, however, full temporal logic

and *PREFIX* given in Sections 4.4.1 and 4.4.2, however, full temporal logic is more powerful than is needed to express  $K(\phi)$ . By considering equivalence of TL formulas and formula  $K_{\text{ETL}}(\phi)$ , we isolate a subset of TL that is necessary and sufficient to express  $K(\phi)$ .

We begin with a version of TL that extends the trace logic of SNL with three standard temporal operators: the *Always* operator  $\square$ , the *Next* operator  $\circ$ , and the *Until* operator  $\mathcal{U}$ . We omit operator  $\diamond$  since it is the dual of  $\square$ . (For definition and further discussion of these temporal operators, see, e.g., [11] and [16].) Our original definition of  $\square$ , in Section 4.4, is for trace logic formulas that do not contain temporal operators. In Section 4.4.2 we considered a version for formulas containing other  $\square$ s. The weaker version of  $\square$  is used to define the Temporal Ordering Axiom, while the stronger version is needed to define the Prefix Axiom. In general, allowing nested temporal operators yields significantly more expressive power than restricting temporal operators to operate over nontemporal formulas.<sup>13</sup> Hence, we also consider use of an additional set of temporal operators that operate over trace logic formulas only:

- the *Restricted Always* operator,  $\overline{\square}$ ;
- the *Restricted Next* operator,  $\overline{\circ}$ ; and
- the *Restricted Until* operator,  $\overline{\mathcal{U}}$ .

It is easy to show that these operators are strictly weaker than their fully temporal counterparts, which we refer to as *Unrestricted* operators.

Various subsets of TL can be constructed by choosing different subsets of the six temporal operators. For example, trace logic with  $\overline{\square}$  and  $\overline{\circ}$  is a (strict) subset of TL. It is easily seen that operators  $\mathcal{U}$  and  $\overline{\mathcal{U}}$  are not needed to encode  $K_{\text{ETL}}(\phi)$ . Hence, the TL subsets of possible interest in expressing a formula equivalent to  $K_{\text{ETL}}(\phi)$  correspond to the subsets of  $\{\square, \circ, \overline{\square}, \overline{\circ}\}$ . The partial ordering of the expressive power of these subsets is given in Figure 4. For each nonempty subset  $S$  in Figure 4, let  $\text{Tr}_S$  denote the trace logic of SNL extended to include the operators in  $S$ . For example, subset 4 of Figure 4 is denoted by  $\text{Tr}_{\square}$  and subset 8 by  $\text{Tr}_{\square\circ}$ . We have proved that  $\text{Tr}_{\square}$  has the necessary and sufficient expressive power to encode  $K_{\text{ETL}}(\phi)$ .

Sufficiency of  $\text{Tr}_{\square}$  is shown by exhibiting a formula in  $\text{Tr}_{\square}$  that is equivalent to  $K_{\text{ETL}}(\phi)$ :

*Definition 5.4. Formula  $K(\phi)$  in  $\text{Tr}_{\square}$*

$$\begin{aligned}
 K_{\square}(\phi) \equiv & \square \phi \wedge \\
 & (\forall i, j, x, y: 1 \leq i \leq m, 1 \leq j \leq m, 1 \leq x, 0 \leq y, i \neq j \vee x \neq y: \\
 & \quad \square(|c_i| < x \Leftrightarrow |c_j| < y) \Rightarrow \square(|c_i| < x \wedge |c_j| < y)) \wedge \\
 & (\forall i, x, v: 1 \leq i \leq m, i \leq x, v \in V: \\
 & \quad \square(c_i.x = v \Rightarrow \square(c_i.x = v)))
 \end{aligned}$$

<sup>13</sup>For example, consider TL formula  $\square(\psi_1 \Rightarrow (\square\psi_2 \vee \circ\psi_3))$ , which asserts that whenever  $\psi_1$  is valid, either  $\psi_2$  is valid thereafter or  $\psi_3$  is valid at the next point in time. This property cannot be expressed using temporal operators only over first-order formulas.

