

PORT AUTOMATA AND THE ALGEBRA OF
CONCURRENT PROCESSES¹

Martha Steenstrup
Michael A. Arbib
Computer and Information Science Department,
Ernest G. Manes
Mathematics and Statistics Department
COINS Technical Report 81-25

(October 1981)

Abstract

We model a process -- the unit of concurrent processing -- as a port automaton, a certain type of nondeterministic sequential machine. We show how to compute the response of these automata, and study the properties of their port interconnections. We are motivated by the work of Milne and Milner, but show that their use of domains, continuous functions, and powerdomains is unnecessary. Instead we can use sets, maps and subsets by applying our theory of greatest fixpoints to the study of the response of nondeterministic machines.

1 The research reported in this paper was supported in part by a grant from the National Science Foundation, No. MCS-8005112, with co-principal investigators M.A. Arbib and E.G. Manes.

Introduction -

Within the past twenty-five years, concurrent computation has moved from the realm of theoretical speculation to that of practical implementation. During this period of time, numerous theoretical models of concurrency have appeared in the literature, [Dijkstra 68b], [Petri 73], [Campbell and Habermann 74], [Hewitt 77], [Brinch Hansen 78], and [Hoare 78] to mention a few. The purposes for which the various models have been developed include the following: to facilitate the investigation of the behaviour of concurrent processes, to aid in the designing of concurrent systems, to illustrate specific process synchronization problems, and to verify the correctness of parallel programs.

Perhaps the simplest yet most versatile model of concurrency is that introduced by Hoare, Communicating Sequential Processes (CSP). Many of the extant concurrency models are based upon CSP, and several of them attempt to formalize the description of the behaviour of these processes. One such model is that of [Milne and Milner 79], in which a process is defined as a certain fixed point of a particular powerdomain equation. The powerdomain construction, utilized by several authors as a formal means for capturing the behaviour of nondeterministic processes, is rather complicated and does not afford a particularly intuitive view of communicating sequential processes. We will discuss this issue in more detail in a later section. It was this aspect of the Milne-Milner model that motivated us to search for an alternative representation.

The Model -

Our formal model of concurrent computation is a derivative of CSP, along the lines of Milne and Milner's, and was originally designed to provide a simple theoretical framework in which to study concurrency. We recognize that a process can be represented as a particular type of nondeterministic sequential machine, which we call a port automaton. Before we present the formal definitions, however, we introduce our informal notion of a process, in order to provide the reader with an intuitive basis from which to examine the formal model.

We consider a process to be a formal machine which is capable of communicating with its environment. Each process possesses a set of ports through which communications take place. Communication consists of a simultaneous exchange of values through a port, that is, a process receives a value from the environment while at the same time the process sends a value to the environment, both the sending and the receiving occurring at the same port. (Actually, in all instances that we consider, the flow of communication is essentially unidirectional, in that one of the exchanged values is always the trivial value, which we denote by "#".)

Each port of a process may be in one of two modes at any given instant: inactivated - not capable of communication, or activated - capable of communication; an activated port may be engaged in communication. Although more than one port of a process may be activated at one time, at most one port may

actually be communicating. The receipt of an input value during communication effects a state transition in the process.

To achieve concurrency, we construct networks of these processes by connecting together ports belonging to separate processes. A network of processes may be regarded both as the set of its component processes and as a single process itself. When viewed as the latter, the connected ports are no longer visible to the environment. Let us consider, for this case, the state transition resulting from an external communication. This state change may involve several state transitions resulting from internal communications. The manner in which these internal transitions are effected is not necessarily deterministic, since the order in which the internal communications occur is not necessarily predetermined. Thus, the state transition caused by an external communication in a network is inherently nondeterministic.

We now proceed to formalize our notion of a process by giving the definition of a port automaton.

Definition - A port automaton P is a collection of objects and maps $(L, Q, \tau, \delta, \beta, X, Y)$ where

L is the set of ports,

Q is the set of states,

$\tau \in 2^Q$ is the set of initial states,

$X = (X_i : i \in L)$, where X_i is the input set for port i ,

$Y = (Y_i : i \in L)$, where Y_i is the output set for port i ,

$\delta : Q \times \bigsqcup_{i \in L} X_i \rightarrow 2^Q$ is the transition map,
where $\bigsqcup_{i \in L} X_i = \{(x, i) : x \in X_i\}$ is the disjoint union of the X_i 's,

$\beta = (\beta_i : i \in L)$, where $\beta_i : Q \rightarrow Y_i$ is the output map for port i ,

all subject to the axiom that for each $q \in Q$

$$\{x \in X_i : \delta(q, (x, i)) \neq \emptyset\} = \emptyset \text{ or } X_i.$$

Definition - We say that port i of P is activated in state q if for all $x \in X_i$, $\delta(q, (x, i)) \neq \emptyset$. Otherwise, we say port i is inactivated.

Now that the reader has been acquainted with both the intuitive and the formal descriptions of a process, we present a representative example of a process modelled as a port automaton. The example has been adapted from [Milner 80]. We will study this example again in more detail when we discuss networks of processes; it is presented here in a somewhat simplified form in order to aid the reader in relating it to the formal definition. The following figure is a schematic depiction of the port automaton, the labels coupled with the descriptions indicating the port names and the nature of the communication passing in the given direction, respectively.

