The Expressive Power of Indeterminate Dataflow Primitives

PRAKASH PANANGADEN* AND VASANT SHANBHOGUE[†]

Computer Science Department, Cornell University, Ithaca, New York 14853

We analyze the relative expressive power of variants of the indeterminate fair merge operator in the context of static dataflow. We establish that there are three different, provably inequivalent, forms of unbounded indeterminacy. In particular, we show that the well-known fair merge primitive cannot be expressed with just unbounded indeterminacy. Our proofs are based on a simple trace semantics and on identifying properties of the behaviors of networks that are invariant under network composition. The properties we consider in this paper are all generalizations of monotonicity. © 1992 Academic Press, Inc.

1. INTRODUCTION

The study of indeterminate computing systems is motivated by seeking to understand the theory underlying what is commonly called "distributed systems" programming. In such systems indeterminacy is introduced when one abstracts away from low level hardware or timing details. The subject receives further impetus from recent interest in parallel and distributed computing. From the purely theoretical point of view the subject is interesting because of the presence of new fundamental concerns, for example, deadlock and fairness, that have no analogue in sequential, determinate programming.

The foundational work in the theory of determinate distributed systems is due to Gilles Kahn (1977). His model of distributed systems is as follows. A system is viewed as a collection of autonomous computing agents communicating by one-way, unbounded data channels. Each computing agent executes its own program and, from time to time, may communicate with other agents by sending data tokens along the channels or reading tokens from incoming channels. If an agent attempts to read a token from an incoming channel that has no data it suspends until data becomes

---

* Present address: School of Computer Science, McGill University, Montreal, Quebec, Canada H3A 2A7.

† Present address: Department of Computer Science, Wichita State University, Wichita, Kansas 67208.

available. Under these conditions an agent computes a continuous function from input token streams to output token streams. One can thus model agents as functions. If one has a network of such agents then one can write down a system of equations to model the behavior of the network. If the network has feedback loops the system of equations will be recursive. Such systems of equations can be solved using elementary fixed-point theory.

Much of the research in indeterminate dataflow is aimed at developing a theory of comparable elegance and utility as Kahn's was for determinate dataflow. The contribution of this paper can be seen as a step towards such a goal as it provides a classification of the types of indeterminate primitives available and shows that monotonicity properties may be violated by common indeterminate primitives.

In this paper we analyze fairness properties of indeterminate dataflow networks. Fairness has been the paradigmatic liveness property and has been studied extensively in a variety of formalisms with several different definitions; see, for example, the recent book by Francez (1986). It is known that fairness introduces unbounded indeterminacy (or countable nondeterminism); the proof is by a simple Koenig's lemma argument. Plotkin's pioneering study of powerdomains for indeterminacy included the observation that the powerdomain that he introduced had been specifically designed for bounded indeterminacy and therefore excluded the study of fair systems (1976). Several people have worked on generalizations of powerdomain techniques that would apply to unbounded indeterminacy (Abramsky, 1983; Apt and Plotkin, 1986; Broy, 1983; Park, 1980, 1982; Plotkin, 1982).

Till recently fairness had been identified with unbounded indeterminacy (Apt and Olderog, 1983). Considerable effort had been expended in formalizing semantics of dataflow networks that would include a satisfactory treatment of the fair merge primitive (Broy, 1983; de Bakker et al., 1989; Keller and Panangaden, 1985; Panangaden, 1985; Park, 1982). All these efforts seemed to take the view that this was the next natural step after Plotkin's work. In the present paper we demonstrate that there are in fact three "levels" of expressive power all embodying unbounded indeterminacy. (Recent work has shown that there are, in fact, even more (McAllester, Panangaden, and Shanbhogue, 1988).) More importantly, from a semanticist's point of view, we show that there are fundamentally different order-theoretic properties corresponding to these levels of expressive power. Thus different semantic theories may be appropriate at these levels. Fair merge, in particular, violates basic monotonicity properties. An important consequence of our proof is that fair merge *cannot* be viewed as being a determinate process with some inputs concealed—a so-called "oracle driven" primitive. Some of the expressiveness results reported here were announced earlier by Panangaden and Stark (1988) using an operational approach.

Many of the results, particularly those relating to monotonicity problems, were observed by Broy (1983), though he did not provide proofs. Smyth had earlier observed that McCarthy's **amb** is not monotone when viewed as a function on the Plotkin powerdomain (Plotkin, 1990).

The three levels that we consider in this paper are expressed as fairness properties of dataflow merges. A merge consists of two input ports and a single output port. Tokens arriving at the input ports are transmitted unaltered to the output port; the relative ordering of tokens on the input port is preserved on the output port but tokens arriving at different input ports may be output in either order at the output port. A *fair* merge will transmit all the tokens that appear at its input ports. An *angelic* merge will transmit all the tokens at a given port if the sequence of tokens appearing at the *other* port is finite. Thus an angelic merge will behave like a fair merge if both input sequences are finite. An *infinity-fair* merge has the dual property. If the sequence of tokens at one of the input ports is infinite all the tokens at the other port will be transmitted. Thus an infinity-fair merge will behave like a fair merge if both input sequences are infinite. The main expressiveness result is that fair merge cannot be implemented by angelic merge and that angelic merge cannot be implemented by infinity-fair merge. It is easy to implement angelic merge with fair merge and a construction by Eugene Stark (1990) shows how to implement infinity-fair merge with angelic merge. Precise definitions of all these primitives and of "implements" will be given in later sections. The terminology for the merge primitives is due to Park (1982).

Our entire analysis is based on the description of a process as a set of possible behaviors. We use *traces* as abstractions of process behavior. By traces we simply mean sequences of events on the ports of the processes forming a dataflow network. Traces have been studied at various levels of mathematical rigor and have a rich mathematical theory (Aalbersberg and Rozenberg, 1988). The order-theoretic properties that we alluded to above are formulated in terms of traces rather than in terms of relations. It would have been preferable if we could have used purely extensional properties such as the input–output relation, but the well-known Brock–Ackerman example shows that the input–output relation computed by a network is not compositional (Brock and Ackerman, 1981).

It turns out that traces not only are compositional but are *fully abstract*. In the first half of the paper we describe an operational semantics for networks and use it to show that traces do indeed give rise to a compositional semantics of dataflow networks. The operational semantics that we describe is closely modeled on work of Stark (1989a, 1989b, 1990; see also Panangaden and Stark, 1988) and also on the IO automata formalism of Lynch and Tuttle (1987). Our expressiveness results do not require that traces be fully abstract, so we do not present a proof here, but a proof of

this result for a slightly different operational semantics may be found in Jonsson's paper that used the suggestions (Jonsson, 1989) of Joost Kok (1987). It is important to be aware that we (and Kok and Jonsson) view the complete output in response to a given (possibly infinite) input stream as observable. This is, of course, a rather liberal view of observability. It is only with this rather strong notion of observability that we can say interesting things about fairness. There is a fascinating study, by Rabinovich and Trakhtenbrot (1988) of the semantics of dataflow networks with only finite observations permitted.

We emphasize that three abstractions of processes are used in this paper—(a) the automata, which describe the operational semantics, (b) trace sets, which are compositional and on which the monotonicity properties are defined, and (c) input–output relations, that fail to be compositional. In fact, our proofs based on trace sets do not rely on the specific operational semantics, once the connection between automata and trace sets has been established. We have included the details of the operational semantics here in order to make the paper more self-contained.

2. Operational Semantics of Dataflow Processes

The operational semantics of dataflow networks is given in an automata theoretic formalism in which the notion of causal independence between concurrent events is taken to be primitive. This is based on the work of Lynch and Tuttle (1987), Lynch and Stark (1987), and Stark (1987, 1989a). Each process has a set of events associated with it. We represent concurrency by a binary relation on events, that tells us when two events are independent. This is the abstraction of describing when two commands in a program can be executed in parallel. This binary relation, which is called the *concurrency relation*, is axiomatized via equations that express the fact that the order of execution of concurrent events can be permuted, thus capturing the causal independence of concurrent events.

2.1. *Port Automata*

We first describe computing agents as automata, that can receive values at "input ports" and transmit values at "output ports." We use the term "port" instead if "channel" to emphasize that this is where an automaton interfaces with its environment. The set of events of an automaton comes equipped with a concurrency relation as mentioned above.

Definition 1. A *concurrent alphabet* is a set E, equipped with a symmetric, irreflexive binary relation $\|_E$, called the *concurrency relation*.

This concept is used in trace theory (Aalbersberg and Rozenberg, 1988; Mazurkiewicz, 1986) to obtain an algebraic theory of traces. We call events related by the concurrency relation *concurrent*. Let V be a set of *data values* called the *value alphabet*. Throughout this paper, we assume a fixed countable value alphabet. We refer to V^{∞} as the domain of streams. We use the term "stream" interchangeably with the term "value sequence."

We now describe the notion of an automaton that can execute events. The input and output events are described as (port, value) pairs. The rest of the events need not be of this form.

DEFINITION 2. A *port automaton* is a tuple

$$M = (E, Q, A),$$

where

• E is a concurrent alphabet of *events*, and Inp and Out are disjoint subsets of E, called the sets of *input events* and *output events*, respectively. Inp $= P^{in} \times V$, and Out $= P^{out} \times V$, for some disjoint finite sets P^{in} and P^{out}. The elements of P^{in} are called *input ports*, and the elements of P^{out} are called *output ports*. The elements of $E \backslash (\text{Inp} \cup \text{Out})$ are called *internal events*.

• Q is a set of *states*, and $q^i \in Q$ is a distinguished *initial state*.