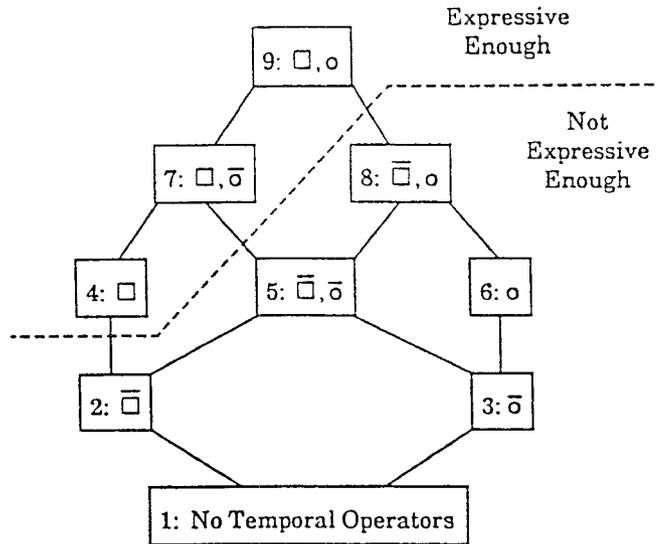


Fig. 4. Temporal logic subsets and expressiveness bounds.

The first conjunct of  $K_{\square}(\phi)$  restricts state-sequences satisfying  $K_{\square}(\phi)$  to always satisfy  $\phi$ . The second and third conjuncts encode the Temporal Ordering and Prefix Properties, restricting state-sequences satisfying  $K_{\square}(\phi)$  to represent computations. Equivalence of  $K_{\square}(\phi)$  and  $K_{\text{ETL}}(\phi)$  is shown using a semantics-preserving mapping from Temporal Logic to Extended Trace Logic [16]. Necessity of  $\text{Tr}_{\square}$  is proved by showing that no formula in  $\text{Tr}_{\square\circ}$  can be equivalent (through the mapping) to  $K_{\text{ETL}}(\phi)$ . Hence, all subsets except 4, 7, and 9 are insufficient.

Note that our results can be strengthened by refining the subset hierarchy. Consider an infinite set of temporal operators based on allowable nesting depth. For any  $x, 0 \leq x$ , let  $\square_x$  denote a version of  $\square$  restricted to operate over formulas with at most  $x$  nested  $\square$ s; similarly define operator  $\circ_x$ . We then obtain an infinite hierarchy of TL subsets. Given the previous results, it is easy to show that, with respect to this refined subset hierarchy,  $\text{Tr}_{\square_1}$  is necessary and sufficient to express a formula equivalent to  $K_{\text{ETL}}(\phi)$ .

## 6. CONCLUSIONS

We have considered a simple trace-based proof system for networks of processes, SNL, with a specification language and inference rules similar to those in most trace-based systems [3, 4, 6, 7, 8, 12, 14, 18]. Through examples that are single-process networks, we showed that SNL is incomplete because it is not expressive enough to encode properties of computation that are needed to verify certain valid network specifications. We then showed that axiomatization of the Temporal Ordering and Prefix Properties is necessary and sufficient to achieve relative completeness. Finally, we showed that the

power of an unrestricted temporal logic *Always* operator is necessary and sufficient to express these axioms.

Our extensions to the consequence rule of SNL could have been encoded elsewhere in the proof system. For example, the network composition rule could have been extended so that Temporal Ordering and Prefix Axioms would be added whenever the rule is applied [16]; then, however, the network composition rule would have to be used even in the case of one-process networks (such as the examples in this paper). Another approach to obtaining relative completeness is to use a different interpretation for logical implication in the consequence rule, treating  $\phi_1 \Rightarrow \phi_2$  as valid iff every legal computation (rather than every state) that always satisfies  $\phi_1$  also always satisfies  $\phi_2$ . This has the disadvantage of requiring a new underlying logic, rather than relying on a first-order logic of states, as we have done in this paper.

A further question of expressiveness arises: Using SNL, can precise specifications be given for all processes of interest? There are some processes for which precise specifications cannot be given, for the same reason that SNL is incomplete: SNL specifications for such processes would require expressing Temporal Ordering or Prefix Properties. However, since we have proved that extension of SNL is necessary for completeness of the proof system, the same extension can be made to the specification language, allowing a wider class of precise specifications to be formulated. It would be interesting to characterize those processes for which precise specifications can be given in various trace-based formalisms. In addition, it might be interesting to consider how our results can be adapted to a network logic that allows *mixed terms* [18], a logic in which processes and specifications are interchangeable.

