

## Semantics of Nondeterminism, Concurrency, and Communication\*

NISSIM FRANCEZ

*Department of Computer Science, The Technion, Haifa, Israel*

C. A. R. HOARE

*Programming Research Group, Oxford University, Oxford, OX2 6PE, England*

DANIEL J. LEHMANN

*Department of Computer Science, the Hebrew University, Jerusalem, Israel*

AND

WILLEM P. DE ROEVER

*Department of Computer Science, University of Utrecht, Utrecht, Holland*

Received September 13, 1978; revised July 25, 1979

### I. INTRODUCTION

#### 1. *Background and Motivation*

One of the more important and active areas in the theory of programming languages is that of concurrent programs, specifically their design, definition, analysis, and verification. Due to recent developments in the technology of microprocessors, there is a trend toward languages supporting distributed activities involving communication rather than concurrent activities on some shared resources, mainly memory. Thus, it becomes very important to supply adequate tools for the definition and analysis of such programs and programming languages.

One recent attempt to design such a language was done by Hoare [9], where the language CSP (communicating sequential processes) was presented informally. This is a language for the expression of nondeterministic, concurrent, and communicating programs.

\* Based on work carried out by Francez, Hoare, and de Roever at Queen's University, Belfast, during the academic year 1976-1977, supported by SRC Grant B/RG/74082.

The completion of the paper was supported by NSF Grants MCS-78-673 and MCS-78-7461, during the stay of Francez and Lehmann at the University of Southern California, Los Angeles, and of de Roever at the University of California, Berkeley, 1977-1978.

The main features which distinguish CSP are:

a. Processes are disjoint, do not have any shared variables. The only contact between processes is by means of communication. Concurrency is explicit on the process level.

b. Communication is achieved by means of input and output operations, which are expressed by primitives of the language. Communication plays a double role of both message passing and synchronization. Communication is always directed, having syntactically specified source and target processes, and a strongly typed message.

c. Processes are nondeterministic, and the language allows one to distinguish between two kinds of nondeterminism, discussed in the sequel.

Syntax and informal meaning of CSP are described by way of example:

$$P::[P_1 \parallel P_2 \parallel P_3],$$

where

$$\begin{aligned} P_1::A_1;[P_2?x \rightarrow T_1 \square P_3!y \rightarrow T_2]; \\ * [P_2!u \rightarrow T_3 \square P_3?v \rightarrow T_4]; \\ P_2::*[P_1?s \rightarrow T_5 \square P_1!t \rightarrow T_6 \square P_3?w \rightarrow T_7]; \\ P_3::A_2;*[B_1 \rightarrow T_8 \square B_2 \rightarrow T_9]; \end{aligned}$$

—“ $\parallel$ ” is the parallel composition operator.

— $A_i$ 's are elementary operations such as assignment, skip, abort, etc.

— $T_i$ 's are unspecified (for abbreviating the example) program sections.

— $P_j?x$  (in  $P_i$ ) is an input command, expressing an input request of  $P_i$  from  $P_j$ , and assignment of the input value to the (local) variable  $x$ . Such a command is to be executed only when  $P_j$  is ready to execute a corresponding output command  $P_i!y$ , meaning a request to output the value of  $y$  to  $P_i$ . Either i/o command waits until the corresponding one is ready.

—“ $\square$ ” is the guard separator. Guards may be Boolean ( $B_i$ 's), passable when true, or i/o commands, passable when a corresponding i/o command in the addressed process is ready.

—“ $*$ ” denotes repetition as long as there exists a passable guard.

—“ $;$ ” is sequential composition.

—All processes have disjoint sets of local variables, the only ones to be assigned.

Thus the language is essentially different from various other attempts to consider concurrent programs, e.g., Concurrent Pascal (Brinch-Hansen [2]), which uses monitors [10] to control access to shared variables and procedures, or the language used by Owicki [15] with critical sections.

In this paper, we

(i) Define a formal (denotational) semantics for the main constructs of CSP.

(ii) Clarify, by means of semantical analysis, the relationship between nondeterminism, concurrency, and communication.

(iii) Suggest a rigorous framework for dealing with termination and deadlock of communicating processes.

(In this paper communication is understood as interaction between disjoint processes.)

A natural question that arises is: What can be learned from the formal semantics (of CSP), that would not be apparent already from any informal semantics? To the rather stale debate concerning this issue, we add an observation arising from our personal experience on struggling with the CSP language. During the process of work toward proof rules for this language, to our own amazement we did not succeed in devising a satisfactory intuitive semantics (we tried as hard as we could). It became clear that first a firm, formal foundation for the semantics of CSP had to be found, before any progress on an operational semantics could be made. This was caused by the unrestricted use of *i/o* guards, which resulted in subtleties concerning nondeterminism, termination, scheduling, waiting, and other operational phenomena. Once the basic intuition was captured formally, the operational semantics became clear as well.

The denotational approach to the definition of the semantics of programming languages originated from a pioneering paper by Scott and Strachey [18], who have shown that a relatively small number of basic semantic constructions are needed for an adequate modeling of the realm of meanings of sequential, deterministic programs. The main idea in this approach is to attach to each program some mathematical object as its meaning, or denotation; see Strachey and Milne [22] for a survey of such a characterization of various programming language constructs. The domain of these objects is called a semantic domain. This attachment enables a mathematical proof of properties of the program, and supplies a justification for various inductive proof rules. The process of attaching a meaning to a program uses induction on the syntactic structure of the program. Because of the presence of circular definitions, e.g., recursive procedures, a mathematical theory had to be developed (Scott [19]) in order to prove the existence of the required denotations. According to this theory, a program denotes a partial function from one domain to another. In case of circular definitions (e.g., recursive ones), it can be shown that there exists a unique partial function which satisfies this definition and is a limit of a sequence of partial functions, each of which is at least as well defined as its predecessor. Using tools from lattice algebra and topology, Scott was able to give the appropriate foundations needed for this approach.

More recently, Plotkin [16] extended this approach to cover also nondeterministic programs. By a construction of power domains, which are domains of certain sets of elements from the base domain ordered in an appropriate approximation ordering (see also Egli [6]), he was able to supply denotations to nondeterministic programs by using set-valued partial functions. Another power domain construction appears in Smyth [20].

Milner [14] suggests a construction, called renewals (or resumptions), to give denotations to concurrent programs. However, the programs he has in mind in [14] involve highly interleaved actions on shared variables. Thus, an action by a process may either deliver a result, or give rise to a new process yet to be interleaved with other processes. Since process interaction in CSP is by means of communication rather than by means

of sharing memory, a different description of the semantics is enabled. We suggest another construction which is adequate for communicating processes in CSP (compare also Milne and Milner [13]) and which reduces the degree of interleaving.