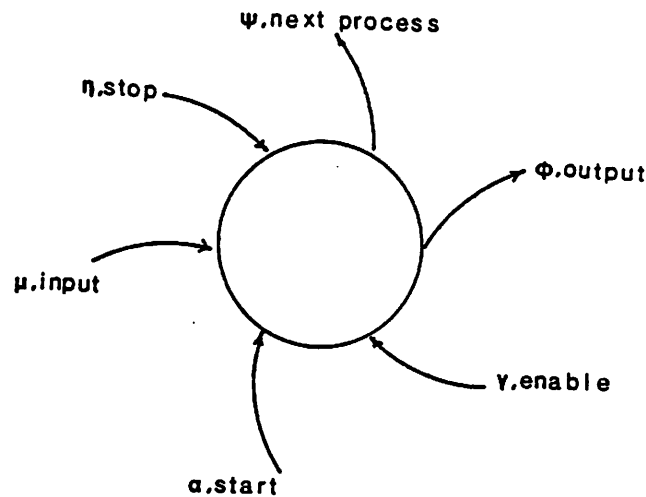


Figure 1

The port automaton in Figure 1 functions as follows. Input arriving through γ enables the automaton, which then waits for a start signal through α . Once this is received, the automaton waits for an integer input through μ and begins processing. At this phase, the automaton may either send an enable signal through ψ or receive a stop processing signal through η . If ψ is sent first, then the automaton waits for the stop signal through η . When a stop signal is received, integer output is sent through ϕ . Then a signal to enable the next automaton is sent through ψ , provided it has not already been sent previously. Pictorially, we may represent this using nondeterministic state graph notation, where each node corresponds to a particular state of the automaton, and each arc corresponds to the port and input value responsible for the state change.

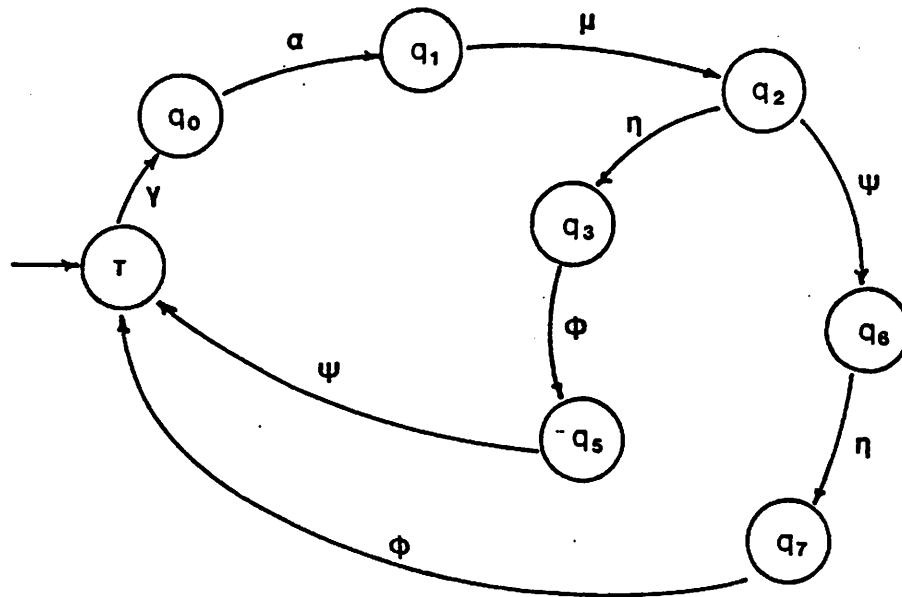


Figure 2

(Note that to be precise, the arc labelled μ really should be represented by multiple arcs leading to multiple states, one for each input value.)

Our next goal is to deduce the reachability map and the response map; in reference to formal machines, the reachability map determines, given an initial state and a set of input values, the state of the machine after receiving these values, while the response map determines, given an initial state and a set of input values, the output of the machine after receiving these values. To accomplish this goal, we must first indicate the order in which inputs from various ports are received. Inputs to a port automaton can be thought of as a set of queues, each queue corresponding to the values to be input at a specific port. Using our automaton of Figure 1, we might have the following arrangement of input queues.

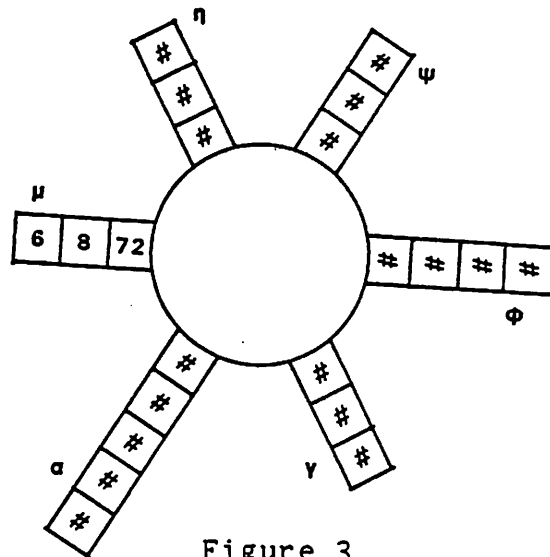


Figure 3

The order in which values are stripped from these queues is, in one sense, dependent upon the order in which the ports become activated. Following our concept of a process, a port automaton may receive an input at any activated port of the automaton, and if more than one port is activated at a given instant, then it is not possible to predetermine which port will actually engage in communication. Formally, inputs to a port automaton are of the form

$$w = (w_1, \dots, w_n), \text{ where } L = \{1, \dots, n\} \text{ and for } 1 \leq i \leq n, w_i \in X_i^*. \text{ This implies that } w \in \prod_{i \in L} X_i^*.$$

We refer to a specific ordering of the ports at which the input values arrive as an interleaving. The order in which inputs are stripped from the various X_i^* 's can also depend upon the interleaving specified. Note that the ordering thus specified may not coincide with the ordering of the port activations; this discrepancy may lead to undefined states in

the computation sequences. An interleaving s may be regarded as an element of L^* . If we pluck input values from the various queues and concatenate them according to the order specified by s , we obtain a string in $(\coprod_{i \in L} X_i)^*$. Thus, each interleaving s induces a mapping from $\prod_{i \in L} (X_i^*)$ to $(\coprod_{i \in L} X_i)^*$, formally described as follows.

For $s \in L^*$, $\phi_s : \prod_{i \in L} (X_i^*) \rightarrow (\coprod_{i \in L} X_i)^*$, where

$$\phi_\Lambda(w) = \Lambda, \text{ for all } w \in \prod_{i \in L} (X_i^*),$$

$$\phi_{j \cdot t}(w) = \begin{cases} \phi_t(w), & \text{if } w_j = \Lambda, \\ x \cdot \phi_t(w_1, \dots, w_{j-1}, v, w_{j+1}, \dots, w_n), & \\ & \text{if } w_j = x \cdot v. \end{cases}$$

At last, we are ready to define the reachability map and the response map.