Since the expressive power of the *Always* operator, or of Temporal Ordering and Prefix Axioms, is an essential component of a relatively complete proof system, we briefly review existing complete systems and identify how this expressive power is encoded. No encoding is needed in [14], since the proof system is based directly on temporal logic. We have shown, however, that the full power of temporal logic used in [14] is unnecessary for proving safety properties. Several proof systems allow explicit reasoning over all possible computations [3, 6, 19]. As we have seen, this gives at least the expressive power of temporal logic, since the states of every computation can be directly and individually referenced.

In [19] the incompleteness of the proof system in [12] is discussed, and a rule is suggested that would render it relatively complete. (A similar rule is proposed in [13].) Informally, the rule asserts the following: Let  $\phi$  be a valid specification for a network  $N$ , and let  $\tau$  be an interleaved trace of all communication events during any computation of  $N$ . Then every prefix of  $\tau$  satisfies  $\phi$ . This rule certainly captures our Prefix Property. The Temporal Ordering Property is encoded as well. To see this, suppose specification  $\phi$  constrains two communication events  $c1.x$  and  $c2.y$  to occur simultaneously. Any trace  $\tau$  including only one of  $c1.x$  and  $c2.y$  will not satisfy  $\phi$  and, thus, cannot correspond to a computation of  $N$ . Suppose, then, that both events are included in  $\tau$ . Consider any prefix  $\tau' \subseteq \tau$  that contains one event but not the

$c1.x$  and  $c2.y$  appears in  $\tau'$ . Hence, no computation of  $N$  can include either event.

In [8] the fact that valid specifications do not always follow from precise specifications is identified, and the author suggests adding a proof rule of the form

$$\frac{N \text{ sat } \phi_1}{N \text{ sat } \phi_2},$$

which can be applied whenever  $\phi_1$  and  $\phi_2$  are such that any network that always satisfies  $\phi_1$  always satisfies  $\phi_2$ . No formal method is given, however, for determining when a pair of specifications is a candidate for application of the rule. Our work has exactly characterized those pairs that qualify and has isolated the expressiveness required of a logic that can recognize them.

## APPENDIX

**THEOREM 3.4.** Soundness of SNL. *Let  $N$  be a network and  $\phi$  a specification such that  $N \text{ sat } \phi$  is a theorem of SNL. Then  $\phi$  is valid for  $N$ .*

**PROOF.** Since we are assuming validity of process specifications, soundness requires showing that whenever the hypothesis of an SNL inference rule is valid, so is the conclusion.

*Network Composition Rule 3.1*

$$\frac{P_1 \text{ sat } \phi_1, \dots, P_n \text{ sat } \phi_n}{P_1 \parallel \dots \parallel P_n \text{ sat } \bigwedge_i \phi_i}$$

Assume that each  $\phi_i$  is valid for  $P_i$ , so  $\kappa.j \models \phi_i$  for all  $\kappa \in CS(P_i)$  and  $j \geq 0$ . We must show that  $\kappa.j \models \bigwedge_i \phi_i$  for all  $\kappa \in CS(N)$  and  $j \geq 0$ , where  $N = P_1 \parallel \dots \parallel P_n$ . Recall that  $Proj(\kappa, P_i)$  denotes the projection of  $\kappa$  onto those channels of  $N$  incident to  $P_i$ . Consider an arbitrary conjunct  $\phi_i$ , an arbitrary  $\kappa \in CS(N)$ , and an arbitrary  $j \geq 0$ . By the definition of  $CS(N)$  (Section 2.1),  $Proj(\kappa, P_i) \in CS(P_i)$ ; hence, by assumption,  $Proj(\kappa, P_i).j \models \phi_i$ . Therefore,  $\kappa.j \models \phi_i$  as well, since the traces that are removed from  $\kappa$  in the projection cannot be for channels mentioned in  $\phi_i$ . Since  $\kappa$ ,  $j$ , and  $\phi_i$  were chosen arbitrarily, we conclude that  $\kappa.j \models \bigwedge_i \phi_i$  for all  $\kappa \in CS(N)$  and  $j \geq 0$ . Thus,  $\bigwedge_i \phi_i$  is valid for  $N$ .