• A is a transition function that maps each pair of states q, r in Q to a subset $A(q, r)$ of $E \cup \{\varepsilon\}$. ε, a special event not in E, is called the *identity event*.

satisfying the following conditions:

-   (*Disambiguation*) $r \neq r'$ implies $A(q, r) \cap A(q, r') = \varnothing$.

    (*Identity*) $\varepsilon \in A(q, r)$ iff $q = r$.

    (*Receptivity*)   For any state q and any input event a, there exists a state r such that $a \in A(q, r)$.

    (*Commutativity*)   For any state q and any events a, b, if $a \| b$, $a \in A(q, r)$ and $b \in A(q, s)$, then there exists a state p such that $a \in A(q, r)$ and $b \in A(r, p)$.

This definition is similar to the definitions of a *port automaton* and an *input–output automaton* due to Lynch and Stark (1989) and Panangaden and Stark (1988), and is closely related to the *input–output automata* of Lynch and Tuttle (1987).

The intuitive significance of these conditions is as follows. Disambiguation states that from a particular state, an event cannot take the automaton to two different states. A basic property of systems is that they cannot

control what their inputs are. They may, of course, ignore their inputs but they cannot determine their inputs, which are supplied by the external environment. The receptivity property makes this precise. If two events are concurrent, and if both of them are enabled in a particular state, then the execution of any one of these two events does not disable the other, and moreover, the execution of these events in either order results in the same final state. This is captured by *commutativity*.

The *transitions* of an automaton are the tiples (q, a, r) with $a \in A(q, r)$. We denote the transition (q, a, r) by $q \xrightarrow{a} r$. The transition $q \xrightarrow{\varepsilon} q$ is called an *identity transition*, and is denoted by id_q.

DEFINITION 3. A *computation sequence* γ is a finite or infinite sequence of transitions of the form

$$q_1 \xrightarrow{a_1} q_2 \xrightarrow{a_2} \cdots$$

The *domain* $\mathrm{dom}(\gamma)$ of γ is the state q_1. A computation sequence is said to be *initial* if $\mathrm{dom}(\gamma)$ is the distinguished start state q'. Two computation sequences γ and δ are *coinitial* if $\mathrm{dom}(\gamma) = \mathrm{dom}(\delta)$.

In earlier work (Panangaden and Stark, 1988), automata in which inputs cannot disable other events were considered. These are called *monotone* automata, and these are port automata satisfying the following additional property.

DEFINITION 4 (Non-Disabling Inputs). If e is an input event at an input port, then $e \| e'$ for any event e' that is not an input event at the same port.

As Example 2 below shows, not all automata satisfy this property, and since we would like to be able to represent automata as in Example 2, we will not restrict our discussion to monotone automata in this paper. The fact that networks of monotone automata cannot implement non-monotone automata in general is the essential content of one of the expressiveness results stated operationally.

We now give two examples of automata. We use $^\wedge$ as an infix operator for representing concatenation of sequences.

EXAMPLE 1 (Buffer). This automaton has one input port and one output port. It reads values and outputs them, guaranteeing to read and output all values that arrive on the input port.

Let the set of states Q be V^*. A state represents the contents of the input port. The initial state is Λ. Let the set of input events Inp be $\{i\} \times V$ and the set of output events Out be $\{o\} \times V$. Then the set of events E is Inp $\cup$ Out.

We now define the transition relation, using v to represent a member of V. $A(q, r) = \{(i, v)\}$ iff $r = q \,\hat{}\, v$. $A(q, r) = \{(o, v)\}$ iff $q = v \,\hat{}\, r$. $A(q, q) = \{\varepsilon\}$.

Every event in Inp is concurrent with every event in Out, and ε is concurrent with any other event.

EXAMPLE 2 (Poll).   This automaton has one input port and one output port. It repeatedly checks its input port for data. If a data value is present, then it is read and output. If not, a special value $\star$ is output.

Let Q and Inp be the same as for the previous example. Let the set of output events Out be the set in the previous example, together with an extra event $(o, \star)$.

Besides the transitions in the previous example, there is an extra transition, and we redefine $A(q, q)$. $A(q, q) = \{\varepsilon, (o, \star)\}$ if $q = A$, and $A(q, q) = \{\varepsilon\}$ otherwise. The only new thing added to the concurrency relation of the previous example is that ε is concurrent with the new event $(o, \star)$.

Note that input events are not concurrent with $(o, \star)$, so that arrival of input disables the output of a $\star$. So the input here has the power to interrupt already enabled output. This automaton is therefore not monotone. It turns out that poll has the expressive power of fair merge.

We end this subsection with some notation. We refer to events of the form (p, v) as p-events. We denote the value component v of an event $e = (p, v)$ by value(e). We also extend the definition of value to sequences—value$((p, v_1)(p, v_2) \cdots)$ is $v_1 v_2 \cdots$. For any computation sequence

$$\sigma = q_1 \xrightarrow{\;a_1\;} q_2 \xrightarrow{\;a_2\;} \cdots$$

we define $ev(\sigma)$, the *sequence of events of* σ, to be $a_1 a_2 \cdots$. We use the symbol Π as a projection operator on sequences of events. Given any sequence t of events and a set S of ports, we use $\Pi_S(t)$ to represent the subsequence of t consisting of the p-events in t for all ports p in S. If S is a singleton set $\{p\}$, then we use the notation $\Pi_p(t)$ instead of $\Pi_{\{p\}}(t)$. If t is the sequence of events of a computation sequence σ, then we also write $\Pi_S(\sigma)$ to mean the same thing as $\Pi_S(t)$, and $\Pi_p(\sigma)$ to mean the same thing as $\Pi_p(t)$. To project out the p-events of a sequence t, we define $\Pi_{\sim p}(t)$ to be the subsequence of t consisting of all the events of t except the p-events. When we compare the projections of a sequence of events onto different ports, we follow the convention of implicitly applying *value* to the projections. We use the notation $t[i]$ to represent the ith event in a sequence t of events, and the notation $\gamma[i]$ to represent the ith transition in a computation sequence γ.

2.2. *Completed Computations and the Input–Output Relation*

In this subsection, we describe which computation sequences of automata we view as "completed," i.e., cannot be extended further. Once we establish this, we then show how we can abstract the input–output behaviour of an automaton from its completed computation sequences.

We first describe the notion of a *history*. Let P be the set of input ports and output ports of an automaton. A *history over P* is defined to be a function from P to V^∞. Then for any computation sequence σ, we can define a history H_σ by letting $H_\sigma(p)$ be value($\Pi_p(\sigma)$). Similarly, for any sequence $t \in (P \times V)^\infty$, we can define a history H_t by letting $H_t(p)$ be value($\Pi_p(t)$). We denote the restriction of H_σ to the input ports by H_σ^{in}, and call it the *input port history* corresponding to σ. We denote the restriction of H_σ to the output ports by H_σ^{out}, and call it the *output port history* corresponding to σ.

We now describe the computation sequences that we consider as "completed." To do this, we extend the prefix ordering on computation sequences to include the concurrency information in the concurrency relation. A finite computation sequence γ is a *prefix* of a computation sequence δ, and we write $\gamma \leqslant \delta$, iff there exists a computation sequence ξ with $\gamma\xi = \delta$. We define *permutation equivalence* to the least congruence $\sim$, respecting concatenation, on the set of finite computation sequences of an automaton such that whenever $a \| b$, the computation sequences $q \overset{a}{\longrightarrow} r \overset{b}{\longrightarrow} p$ and $q \overset{b}{\longrightarrow} s \overset{a}{\longrightarrow} p$ are $\sim$-related. We define the *permutation preorder* relation $\sqsubseteq$ on finite computation sequences of A as the transitive closure of $\leqslant \cup \sim$. Define $\simeq\, =\, \sqsubseteq \cap \sqsupseteq$. We can now extend the permutation preorder relation to infinite computation sequences by defining $\gamma \sqsubseteq \delta$ iff for every finite $\gamma' \leqslant \gamma$, there exists a finite $\delta' \leqslant \delta$, such that $\gamma' \sqsubseteq \delta'$. We define $\simeq\, =\, \sqsubseteq \cap \sqsupseteq$ for infinite computation sequences also.

We would like a notion "completed" computation sequence, in which all events that could happen at any state have either happened or been disabled. This is a *weak fairness* condition, and we formalize it as follows.

DEFINITION 5. A computation sequence γ is called *completed* if either it is finite and no non-input event is enabled at its end, or it is infinite and there is no event e enabled at every state in $\gamma[i..]$ and concurrent with every event in $\gamma[i..]$ for some i.

This turns out to be identical to $\sqsubseteq$-maximality for a particular input (Shanbhogue, 1990).

For the port automata considered by Panangaden and Stark (1988), no two non-input events were concurrent. In this paper, the notion of a *fair initial computation sequence* was introduced. For single automata, this

notion is equivalent to our notion of completion. We will consider
networks of automata presently, and for the networks of monotone
automata considered by Panangaden and Stark (1988), the fair computa-
tion sequences coincide with the completed computation sequences, which
coincide with the $\sqsubseteq$-maximal computation sequences. Moreover, the
projections of maximal computation sequences onto individual automata
in the network are themselves maximal. When we consider networks of
general port automata as in this paper, this is no longer the case. So, as we
discuss in the next subsection, we give a different notion of completion for
networks of automata. For networks of monotone automata, however
(Panangaden and Stark, 1988; Stark, 1990), the projections of a maximal
computation sequence onto individual automata are again maximal. This
means that even though our above notion of completion is equivalent to
maximality for general automata, this notion can be extended to networks
of monotone automata but not to networks of general automata. This is
unfortunate, but unavoidable.

We could think of the preorder $\sqsubseteq$ as the prefix ordering in which con-
current information has been encoded. It is quite pleasant to be able to
state completedness as a maximality property of computation sequences.