Definition - The reachability map for a port automaton P is $r = (r_s : s \in L^*)$, where $r_s : \prod_{i \in L} (X_i^*) \rightarrow 2^Q$ is the map $\tilde{r} \cdot \phi_s$, and where $\tilde{r} : (\coprod_{i \in L} X_i)^* \rightarrow 2^Q$ is the conventional reachability map:

$$\tilde{r}(\Lambda) = \tau,$$

$$\tilde{r}(v \cdot (x, i)) = \delta^\#(\tilde{r}(v), (x, i)), \text{ where } i \in L, \\ x \in X_i, v \in (\coprod_{i \in L} X_i)^*, \text{ and } ()^\#$$

is extension by union, (ie.

$$\delta^\#(\tilde{r}(v), (x, i)) =$$

$$\bigcup_{r_j \in \tilde{r}(v)} \delta(r_j, (x, i)).$$

Definition - The response map for a port automaton P is $f = (f_s : s \in L^*)$, where $f_s : \prod_{i \in L} (X_i^*) \rightarrow (\prod_{i \in L} 2^{Y_i}) / \emptyset$ (see footnote) is the map $\tilde{f} \cdot \phi_s$, where $\tilde{f} : (\prod_{i \in L} X_i^*) \rightarrow (\prod_{i \in L} 2^{Y_i}) / \emptyset$ is the map

$$\tilde{f}(\Lambda) = \emptyset,$$

$$\tilde{f}(v \cdot (x, i)) = \beta_i^{\#}(\tilde{r}(v)), \text{ where } i \in L, x \in X_i, \\ v \in (\prod_{i \in L} X_i)^*, \text{ and } ()^{\#} \text{ is} \\ \text{extension by union.}$$

Processes Do Not Need Powerdomains -

Milne and Milner require a great deal of algebraic structure for their definition of process. We argue in this section that such structure is unnecessary, but first we recall their setting. (The exact definitions are not required for the present discussion; they may be found in [Milne and Milner 79], pp. 304-305.) A domain D is an ω -algebraic complete partial order (cpo). The cartesian product of two domains is a domain. For any denumerable indexing set L and any family $(D_i : i \in L)$ of domains, the indexed sum

$$D_i = \{(i, d) : i \in L, d \in D_i\} \cup \{\perp_s\}$$

is also a domain, with $s \sqsubseteq s'$ iff $s = \perp_s$, or $s = (i, d)$, $s' = (i, d')$, and $d \sqsubseteq d'$. For any pair D, E of domains, the function domain $D \rightarrow E$ of continuous functions from D to E is a domain

The quotient set $(\prod_{i \in L} 2^{Y_i}) / \emptyset$ means that we identify the empty sets from the various 2^{Y_i} to a single set. Thus, $(\prod_{i \in L} 2^{Y_i}) / \emptyset$ is equivalent to $(\prod_{i \in L} (2^{Y_i} - \{\emptyset\}) + \{\emptyset\})$.

with the ordering

$$f \sqsubseteq f' \iff (\forall d \in D)(fd \sqsubseteq f'd).$$

Finally, for any domain D , there is a domain $\mathcal{P}(D)$, the powerdomain of D , whose members are a certain subset of the power set of D and whose construction is a slight variant of that given by [Smyth 76].

Given a finite set L of ports, Milne and Milner would (with slightly different notation) introduce a pair of domains X_i and Y_i for each i in L , where they call L a sort. They then define the domain P of processes of sort L as a particular solution of the isomorphism

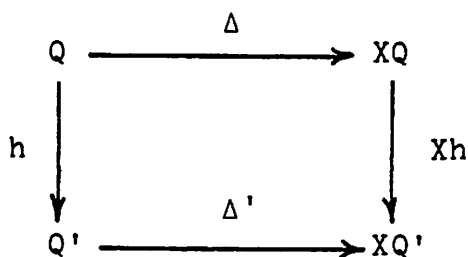
$$P_L = \mathcal{P}\left(\coprod_{i \in L} (Y_i \times [X_i \rightarrow P_L])\right) \quad (1)$$

They note (p. 305) that "the existence of such isomorphisms (without using \mathcal{P}) is due to Scott; a major purpose in the work of Plotkin and Smyth was to justify the use of the powerdomain construction in such isomorphisms." Scott was motivated to use domains and continuous functions because one can then solve the isomorphism $D = [D \rightarrow D]$, but one cannot solve $D = D^D$ in the category Set of sets and maps, where D comprises all maps from set D to itself. Similarly, Smyth and Plotkin introduced the powerdomain construct so that they could solve the isomorphism $D = \mathcal{P}(D)$, since $D = 2^D$ is not solvable.

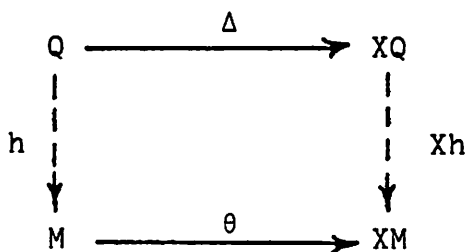
Our critique of (1) is two-fold. (a) It provides an unwieldy conceptual apparatus which makes the specification of even a simple process far less intuitive than the formal machine

approach presented above. (b) The job that domains, continuous functions, and powerdomains were introduced to handle can in fact be accomplished in the category of sets. Claim (a) we take to be clear to any reader of Milne and Milner; for claim (b) we must introduce the notion of greatest fixpoint from [Arbib and Manes [80]].