*Consequence Rule 3.2*

$$\frac{N \text{ sat } \phi_1, \models \phi_1 \Rightarrow \phi_2}{N \text{ sat } \phi_2}$$

Let  $\phi_1$  be valid for  $N$ , so  $\kappa.j \models \phi_1$  for all  $\kappa \in CS(N)$  and  $j \geq 0$ . Then, by  $\models \phi_1 \Rightarrow \phi_2$  and predicate logic,  $\kappa.j \models \phi_2$  for all  $\kappa \in CS(N)$  and  $j \geq 0$ . Thus,  $\phi_2$  is valid for  $N$ .  $\square$

**THEOREM 4.3. Preciseness-Preservation.** *Let  $\phi_i$  be a precise specification for  $P_i$ ,  $1 \leq i \leq n$ , and let  $N = P_1 \parallel \cdots \parallel P_n$ . Then  $\bigwedge_i \phi_i$  is a precise specification for  $N$ .*

**PROOF.** We must show that  $\bigwedge_i \phi_i$  satisfies both parts of Definition 3.5 of precise specifications.

- (1) ( $\bigwedge_i \phi_i$  is valid for  $N$ .) Since the  $\phi_i$  are precise specifications for their respective  $P_i$ , they are valid. That  $\bigwedge_i \phi_i$  is then valid for  $N$  was proved in Soundness Theorem 3.4.
- (2) (If  $\kappa$  is any computation containing traces for the channels in  $N$  and if  $\kappa.j \models \bigwedge_i \phi_i$  for all  $j \geq 0$ , then  $\kappa \in CS(N$ .) Consider any  $\kappa$  containing traces for the channels in  $N$  in which  $\kappa.j \models \bigwedge_i \phi_i$  for all  $j \geq 0$ . Recall that  $Proj(\kappa, P_i)$  denotes the projection of  $\kappa$  onto those channels of  $N$  incident to  $P_i$ . Since  $\kappa.j \models \bigwedge_i \phi_i$  for all  $j \geq 0$ ,  $Proj(\kappa, P_i).j \models \phi_i$  for all  $j \geq 0$  and  $1 \leq i \leq n$ . Thus, by the preciseness of the  $\phi_i$ ,  $Proj(\kappa, P_i) \in CS(P_i)$ ,  $1 \leq i \leq n$ . Then, by the definition of  $CS(P_1 \parallel \cdots \parallel P_n)$  (Section 2.1),  $\kappa \in CS(N)$ .  $\square$

**THEOREM 4.6. Soundness of ORDERING.** *If  $\kappa$  is a computation, then  $\kappa \models ORDERING$ .*

**PROOF.** Let  $\kappa$  be any computation, and suppose that the antecedent of *ORDERING*, holds:  $\kappa \models \square(|c1| < x \Leftrightarrow |c2| < y)$  for some  $x \geq 1$ ,  $y \geq 0$ . We prove by induction that each state of  $\kappa$  satisfies  $(|c1| < x \wedge |c2| < y)$ , so  $\kappa \models \square(|c1| < x \wedge |c2| < y)$  and, therefore,  $\kappa \models ORDERING$ .

*Base Case.* We show that  $\kappa.0$  satisfies  $(|c1| < x \wedge |c2| < y)$ . From  $x \geq 1$  and computation condition (CC1) (Section 2.1),  $\kappa.0$  satisfies  $|c1| < x$ . Then, by the antecedent of *ORDERING*,  $\kappa.0$  also satisfies  $|c2| < y$ . Thus,  $\kappa.0$  satisfies  $(|c1| < x \wedge |c2| < y)$ .