EXAMPLE 3. Consider an automaton with one output port o and only
two internal events e_1 and e_2. There are only three states, q_1, q_2, q_3. q_1 is
the initial state. The transitions are completely described by $A(q_1, q_1) =$
$\{\varepsilon, e_1\}$, $A(q_1, q_2) = \{e_2\}$, $A(q_2, q_3) = \{(o, v)\}$, where v is some fixed value,
$A(q_2, q_2) = \{\varepsilon\}$, and $A(q_3, q_3) = \{\varepsilon\}$. There is an infinite sequence γ of
e_1-transitions from the initial state, and the event e_2 is continuously
enabled, following which an output would be enabled. But the event e_2 is
not concurrent with e_1, and the automaton is really making an indeter-
minate choice at each step without any fairness constraints. So we would
accept γ as a completed computation sequence.

This example makes clear the distinction between completion and
continuous enabling.

Now we can describe the *input–output relation* of an automaton. This
describes the input–output behaviour and corresponds to the *observable*
aspect of network behavior.

DEFINITION 6. The *input–output relation* of an automaton is the set of
all pairs $(H_\sigma^{\text{in}}, H_\sigma^{\text{out}})$ with σ being a completed computation sequence of the
automaton.

We can also equivalently consider the input–output relation to be a set
of pairs of tuples of streams, the first tuple of each pair consisting of

streams at the input ports and the second tuple consisting of streams at the output ports. We refer to the input–output relation as the *IO-relation*.

If the IO-relation turns out to be the graph of a function then the automaton is said to *compute* that function. We now describe a subclass of the class of automata that compute functions. The following definition is due to Stark (1987).

DEFINITION 7. An automaton is *deterministic* if it satisfies the following condition: $b \parallel b'$ whenever b, b' are distinct non-input events both enabled at some state.

Intuitively, a determinate automaton does not exhibit "internal indeterminacy." These are the automata that correspond to the computing agents that Kahn considered in his treatment of networks, in view of the following theorem of Stark (1987).

THEOREM 1. *A function f is computed by a determinate automaton iff f is a continuous function.*

2.3. Networks of Automata

We now describe how we can build networks of automata. We consider three operations—aggregation, feedback, and output hiding. Ordinary sequential composition of programs can be described in terms of these operations. Thus these operations cover the normal intuitions of forming networks from individual components.

Aggregation involves taking two networks with disjoint sets of ports and "keeping them side by side" to obtain a new network. Feedback (see Fig. 1) involves identifying an output port p_1 of the network with an input port p_2 of that network, so that the port p_2 can no longer be observed by the environment and can no longer serve as an input port. The values appearing at p_2 will be those that are output at p_1. One stipulation that


FIG. 1. Feedback.

we make is that the output port p_1 and the input port p_2 not be ports of the same automaton in the network. Our last operation, output hiding, removes an output port, so that the environment can no longer observe values at this hidden port.

We can define the first two operations as special cases of the *composition* of a finite *compatible* set of automata.

DEFINITION 8.   If I is a finite index set, then a set $\mathscr{S} = \{M_i : i \in I\}$ of automata is said to be *compatible* if

• for all i, $j \in I$ such that $i \neq j$ we have $(E_i \backslash (\mathrm{Inp}_i \cup \mathrm{Out}_i)) \cap (E_j \backslash (\mathrm{Inp}_j \cup \mathrm{Out}_j)) = \varnothing$, that is, the sets of internal events of any pair of automata are disjoint, and,

• for any port name, at most two automata may have that name in common, and in that case, it must be the name of an output port of one automaton and an input port of the other automaton,

where $M_i = (E_i, Q_i, A_i)$, and Inp_i is the set of input events of M_i, and Out_i is the set of output events of M_i.

The shared port names represent ports that get connected when the set of automata are composed together. We then obtain a *network automaton*. The input ports of the network are all the input ports of the M_i's, excluding those that are shared. The output ports of the network are all the output ports of the M_i's.

DEFINITION 9.   The *composition* of a compatible set $\mathscr{S}$ of automata is the automaton $\prod M_i = (E, Q, A)$, where

• $E = \bigcup E_i$, with $a \| b$ iff $a \|_i b$ for all $i \in I$ such that both a and b are in E_i.
• $\mathrm{Out} = (\bigcup \mathrm{Out})$, and $\mathrm{Inp} = (\bigcup \mathrm{Inp}_i) \backslash (\bigcup \mathrm{Out}_i)$,
• $Q = \prod_{i \in I} Q_i$
• $q^i = (q_i^i : i \in I)$
• $e \in A((q_i : i \in I), (r_i : i \in I))$ iff for all $i \in I$, either $e \notin E_i$ and $r_i = q_i$, or else $e \in E_i$ and $e \in A_i(q_i, r_i)$.

The definition of $\|$ above implies that events of distinct automata, that do not share any ports, are concurrent, because they are not both in the event set of any single automaton. We now note that the above definition includes the definition of *aggregation*, which is what happens when two automata do not share any ports, and the definition of *feedback*, which is what happens when two automata share a port name and it is the output

port of one automaton and an input port of the other automaton. We now explicitly define *output hiding*.

DEFINITION 10. If $\mathscr{A}$ is an automaton with input ports P^{in} and output ports P^{out}, and $\mathscr{S}$ is a subset of P^{out}, then the *output hiding* of $\mathscr{S}$ in $\mathscr{A}$ results in the automaton with input ports P^{in} and output ports $P^{out}\backslash\mathscr{S}$ with exactly the same sets of events and states and the same transition relation as $\mathscr{A}$.

When we compose two automata with a shared port name p, p must be an output port of one automaton and an input port of the other automaton. The two automata connected in this manner may execute a single event in the composed automaton. It is important to remember that this corresponds to two different events in the original collection of automata. By defining composition in this way, we do not have to worry about liveness conditions, to ensure that values output onto p will eventually arrive at the input port of the other automaton.

We can define *history*, *input port history*, and *output port history*, corresponding to computation sequences of networks, just as we did for computation sequences of single automata. We recall that, for a single automaton, the completed computation sequences were the $\sqsubseteq$-maximal ones. We could now define *permutation equivalence* and the *permutation preorder* $\sqsubseteq$ as we did for single automata, and then talk about the $\sqsubseteq$-maximal computation sequences. But we argue that this maximality condition is not a reasonable definition of completion for networks in general. We illustrate this point of view by an example.

EXAMPLE 4. Consider the network in Fig. 2, consisting of a buffer automaton, as in Example 1, and a poll automaton, as in Example 2. Suppose we took maximality of computation sequences as defining completion. Then the computation sequence starting with an input event at the input port of the buffer, followed by an infinite sequence of ★ outputs by the poll, is a maximal computation sequence, because even though the output event for buffer (input event for poll) is enabled at every state following the read event of the buffer, it does not commute with a ★ output. But this computation sequence does not really make sense as a completed sequence, because we want the buffer to output all its input values. In other words,


FIG. 2. A buffer and a poll.

the projection of the network computation sequence onto the ports of the buffer is not a completed computation sequence for the buffer.

The following definition, however, captures the proper notion of completion.

DEFINITION 11.   A computation sequence γ for a network of automata is *completed* if, for every automaton M in the network, the restriction of γ to M is a completed computation sequence for M.

Intuitively, this definition makes sense, because we would like to call those computation sequences of the network "completed," in which every component automaton gets chances to execute and exhibits a completed computation sequence in every computation of the network.

Just as we defined the input–output relation for a single automaton earlier, we can now define the following:

DEFINITION 12.   The *input–output relation* of a network of automata is the set of all pairs $(H_\sigma^{in}, H_\sigma^{out})$ with σ being a completed computation sequence of the network, H_σ^{in} being the input port history corresponding to σ, and H_σ^{out} being the output port history corresponding to σ.

2.4. Observability, Buffering, and Processes

It is clear that, for an "external" observer of an automaton, all that can be observed is a sequence of input events at each input port of the automaton, and a sequence of output events at each output port of the automaton. Due to the asynchronous nature of communication between automata, there may be an arbitrary delay between the emission of a value at an output port and the observation of that value. So it is not possible for an observer to determine the order of emission of values on *different* ports.

To reason about processes, we abstract away from computation sequences by considering only the input events and output events, as these are the events that interface the process with the external world. But the definition of automata do not allow any two events on different ports to be commuted. That means that if there is a computation sequence γ with event e on output port p occurring before an event e' on some other output port p', then there may be no computation sequence with e' preceding e and the rest of γ unchanged. But, as far as the external observer is concerned, he should not be able to distinguish the automaton in which e always occurs before e' from the automaton in which e may precede e', or e' may precede e. So, as far as the external observer is concerned, these two automata

should have the same set of "traces"—abstractions from computation sequences.

For this reason, we describe processes as automata with buffers attached to each of the input and output ports. As we shall see, this will allow us to commute events on different ports. A buffer was described as a specific automaton in Example 1. Having buffers in the description of a process will also make the arbitrary delay between the emission of a value and its observation explicit in computation sequences of the network.

We now formally describe what we mean by a "buffered automaton."

DEFINITION 13. A *process* is the composition of

(i) an automaton, called the *central automaton* of the process, with input ports $i_1, i_2, \ldots i_n$ and output ports $o_1, o_2, \ldots o_m$, and

(ii) $m + n$ buffers, called the *process buffers*, n of them having output ports $i_1, i_2, \ldots i_n$, respectively, and the remaining m of them having input ports $o_1, o_2, \ldots o_m$, respectively. The rest of the ports of the buffers are disjoint from the set $\{i_1, i_2, \ldots i_n, o_1, o_2, \ldots o_m\}$,

with the ports $\{i_1, i_2, \ldots i_n, o_1, o_2, \ldots o_m\}$ hidden.