Let X be any endofunctor of the category of sets. We introduce a category CoDyn(X) whose objects (Q, Δ) are X -codynamics, that is maps $\Delta: Q \rightarrow XQ$, and whose morphisms $h: (Q, \Delta) \rightarrow (Q', \Delta')$ are maps $h: Q \rightarrow Q'$ for which we have commutativity of



We say that X has a greatest fixpoint (M, θ) if (M, θ) is terminal in the category CoDyn(X), that is, if for each X -codynamics (Q, Δ) , there is a unique $h: Q \rightarrow M$ such that



Such a θ , if it exists, is an isomorphism $M \cong XM$. Moreover, the

solution (M, θ) is itself unique up to isomorphism in CoDyn(X).

The crucial point for our discussion is the first example given by Arbib and Manes. They note that a finite automaton (without specified initial state) can be given by a pair of maps $\delta : Q \times A \rightarrow Q$ (the next-state map, transition map, or dynamics) and $\beta : Q \rightarrow Y$ (the output map) or, equivalently, by a single map

$$\Delta_{\beta, \delta} : Q \rightarrow (Y \times [A \rightarrow Q]), q \mapsto (\beta(q), \delta(q, \cdot)) \quad (2)$$

In other words, such an automaton can be presented as a codynamics of the functor X defined by

$$QX = Y \times [A \rightarrow Q].$$

They then observe that the greatest fixpoint for this functor is the set of all response maps $[A^* \rightarrow Y]$, with the isomorphism

$$[A^* \rightarrow Y] \xrightarrow{\cong} Y \times (A \rightarrow [A^* \rightarrow Y]) \quad (3)$$

given by $h \rightarrow (h(\Lambda), a \mapsto hL_a)$, where $hL_a(w) = h(aw)$. The verification that (3) is the greatest fixpoint of our X is given by checking (corresponding to (2)) that there is a unique σ satisfying

$$\begin{array}{ccc} Q & \xrightarrow{\Delta_{\beta, \delta}} & XQ \\ \sigma \downarrow & & \downarrow X\sigma \\ [A^* \rightarrow Y] & \xrightarrow{\cong} & X[A^* \rightarrow Y] \end{array} \quad (4)$$

In fact, we can easily check that σ is the observability map

which sends each q to its response map $\sigma(q) : A^* \rightarrow Y$ which specifies, for each input sequence w , the output $\sigma(q)(w) = \beta(\delta^*(q,w))$ that the automaton would emit after reading in w from q as initial state.

Now compare Milne and Milner's fixpoint definition (1) of the domain of processes with our definition (3) and (4) of the set of response maps as a greatest fixpoint. We may say that the Milne-Milner approach to processes is like a theory of automata which only talks of response maps $[A^* \rightarrow Y]$, and never discusses finite-state implementations ($\delta : Q \times A \rightarrow Q$, $\beta : Q \rightarrow Y$). Our approach allows us to exploit the benefits of both interpretations, adopting whichever is more convenient to the purpose at hand, and moving back and forth where necessary. Moreover, it does so using sets together with ordinary subsets and functions, rather than introducing the more complex apparatus of domains, powerdomains, and continuous functions.

With this background, let us re-analyze our definition of a port automaton, and of the response map of such an automaton. The port automaton $(L, Q, \tau, \delta, \beta, X, Y)$ may be re-presented as a codynamics

$$\Delta_{\beta, \delta} : Q \rightarrow \prod_{i \in L} (Y_i \times [X_i \rightarrow 2^Q]) \quad (5)$$

which makes explicit the linkage which requires that the input and output must use the same port i in any activation. The form corresponding to (1) would be

$$Q \rightarrow \mathcal{P} \left(\prod_{i \in L} (Y_i \times [X_i \rightarrow Q]) \right) \quad (6)$$

which differs from (5) in the use of \mathcal{P} rather than the subset operator, and in allowing nondeterminism to couple a choice of output with a choice of deterministic next-state map. We shall not discuss the latter difference further. The crucial question is "Can we afford to work with sets and 2^Q , or must we work with domains and replace 2^Q by $\mathcal{P}(Q)$ in (5)?"

For the purpose of the present section (but not for the study of interconnection), we limit ourselves to the case in which L is a singleton, and consider

$$\Delta_{\beta, \delta} : Q \rightarrow Y \times [A \rightarrow 2^Q] \quad (7)$$

Note that the functor implicit here cannot have a greatest fixpoint, since we can solve

$$Q = Y \times [A \rightarrow 2^Q]$$

in Set only if Y is empty, as we see by a simple cardinality argument. However, all is not lost, for we may replace the $\Delta_{\beta, \delta}$ of (7) by the map

$$\hat{\Delta}_{\beta, \delta} : \bar{Q} \rightarrow 2^Y \times [A \rightarrow \bar{Q}] \quad (8)$$

where $\bar{Q} = 2^Q$, and we have

$$\text{pr}_1 \cdot \hat{\Delta}_{\beta, \delta}(p) = \{\text{pr}_1 \cdot \Delta_{\beta, \delta}(q) : q \in p\}$$

while

$$(\text{pr}_2 \cdot \hat{\Delta}_{\beta, \delta}(p))(a) = \cup \{(\text{pr}_2 \cdot \Delta_{\beta, \delta}(q))(a) : q \in p\}$$

[Arbib and Manes 75] have given a general theory of nondeterministic dynamics $QX \rightarrow QT$, where QT may be 2^Q ,

probability distributions over Q , fuzzy sets on Q , (or even $\mathcal{P}(Q)$ [D. Lehmann, personal communication, 1976]). The correct setting for studying observability maps for nondeterministic processes ($Q \times A \rightarrow 2^Q$) then proves to be at the level of (8) rather than that of (7). We have chosen not to develop the general theory here, but we note that in (8) both $\bar{Q} = 2^Q$ and 2^Y are semilattices under set inclusion; that the collection of all set-theoretic maps, $[A \rightarrow 2^Q]$, then inherits a semilattice structure; and then that the cartesian product $2^Y \times [A \rightarrow \bar{Q}]$ is a semilattice. (The general theory requires that T be a fuzzy theory (= monad = algebraic theory), and that semilattices be replaced by T -algebras.) The greatest fixpoint we seek, then, is not for the endofunctor $QX = Y \times [A \rightarrow 2^Q]$ of Set (we have already seen that $QX \cong Q$ cannot be satisfied for this X), but rather for the endofunctor $\hat{Q}X = 2^Y \times [A \rightarrow \bar{Q}]$ in the category of semilattices and their homomorphisms. The general theory of [Arbib and Manes 75] then specializes to the observation that X does have a greatest fixpoint, namely $[A^* \rightarrow 2^Y]$, equipped with the semilattice isomorphism