*Induction.* Suppose that  $\kappa.(i-1)$  satisfies  $(|c1| < x \wedge |c2| < y)$  for some  $i > 0$ . We show that  $\kappa.i$  satisfies  $(|c1| < x \wedge |c2| < y)$ . Assume, for the sake of a contradiction, that  $\kappa.i$  satisfies  $(|c1| \geq x \vee |c2| \geq y)$ . Then, by the antecedent of *ORDERING*,  $\kappa.i$  also satisfies  $(|c1| \geq x \wedge |c2| \geq y)$ . Since  $\kappa.(i-1)$  satisfies  $(|c1| < x \wedge |c2| < y)$ , if  $c1$  and  $c2$  are distinct then two channel traces change between  $\kappa.(i-1)$  and  $\kappa.i$ , contradicting computation condition (CC4) (Section 2.1). If  $c1$  and  $c2$  are the same channel, then  $x \neq y$ , a channel trace increases in length by more than one between  $\kappa.(i-1)$  and  $\kappa.i$ , and condition (CC3) is contradicted. Hence,  $\kappa.i$  satisfies  $(|c1| < x \wedge |c2| < y)$ , and the induction is complete.  $\square$

**LEMMA 4.9.** *Let  $\kappa$  be any infinite sequence of states.  $\kappa$  represents a computation iff  $\kappa \models ORDERING \wedge PREFIX$ .*

**PROOF.** [ $\Rightarrow$ ] (If  $\kappa$  represents a computation, then  $\kappa \models ORDERING \wedge PREFIX$ .) This follows directly from soundness of *ORDERING* and *PREFIX* (Theorems 4.6 and 4.8).

[ $\Leftarrow$ ] (If  $\kappa \models \text{ORDERING} \wedge \text{PREFIX}$ , then  $\kappa$  represents a computation.) We prove the contrapositive: If  $\kappa$  does not represent a computation, then  $\kappa \not\models \text{ORDERING} \wedge \text{PREFIX}$ . By the definition of a computation (Section 2.1), if  $\kappa$  does not represent a computation then  $\kappa$  satisfies  $\neg(\text{CC1} \wedge \text{CC2} \wedge \text{CC3} \wedge \text{CC4})$ . Formula  $\neg(\text{CC1} \wedge \text{CC2} \wedge \text{CC3} \wedge \text{CC4})$  can be rewritten as

$$(\neg\text{CC2}) \vee (\text{CC2} \wedge \neg\text{CC1}) \vee (\text{CC2} \wedge \neg\text{CC3}) \vee (\text{CC2} \wedge \text{CC3} \wedge \neg\text{CC4}). \quad (6)$$

We show that if  $\kappa$  satisfies any of the disjuncts in (6), then  $\kappa$  does not satisfy both *ORDERING* and *PREFIX*.

- (1)  $\kappa$  does not satisfy (CC2): Some trace in some state is not a prefix of the corresponding trace in the subsequent state. Therefore, *PREFIX* does not hold.
- (2)  $\kappa$  satisfies (CC2) but not (CC1): Some trace is nonempty in the initial state, so let  $|c| = x$  in  $\kappa.0$  with  $x \geq 1$ . Since (CC2) holds,  $\kappa$  satisfies  $\Box(|c| \geq x)$ , and therefore,  $\Box(|c| < x \Leftrightarrow |c| < 0)$ . However,  $\kappa$  does not satisfy  $\Box(|c| < x \wedge |c| < 0)$ , so *ORDERING* does not hold.
- (3)  $\kappa$  satisfies (CC2) but not (CC3): Some trace increases in length by more than one between states, so let  $|c| = x$  in some  $\kappa.i$  and  $|c| = x + y$  in  $\kappa.(i + 1)$ , with  $y > 1$ . Since (CC2) holds,  $\kappa$  satisfies  $\Box(|c| < x + 1 \Leftrightarrow |c| < x + y)$ . However,  $\kappa$  does not satisfy  $\Box(|c| < x + 1 \wedge |c| < x + y)$  so *ORDERING* does not hold.
- (4)  $\kappa$  satisfies (CC2) and (CC3) but not (CC4): More than one trace changes between some state and its subsequent state, so let  $|c1| = x$  and  $|c2| = y$  in some  $\kappa.i$ , and let  $|c1| = x + 1$  and  $|c2| = y + 1$  in  $\kappa.(i + 1)$ . Since (CC2) holds,  $\kappa$  satisfies  $\Box(|c1| < x + 1 \Leftrightarrow |c2| < y + 1)$ . However,  $\kappa$  does not satisfy  $\Box(|c1| < x + 1 \wedge |c2| < y + 1)$ , so *ORDERING* does not hold.  $\square$

#### ACKNOWLEDGMENTS

We are grateful to Prakash Panangaden for many useful discussions regarding the results described in Section 5, and to Stephen Brookes for his comments on an earlier version of the paper.