Figure 3 shows a process; we may also refer to this as a buffered automaton.

By our definition of completed computation sequences, we ensure that every value input to the process will arrive at an input port of the central automaton, and every value output by the central automaton will be output by the process.

Just as we defined a compatible set of automata in the previous section, we can similarly define a compatible set of processes—the set consisting of the automata that make up these processes must be compatible. The *composition* of a compatible set of processes is the network obtained by


FIG. 3.   The structure of a process.

composing the set of network automata representing the individual processes. A network of processes may also be obtained by hiding some of the output ports of the composition of a compatible set of processes.

We can now talk about notions of implementability and non-implementability.

DEFINITION 14. A set S of processes can *implement* a relation R if there is a finite network N, built out of copies of processes in S, such that R is the input–output relation of N.

DEFINITION 15. A set S of processes can *implement* a process or a network M if there is a finite network N, built out of copies of processes in S, such that N and M have the same input–output relation.

In this paper, we differentiate between classes of processes, as exemplified by the merge processes, by proving non-implementability results between these classes.

3. TRACE SETS

We can describe the behaviour of a network by describing its input–output behaviour. The input–output relation of a network of processes is the input–output relation of the associated network of automata. Since "buffering" an automaton does not affect the input–output relation, we have a process with the input–output relation of poll (Example 2). But note that we do not consider "arbiters," which can distinguish between the order of arrival of values at distinct input ports. This is because, for the processes we consider, we cannot distinguish between the order of arrival of values at distinct input ports, because the central automaton of the process may see the values in either order, as the buffers at the input ports of the process may produce the corresponding output events in either order.

The input–output relation semantics for processes and networks fails to be compositional. The reason for this is that the input–output relation describes which sequences of values can be obtained at the output ports for given sequences of values at the input ports. So this view describes the *entire* output when given the *entire* input. It does not describe which particular values in the output sequences *really depended* on the presence of which particular values in the input sequences. Briefly, there is no causality information between events encoded in the input–output relation. One way to express causality information is to consider sequences of input and output events.

DEFINITION 16.   If γ is a computation sequence, then we define $\mathrm{tr}(\gamma)$ to be the subsequence of $\mathrm{ev}(\gamma)$, consisting of all the input and output events in $\mathrm{ev}(\gamma)$.

DEFINITION 17.   A *trace* of a network N of processes is a sequence t of input events and output events of N, such that $t = \mathrm{tr}(\gamma)$ for some completed computation sequence γ of N. We write $\mathrm{Trset}(N)$ for the set of traces of a network N.

In an earlier discussion of network semantics, trace sets were called archives (Keller and Panangaden, 1986). Presented in this way, the traces appear as an abstraction of computation sequences that were defined using an operational formalism. The important point is that we can define composition rules directly on trace sets, and this allows us to build up trace sets of complex networks structurally. If t is a sequence of events and P is a process or a network, $\Pi_P(t)$ will represent the subsequence of t consisting of all the events on the input and output ports of P in t.

We emphasize here that due to the description of processes as automata with buffers, the trace set of a process satisfies natural closure conditions —if two events e_1 and e_2 are adjacent in a trace t, and either (i) both are input events at different ports, or (ii) both are output events at different ports, or (iii) e_1 is an output event and e_2 is an input event, then flipping these two events gives us a trace of the process. This property of trace sets of processes is natural, as discussed in the subsection on observability earlier.

We would like to prove a theorem that the trace set of a network may be described in terms of the trace sets of its subnetworks. In order to do this we need to formalize some preliminary notions. We define the notion of corresponding events in a trace or a computation sequence.

DEFINITION 18.   In a trace or computation sequence of a network with two of its ports being p_1 and p_2, a p_1-event occurrence e_1 and a p_2-event occurrence e_2 are said to be *corresponding events* if e_1 is the ith event on port p_1, and e_2 is the ith event on port p_2 for some $i > 0$.

We use the predicate "precedes" to indicate precedence between events in a trace.

DEFINITION 19. $\mathrm{precedes}(p_1, p_2, t) =_{\mathrm{def}}$ for every finite prefix of t, value($\Pi_{p_2}(t)$) is a prefix of value $\Pi_{p_1}(t)$).

Informally, this means that every p_2-event occurrence is preceded by a corresponding p_1-event occurrence with the same value. This predicate is useful for defining the trace sets of networks built using feedback.

The following theorem states that traces are compositional.

THEOREM 2. Consider any decomposition of a network N into sub-networks $\{N_i\}$ for i in some finite index set I, where the N_i are not necessarily individual processes. Then the set of traces $\mathrm{Trset}(N)$ of the network is exactly the set of sequences $\Pi_N(t)$, such that

(i) t is a sequence of events at the input and output ports of the N_i's distinguishing between an output port and an input port even if they have the same name p by calling them p_{out} and p_{in} respectively,

(ii) for all $i \in I$, $\Pi_{N_i}(t)$ is a trace of N_i, and

(iii) for every port p that is an output port of N_i, as well as an input port of N_j, i not necessarily distinct from j, $\mathrm{precedes}(p_{\mathrm{out}}, p_{\mathrm{in}}, t)$ and $\Pi_{p_{\mathrm{out}}}(t) = \Pi_{p_{\mathrm{in}}}(t)$.

Proof. Let t be a sequence of events, satisfying conditions (i), (ii), and (iii) above. We show that $\Pi_N(t)$ is a trace of N. Let $t_i = \Pi_{N_i}(t)$. For each t_i, there is a completed computation sequence γ_i of N_i, such that $t_i = \mathrm{tr}(\gamma_i)$. Then we can dovetail among the computation sequences $\{\gamma_i : i \in I\}$ to obtain a computation sequence γ, such that t is the subsequence of $\mathrm{ev}(\gamma)$ consisting of all the input and output events of the N_i. Now, using condition (iii) in the statement of the theorem, and by the property of non-disabling inputs of processes, we obtain a new computation squence γ'' in which for every shared port p, the output events at p_{out} immediately precede the corresponding input events at p_{in}. Replacing each such adjacent pair of events at p_{out} and p_{in} by a single p-event, we obtain a computation sequence γ' of the network N in which the output events on the shared port are also the input events on the shared port. Any automaton M in the network is in one of the subnetworks, say N_i, and so the restriction of γ' to M is the same as the restriction of γ_i to M. Since the latter is completed, because γ_i is a completed computation sequence of N_i, so is the restriction of γ' to M. Therefore γ' is a completed computation sequence of the whole network, and so $\mathrm{tr}(\gamma')$ is a trace and it is exactly the restriction of t to the events on the input and output ports of the network.

Let t' be a trace of N corresponding to the completed computation sequence γ of the network, composed of subnetworks N_i. The restriction γ_i of γ onto any subnetwork N_i is then a completed computation sequence of N_i by the definition of a completed computation sequence of a network. Let t'' be the subsequence of $\mathrm{ev}(\gamma)$ consisting of all the input and output events of the N_i in $\mathrm{ev}(\gamma)$. Now, for every port that is shared by an N_i and an N_j, i not necessarily distinct from j, replace an event on that port in t'' by two events—an output event on the output port followed by an input event on the input port. The new event sequence t is a sequence of events such that its restriction to the events of any subnetwork N_i is a trace of

that subnetwork, and moreover, the restriction of t to events on the input and output ports of the network is the same as the trace $t' = \mathrm{tr}(\gamma)$. ∎

It is possible to have processes with different sets of traces, but the same IO-relation. Brock and Ackerman (1981) have such an example, but their example uses a powerful primitive, *fair merge*. There are other examples using only finite indeterminacy (Russell, 1989).

We use vector notation to talk about collections of ports. We project traces onto vectors of ports and assume that we get a vector of event sequences. Thus, we may write, "a trace t when projected onto ports **a** produces the sequences **α**." This notation allows us to write many projections at once and also allows us to name the components of the vectors as the need arises.

We described the construction of networks in the previous section. We use the notation $N \| M$ to represent the network built by aggregation from the two subnetworks N and M. We use the notation $\mathrm{loop}(p_1, p_2, N)$ to represent network built by identifying an output port p_1 of the network N with the input port p_2 of N. (See Fig. 4.) Last the traces of a network built by hiding some output ports of another network are obtained by simply taking the traces of the original network and projecting out the events on the hidden output ports. We use the notation $N \backslash S$ for the network N with the ports S of N hidden.

THEOREM 3.

$$\mathrm{Trset}(N \| M) = \{ t \in (P \times V)^{\infty} \mid \Pi_N(t) \in \mathrm{Trset}(N) \wedge \Pi_M(t) \in \mathrm{Trset}(M) \},$$

where P is the set of input and output ports of N and M.

THEOREM 4. $\mathrm{Trset}(\mathrm{loop}(p_1, p_2, N))$ *consists of every sequence* $\Pi_{\sim p_2}(t)$ *such that* $t \in \mathrm{Trset}(N)$ *and* $\mathrm{precedes}(p_1, p_2, t)$ *and* $(\Pi_{p_1}(t) = \Pi_{p_2}(t))$.


FIG. 4. Loop formation.

The statement of the theorem expresses the idea that in the network with the feedback, the events on p_2 arise from events on p_1.

Let us define a predicate similar to precedes, that captures "immediate precedence" of events.

DEFINITION 20.   immprecedes$(p_1, p_2, t) =_{\text{def}}$ precedes$(p_1, p_2, t) \land$ (every p_2-event is immediately preceded by the corresponding p_1-event in t).

Then it follows by monotonicity (Definition 4) that for any trace t' of loop(p_1, p_2, N), there is a trace $t \in \text{Trset}(N)$ such that immprecedes$(p_1, p_2, t) \land (\Pi_{p_1}(t) = \Pi_{p_2}(t))$.