$$L : [A^* \rightarrow 2^Y] \rightarrow 2^Y \times [A \rightarrow [A^* \rightarrow 2^Y]]$$

which pairs each $h \in [A^* \rightarrow 2^Y]$ with $h(\Lambda) \in 2^Y$ and $(a \mapsto hL_a)$ in $[A \rightarrow [A^* \rightarrow 2^Y]]$. The unique semilattice homomorphism $\sigma : 2^Q \rightarrow 2^{[A^* \rightarrow Y]}$ satisfying

$$\begin{array}{ccc}
 2^Q & \xrightarrow{\hat{\Delta}_{\beta, \delta}} & 2^Y \times [A \rightarrow 2^Q] \\
 \sigma \downarrow & & \downarrow \hat{\chi}\sigma \\
 [A^* \rightarrow 2^Y] & \xrightarrow{L} & 2^Y \times [A \rightarrow [A^* \rightarrow 2^Y]]
 \end{array}$$

is indeed the nondeterministic observability map.

The key point in all this is that we move from the unsolvable $Q = 2^Y \times [A \rightarrow 2^Q]$ to the solvable $\bar{Q} = 2^Y \times [A \rightarrow \bar{Q}]$ subject to the restriction that this be a semilattice isomorphism. We note the following:

(i) We use arbitrary subsets 2^Y and maps $A \rightarrow \bar{Q}$, with no restriction to powerdomains or continuous maps.

(ii) The greatest fixpoint $[A^* \rightarrow 2^Y]$ is not of the form 2^B for any B.

(iii) The use of semilattices is the result of the choice $QT = 2^Q$ to express the nondeterminism in our processes. We can just as easily handle stochastic nondeterminism by replacing semilattices by T-algebras (see [Arbib and Manes 75] for the general theory) for $QT =$ probability distributions over Q.

Machine Interconnection -

Up until now, we have been laying a formal foundation upon which to build our understanding of concurrent computation. We now introduce concurrency in the form of networks of communicating processes, realized as connections of port automata.

We make distinctions between types of port automata depending upon what properties each possesses. The kind of port automaton we define below permits no nondeterminism when viewed as a single entity.

Definition - A port automaton P is said to be deterministic if τ is a singleton and δ is a partial function.

We wish to regard deterministic port automata as the atomic elements in a network of port automata, and we shall observe later on that our intuition is indeed correct. Before we begin discussing machine interconnection, however, we illustrate two examples of deterministic port automata.

(Note: We have selected a state graph depiction of port automata in an effort to facilitate comprehension. The state graph scheme is to be interpreted as follows:

(1) Each node is labelled in the following format:
 state | set of activated ports : output sets | set of inactivated ports : output sets. Thus, $q | \alpha_1:T_1, \alpha_2:T_2 | \gamma_1:U_1$ is a typical

node labelling.

(2) Each arc is labelled in the following format: activated port in previous state : input value. Thus, $\mu_3:V_3$ is a typical arc labelling.

(3) If T_j , U_j , or $V_j = 1$, then the labelling is abbreviated, using only the port label.)

The first example we consider is that of a binary semaphore [Dijkstra 68a]. Two ports, α and γ , are necessary; the P operation is performed through port α , and the V operation is performed through port γ . (Were we to simulate the situation in which n processes access a semaphore, then our port automaton model would have to have $2n$ ports, one α and one γ dedicated to each of the n processes.) Initially, the semaphore is in a state in which a P operation may be performed. A V operation may be performed at any time.

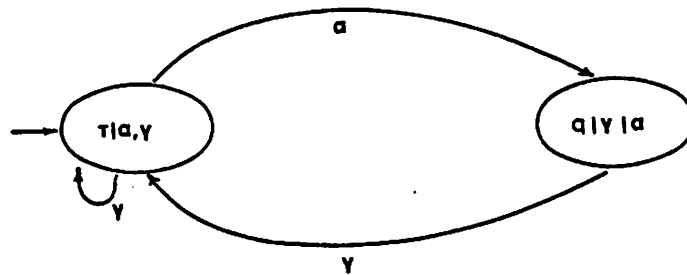


Figure 4

The second example we have chosen is a register capable of holding values from a set $A = \{a_1, a_2, \dots\}$, which may or may not be finite. The register has two ports, α and γ ; a value is written in at port α , and is read out (nondestructively) at port

γ . In the initial state, the register is assumed to be empty, and if a read is attempted the empty string is returned. For each different $a_j \in A$ written into the register, we regard the register as being in a separate state, since intuitively the contents of the register determines its state.

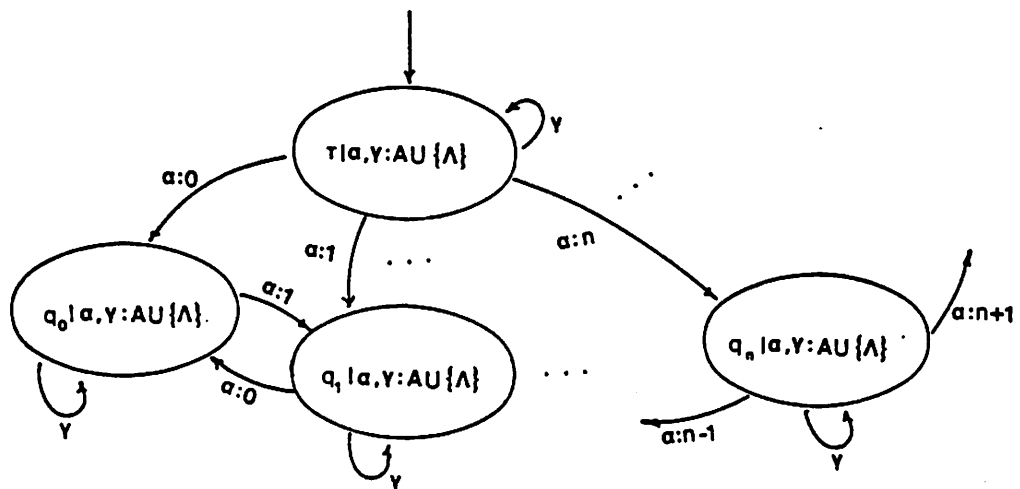


Figure 5

Recall that we had previously informally claimed that a network of processes could be regarded as a single process in itself. We will show formally that the interconnection of two port automata is indeed a port automaton. For the immediate discussion, we will adopt the following convention:

Let P^1 and P^2 be two port automata, where $P^j = (L^j, Q^j, \tau^j, \delta^j, \beta^j, x^j, \gamma^j)$ for $j = 1, 2$.

In order to interconnect two port automata, one or more ports from each automaton must be joined, under the constraint

that the input set for each such port is the same as the output set for the corresponding port. We thus have

Definition - Let $S^j \subseteq L^j$. The map $c : S^1 \rightarrow S^2$ is a port connection map if c is a bijection such that for each $i \in S^1$, $X_i = Y_{c(i)}$ and $Y_i = X_{c(i)}$. We call $c(i)$ the complement of i , and we say that i and $c(i)$ are complementary ports.

A port connection of two port automata yields a port automaton, such that the connected ports are subsumed and no longer visible to the environment. For the most part, the formal definition of a port connection is straight forward. As we shall see, it is the internal communications via the connected ports that make the behaviour of the composite port automaton nontrivial.

Definition - The port connection of P^1 and P^2 for a port connection map c , $P^1 \parallel_c P^2$, is a port automaton $P = (L, Q, \tau, \delta, \beta, X, Y)$ where

$$L = (L^1 - S^1) \cup (L^2 - S^2),$$

$$Q = Q^1 \times Q^2,$$

$$\tau = \tau^1 \times \tau^2 \in 2^Q,$$

$$X = (X_i : i \in L), \text{ where}$$

$$X_i = \begin{cases} X_i^1, & \text{for } i \in L^1 - S^1, \\ X_i^2, & \text{for } i \in L^2 - S^2, \end{cases}$$

$$Y = (Y_i : i \in L), \text{ where}$$

$$Y_i = \begin{cases} Y_i^1, & \text{for } i \in L^1 - S^1, \\ Y_i^2, & \text{for } i \in L^2 - S^2, \end{cases}$$

$\beta = (\beta_i : i \in L)$, where

$$\beta_i = \begin{cases} Q \xrightarrow{pr_1} Q^1 \xrightarrow{\beta_i^1} Y_i^1, & \text{for } i \in L^1 - S^1, \\ Q \xrightarrow{pr_2} Q^2 \xrightarrow{\beta_i^2} Y_i^2, & \text{for } i \in L^2 - S^2. \end{cases}$$

We have left the definition of δ for last since it is somewhat involved and must be unravelled in stages. We desire to find an expression for $\delta(q, (x, i))$ which captures the interactions of P^1 and P^2 . Since communication between port automata is symmetric, we may assume that $i \in L^1 - S^1$. The complementary case may be done similarly.

Recall that state transitions in a port automaton are to occur as the result of communication with the environment. In a network of port automata, internal communications may take place in the interval between external communications. To determine what internal communications are possible between P^1 and P^2 , we must know which of the connected ports are activated in a state $q \in Q$. We describe these in a set Z_q , the set of S^1 -components of the set of simultaneously activated pairs of complementary ports. Let $q = (q^1, q^2)$.

$$Z_q = \{k \in S^1 : \exists x \in X_k^1 \text{ such that } \delta^1(q^1, (x, k)) \neq \emptyset \\ \text{and } \exists \bar{x} \in X_{c(k)}^2 \text{ such that } \delta^2(q^2, (\bar{x}, c(k))) \neq \emptyset\}.$$

The internal communications are responsible for state transitions in both P^1 and P^2 , owing to the symmetric nature of the value exchanges involved. Let us observe what happens to P^1

and p^2 as a result of one internal communication via complementary ports k and $c(k)$, where $k \in Z_q$. We define a function $\hat{\delta}: Q \rightarrow 2^Q$ as follows:

$$\begin{aligned}\hat{\delta}(q) &= \hat{\delta}(q^1, q^2) \\ &= \{(\bar{q}^1, \bar{q}^2) : \bar{q}^1 \in \delta^1(q^1, \beta_{c(k)}^2(q^2)) \text{ and} \\ &\quad \bar{q}^2 \in \delta^2(q^2, \beta_k^1(q^1)), \text{ where } k \in Z_q\}.\end{aligned}$$

We must not discount the possibility of multiple internal communications occurring between two external communications. Therefore, we extend $\hat{\delta}$ to its transitive closure $\hat{\delta}^*$ to include this possibility. For $j \geq 0$, $\hat{\delta}^j(q)$ is equivalent to the set of states reachable from q by j internal state transitions.

$$\begin{aligned}\hat{\delta}^0(q) &= \{q\}, \\ \hat{\delta}^j(q) &= \hat{\delta}^\#(\hat{\delta}^{j-1}(q)), \text{ where } ()^\# \text{ is extension by union,} \\ \hat{\delta}^*(q) &= \bigcup_{j=0}^{\infty} \hat{\delta}^j(q).\end{aligned}$$

Having represented by $\hat{\delta}^*(q)$ all states reachable from state q as a result of exclusively internal state transitions, we are now in a position to define the state transition directly effected by the input at port i . We first define this transition function for a single state $q \in Q$ as follows:

$$\begin{aligned}\tilde{\delta}: Q \times \prod_{i \in L} X_i &\rightarrow 2^Q, \text{ where} \\ \tilde{\delta}(q, (x, i)) &= \tilde{\delta}((q^1, q^2), (x, i)) \\ &= \{(\bar{q}^1, q^2) : \bar{q}^1 \in \delta^1(q^1, (x, i))\}.\end{aligned}$$

By extending $\tilde{\delta}$ to the set of states obtained from internal transitions, we finally arrive at the completed definition of the transition function.

$$\delta(q, (x, i)) = \tilde{\delta}^{\#}(\hat{\delta}^*(q), (x, i)), \text{ where } ()^{\#} \text{ is extension by union.}$$

Thus, we have formalized our conjecture that a port connection of two port automata should again be a port automaton.