The importance of this work is twofold:

- (1) The formal semantics for CSP clarifies many of the complex issues which are needed in order to formulate and validate proof rules for CSP; compare [7].
- (2) The clearer relationship between concurrency, nondeterminism, and communication suggests a way for both the design of language constructs which diminish the danger of deadlock, and the construction of terminating programs.

In the next section, we specify some new contributions of the paper, which expand on (1) and (2) above.

## 2. *What is New?*

### 2.1. *A Priori Semantics*

We regard a single process (taken out of a set of communicating processes) by itself to be a *semantically meaningful entity*, which deserves a denotation of its own. Therefore, we are led to a definition of semantics which attributes a separate meaning to each component  $P_i$  of  $P::[P_1 \parallel \dots \parallel P_n]$ . This kind of semantics is called *a priori* because it denotes *all* the communication capabilities of  $P_i$  when confronted with *any* environment, i.e., all other processes in  $P$ . At the next level we introduce a binding operator  $\mathcal{B}$ , which combines the set of all separate a priori meanings of all  $P_i$ 's to a joint meaning of  $P$ . This  $n$ -ary binding will be compared with Milne and Milner's binary binding in Section III.1.

Since every process has its own (disjoint) local memory, the only contact with other processes being via communication, the degree of interleaving is much smaller than with shared variables, as in each process local computation, involving no communication, does not influence in any way similar computations in other processes. Therefore, the semantics of each separate process is determined by (1) providing an initial state for the local variables and (2) describing its reaction to every possible message requested. Since the values of these messages may vary, each value specifies a different possibility for continuing the computation. These possibilities will be expressed as branches of a tree. This leads to the construction of a new semantic domain, which we call the domain of *history trees*. The histories in question are histories of *communication* (i.e., traces of records of communications that might have taken place). We show that these histories are sufficient for the description of deadlock since deadlock can be caused *only* by some communications failing to happen. With these histories we provide a uniform alternative to mythical ("ghost") variables (e.g., Clint [3] and Owicki [15]), since these variables are used to capture parts of such histories.

In this context, communication involves a message passing from a source process to a target process. Other approaches are, e.g., Milne and Milner [13], where emphasis

is put on *exchange* of values, or Kahn [11], where message transmission is *buffered*, not synchronizing.

## 2.2. Nondeterminism, Concurrency, and Communication

In our semantic domain we distinguish between *two kinds of nondeterminism*, which are expressed in CSP by means of two kinds of guarded commands [4] (this may be one of the innovations of CSP). These two types differ in the way nondeterminism is resolved, and have a different impact on achieving successful communication and deadlock freedom.

The one kind, using Boolean guards, we call local nondeterminism, and is the “old” notion introduced by Dijkstra [4]. Examined in connection with communication, it occurs when a process  $P_i$  can communicate with any of  $P_{i_1} \cdots P_{i_n}$ , and decides *on its own* for which communication to wait, i.e., independent of any consultation with the other processes.

The second kind, using i/o guards, we call global nondeterminism, and is resolved by inspecting the other processes w.r.t. to their willingness to communicate. Only mutual willingness to communicate may result in a de facto communication.

Mixtures of Boolean guards and i/o guards are not considered in this paper. As a simple example, consider the difference between the following two programs:

$$[P_1::[\text{true} \rightarrow P_2 ?x \square \text{true} \rightarrow P_2 !0] \parallel P_2::P_1 !1] \quad (1)$$

and

$$[P_1::[P_2 ?x \rightarrow \text{skip} \square P_2 !0 \rightarrow \text{skip}] \parallel P_2::P_1 !1]. \quad (2)$$

In the first,  $P_1$  may choose the second alternative and cause a deadlock. In the second, successful termination is guaranteed.

In our semantic domain of history trees, this distinction is reflected by letting the history trees have two kinds of nonleaves, on which the binding operator operates differently. This difference reflects also the different implications of the presence of the two kinds of nondeterminism on freedom from deadlock, and will underlie a future proof rule [7].

## 2.3. The Use of End-Signaling for Termination

CSP [9] enables a neat handling of loops which depend on communication guards, and also enables the abortion of the corresponding selection. Upon termination, a process  $P_i$  reaches a final state, which may be sensed by all processes communicating with  $P_i$ . A guard consisting of a communication request is regarded as false iff the target process has already terminated! Otherwise, waiting occurs, because the communication may take place in the future. Thus, a loop depending on guards communicating with  $P_{i_1} \cdots P_{i_n}$  is exited only if *all* of these processes have terminated. Correspondingly, a selection depending on such guards aborts.

Note that termination is in general *not* a property of a single process. As a typical example, consider a "service process" which responds to requests until it receives a signal meaning "terminate!" If presented with an infinite sequence of requests, it should produce an infinite sequence of responses. Only the pair consisting of user-process and service-process may provably terminate. For a more interesting example, displaying how intricate the termination of such programs may be, see Dijkstra [5]. Thus, although every possibility for (non)termination is already present in the a priori meaning of a single process, actual (non)termination of the combination of all processes is determined only on the level of the binding operator  $\mathcal{B}$ .

Dealing with *terminating* processes is an essential feature of the conception of CSP. This may be a distinguishing trait of our formalism compared to Milne and Milner's, in its present form [13]. They consider nonterminating processes, and the problem of termination has never arisen in their work.

## II. A DOMAIN $\mathcal{T}_i$ AND SEMANTIC EQUATIONS FOR $\mathcal{M}[P_i]$ , THE A PRIORI SEMANTICS OF $P_i$

The distinction between local and global nondeterminism implies that a domain of ordinary trees (whose arcs are possibly labeled by records of communication) is not sufficiently structured to reflect this distinction. For, whenever local nondeterminism is resolved, a number of independent alternatives are created, each of which has to be independently confronted with the environment. However, global nondeterminism postpones resolution until the moment of binding since it looks for mutual consent with the environment and can therefore only be resolved during binding.

Therefore, a more refined structure of (finite and infinite) trees with *two* kinds of nodes, called *local* nodes and *global* nodes, is needed.

Since at any stage in its computation, at the level of elementary statements and operations, a process can make only a finite number of nondeterministic choices, any local node has only a finite, positive, number of outgoing arcs.

A global node signals willingness to communicate. Therefore, any arc outgoing a global node is labeled by a target process identifier. Willingness to communicate means either willingness to output a value or to input one of the appropriate type.

At any instant, a process may have a global nondeterministic choice to communicate with a finite number of processes, and therefore the corresponding node will have a finite number of outgoing edges, specifying these processes.

There will be also nodes corresponding to a single input command, and these may have an infinite number of outgoing branches, each labeled by a record of communication corresponding to one of the possible input values. An output command will be represented by a node with a single outgoing edge, labeled by the record of communication corresponding to the output value.

We now proceed with the formal definition of this domain.

DEFINITION. Let  $A$  be any nonempty set, called the *Communication Alphabet*.

Members of  $A$  represent messages passed via i/o from one process to other. In this paper we shall assume that  $A = \{n \mid n \geq 0\}$ .