THEOREM 5.   *Suppose S_O is a subset of the set of output ports of a network N. Then*

$$\text{Trset}(N \backslash S_O) = \{ \Pi_{S_O}(t) \mid t \in \text{Trset}(N) \}$$

All these proofs follow directly from Theorem 2 that establishes compositionality of trace sets.

The fact that we have a compositional description of network behavior allows one to prove properties of networks by structural induction on the network.

4. THE MERGE PRIMITIVES

In this section, we describe the merge primitives fair merge, angelic merge, and infinity-fair merge. We describe both their input–output behavior and their trace sets. The point of introducing these primitives is to study and compare the interesting classes of processes in which they lie. We emphasize that we wish to make interesting distinctions between the input–output relations of these primitives. We use the compositionality of trace sets to make these distinctions. Informal descriptions of the merge primitives were given in the introduction. Now we formally describe their input–output relations, and specific trace sets.

We first describe angelic merge. Intuitively, angelic merge can avoid getting stuck at input ports with no input, but need not be fair on infinite input streams. The input–output relation of angelic merge consists of triples $(s_i, s_{i'}, s_o)$ of value sequences, such that s_o can be broken up into two subsequences s_1, s_2 such that

   (i) s_1 is a prefix of s_i and s_2 is a prefix of $s_{i'}$,
   (ii)   if s_1 is finite, then $s_2 = s_{i'}$, and if s_2 is finite, then $s_1 = s_i$.

We now give a specific trace set that has the above IO-relation associated with it. If the angelic merge process has input ports i, i' and output port o, then the trace set consists of all event sequences $t \in (\{i, i', o\} \times V)^\infty$ such that $\Pi_o(t)$ can be broken up into two subsequences s_1, s_2 such that

(i)   value(s_1) is a prefix of value($\Pi_i(t)$) and value(s_2) is a prefix of value($\Pi_{i'}(t)$),

(ii)   if $\Pi_i(t)$ is finite, then value(s_2) = value($\Pi_{i'}(t)$),

(iii)   if $\Pi_{i'}(t)$ is finite, then value(s_1) = value($\Pi_i(t)$), and

(iv)   for every prefix t' of t, the prefix of s_1 in t' is a prefix of $\Pi_i(t')$, and the prefix of s_2 in t' is a prefix of $\Pi_{i'}(t)$.

Before describing infinity-fair merge, we first describe a primitive *anyints* that embodies unbounded indeterminacy. It has one input port and one output port. It consumes any value on the input port and produces any positive integer on the output port. In particular, we could feed it an infinite stream of zeros and have it produce an infinite stream of arbitrary positive integers on its output port. With *anyints* we can easily program a merge process that is fair if both its input streams are infinite. We use the *anyints* to produce an arbitrary stream of integers. We can then use this as an "oracle" for a determinate process that has three input ports and a single output port. The determinate process uses the arbitrary stream of integers to determine how many tokens to read off each of its input ports before switching to the other input port. The adjective infinity-fair was coined by Park (1982) for such a process.

The intuition behind infinity-fair merge is captured by the above simulation. If both input sequences are infinite, then it reads both of them completely and outputs an interleaving of them. The process will however get stuck at an input port with no more values left to read, if either of the input sequences is finite. The input–output relation of infinity-fair merge consists of triples $(s_i, s_{i'}, s_o)$ of value sequences, such that s_o can be broken up into two subsequences s_1, s_2 such that

(i) s_1 is a prefix of s_i and s_2 is a prefix of $s_{i'}$,

(ii)   if s_1 is infinite, then $s_2 = s_{i'}$, and if s_2 is infinite, then $s_1 = s_i$.

We now give a specific trace set that has the above IO-relation associated with it. If the infinity-fair merge process has input ports i, i' and output port o, then the trace set of the process consists of all sequences $t \in (\{i, i', o\} \times V)^\infty$ such that $\Pi_o(t)$ can be broken up into two subsequences s_1, s_2 such that

(i)   value(s_1) is a prefix of value($\Pi_i(t)$) and value(s_2) is a prefix of value($\Pi_{i'}(t)$),

(ii)   if $\Pi_i(t)$ is infinite, then value(s_2) = value($\Pi_{i'}(t)$),

(iii)   if $\Pi_{i'}(t)$ is infinite, then value(s_1) = value($\Pi_i(t)$),

(iv)   if either of $\Pi_i(t)$ or $\Pi_{i'}(t)$ is finite, then $\Pi_o(t)$ is finite, and either value(s_1) = value($\Pi_i(t)$) or value(s_2) = value($\Pi_{i'}(t)$), and

(iv)   for every prefix t' of t, the prefix of s_1 in t' is a prefix of $\Pi_i(t')$, and the prefix of s_2 in t' is a prefix of $\Pi_{i'}(t)$.

It is also true that infinity-fair merge can implement anyints. To exhibit this, we use an infinity-fair merge with fixed input sequences 0^∞ and 1^∞. Since the output must contain infinitely many 0's and infinitely many 1's, we can read off the lengths of the blocks of 1's and the blocks of 0's, and make the output of *anyints* be the sequence consisting of these lengths. Therefore *anyints* and infinity-fair merge can implement each other, and we can use either one interchangeably in implementability and non-implementability proofs.

Intuitively, fair merge requires sensitivity to the presence or absence of values at input ports. Keller (1978) considered the addition of a command with syntax *poll a*, where a is an input port, to the usual language of determinate node programs consisting of read, write, and internal commands. The command has the effect of indicating whether there is a value available to be read on port a. Thus it is sensitive to the presence or absence of values at input ports. With this command, it is easy to program a fair merge. We can repeatedly poll the two input ports and read a value each time if it is present. This process never waits to read data are not going to appear nor will it consistently favour a particular port. We introduce a primitive called poll that has similar behaviour. A poll primitive has one input port andone output port. It outputs a shuffle of the input sequence and an infinite sequence of ★'s, where ★ is a special value that cannot be produced by other processes. With poll, one can easily implement fair merge. All one needs is a determinate zipper. This is a process with two input ports and one output port that reads values from each input port in strict alternation, and copies the values onto its output port. To implement a fair merge one·puts a poll in front of each input port to the zipper. Then the resulting output can be filtered to remove the ★'s. It is trivial to implement poll with a fair merge—just merge the input stream with an infinite sequence of ★'s. Therefore poll and fair merge can implement each other.

Formally, the input–output relation of fair merge consists of triples $(s_i, s_{i'}, s_o)$ of value sequences, such that s_o is an interleaving of s_i and $s_{i'}$. We can give a specific trace set with this IO-relation associated with it. If the fair merge process has input ports i, i' and output port o, then the trace

set consists of all event sequences $t \in (\{i, i', o\} \times V)^{\infty}$ such that $\Pi_o(t)$ can be broken up into two subsequences s_1, s_2 such that

    (i)   value(s_1) = value$(\Pi_i(t))$, and

    (ii) value(s_2) = value$(\Pi_{i'}(t))$, and

In the next section, we use properties of trace sets of these merge processes to distinguish between them.

5. INEXPRESSIVENESS RESULTS

In this section, we study properties of the merge processes defined in the previous section, and prove that the different merge processes actually form different levels in a hierarchy of expressiveness. The properties that we define and study in this section are two different notions of "monotonicity." We refer to these as "Hoare-monotonicity" and "Smyth-monotonicity," and these are properties of trace sets. We prove that Smyth-monotonicity of trace sets is preserved under network composition, but it turns out to be exceedingly difficult to prove that Hoare-monotonicity of trace sets is preserved under network composition. So a stronger property of trace sets, that implies Hoare-monotonicity, needs to be proved. This is a "continuity" kind of property.

5.1. *Hoare-Monotonicity*

The intuition behind the definition of Hoare-monotonicity is that, in a monotone network, i.e., satisfying the non-disabling input property (Definition 4), arrival of values at input ports cannot disable output of values that were akready enabled before the arrival of the values at the input ports. We formulate this property as a property of trace sets and also describe a stronger property—limit-closure—that is preserved by network composition.

DEFINITION 21. Let t_1 and t_2 be two traces of a network N. Then the *relative order of events* in t_1 is said to be preserved in t_2 if the following two conditions hold:

    (i)   for every port p of N, $\Pi_p(t_1) \sqsubseteq \Pi_p(t_2)$, and

    (ii) for every pair e_1, e_2 of event occurrences in t_1, if e_1 is the ith event on port p_1, and e_2 is the jth event on port p_2, then e_1 precedes e_2 in t_1 iff the ith event on port p_1 in t_2 precedes the jth event on port p_2 in t_2.

DEFINITION 22. Let N be a network with input ports **a** and output ports **b**. We say that Trset(N) is *Hoare-monotone* if, for any trace

$t \in \text{Trset}(N)$ and any finite prefix t' of t such that $\Pi_a(t) = \alpha$, and for any $\alpha' \sqsupseteq \alpha$, there is a trace $t'' \in Trset(N)$ with $\Pi_a(t'') = \alpha'$, with $\Pi_b(t) \underset{\sim}{\sqsubseteq} \Pi_b(t'')$, with t' a prefix of t'', and with the relative order of the events in t preserved in t''.

This says that, given a trace, we can always extend the input, and there will be some trace that represents the response of the network to the new input and in this new trace, the sequence of values on each port will be an extension of, or equal to, the sequence of values seen before. Clearly we cannot expect that every response will be an extension (in the above sense) of, or equal to, the old response since the networks are indeterminate. This definition captures the idea that if there were an enabled output then adding new input would not disable this output.

The definition of the Hoare-monotonicity of a trace set is quite involved since it deals with the orderings of events in traces, and not just with the input and output sequences. We now define a predicate on the input–output relations of networks, such that Hoare-monotonicity of the trace set of a network implies the truth of this predicate on the input–output relation of that network. We give the same name, Hoare-monotonicity, to this predicate.