Machine Decomposition -

We would like to show that the intuitive notion that any port automaton may be simulated by a port connection of deterministic port automata is indeed correct. However, the way in which we decompose a port automaton, though simple, is perhaps not obvious. That is, the port automata which comprise the port connection are not necessarily the deterministic port automata with which we are familiar, for example registers and semaphores.

Remember that arbitrary port automata needn't be deterministic. We must, somehow, incorporate nondeterminism into our simulation, and we do this by a port connection of our deterministic automata. In fact, we shall find that almost any non-deterministic port automaton can be simulated by as few as two deterministic port automata connected appropriately.

Informally, the argument proceeds as follows. Let us say that we wish to simulate an arbitrary port automaton $P = (L, Q, \tau, \delta, \beta, X, Y)$. We will define a deterministic port automaton

P^1 , which is very similar to P in that it possesses P 's port label set plus one additional port α through which communications pass to the other deterministic automaton P^2 , and P^1 's state set essentially "includes" P 's state set plus an extra set of states to simulate the nondeterminism. This latter clause will be clarified in a moment. P^2 is a simple automaton consisting of one state and one port $c(\alpha)$, (where c is the port connection map). The state transition map simply takes the single state to itself.

$$\begin{aligned}
 P^2 \text{ ----- } L^2 &= \{c(\alpha)\}, \\
 Q^2 &= \tau^2, \\
 X_{c(\alpha)}^2 &= 1 = Y_{c(\alpha)}^2, \\
 \delta^2(\tau^2, (\#, c(\alpha))) &= \tau^2, \\
 \beta^2(\tau^2) &= \#.
 \end{aligned}$$

The key idea in the design of P^1 is that it is the exchange with P^2 that will enable the deterministic P^1 to simulate the nondeterministic P . We shall enumerate possible next states, and then let the number of exchanges between P^1 and P^2 determine the new state of P^1 , which when followed by an external transition, corresponds to the new state of P , following the identical external transition. Thus far, we have

$$\begin{aligned}
 P^1 \text{ ----- } L^1 &= L \cup \{\alpha\}, \\
 Q^1 &= \tilde{Q} \cup Z \text{ (explained later on)}, \\
 X_i^1 &= X_i, \quad Y_i^1 = Y_i, \text{ for } i \in L, \\
 X_\alpha^1 &= 1 = Y_\alpha^1.
 \end{aligned}$$

We now begin to construct, in stages, $Q^1 = \tilde{Q} \cup Z$. For $q \in Q$, $x \in X_i$, $i \in L$, define $\delta(q, (x, i)) = \{q_1, q_2, \dots\} = S$, a possibly infinite set. First, let us describe \tilde{Q} . For each $q \in Q$, we define $\tilde{q} \in \tilde{Q}$ such that if port i is activated in q , then port i is not activated in \tilde{q} . However, port α is activated in \tilde{q} to allow communication between P^1 and P^2 , in order to simulate the states contained in the set S . We will return to \tilde{Q} later on.

We now proceed to define Z , the state set associated with the internal communications between P^1 and P^2 . We decompose Z as follows,

$$Z = \bigcup_{\tilde{q} \in \tilde{Q}} Z_{\tilde{q}}, \text{ where } Z_{\tilde{q}} \text{ is the set of states generated by internal transitions starting from state } \tilde{q}.$$

Note that since the internal communications are discrete, then the maximum number of states generated by internal communications beginning in state $\tilde{q} \in \tilde{Q}$ is at most countably infinite, the maximum being achieved when a different state transition occurs as the result of each internal communication. Thus, $|Z_{\tilde{q}}| \leq \infty$. In fact, we define the cardinality of $Z_{\tilde{q}}$ to be equal to N_q , where N_q is the maximum of the cardinalities of the sets $\delta(q, (x, i))$ over all $q \in Q$, $x \in X_i$, and $i \in L$. This particular definition is made, since we shall see that each state in $Z_{\tilde{q}}$ corresponds to a state in $S = \delta(q, (x, i))$. We will, therefore, require that for all $q \in Q$, $x \in X_i$, and $i \in L$, $\delta(q, (x, i))$ must be at most countably infinite.

Thus, we may enumerate the elements of $Z_{\tilde{q}} = \{z_1, z_2, \dots\}$, where

$$\delta^1(\tilde{q}, (\#, \alpha)) = z_1,$$

$$\delta^1(z_j, (\#, \alpha)) = z_{j+1}, \text{ for } 1 \leq j < N_q.$$

This explicitly shows that communication of P^1 with P^2 generates in P^1 the state set $Z_{\tilde{q}}$ from which the sets corresponding to S for each $x \in X_i$ and $i \in L$ will be created.

Let us now turn to the final phase of our definition, that is, the effect of the external communication through port i . We can enumerate the elements of $S = \delta(q, (x, i)) = \{q_1, q_2, \dots\}$. We then define the value of the transition function for the specified external communication to be

$$\delta^1(z_j, (x, i)) = \tilde{q}_j \in \tilde{Q}, \text{ if } q_j \in \delta(q, (x, i)).$$

Note that for $k = |\delta(q, (x, i))| \leq j \leq N_q$, we define

$$\delta^1(z_j, (x, i)) = \tilde{q}_k .$$

To recapitulate, we see that our set of next states $\{\tilde{q}_1, \tilde{q}_2, \dots\}$ for $\delta^1(\tilde{q}, (x, i))$ is generated from an intermediate state set of next states, $Z_{\tilde{q}}$, which in turn is generated by communication with P^2 . We have completed our definition of the state set and the transition map for P^1 ; (the output map β^1 is defined in the obvious way.)