DEFINITION. A record of communication (roc) is a triple  $\sigma = \langle a, i, j \rangle, a \in A$ .

The intended interpretation of  $\sigma$  is that of the message  $a$  passed from  $P_i$  to  $P_j$ . Let  $1 \leq i, j \leq n$ . Then  $\Sigma_i^j = \{\langle a, i, j \rangle \mid a \in A\}$ , where  $i \neq j$ , and  $\Sigma_i^i = \emptyset$ . Also, let  $\Gamma_i = \{1, \dots, n\} - \{i\}$ .  $\Sigma_i = \bigcup_{j \in \Gamma_i} \Sigma_i^j$  and  $\Sigma^j = \bigcup_{i \in \Gamma_j} \Sigma_i^j$ .

DEFINITION. Let  $V_i$  denote the set of (local) variables of  $P_i$ .  $S_i = [V_i \rightarrow A] \cup \{fail\}$  is the set of states of  $P_i$ .

We consider a state to be mapping from variables to values. We avoid the consideration of "environments" [18] since these do not change in the restricted language we consider; *fail* is a special state denoting a failing computation.

For  $s \in S_i, s \neq fail, x \in V_i,$  and  $a \in A, s_a^x = \lambda y. \text{ if } y = x \text{ then } a \text{ else } s(y).$

Next, we define the complete partial order (cpo)  $\mathcal{F}_i$  as the least solution (in the category of cpo's) of a domain equation. The reader unfamiliar with the technicalities of this kind of equations could consult [21, 22].  $\mathcal{F}_i$  is the domain of history trees corresponding to  $P_i$ , and will be used as the range of the semantic function  $\mathcal{M}[[P_i]]$ , characterizing the a priori semantics of  $P_i$ .

$$\mathcal{F}_i = \left( S_i \cup \left( \bigcup_{j \in \Gamma_i} [\Sigma_j^i \rightarrow \mathcal{F}_i] \right) \cup \left( \bigcup_{j \in \Gamma_i} (\Sigma_i^j \times \mathcal{F}_i) \right) \cup (\Gamma_i \times \mathcal{F}_i)^+ \times S_i \cup \mathcal{F}_i^+ \right)_{\perp}. \quad (1)$$

Here:

$$X^+ = \mu F(X) \quad \text{and} \quad F(X) = \lambda Y \cdot X \oplus X \otimes Y,$$

i.e.,  $X^+$  is the domain of finite (nonempty) sequences over  $X - \{\perp\}$  with the following ordering: there is one bottom element, sequences of different lengths are not comparable, and sequences of the same length are ordered coordinatewise by the ordering inherited from  $X$ .

Equation (1) defines a (finite or infinite) tree in  $\mathcal{F}_i$  to be either bottom or belonging to one of five addends.

Formally: If  $A$  is a partially ordered set then  $A_{\perp}$  is obtained by adding to  $A$  a new bottom element. The union symbol  $\cup$  denotes disjoint union of partially ordered sets (no bottom element is added); union is thus associative. (For instance, an element  $A \cup B \cup C$  is either in  $A$  or in  $B$  or in  $C$  and corresponds therefore with three cases; this is in contrast to the customary usage of the disjoint sum  $A + B + C$  which adds more bottom elements and therefore creates more partially defined objects (and is not associative)).

If  $A$  is a set and  $B$  a cpo,  $[A \rightarrow B]$  is the cpo of all total functions from  $A$  to  $B$  with the obvious ordering.

The symbol  $\times$  denotes Cartesian product  $\mu$  denotes the least fixed point operator.

The coalesced sum  $\oplus$  and the coalesced product  $\otimes$  have been defined in [12, 21]. They are used to avoid the introduction in  $X^+$  of partially undefined and infinite objects, as explained in [12] (warning: In [12] the coalesced sum is denoted by  $\cdot$ ).


Intuitive explanation: The solution to Eq. (1) can be thought of as the domain of all finite and infinite trees, which have the following nodes:

- a. *Leaves* are labeled by  $S_i \cup \{\perp\}$  and have no outgoing arcs.
- b. *Input nodes* have a (possibly infinite) number of outgoing arcs, each labeled by some  $\sigma \in \Sigma^i$ .
- c. *Output nodes* have one outgoing arc, labeled by some  $\sigma \in \Sigma_i$ .
- d. *Global nodes* have a finite, positive, number of outgoing arcs, labeled by  $\Gamma_i$  and an additional unlabeled arc (will be denoted in figures as  $\square$ ).
- e. *Local nodes* have a finite, positive, number of unlabeled arcs (will be denoted in figures as  $\bullet$ ).


Note that each kind of node corresponds to a particular addend in (1). Equation (1) has been designed so as to induce the following ordering on  $\mathcal{T}_i$  (see [12, 21]):

$$T_1 \sqsubseteq T_2 \text{ iff } T_2 \text{ may be obtained from } T_1 \text{ by replacing some } \perp\text{-labeled leaf by some } T' \in \mathcal{T}_i .$$

Thus,



This is very refined ordering, which, e.g., distinguishes between



Such distinctions are natural in a context in which histories (and not only terminal values) matter. One could envisage a semantics of the whole program reflecting the difference of such incomparable trees, e.g., by counting nondeterministic choices involved in producing a given final value. Our specific semantics ignored this difference at the binding level. Both trees will produce the same results while bound to the same environment. The smallest element in this ordering is  $\perp$ . We shall denote by  $i_S, i_I, i_O, i_G$ , and  $i_L$  the corresponding injections from the components to their corresponding copies in  $\mathcal{T}_i$ ; cf. [18].



Thus,

$$i_S: S_i \rightarrow \mathcal{F}_i, \quad i_I: \bigcup_j [\Sigma_j^i \rightarrow \mathcal{F}_i] \rightarrow \mathcal{F}_i, \quad \text{etc.}$$

Before defining the semantic function  $\mathcal{M}[P_i]$ , we define an auxiliary function  $\mathcal{R}$  (replacement), which generalizes functional composition from the sequential case.

When defining the meaning of  $S_1 ; S_2$ ,  $S_1$  has already produced an (intermediate) history tree. Thus, the meaning of  $S_2$  has to be applied to *all* possible leaves of that tree, and will in general depend on the states labeling the leaves of this tree; e.g., if  $S_2$  starts with some Boolean selection, the state will determine the selected branch(es).

$$\mathcal{R}: \mathcal{F}_i \times [S_i \rightarrow \mathcal{F}_i] \rightarrow \mathcal{F}_i$$

and the meaning of  $\mathcal{R}[T, F]$  is the tree obtained by replacing every leaf labeled  $s$  by the tree  $F(s)$ . (We assume  $F(\text{fail}) = \text{fail}$ .)