DEFINITION 23.   The input–output relation of a network N is said to be *Hoare-monotone*, if whenever for some input α, there is an output β, and $\alpha' \sqsupseteq \alpha$, then there is a possible output $\beta' \sqsupseteq \beta$.

Note that since we are only talking of the input–output relation here, we do not have requirements as for Hoare-monotonicity of trace sets, such as the preservation of relative order of events in traces.

The reason that we could not work with Hoare-monotonicity of the input–output relation is that this does not imply Hoare-monotonicity of the trace set, as the example below shows.

EXAMPLE 5.   We describe a process with input ports a', b' and output port c', and with the central automaton of the process having input ports a, b, and output port c. The central automaton repeatedly executes the following conditional statement: if a has no values to be read, then write a 1 on c and then read from b, otherwise first read a value from a, then read a value from b and then write a 1 on c. The input–output relation of this process has the property that when there are no values at a', then the output at c' is of the same length as the input at b'. Moreover, when there is input at a', then the length of the output at c' is either the length of the input at b' or 1 more than the length of the input at b'. It can easily be checked that this is a Hoare-monotone input–output relation. But the trace

set of this process is not Hoare-monotone, as the sequence $(c', 1)$, $(b', 1)$, $(c', 1)$, ... is a trace, but when the input at a' is extended to a non-empty sequence, then the relative order of events in the above trace cannot be preserved, and for some $i > 0$, the ith c'-event follows the ith b'-event.

However, we justify the use of our definition of Hoare-monotonicity of trace sets by arguing that all processes whose central automata are monotone have a Hoare-monotone trace set. The point here is that for such processes, input events commute with all other events. So if t is a trace and γ is the corresponding completed computation sequence, then adding new input events after any prefix γ_p of γ gives us a new computation sequence extending γ_p and with the relative order of events in γ preserved. This can now be $\sqsubseteq$-extended to a completed computation sequence (Panangaden and Stark, 1988). It is of course possible that one could have a process with a non-monotone central automaton that has a Hoare-monotone trace set. Thus our proof is more general than that of Panangaden and Stark (1988).

To prove the preservation of Hoare-monotonicity of trace sets for network composition, we need to consider aggregation and feedback. Of these, aggregation is more or less trivial whereas feedback, not surprisingly, involves an inductive proof. In order for the inductive proof to go through, we need to introduce limit-closure.

DEFINITION 24. Trset(N) is said to be *limit-closed* if it satisfies the following property. Suppose that $\{t_i \mid i < \omega\}$ is a set of traces of N such that for any port p, $\Pi_p(t_i)$ is a chain in the prefix ordering, and the relative order of events is preserved along this chain. Then there is a trace $t \in \text{Trset}(N)$ such that, for any port p, we have $\Pi_p(t) = \bigsqcup \Pi_p(t_i)$ and the relative order of events in t_i is preserved in t for every i.

We now show that a network composed of components with Hoare-monotone and limit-closed trace sets has to have a Hoare-monotone and limit-closed trace set. This is the fundamental theorem on which our first expressiveness result rests.

THEOREM 6. *The aggregate of two networks, N and M, with Hoare-monotone and limit-closed trace sets has a Hoare-monotone and limit-closed trace set.*

Proof. Let M and N be networks with Hoare-monotone and limit-closed trace sets. Let t be a trace of $N \| M$. Then t is a shuffle of t_1 and t_2, where t_1 is a trace of N and t_2 is a trace of M. Extending the input streams of $N \| M$ extends the input streams of N and M. By Hoare-monotonicity of N and M, there are traces t_1' and t_2' of N and M, respectively, with the

extended input and output and with relative order of events in t_1 and t_2 preserved in t'_1 and t'_2 respectively. Then t', an appropriate shuffle of t'_1 and t'_2, is the desired trace of $N \| M$.

The proof that limit-closure is preserved is very similar. One can decompose traces of $N \| M$ to obtain traces of N and M. One can obtain the required limiting vector by using the limit-closure of the trace sets of the components and appropriately shuffling the corresponding traces. ∎

Now we consider the case of feedback. For definiteness, let us suppose that the network in question is M and that the output port b of M is connected to the input port c forming a feedback loop. Let the network M without this feedback loop be called N. Note that N has an extra input port c. Recall that the relationship between the trace sets of the two networks is

$$\mathrm{Trset}(M) = \{ \Pi_{\sim c}(t) \mid t \in \mathrm{Trset}(N) \wedge \mathrm{precedes}(b, c, t) \wedge (\Pi_c(t) = \Pi_b(t)) \}.$$

We will use the following convenient notation: suppose that $\boldsymbol{\alpha}$ is a vector of n sequences, then $\boldsymbol{\alpha}; \gamma$ will mean the vector consisting of the sequences in $\boldsymbol{\alpha}$ and the sequence γ in the $(n+1)$th position.

THEOREM 7. *Suppose the networks M and N are as above. If* $\mathrm{Trset}(N)$ *has a Hoare-monotone and limit-closed trace set, then does* $\mathrm{Trset}(M)$.

Proof. Suppose $t \in \mathrm{Trset}(M)$, $\Pi_I(t) = \boldsymbol{\alpha}$, and $\Pi_O(t) = \boldsymbol{\beta}$, where I and O are the sets of input ports and output ports, respectively, of M. Now we consider extending the input to $\boldsymbol{\alpha}'$, and we pick a finite prefix t_p of t. By Theorem 3, there is some trace $t_0 \in \mathrm{Trset}(N)$ such that $\Pi_I(t_0) = \boldsymbol{\alpha}; \Pi_c(t_0)$ and $\Pi_O(t_0) = \boldsymbol{\beta}$ and $\mathrm{precedes}(b, c, t_0)$ and $\Pi_c(t_0) = \Pi_b(t_0)$. Let t'_p be the finite prefix of t_0 such that $\Pi_{\sim c}(t'_p) = t_p$. Since the original network N is Hoare-monotone, we can extend the input to $\boldsymbol{\alpha}; \Pi_c(t_0)$ and there will be some new trace of N, call it t_1, such that $t_1 \sqsupseteq t'_p$ and $\Pi_I(t_1) = \boldsymbol{\alpha}; \Pi_c(t_0)$ and $\Pi_O(t_1) = \boldsymbol{\beta}_1 \sqsupseteq \boldsymbol{\beta}$ and with the relative ordering of events of t_0 preserved in t_1, and hence $\mathrm{precedes}(b, c, t)$. If $\Pi_c(t_0)$ is infinite, then $\Pi_c(t_1) = \Pi_b(t_1)$, but if $\Pi_c(t_0)$ is finite, then we have no guarantee that $\Pi_c(t_1) = \Pi_b(t_1)$. In the rest of the discussion we describe how to build a trace $t' \in \mathrm{Trset}(N)$, such that $t' \sqsupseteq t'_p$, $\Pi_O(t') = \boldsymbol{\beta}'$, $\Pi_I(t') = \boldsymbol{\alpha}'; \Pi_c(t')$, $\Pi_b(t') = \Pi_c(t')$ and with $\boldsymbol{\beta} \sqsubseteq \boldsymbol{\beta}'$ and $\mathrm{precedes}(b, c, t')$. $\Pi_{\sim c}(t')$ will then be present in $\mathrm{Trset}(M)$ and will demonstrate that M is also Hoare-monotone.

The construction of t' is carried out using traces from $\mathrm{Trset}(N)$. Since we know that $\mathrm{Trset}(N)$ is Hoare-monotone, we can construct t' in successive approximations knowing that at each stage we can do it in such a way that the sequences of events on the ports increase at each stage of the construction. Let t'_1 be a prefix of t_1 upto the first new b-event in t_1. Then, by

Hoare-monotonicity, there is a trace t_2, extending t_1', with the input on input port c extended by a new c-event that matches the new b-event in h_1, and with an extended output.

We iterate this process by successively considering all the b-events, and obtain, in this way, a sequence of traces $t_0, t_1 \cdots t_n \cdots$ such that the event sequences on all the ports are increasing and for which value($\Pi_c(t_{n+1})$) is the prefix of value($\Pi_b(t_n)$) of length $n + 1 + $ (length of $\Pi_c(t_0)$). At any stage, if $\Pi_b(t_n) = \Pi_c(t_n)$, then the trace t_n has the required property. However, it is possible that, when we look at the sequence of traces projected onto c, we obtain an infinite properly increasing chain of finite sequences. In this case, we have a sequence of inputs $\alpha'; \Pi_c(t_i)$ and a corresponding sequence of outputs β_i. By the limit-closure of N, we have a trace t_∞ with input $\bigsqcup (\alpha'; \Pi_c(t_i)) = \alpha'; \bigsqcup \Pi_c(t_i)$ and outputs $\bigsqcup \beta_i$ with the relative order of events from each t_n being preserved. The sequence corresponding to port b is $\bigsqcup \Pi_b(t_i)$ in $\bigsqcup \beta_i$.

Any event in $\bigsqcup \Pi_c(t_i)$ is in some $\Pi_c(t_i)$, hence preceded by a corresponding b-event in the trace t_i, and hence preceded by a corresponding b-event in $\bigsqcup \Pi_b(t_i)$ (because relative order is preserved). Therefore, precedes(b, c, t_∞) holds and so t_∞ is the required trace.