#### REFERENCES

1. APT, K. R. Ten years of Hoare's logic: A survey—Part I. *ACM Trans. Program. Lang. Syst.* 3, 4 (Oct. 1981), 431–483.
2. BROCK, J. D., AND ACKERMAN, W. B. Scenarios: A model of non-determinate computation. In *Formalization of Programming Concepts*. Lecture Notes in Computer Science, vol. 107. Springer-Verlag, New York, 1981, pp. 252–259.
3. BROOKES, S. D. A semantics and proof system for communicating processes. In *Logics of Programs*. Lecture Notes in Computer Science, vol. 164. Springer-Verlag, New York, 1984, pp. 68–85.
4. CHEN, Z. C., AND HOARE, C. A. R. Partial correctness of communicating sequential processes. In *Proceedings of the IEEE International Conference on Distributed Computing Systems* (Paris, Apr. 1981), IEEE, New York, pp. 1–12.

5. COOK, S. A. Soundness and completeness of an axiom system for program verification. *SIAM J. Comput.* 7, 1 (Feb. 1978), 70-90.
6. HEHNER, E. C. R., AND HOARE, C. A. R. A more complete model of communicating processes. *Theor. Comput. Sci.* 26 (Sept 1983), 105-120.
7. HOARE, C. A. R. *Communicating Sequential Processes*. Prentice-Hall, Englewood Cliffs, N.J., 1985.
8. JONSSON, B. A model and proof system for asynchronous networks. In *Proceedings of the 4th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing* (Aug. 1985) ACM, New York, pp. 49-58.
9. LAMPORT, L. Proving the correctness of multiprocess programs. *IEEE Trans. Softw. Eng. SE-3*, 2 (Mar. 1977), 125-143.
10. MANNA, Z., AND PNUELI, A. Verification of concurrent programs: The temporal framework. In *The Correctness Problem in Computer Science*, R. S. Boyer and J. S. Moore, Eds. International Lecture Series in Computer Science. Academic Press, London, 1981, pp. 215-273
11. MANNA, Z., AND PNUELI, A. Verification of concurrent programs: A temporal proof system. In *Proceedings of the 4th School on Advanced Programming* (Amsterdam, June 1982), pp. 163-255.
12. MISRA, J., AND CHANDY, K. M. Proofs of networks of processes. *IEEE Trans. Softw. Eng.* 7, 7 (July 1981), 417-426.
13. NGUYEN, V. The incompleteness of Misra and Chandy's proof systems. *Inf. Process. Lett.* 21 (Aug. 1985), 93-96.
14. NGUYEN, V., DEMERS, A., GRIES, D., AND OWICKI, S. A model and temporal proof system for networks of processes. *Distrib. Comput.* 1, 1 (Jan. 1986), 7-25.
15. SCHOENFIELD, J. R. *Mathematical Logic*. Addison-Wesley, Reading, Mass., 1967.
16. WIDOM, J. Trace-based network proof systems: Expressiveness and completeness. Ph.D. thesis, Computer Science Dept. Cornell Univ., Ithaca, N.Y., May 1987.
17. WIDOM, J., GRIES, D., AND SCHNEIDER, F. B. Completeness and incompleteness of trace-based network proof systems. In *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (Jan. 1987), ACM, New York, pp. 27-38.
18. ZWIERS, J. *Compositionality, Concurrency, and Partial Correctness*. Lecture Notes in Computer Science, vol 321 Springer-Verlag, Berlin, 1989.
19. ZWIERS, J., DE ROEVER, W. P., AND VAN EMDE BOAS, P. Compositionality and concurrent networks: Soundness and completeness of a proofs system. In *Proceedings of the 12th International Colloquium on Automata, Languages, and Programming*. Lecture Notes in Computer Science, vol. 194. Springer-Verlag, Berlin, 1985, pp. 509-519.

Received February 1989; revised January 1990, November 1990, and August 1991; accepted August 1991