$\mathcal{R}$  is defined recursively by structural induction on  $T$ :

$$\begin{aligned} \mathcal{R}(T, F) &= \perp, & \text{if } T &= \perp, \\ &= F(T), & \text{if } T &\in i_S(S_i), \\ &= i_I(\lambda\sigma \cdot \mathcal{R}[T(\sigma), F]), & \text{if } T &\in i_I([\Sigma_j^i \rightarrow \mathcal{F}_i]), \\ &= i_O(\langle T \downarrow 1, \mathcal{R}[T \downarrow 2, F] \rangle), & \text{if } T &\in i_O(\Sigma_i^j \times \mathcal{F}_i), \\ &= i_G(\langle \langle \langle T' \downarrow 1 \downarrow 1, \mathcal{R}[T' \downarrow 1 \downarrow 2, F] \rangle, \dots, \langle T' \downarrow k \downarrow 1, \mathcal{R}[T' \downarrow k \downarrow 2, F] \rangle \rangle, F(s) \rangle), \\ & \quad \text{where } s = T \downarrow 2 \text{ and } T' = T \downarrow 1, & \text{if } T &\in i_G(\Gamma_i \times \mathcal{F}_i^k \times s), \\ &= i_L(\langle \mathcal{R}[T \downarrow 1, F], \dots, \mathcal{R}[T \downarrow k, F] \rangle), & \text{if } T &\in i_L(\mathcal{F}_i^k) \end{aligned} \quad (2)$$

( $\downarrow i$  denotes the projection to the  $i$ th component of an  $n$ -tuple).

Note that  $\mathcal{R}$  does not affect infinite paths in  $T$ . Clearly,  $\mathcal{R}$  is continuous. The apparent (notational) complexity of the definition of  $\mathcal{R}$  is due to the fact that a domain equation for denoting our history trees had to be employed. Had we considered such trees as self-explanatory, a notationally simpler definition of  $\mathcal{R}$  could have arisen; compare Hoare [23].


Next, we proceed with the formal definition of  $\mathcal{M}[P_i]: S_i \rightarrow \mathcal{F}_i$ , and the informal explanation of each clause in the definition.

- (1)  $\mathcal{M}[Q](\text{fail}) = \text{fail}$  for all  $Q$ . In the sequel, we assume  $s \neq \text{fail}$ !
- (2) *Input*

$$\mathcal{M}[P_j ?x] = \lambda s \cdot i_I(\lambda\sigma \cdot i_S(s_{\sigma \downarrow 1}^x)).$$

We get a new function, returning for each  $s$  the modified state, which records the side effect of input the value component  $\sigma \downarrow 1$  of the input (roc)  $\sigma$ .


Thus,  $P_j ?x$  creates the following tree, to be called an input node.



(3) *Output*

$$\mathcal{M}[P_j!x] = \lambda s \cdot i_O(\langle\langle s(x), i, j \rangle, i_S(s) \rangle).$$

We get a tree with one arc to be called an output node.



The subtree labeled  $s$  indicates that output has no side effect, and the output value will have to be matched by  $\mathcal{B}$  to an input arc in the tree corresponding to  $P_j$ .

(4) *Assignment*

$$\begin{aligned} \mathcal{M}[x := e] = \lambda s. \text{ if } \mathcal{V}(e, s) = \text{fail then fail} \\ \text{else } i_S(s_{\mathcal{V}(e, s)}^x). \end{aligned}$$

$\mathcal{V}(e, s)$  is an auxiliary function which computes the value of an expression  $e$  in state  $s$ . We assume  $\mathcal{V}(e, s) = \text{fail}$  if  $e$  is undefined. We do not consider recursive functions here, so the evaluation of expressions always terminates, and yields *fail* in cases like division by 0. The meaning of assignment is to update the state  $s$ . Note that it does not create any new arcs in the tree.

(5) *Skipping*

$$\mathcal{M}[\text{skip}] = \lambda s \cdot i_S(s). \quad \text{Obvious.}$$

(6) *Sequential composition*

$$\mathcal{M}[S_1 ; S_2] = \lambda s \cdot \mathcal{B}[\mathcal{M}[S_1](s), \mathcal{M}[S_2]].$$

For a given state  $s$ , we first apply  $\mathcal{M}[S_1]$  to  $s$ , obtaining a tree, say  $T_s$ . Then, we apply

the replacement operator  $\mathcal{R}$  to  $T_s$  and the function  $\mathcal{M}[[S_2]]$ . This reflects the fact that the operation of  $S_2$  depends upon the final state of  $\mathcal{M}[[S_1]]$ , which it continues.

Note that  $\mathcal{M}[[S_1]](\text{fail}) = \text{fail}$  by assumption (case (1)), and  $\mathcal{R}$  may therefore be applied with  $\mathcal{M}$  as an argument. Also, if  $S_1$  has a nonterminating computation,  $\mathcal{M}[[S_1]](s)$  will have an infinite path, which will not be affected by  $\mathcal{R}$ .

(7) *Boolean selection.* We treat here the case of two guards only. The extension to any number of guards should be clear. We assume guards are always defined.

$$\mathcal{M}[[B_1 \rightarrow S_1 \square B_2 \rightarrow S_2]] = \lambda s \cdot T_s,$$


where

$$\begin{aligned} T_s = \text{case } \langle \mathcal{V}(B_1, s), \mathcal{V}(B_2, s) \rangle \text{ of} \\ \langle ff, ff \rangle: i_s(\text{fail}); \\ \langle ff, tt \rangle: \mathcal{M}[[S_2]](s); \\ \langle tt, ff \rangle: \mathcal{M}[[S_1]](s); \\ \langle tt, tt \rangle: i_1(\langle \mathcal{M}[[S_1]](s), \mathcal{M}[[S_2]](s) \rangle). \end{aligned}$$

In case both guards are false, computation is aborted. In case exactly one guard is true, then  $\mathcal{M}$  of the corresponding guarded statement is applied to  $s$ .

In case both guards are true, a *local* node is created, reflecting in its two unlabeled subtrees the two independent continuations, thus recording local nondeterminism, which will cause independent binding of each subtree.

The picture for this case is:



(8) *i/o directed selection.* Again, we shall describe the semantics of a particular case, involving only two guards, both being input guards. The description of  $\mathcal{M}$  for more (or less) than two guards, and for output guards, should be clear.

$$\mathcal{M}[[P_j ?x \rightarrow S_1 \square P_k ?y \rightarrow S_2]] = \lambda s \cdot T_s,$$


where

$$\begin{aligned} T_s = i_G(\langle \langle \langle j, i_1(\lambda \sigma \cdot \mathcal{M}[[S_1]](s_{\sigma_{11}}^x)) \rangle, \\ \langle k, i_1(\lambda \sigma \cdot \mathcal{M}[[S_2]](s_{\sigma_{11}}^y)) \rangle \rangle \rangle, \\ i_s(\text{fail}) \rangle). \end{aligned}$$

An input can be accepted from either  $P_j$  or from  $P_k$ , and then the corresponding

guarded statement will be executed. The decision as to which continuation to take is postponed to the binding time, when the state of  $P_j$  and  $P_k$  will be available, thus reflecting the global nondeterminism.