Now we prove that Trset(M) is limit-closed. Suppose $(\alpha_1, \beta_1), \ldots$ are input-output pairs corresponding to traces $t_1, \ldots$ in Trset(M), such that the relative ordering of events is preserved. Let $t_1', \ldots$ be the corresponding traces in Trset(N) such that immprecedes(b, c, t_i') holds for every i. It is then easy to see that the relative ordering of events is preserved along $t_1', \ldots$ By limit-closure of Trset(N), $\exists$ a trace t_∞ with input $\bigsqcup (\alpha_i; \Pi_c(t_i))$ and output $\bigsqcup \beta_i$ with relative order of events in every t_i being preserved in t_∞. Since $\Pi_c(t_i) = \Pi_b(t_i)$ for every i, $\bigsqcup \Pi_c(t_i) = \bigsqcup \Pi_b(t_i)$ and as before, precedes(b, c, t) holds. So $\Pi_{\sim c}(t)$ is a trace of M with the required property. ∎

We now deal with output hiding.

THEOREM 8. Any network obtained by hiding some output ports of a network with a Hoare-monotone and limit-closed trace set has a Hoare-monotone trace set.

Proof. Let N be a network with a Hoare-monotone and limit-closed trace set, and let M be the network obtained by hiding some set S of output ports of N. Let t be a trace of M, and let t_p be a finite prefix of t. Let the input be α and the output be β. By Theorem 5, there is a trace t' of N such that $t = \Pi_{\sim S}(t')$. Let t_p' be the prefix of t' such that $\Pi_{\sim S}(t_p') = t_p$. The output in t' is $\beta; \gamma$ where γ consist of the sequences of events at the ports in S. When α is extended to α', then by Hoare-monotonicity of N, there is a trace $s' \sqsupseteq t_p'$ with input sequences α' and output sequences $\beta'; \gamma'$, $\beta \sqsubseteq \beta'$

and $\gamma \sqsubseteq \gamma'$. Therefore $\Pi_{\sim S}(s')$ is a trace of M with the desired property. ∎

COROLLARY 1. *Any network consisting of components with Hoare-monotone and limit-closed trace sets has a Hoare-monotone trace set.*

Proof. Any network implemented by components with Hoare-monotone and limit-closed trace sets is obtained by composing the components and then hiding some output ports. As proved earlier, any network obtained by composing components with Hoare-monotone and limit-closed components has a Hoare-monotone and limit-closed trace set. Hiding some output ports preserves Hoare-monotonicity of trace sets. Hence the resulting network has a Hoare-monotone trace set. ∎

THEOREM 9. *No total subset of the input–output relation of fair merge can be implemented by any finite network of Hoare-monotone and limit-closed processes.*

Proof. The trace set of any finite network built out of Hoare-monotone and limit-closed processes by aggregation, feedback and output hiding is Hoare-monotone, and hence the input–output relation is also Hoare-monotone. But no total subset of the IO-relation of fair merge is Hoare-monotone. To see this, consider the case where the input sequences of fair merge are 1^{∞} and Λ, the empty sequence. Then 1^{∞} is the only possible output sequence. Now if the sequence on the second input port is extended to 2, then the output sequence must include the value 2, and no such output sequence is an extension of 1^{∞}. ∎

We can therefore distinguish between angelic merge and fair merge, since the following corollary is immediate.

COROLLARY 2. *Angelic merge cannot implement the* IO-*relation of fair merge.*

Note that, although fair merge is also not limit-closed, and limit-closure is preserved under composition, it might seem that we could prove the expressiveness theorem using the limit-closure property instead of Hoare-monotonicity. This is not possible, because infinity-fair merge can be obtained, as in the previous section, by composing some networks with Hoare-monotone and limit-closed trace sets and then hiding some output ports. But it is not limit-closed. Consider the inputs $(1, 2), (1^2, 2), (1^3, 2), \ldots$ to the infinity-fair merge process. $1, 1^2, 1^3, \ldots$ are possible outputs, respectively. The limit of the inputs is $(1^{\infty}, 2)$ and the limit of the outputs is 1^{∞}, which is not a valid output sequence because 2 has not been output. This

shows that we cannot just consider limit-closure, and we really need to talk about Hoare-monotonicity.

5.2. *Smyth-Monotonicity*

Hoare-monotonicity described the property that given a trace t, if we consider an increased input, then there is a trace t' with that increased input and an equal or increased output, and further the relative ordering of events in t is preserved in t'. We now define a "similar" property. Informally, this property says that given a trace t, if we consider a decreased input, then there is a trace t' with that decreased input and an equal or decreased output, and further the relative ordering of events in t' is preserved in t.

DEFINITION 25. Trset(N) is said to be *Smyth-monotone* if for any trace t with input sequences α and output sequences β, if $\alpha' \sqsubseteq \alpha$, then there is a trace t' with input α' and output β' and $\beta' \sqsubseteq \beta$ and the relative order of events in t' is preserved in t.

The next lemmas establish that Smyth-monotonicity is preserved by network composition.

LEMMA 1. *Suppose N_1 and N_2 are networks with Smyth-monotone trace sets. Then their aggregate N also has a Smyth-monotone trace set.*

Proof. Let t be a trace of N. Then t is an interleaving of a trace t_1 of N_1 and a trace t_2 of N_2. Let α and β be the input sequences and output sequences, respectively, in t_1. Let γ and δ be the input sequences and output sequences, respectively, in t_2. Let $\alpha' \sqsubseteq \alpha$ and $\gamma' \sqsubseteq \gamma$. Then, by Smyth-monotonicity of N_1 and N_2, there are traces t_3 and t_4, of N_1 and N_2 respectively, such that the relative order of events in t_3 and t_4 are preserved in t_1 and t_2 respectively. Then an appropriate interleaving of t_3 and t_4 is a trace of N, the relative order of events of which are preserved in t. ∎

LEMMA 2. *If N is a network with a Smyth-monotone trace set, and has an output port p_1 and an input port p_2, then the network $\mathrm{loop}(p_1, p_2, N)$ has a Smyth-monotone trace set.*

Proof. Let M be the network $\mathrm{loop}(p_1, p_2, N)$. We recall that

$$\mathrm{Trset}(M) = \{\Pi_{\sim p_2}(t) \mid t \in \mathrm{Trset}(N) \wedge \mathrm{precedes}(p_1, p_2, t) \wedge (\Pi_{p_1}(t) = \Pi_{p_2}(t))\}$$

Let t' be a trace of M with α and β as the input sequences and output sequences respectively. There is a trace t of N, such that $t' = \Pi_{\sim p_2} t$ and $\mathrm{precedes}(p_1, p_2, t)$ and $\Pi_{p_1}(t) = \Pi_{p_2}(t)$. Let $\delta = \mathrm{value}(\Pi_{p_1}(t))$. Then the

input sequences in t are $\alpha; \delta$. Let $\alpha' \sqsubseteq \alpha$. By Smyth-monotonicity of N, there is a trace r_1 of N, such that the input sequences are $\alpha'; \delta$ and value($\Pi_{p_1}(r_1)) = \delta_1 \sqsubseteq \delta$ and the relative order of events in r_1 are preserved in t.

If $\delta = \delta_1$ then we are done because then $\Pi_{\sim p_2} r_1$ is the desired trace of M. If not, we iterate as outlined below until we finally obtain the desired trace. At each stage of the iteration, we construct new traces such that the relative order of events in newer traces is preserved in the older traces.

The iteration proceeds as follows: If $\delta \neq \delta_1$ then δ_1 is a proper prefix of δ. This implies that $\alpha'; \delta_1$ is a prefix of $\alpha'; \delta$. So, by Smyth-monotonicity of N, there is a trace r_2 with input sequences $\alpha'; \delta_1$ and value($\Pi_{p_1}(r_2)) = \delta_2 \sqsubseteq \delta_1$ and precedes(p_1, p_2, p_2).

If we repeat he above step, we will reach $\delta_i = \delta_{i+1}$ and then we will be done. This procedure will terminate because the prefix ordering is a well-ordering and δ_i must certainly hit $\varnothing$ and stabilize if it does not stabilize earlier. ∎

LEMMA 3. *Any network obtained by hiding some output ports of a network with a Smyth-monotone trace set has a Smyth-monotone trace set.*

Proof. Let N be a network with a Smyth-monotone trace set, and let M be the network obtained by hiding some set S of output ports of N. Let t be a trace of M. Let the input be α and the output be β. By Theorem 5, there is a trace t' of N such that $t = \Pi_{\sim S}(t')$. The output in t' is $\beta; \gamma$ where γ consist of the sequences of events at the ports in S. Let $\alpha' \sqsubseteq \alpha$. Then by the Smyth-monotonicity of N, there is a trace $s' \in \text{Trset}(N)$ with input α' and output $\beta'; \gamma'$, such that the relative order of events in s' is preserved in t'. Then $\Pi_{\sim S}(s')$ is a trace of M with the desired property. ∎

We now define a predicate on input–output relations that is true for any network whose trace set has the Smyth-monotonicity property. We will refer to this predicate as Smyth-monotonicity too.

DEFINITION 26. The input–output relation of a network N is said to be *Smyth-monotone*, if whenever for some input α, there is an output β, and $\alpha' \sqsubseteq \alpha$, then there is a possible output $\beta' \sqsubseteq \beta$.

LEMMA 4. *No total subset of the IO-relation of angelic merge is Smyth-monotone.*

Proof. Consider the following three pairs of sequences as inputs:

1. $\langle 1^\infty, 2^\infty \rangle$,
2. $\langle 1^\infty, \Lambda \rangle$,
3. $\langle \Lambda, 2^\infty \rangle$.