The picture corresponding to this case is:



The *fail* subtree will be used if both  $P_j$  and  $P_k$  have terminated, a fact to be noticed at binding.

Since all the auxiliary functions applied so far are continuous, the definition of  $\mathcal{M}$  for loops by means of least fixed points is justified.

(9) *Boolean repetition*

$$\mathcal{M}[\![*][B_1 \rightarrow S_1 \square B_2 \rightarrow S_2]\!] = \mu(\lambda F \cdot \lambda s \cdot T_s),$$

where

$$\begin{aligned} T_s = & \text{case } \langle \mathcal{V}(B_1, s), \mathcal{V}(B_2, s) \rangle \text{ of} \\ & \langle ff, ff \rangle: i_s(s); \\ & \langle ff, tt \rangle: \mathcal{R}[\mathcal{M}[\![S_2]\!]](s, F); \\ & \langle tt, ff \rangle: \mathcal{R}[\mathcal{M}[\![S_1]\!]](s, F); \\ & \langle tt, tt \rangle: i_L(\langle \mathcal{R}[\mathcal{M}[\![S_1]\!]](s, F), \mathcal{R}[\mathcal{M}[\![S_2]\!]](s, F) \rangle). \end{aligned}$$

If both guards are false, the loop is exited. In case exactly one guard is true, the corresponding guarded statement is executed and the whole guarded command is attempted again.

In case both guards are true, a local node is created, and both continuations are recorded as subtrees of this node.

(10) *i/o directed repetition*

$$\mathcal{M}[\![*][P_j ?x \rightarrow S_1 \square P_k ?y \rightarrow S_2]\!] = \mu(\lambda F \cdot \lambda s \cdot T_s),$$



where

$$\begin{aligned} T_s = & i_G(\langle \langle \langle j, i_1(\lambda \sigma \cdot \mathcal{R}[\mathcal{M}[\![S_1]\!]](s_{\sigma 11}^x), F) \rangle \rangle, \\ & \langle k, i_1(\lambda \sigma \cdot \mathcal{R}[\mathcal{M}[\![S_2]\!]](s_{\sigma 11}^y), F) \rangle \rangle, \\ & i_s(s) \rangle). \end{aligned}$$

If both processes end, the loop is exited. Otherwise an input is selected (again, postponing the decision from which process), the corresponding statement  $S$  executed, and the whole loop attempted again.

This completes the definition on  $\mathcal{M}[[P_i]]$  by means of semantic equations. As a simple example, reconsider the program presented in the Introduction.

Let  $s \in S_1$ ,  $s' \in S_2$  be two initial local states.



Anticipating the binding function, one can see intuitively that the second program can never fail, since it will choose that alternative in  $P_1$  which will match  $P_2$  (i.e., the first). On the other hand, the first program *may* fail, if  $P'_1$  chooses its second guard as a “wrong” independent choice, thereby causing deadlock.

### III. THE BINDING OPERATOR $\mathcal{B}$

1. The purpose of the binding operator is to attach a joint meaning to a concurrent command  $P::[P_1 \parallel \dots \parallel P_n]$  in an initial state  $\langle s_1, \dots, s_n \rangle$  by using the history trees  $\mathcal{M}[[P_1]](s_1), \dots, \mathcal{M}[[P_n]](s_n)$ .

We restrict ourselves to “closed” concurrent commands since the semantics for, and the proper restrictions to be imposed upon, nonclosed concurrent commands are still under scrutiny; cf. Hoare [9]. (By “closed” we mean that no component  $P_i$  of  $P$  communicates with any process not among  $P_1, \dots, P_{i-1}, P_{i+1}, \dots, P_n$ .)

The computation of  $P$  starts in  $\langle s_1, \dots, s_n \rangle$  and may produce a set of such  $n$ -tuples as final states. While binding, all the histories of communication are “forgotten” and only final states are left; cf. also Milne and Milner [13]. Any occurrence of an infinite computation in any process  $P_i$  is recorded by a single formal value  $\perp$  standing for the “undefined”  $n$ -tuple. We also include two other states, *deadlock*, which records a deadlock situation, and *fail*, which records abortion in any component. The meaning of  $P$  is given by  $\mathcal{B}(\mathcal{M}[P_1](s_1), \dots, \mathcal{M}[P_n](s_n))$ .

In Milne and Milner [13], a binary binding of processes is introduced. Thus,  $[P_1 \parallel P_2 \parallel P_3]$  is interpreted as, e.g.,  $[(P_1 \parallel P_2) \parallel P_3]$ , where “(, )” correspond to the binary binding. The binary binding is associative. This operation leads to intermediate trees which describe a mixture of global *and* local nondeterminism on the same level. This mixture arises from internalization of communication between the bound processes. Consider the following example:

$$P::[P_1 \parallel P_2 \parallel P_3],$$

where

$$P_1::[P_2 ?x \rightarrow P_3 ?y \square P_3 ?y \rightarrow P_2 ?x],$$

$$P_2::P_1 !z,$$

$$P_3::P_1 !u.$$

Then,  $[(P_1 \parallel P_2) \parallel P_3]$  is equivalent to

$$[[\text{true} \rightarrow x := z; P_3 ?y \square P_3 ?y \rightarrow x := z] \parallel P_1 !u].$$

Such mixed trees introduce subtleties, which on the level of program constructs, as seen from the above example, introduce guarded commands with a mixture of Boolean and i/o guards. These subtleties are our reason for excluding such mixtures from the present paper, and considering instead  $n$ -ary binding of closed commands, which avoids the introduction of such trees.

### 1.1. On the Egli–Milner Order

Before defining the binding operator  $\mathcal{B}$  we first need to introduce the concept of the Egli–Milner order to describe the value domain of  $\mathcal{B}$ —a certain collection of subsets of  $S_1 \times \dots \times S_n \cup \{\perp, \text{fail}, \text{deadlock}\}$  and its underlying structure (since  $\mathcal{B}$  is defined recursively,  $\mathcal{B}$ 's existence follows from the usual continuity considerations with respect to this order).

The concept of the Egli–Milner order (Egli [6], Plotkin [16]) dates back to 1975, and constituted a breakthrough in the semantics of nondeterminism, and a fortiori, of concurrency; its application in de Roever [17] resulted in the first comprehensive model—i.e., including a characterization of nondeterminism—of Dijkstra’s predicate transformer; and it is based on a powerful intuition which is best explained in Egli’s unpublished paper, extensively cited in, e.g., de Bakker [1].

Let  $D$  denote any nonempty set with  $\perp \notin D$ .  $P_{E-M}(D \cup \{\perp\})$  denotes the collection  $C$  of all nonempty subsets  $V$  of  $D \cup \{\perp\}$  satisfying: If  $V \in C$  and  $V$  is infinite, then  $\perp \in V$ . Order  $P_{E-M}(D \cup \{\perp\})$  as follows: For  $V_1, V_2 \in P_{E-M}(D \cup \{\perp\})$ .

$$\begin{aligned} V_1 \sqsubseteq_{E-M} V_2 \text{ iff } & \textit{either} \perp \in V_1 \text{ and } V_1 - \{\perp\} \subseteq V_2 \\ & \text{(set-theoretical containment)} \\ & \textit{or} \perp \notin V_1 \text{ and } V_1 = V_2. \end{aligned}$$

$\langle P_{E-M}(D \cup \{\perp\}), \sqsubseteq_{E-M} \rangle$  is a complete partial order, called the Egli-Milner order.

We shall use the property that if  $V_i \in P_{E-M}(D \cup \{\perp\})$  and  $V_i \sqsubseteq_{E-M} V_{i+1}$ , for  $i \in \mathbb{N}$ , then  $y \in \text{l.u.b.}_i V_i$  iff  $\exists j \cdot y \in V_j$ . This concept is related to nondeterminism in that functions with values in  $P_{E-M}(D \cup \{\perp\})$ , such as  $\mathcal{B}$ , describe nondeterministic program behavior. The extension of this concept to subsets of any complete partial order  $\langle L, \sqsubseteq_L \rangle$  has been described by Plotkin [16]. Smyth [20] provides a masterful account how this order can be more simply described for the general case.

### 1.2. The Definition of $\mathcal{B}$

The functionality of  $\mathcal{B}$  for any  $n \geq 2$  is given by

$$\mathcal{B}: \mathcal{T}_1 \times \cdots \times \mathcal{T}_n \rightarrow P_{E-M}((S_1 \times \cdots \times S_n) \cup \{\perp, \textit{fail}, \textit{deadlock}\}).$$

First, we give a recursive definition of  $\mathcal{B}(T_1, \dots, T_n)$  accompanied by an informal explanation of the role of each clause in the definition. We shall number the successive clauses in the margin, for convenience.

$$\mathcal{B}(T_1, \dots, T_n) = \{\perp \mid \exists i (1 \leq i \leq n). T_i = \perp\}. \quad (1)$$

$\{\perp\}$  is the bottom element of  $P_{E-M}$  mentioned above, and denotes undefined information. A  $\perp$ -node in  $\mathcal{T}_i$  is used to describe approximations to elements in  $\mathcal{T}_i$ , and will appear in the approximations to the a priori semantics of loops within  $P_i$ . This clause is needed for the continuity of  $\mathcal{B}$ .

(For all subsequent clauses, assume  $\forall i (1 \leq i \leq n). T_i \neq \perp$ .)

$$\bigcup \{\textit{fail} \mid \exists i (1 \leq i \leq n). T_i = i_S(\textit{fail})\}. \quad (2)$$

*Fail* is a formal value denoting some machine-detectable error, such as a selection with false guards only, etc. It is preserved under the a priori semantics (see clause 1 in the definition of  $\mathcal{M}$ ), and could be used to issue an error message (Goguen [8]). Once such an error occurs within any  $\mathcal{M}[[P_i]](s_i)$ , it will be reflected in the value of  $\mathcal{B}$ . Note that  $\perp \sqsubseteq \textit{fail}$ .

(For all subsequent clauses, assume  $\forall i (1 \leq i \leq n). T_i \neq i_S(\textit{fail})$ .)

$$\bigcup \{\langle s_1, \dots, s_n \rangle \mid \forall i (1 \leq i \leq n). T_i = s_i \in i_S(S_i)\}. \quad (3)$$

This is the case of successful termination of all  $P_i$ 's, each reaching a final state  $s_i \in S_i$ . Then we add the tuple  $\langle s_1, \dots, s_n \rangle$  to the set of final values of  $\mathcal{B}$ .

$$\bigcup \mathcal{B}(T_1, \dots, T_{i-1}, T_i(T_j \downarrow 1), \dots, T_{j-1}, T_j \downarrow 2, \dots, T_n) \\ \text{if } T_i \in i_1([\Sigma_j^i \rightarrow \mathcal{F}_i]) \text{ and } T_j \in i_0(\Sigma_j^i \times \mathcal{F}_j). \quad (4)$$

This is the case of successful communication, where  $T_i$  contains an input node (from  $P_j$ ) and  $T_j$  contains a corresponding output node (to  $P_i$ ), with matching roc's. Then,  $T_i$  is replaced by the subtree corresponding to this matching roc, which is obtained by applying  $T_i$ , which is a function, to this roc, and  $T_j$  is replaced by its (unique) subtree; then  $\mathcal{B}$  is called recursively.

$$\bigcup \{fail \mid T_i \in i_1([\Sigma_j^i \rightarrow \mathcal{F}_i]) \text{ and } T_j \in i_5(S_j), \text{ or } T_i \in i_0(\Sigma_j^i \times \mathcal{F}_i) \text{ and } T_j \in i_5(S_j)\} \quad (5)$$

This is the case of unsuccessful communication, where  $T_i$  is an input node or an output node, and  $T_j \in S_j$  is a final state of  $P_j$ , meaning that  $P_j$  has already terminated. A communication attempt with a terminated process is interpreted as failure, and the value is  $\{fail\}$ .

$$\bigcup \mathcal{B}(T_1, \dots, T_{i-1}, T_i^j, \dots, T_n) \quad \text{if } T_i = \langle T_i^1, \dots, T_i^m \rangle \in i_1(\mathcal{F}_i^+), \quad 1 \leq j \leq m, \quad (6)$$

This is the case of local nondeterminism in  $P_i$  involving selection with Boolean guards. As already noted, the meaning is that *any* of the subtrees of  $T_i$  can be chosen, and bound to the other  $T_j$ 's. Thus, we pick an arbitrary  $j$ ,  $1 \leq j \leq m$ , and replace  $T_i$  by its subtree  $T_i^j$  in the recursive call. Since the union is taken over all possibilities, each  $T_i^j$  will be considered.

$$\bigcup \mathcal{B}(T_1, \dots, T_{i-1}, T_i^j, \dots, T_{j-1}, T_j, \dots, T_n) \\ \text{if } T_i = \langle \langle k_1, T_i^1 \rangle, \langle k_2, T_i^2 \rangle, \dots, \langle k_m, T_i^m \rangle, s \rangle, \\ \text{and} \quad (7)$$

$$\text{either } T_i^j \in i_1([\Sigma_{k_j}^j \rightarrow \mathcal{F}_i]), \text{ and } T_{k_j} \in i_0(\Sigma_{k_j}^i \times \mathcal{F}_{k_j})$$

$$\text{or } T_i^j \in i_0(\Sigma_{k_j}^{k_j} \times \mathcal{F}_i) \text{ and } T_{k_j} \in i_1([\Sigma_{k_j}^{k_j} \rightarrow \mathcal{F}_{k_j}]), \quad \text{for } 1 \leq j \leq m,$$