In the second and third cases, the output of the angelic merge is determined. Thus, any total subset of the IO-relation of angelic merge must produce exactly the same output as angelic merge would; namely 1^∞ and 2^∞, respectively. Now we consider the possible outputs in the first case. Note that the input in this case extends the inputs in both cases two and three. Thus if the IO-relation is to be Smyth-monotone *all* the possible outputs must extend both 1^∞ and 2^∞; this is obviously impossible. ∎

THEOREM 10. *No total subset of the* IO-*relation of angelic merge can be implemented by any finite network of processes with Smyth-monotone trace sets.*

Proof. Smyth-monotonicity of trace sets is preserved under aggregation, feedback and output hiding. Therefore the trace set of any finite network of Smyth-monotone processes must be Smyth-monotone, and hence the network must have a Smyth-monotone input-output relation. But no total subset of the IO-relation of angelic merge can have a Smyth-monotone input-output relation. ∎

We can therefore distinguish between infinity-fair merge and angelic merge, since the following corollary is immediate.

COROLLARY 3. *Infinity-fair merge cannot implement the* IO-*relation of angelic merge.*

5.3. Implementability Results

The results in the previous subsection show that infinity-fair merge cannot implement angelic merge, and that angelic merge cannot implement fair merge. Eugene Stark (1990) showed how one could implement infinity-fair merge using angelic merge. We now show how one could implement angelic merge using fair merge, thus providing a straightforward picture of expressibility (Fig. 5).

Figure 6 shows the network that implements angelic merge. The process P is a determinate process whose behaviour can be described as follows: it

fair

angelic

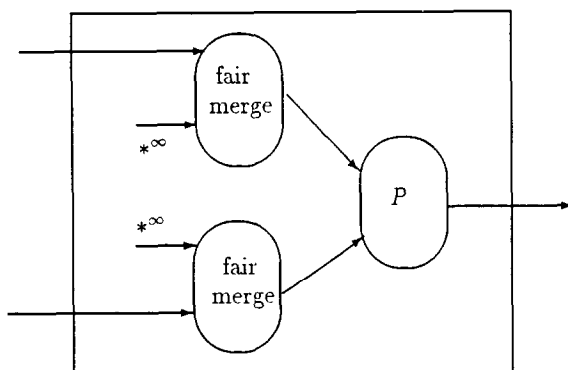infinity-fair

FIG. 5. A straight-line picture of merges.

FIG. 6. Angelic merge from fair merge.

reads values from its first input port and outputs the values that are not *, until it reads two successive *'s, in which case it switches to reading from its second input port and reads and outputs values that are not * until it reads two successive *'s, in which case it switches back to the first input port and continues as above. We could formally describe this as an automaton, but the behaviour should be clear from the given description.

6. CONCLUSIONS

In this paper we have described a trace semantics for indeterminate dataflow networks and used it to prove expressiveness theorems. We have shown that there are three levels of *unbounded indeterminacy* and these are characterized by different monotonicity properties. We make no claims that this is an exhaustive classification, indeed we have discovered related expressiveness results that refine the ones discussed in the present paper (McAllester, Panangaden, and Shanbhogue, 1988). The main significance of the results derives from the significance for order theoretic approaches to fixed-point semantics for dataflow networks. The fact that different monotonicity properties are obeyed suggests that different approaches to fixed-point semantics might be appropriate at the different levels. Thus, for example, one may not need as elaborate an approach if one is only interested in modeling infinity-fair merge and not fair merge.

Recent work by Samson Abramsky (1989) and E. W. Stark (1989b) has in fact shown that there are approaches that give fixed-point semantics for dataflow networks that include infinity-fair merge and angelic merge respectively. The problem of describing a fixed-point principle for dataflow networks with fair merge remains open.

REFERENCES

AALBERSBERG, I. J., AND ROZENBERG, G. (1983), Theory of traces, *Theoret. Comput. Sci.* **60**, No. 1, 1.

ABRAMSKY, S. (1983), On semantic foundations for applicative multiprogramming, *in* "Proceedings of the Tenth International Conference On Automata, Languages And Programming" (J. Diaz, Ed.), pp. 1–14, Springer-Verlag, New York.

ABRAMSKY, S. (1989), A generalized Kahn principle for abstract asynchronous networks, *in* "Proceedings of the 1989 Symposium on Mathematical Foundations of Programming Language Semantics."

APT, K. R., AND OLDEROG, E.-R. (1983), Proof rules and transformations dealing with fairness, *Sci. Comput. Prog.* **3**, 65.

APT, K. R., AND PLOTKIN, G. D. (1986), Countable nondeterminism and random assignment, *J. Assoc. Comput. Mach.* **33**, No. 4, 724.

BROCK, J. D., AND ACKERMAN, W. B. (1981), Scenarios: A model of nondeterminate computation, *in* "Formalization of Programming Concepts" (J. Diaz and I. Ramos, Eds.), pp. 252–259, Lecture Notes in Computer Science, Vol. 107, Springer-Verlag, Berlin/ New York.

BROY, M. (1983), Fixed-point theory for communication and concurrency, *in* "Formal Description of Programming Concepts II" (D. Bjoerner, Ed.), pp. 125–148, North-Holland, Amsterdam.

DEBAKKER, J. W., BERGSTRA, J. A., KLOP, J. W., AND MEYER, J.-J. CH. (1984), Linear time and branching time semantics for recursion with merge, *Theoret. Comput. Sci.* **34**, 134.

FRANCEZ, N. (1986), "Fairness," Springer-Verlag, Berlin/New York.

JONSSON, B. (1989), A fully abstract trace model for dataflow networks, *in* "Proceedings of the Sixteenth Annual ACM Symposium On Principles Of Programming Languages, 1989."

KAHN, G. (1977), The semantics of a simple language for parallel programming, *in* "Information Processing 74," pp. 993–998, North-Holland, Amsterdam.

KELLER, R. M. (1978), Denotational models for parallel programs with indeterminate operators, *in* "Formal Description of Programming Concepts" (E. J. Neuhold, Ed.), pp. 337–366, North-Holland, Amsterdam.

KELLER, R. M., AND PANANGADEN, P. (1985), Semantics of networks containing indeterminate operators, *in* "Proceedings of the 1984 CMU Seminar on Concurrency," pp. 479–496, Lecture Notes in Computer Science, Vol. 197, Springer-Verlag, Berlin/ New York.

KELLER, R. M., AND PANANGADEN, P. (1986), Semantics of digital networks containing indeterminate operators, *Distrib. Comput.* **1**, No. 4, 235.

KOK, J. (1987), A fully abstract semantics for dataflow nets, *in* "Proceedings of Parallel Architectures And Languages Europe 1987," pp. 351–368, Springer-Verlag, Berlin.

LYNCH, N. A., AND STARK, E. W. (1989), A proof of the Kahn principle for input/output automata, *Inform. and Comput.*

LYNCH, N. A., AND TUTTLE, M. (1987), "Hierarchical Correctness Proofs for Distributed Algorithms," Technical Report MIT/LCS/TR-387, M. I. T. Laboratory for Computer Science.

MAZURKIEWICZ, A. (1986), "Advanced Course in Petri Nets," pp. 279–324, Lecture Notes in Computer Science, Vol. 255, Springer-Verlag, Berlin/New York.

MCALLESTER, D., PANANGADEN, P., AND SHANBHOGUE, V. (1988), Nonexpressibility of fairness and signaling, in "Proceedings of the 29th Annual Symposium of Foundations of Computer Science, 1988."

PANANGADEN, P. (1985), Abstract interpretation and indeterminacy, in "Proceedings of the 1984 CMU Seminar on Concurrency," pp. 497–511, Lecture Notes in Computer Science, Vol. 197, Springer-Verlag, Berlin/New York.

PANANGADEN, P., AND STARK, E. W. (1988), Computations, residuals and the power of indeterminacy, in "Proceedings of the Fifteenth ICALP" (Timo Lepisto and Arto Salomaa, Eds.), pp. 439–454, Lecture Notes in Computer Science, Vol. 137, Springer-Verlag, Berlin/New York.

PARK, D. (1989), On the semantics of pair parallelism, in "Proceedings of the Winter School on Formal Software Specification," pp. 504–526, Lecture Notes In Computer Science, Vol. 86, Springer-Verlag, New York.

PARK, D. (1982), The "fairness problem" and non-deterministic computing networks, in "Proceedings of the Fourth Advanced Course on Theoretical Computer Science," pp. 133–161, Mathematisch Centrum.

PLOTKIN, G. D. (1976), A powerdomain construction, SIAM J. Comput. 5, No. 3, 452.

PLOTKIN, G. D. (1982), A powerdomain for countable nondeterminism, in "Proceedings of the Ninth ICALP," Lecture Notes in Computer Science, Vol. 140, Springer-Verlag, Berlin/New York.

PLOTKIN, G. D. (1990), personal communication to Prakash Panangaden (remark attributed to M. Smyth in 1976).

RABINOVICH, A., AND TRAKHTENBROT, B. A. (1988), Nets of processes and dataflow, in "Proceedings of ReX School on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency," Lecture Notes in Computer Science.

RUSSELL, J. (1989), "Full Abstraction and Fixed-Point Principles for Indeterminate Computation," Ph.D. Thesis, Cornell University.

SHANBHOGUE, V. (1990), "The Expressiveness of Indeterminate Dataflow Primitives," Ph.D. Thesis, Cornell University.

STARK, E. W. (1987), Concurrent transition system semantics of process networks, in "Proceedings Of The Fourteenth Annual ACM Symposium On Principles Of Programming Languages," pp. 199–210.

STARK, E. W. (1989a), Concurrent transition systems, Theoret. Comput. Sci., 221.

STARK, E. W. (1989b), "A Simple Generalization of Kahn's Principle to Indeterminate Dataflow Networks," Technical Report TR 89-29, SUNY at Stony Brook.

STARK, E. W. (1990), On the relations computable by a class of concurrent automata, in "Proceedings Of The Fifteenth Annual ACM Symposium On Principles Of Programming Languages," IEEE, New York; also Technical Report TR 88-09, SUNY at Stony Brook.