This is a case of global nondeterminism in  $P_i$ .  $T_i$  is a global node, with subtrees corresponding to communication with  $P_{k_1}, \dots, P_{k_m}$ . For some  $j$ ,  $1 \leq j \leq m$ ,  $T_i^j$  is an input node (corresponding to an input guard), addressing  $P_{k_j}$ , and  $T_{k_j}$  is an output node addressing  $P_i$ . Thus the global nondeterminism can be successfully resolved. Note that this binding step does not reflect the establishment of the corresponding communication. This communication will be detected at the next level of recursion, when the input node  $T_i^j$  is confronted with the output node  $T_{k_j}$ . A similar case arises if  $T_i^j$  is an output node (to  $P_{k_j}$ ) and  $T_{k_j}$  is an input node.



$$\begin{aligned}
& \mathcal{B}(T_1, \dots, T_{i-1}, T_i^p, \dots, T_{j-1}, T_j^q, \dots, T_n) \\
& \text{if } T_i = \langle\langle\langle k_i, T_i^1 \rangle, \dots, \langle k_m, T_i^m \rangle, s \rangle, \\
& \quad T_j = \langle\langle\langle l_1, T_j^1 \rangle, \dots, \langle l_r, T_j^r \rangle, s' \rangle \\
& \text{and for some } p, q, k_p = j, l_q = i, \text{ and} \\
& \text{either } T_i^p \in i_1([\Sigma_j^i \rightarrow \mathcal{F}_i]), \quad T_j^q \in i_0(\Sigma_j^i \times \mathcal{F}_j), \\
& \text{or } T_i^p \in i_0(\Sigma_i^j \times \mathcal{F}_i), \quad T_j^q \in i_1([\Sigma_i^j \rightarrow \mathcal{F}_i]).
\end{aligned} \tag{8}$$

This is another case of resolvable global nondeterminism in both  $P_i$  and  $P_j$ . Each of them has an i/o guard addressing the other with matching roc's. Again, the actual communication will be detected at the next recursive call.

$$\begin{aligned}
& \bigcup \mathcal{B}(T_1, \dots, T_{i-1}, s, T_{i+1}, \dots, T_n) \\
& \text{if } T_i = \langle\langle\langle k_i, T_i \rangle, \dots, \langle k_m, T_i \rangle, s \rangle \\
& \text{and } \forall j (1 \leq j \leq m). T_{k_j} \in i_S(S_{k_j}).
\end{aligned} \tag{9}$$

This is the case of unresolvable global nondeterminism in  $P_i$ , and then  $T_i$  is replaced by the "escape" state  $s$ , which accompanies the guards. The state  $s$  is a proper state if this node was generated in an i/o directed loop (clause 10 in the definition of  $\mathcal{M}$ ) or equals *fail* in case of selection (clause 7 in the definition of  $\mathcal{M}$ ). The global nondeterminism is unresolvable only if all addressed processes  $P_{k_1}, \dots, P_{k_m}$  have terminated.

$$\bigcup \{\text{deadlock} \mid \text{if none of the other clauses is applicable}\}. \tag{10}$$

This case arises when a group of processes are involved in some cyclic communication, while all the rest have terminated. This is a deadlock state, and is recorded as such in the value of  $\mathcal{B}$ .

Note that we are able to detect a *nondeterministically possible* deadlock state. Compare Milne and Milner [13].

In the example at the end of the last section, one would get

$$\mathcal{B}(\mathcal{M}[[P_1]](s), \mathcal{M}[[P_2]](s')) = \{\langle (s)_1^\# \rangle, s'\}$$

whereas

$$\mathcal{B}(\mathcal{M}[[P'_1]](s), \mathcal{M}[[P_2]](s')) = \{\langle (s)_1^\# \rangle, s', \text{deadlock}\}.$$

2. The first thing that has to be done is to show that if the equation is written as  $\mathcal{B} = \tau(\mathcal{B})$ , the functional  $\tau$  is continuous in  $\mathcal{B}$ . Let  $\langle \mathcal{B}^i \rangle_{i=0}^\infty$  be a sequence of partial functions from  $\mathcal{T}_1 \times \dots \times \mathcal{T}_n$  to  $P_{E-M}(S_1 \times \dots \times S_n \cup \{\perp, \text{fail}, \text{deadlock}\})$ ,  $\mathcal{B}_0 \sqsubseteq \mathcal{B}_1 \sqsubseteq \dots \sqsubseteq \mathcal{B}_i \sqsubseteq \dots$ , and  $\mathcal{B}^\infty = \text{l.u.b.}_i \mathcal{B}^i$ ; then  $\tau(\mathcal{B}^\infty)(T_1, \dots, T_n)$  is obtained by replacing all occurrences of  $\mathcal{B}$  in  $\tau$  by  $\mathcal{B}^\infty$ ; from  $\mathcal{B}_0 \sqsubseteq \mathcal{B}_1 \sqsubseteq \dots \sqsubseteq \mathcal{B}_i \sqsubseteq \dots$  it follows that  $\mathcal{B}^\infty(T'_1, \dots, T'_n) = \text{l.u.b.}_i \mathcal{B}^i(T'_1, \dots, T'_n)$  for arbitrary  $T'_i$ ; by continuity of  $\bigcup$  one obtains  $\tau(\mathcal{B}^\infty)(T_1, \dots, T_n) = \text{l.u.b.}_i \tau(\mathcal{B}^i)(T_1, \dots, T_n)$ .

One of the cardinal principles of denotational semantics being that all semantically meaningful functions are continuous, one would like to show next that  $\mathcal{B}$  is continuous in its arguments.  $\mathcal{B}$  being the least upper bound of the sequence  $\Omega \sqsubseteq \tau(\Omega) \sqsubseteq \tau^2(\Omega) \cdots \sqsubseteq \tau^i(\Omega) \cdots$  (where  $\Omega$  is the completely undefined function), it is enough to show that  $\tau^i(\Omega)$  is continuous for every  $i$ ; it suffices to show that if  $A$  is continuous then so is  $\tau(A)$  since  $\Omega$  is obviously continuous.

The first step is to see that if  $A$  is monotone then  $\tau(A)$  is monotone. Suppose  $T_1 \sqsubseteq T_1'$ : Then case analysis shows that if  $a \in (S_1 \times \cdots \times S_n) \cup \{fail, deadlock\}$  and  $a \in \tau(A)(T_1, T_2, \dots, T_n)$ , then  $a \in \tau(A)(T_1', T_2, \dots, T_n)$  and if  $\perp \in \tau(A)(T_1, T_2, \dots, T_n)$  then  $\perp \in \tau(A)(T_1, T_2, \dots, T_n)$ ; the case analysis is tedious but standard, and therefore omitted. The last step is to show that if  $A$  is continuous and  $T_1^0 \sqsubseteq T_1^1 \sqsubseteq \cdots \sqsubseteq T_1^i \sqsubseteq \cdots$  is an ascending sequence whose l.u.b. is  $T_1^\infty$  then, if  $a \in (S_1 \times \cdots \times S_n) \cup \{fail, deadlock\}$  and  $a \in \tau(A)(T_1^\infty, T_2, \dots, T_n)$  then there are  $a, i$  such that  $a \in \tau(A)(T_1^i, T_2, \dots, T_n)$  and that if for every  $i$ ,  $\perp \in \tau(A)(T_1^i, T_2, \dots, T_n)$  then  $\perp \in \tau(A)(T_1^\infty, T_2, \dots, T_n)$ . Both properties are checked by case analysis.

We give a detailed proof of one case.

Assume  $A$  is continuous, and let  $T_1^0 \sqsubseteq T_1^1 \sqsubseteq \cdots \sqsubseteq T_1^i \sqsubseteq \cdots$ ,  $T_1^\infty = \text{l.u.b.}_i(T_1^i)$ . Let  $y \in \tau(A)(T_1^\infty, T_2, \dots, T_n)$ , where  $y = \langle s_1, \dots, s_n \rangle$ . We want to show that  $\exists i$  s.t.  $y \in \tau(A)(T_1^i, T_2, \dots, T_n)$ . From the form of  $\tau$ 's definition, there are six clauses due to which this  $y$  could be generated ( $y$  is a tuple of final states!). These are clauses 3, 4, 6, 7, 8, 9. Since, by assumption,  $A$  is continuous in its arguments (and application and projection are continuous as well), we have that

$$A(T_1^\infty, T_2, \dots, T_n) = \text{l.u.b.}_i A(T_1^i, T_2, \dots, T_n).$$

By a property of  $P_{E-M}$  of a flat domain  $D$ ,  $d \in \text{l.u.b.}_i V_i$  implies that  $\exists j. d \in V_j$ , for  $V_i \in P_{E-M}(D)$ . This implies the claim for clauses 4, 6, 7, 8, 9.

For clause 3, we have that

$$T_1^\infty = s_1 \text{ (} S_1 \text{ has no } \perp \text{!)}, \text{ and therefore } \exists i. T_1^i = s_1; \text{ again the claim follows.}$$

Similar arguments can be given for the remaining arguments  $T_2, \dots, T_n$ .

### A Semantic Variant

According to the semantics presented, a possible outcome of a program is the set  $\{\perp, fail\}$ . This represents a nondeterministic situation, where there is a nonending computation and a failing computation. This could be interpreted operationally as terminating (actually aborting) the *whole* concurrent program once a local failure is detected.

An alternate semantics could be that in the presence of nontermination the result is  $\{\perp\}$ , and any failure is disregarded. In order to achieve this semantics, one has to restrict the application of the "negative" clauses 2 and 5 only if no other, "positive" clause, is applicable.

## ACKNOWLEDGMENTS

We are grateful for helpful remarks from D. Albert, M. Clint, E. W. Dijkstra, D. Harel, R. Milner, G. Plotkin, and A. Pnueli. Special thanks are due to Robert Milne, who discussed in detail previous drafts and helped to improve both contents and presentation. We are grateful to SRC of the United Kingdom and to NSF of the United States for providing funds for living and traveling.

## REFERENCES

1. J. W. DE BAKKER, Semantics and termination of nondeterministic recursive programs, in "Proceedings, 3rd Coll. Automata, Languages and Programming," Edinburgh Univ. Press, Edinburgh, 1976.
2. P. BRINCH-HANSEN, The programming language Concurrent Pascal, *IEEE Trans. Software Engrg.* 1, 2 (1975), 199-207.
3. M. CLINT, Program proving: Coroutines, *Acta Informatica* 2, No. 1 (1973), 50-63.
4. E. W. DIJKSTRA, "A Discipline Programming," Prentice-Hall, Englewood Cliffs, N. J., 1976.
5. E. W. DIJKSTRA *et al.*, An elephant inspired by the Dutch National Flag, EWD 608, Burroughs-Nuenen, 1977; see also EWD 607.
6. H. EGLI, "A Mathematical Model for Nondeterministic Computations," Technological University, Zurich, 1975.
7. K. R. APT, N. FRANCEZ, AND W. P. DE ROEVER, A proof system for communicating sequential processes, *TOPLAS*, submitted for publication.
8. J. GOGUEN, Abstract errors for abstract data types, in "Proceedings, IFIP Working Conference on Formal Description of Programming Concepts, 31 July to 5 August 1977, New Brunswick."
9. C. A. R. HOARE, Communicating sequential processes, *Comm. ACM* 21 (1978).
10. C. A. R. HOARE, Monitors: An operating systems structuring concept, *Comm. ACM* 17 (1974), 549-557.
11. G. KAHN, The semantics of a simple language for parallel programming, *IFIP*, 1974.
12. D. J. LEHMANN AND M. B. SMYTH, Algebraic specifications of data types: A synthetic approach, *Math. Systems Theory*, in press. (Summary in "Proceedings, 18th Annual Symposium on F.O.C.S. Providence, R. I., Oct. 1977," pp. 7-12.)
13. G. MILNE AND R. MILNER, Concurrent processes and their syntax, *J. Assoc. Comput. Mach.* 26 2 (1979).
14. R. MILNER, Processes: A mathematical model of computing agents, in "Logic Colloquium 1973," North-Holland, Amsterdam, 1973.
15. S. OWICKI AND D. GRIES, An axiomatic proof technique for parallel programs, I, *Acta Informatica* 6 (1976), 319-340.
16. G. D. PLOTKIN, A power domain construction, *SIAM J. Comput.* 5, No. 3 (September 1976)
17. W. P. DE ROEVER, Dijkstra's predicate transformer, nondeterminism, recursion, and termination, in "Proceedings, Conference on Mathematical Foundation of Computer Science, 1976," Lecture Notes in Computer Science, Springer-Verlag, New York/Berlin, 1976.
18. D. SCOTT AND C. STRACHEY, Towards a mathematical semantics for computer languages, in "Proceedings, Symposium on Computers and Automata," Microwave Research Institute, 1971.
19. D. SCOTT, Outline of mathematical theory of computation, in "Proceedings, 4th Princeton Conf. on Info. Sci and Sys., 1970."
20. M. SMYTH, Power domains, *J. Comput. System Sci.* 16 (1978), 23-36.
21. J. STOY, "Denotational Semantics of Programming Languages: The Scott-Strachey Approach," MIT Press, Cambridge, Mass., 1977.
22. C. STRACHEY AND R. MILNE, "A Theory of Programming Language Semantics," Chapman & Hall, London, 1977.
23. C. A. R. HOARE, "A Model for Communicating Sequential Processes," Oxford, December 1978.