This concludes the proof of the following theorem:

Theorem - Let P be a port automaton with the following restrictions: τ is single-valued, and $\delta(q, (x, i))$ is countable for all $q \in Q$, $x \in X_i$, and $i \in L$. Then the response of P is identical to the response of $P^1 \parallel_c P^2$, for a port connection c of a deterministic port automaton P^1 and a one-state deterministic automaton P^2 .

We proceed to note a few consequences of the above construction.

(1) The set of reachable states of P is (under the isomorphism " \sim ") contained in the set of reachable states of $P^1 \parallel_c P^2$. However, there is not necessarily any correlation between the unreachable states of P and those of $P^1 \parallel_c P^2$.

(2) Note that although the transition from q to $q_j \in \delta(q, (x, i))$ may require only one time unit, the similar transition from \tilde{q} to \tilde{q}_j , where $\tilde{q}_j = \delta^1(z_j, (x, i))$, requires at least j more time units, since P^1 must pass through j intervening states z_1, \dots, z_j before reaching \tilde{q}_j . (Note that if, for all $\tilde{q} \in \tilde{Q}$, $|Z_{\tilde{q}}|$ is finite and the maximum cardinality of $|Z_{\tilde{q}}| = M \leq \infty$, then $P^1 \parallel_c P^2$ will require at most M more time units at each transition than P . Thus, if, for an input of length l , P requires $t(l)$ time units, then $P^1 \parallel_c P^2$ would require at most $\tilde{t}(l) = (M + 1)t(l)$ time units.)

(3) If $|Z_{\tilde{q}}|$ is infinite for some $\tilde{q} \in \tilde{Q}$, then the possibility of a nonterminating sequence of transitions between P^1 and P^2 exists. As a result, no future external communication may take

take place. This can be remedied by requiring an input to be present at a port in a finite but not necessarily bounded length of time following activation. This does not change the definition of a port connection, but rather is a restriction which may be necessary to add to a particular process description in order for the process to behave in the manner desired. Deadlock may not occur in a network of processes where this constraint is imposed upon the processes. (Note that this measure is necessary only in the situation in which P's inputs behave properly and $Z_{\bar{q}}$ is not finite. In this case, we must ensure that P^1 also behaves properly, and so restrict its time from activation to communication to be finite for each port.)

This concludes our formal discussion of process decomposition.

Examples -

We now return briefly to the port automaton, illustrated in Figure 1, and consider the consequences of connecting together m of these automata in a cycle. The port connection maps for each such automaton P^i , $1 \leq i \leq m$, are as follows:

For $i = 2$ to $m - 1$:

$$S^{i-1} = \{\gamma_{i-1}\}, T^{i-1} = \{\psi_{i-1}\},$$

$$S^i = \{\gamma_i\}, T^i = \{\psi_i\},$$

$$S^{i+1} = \{\gamma_{i+1}\}, T^{i+1} = \{\psi_{i+1}\}.$$

$$c^1 : T^{i-1} \rightarrow S^i.$$

$$c^2 : T^i \rightarrow S^{i+1}.$$

(Note: P^1 and P^m are connected together.)

Each of these processes has no terminating state and will attempt to execute forever. However, if an input at any port of any of the m processes is not received within a finite length of time, then a deadlock situation will result, as described in the following scenario.

Suppose that for some $1 \leq i \leq m$, the process P^i does not receive an external input that it needs in order to continue execution. Then it will be unable to send any more messages through port ψ in order to enable the next process. Process P^{i+1} may at present be executing, but it will soon receive no new enable signal, and then it too will stop execution. Eventually, all processes (except P^i) will be in a state corresponding to τ , that is, each will be waiting for an input through port γ (provided of course that all inputs from the external ports have been received). The similarity of the processes and their cyclic arrangement fosters this type of deadlock.

In the above discussion, we spoke loosely about the notions of deadlock and nontermination; we will now attempt to solidify these notions. Consider a network of port automata.

Definition - We say that a computation of a network is in a deadlocked state if, for each pair of connected ports, at most one port of the pair is activated, and for each activated external port, the expected input value will take an infinite amount of time to arrive.

Definition - The computation of a network is said to be non-terminating if, in every state of the composite automaton, there exists at least one activated external port. A computation is said to have achieved termination if all external ports of all automata in the network are inactivated.

Note that at the level of the component automata of the network, no state transitions may occur in any automaton in a deadlock situation. However, it is possible that some of the automata in the network are themselves port connections of other port automata, and at that level, they may undergo various internal state transitions as the result of internal communications. These internal transitions are, however, invisible to the rest of the network, since they do not change the activation of any of the ports at the network level.

In addition to deadlock and non-termination, we also define lockout which may arise as a side-effect of a port connection of automata. An instance of this occurs in our next example.

Definition - We say that a non-terminating computation of a network is locked out of a set R of reachable states if, after attaining a state q in a finite length of time, the set of states R is reachable from q but the computation will take an infinite amount of time to reach any state in R .

The last example, a critical section problem for two processes, is adapted from [Ladner 80] to fit our process conventions. Each of the processes, realized as a port automaton, may be in one of the following four states: "dormant", "trying", "critical section", and "end" (abbreviated "d", "t", "c", and "e", respectively). Each process has two ports, α and γ , where the status of the process (namely whether or not it is in its critical section) is sent through α , and the status of the other process is received through γ . (We shall adopt the convention that "2" indicates that a process is in its critical section and that "1" indicates otherwise. At each state, the process will attempt to inform the other process of its present state, while concurrently it will be expecting to receive the present state of the other process.

A process may only change state as a result of receiving the state of the other process. The process begins in the "dormant" state and may then proceed to the "trying" state. From the "trying" state, the process may advance to its "critical section" only if the other process is not in its critical section. From the "critical section", the process proceeds to the "end" state, and from "end", back to "dormant", indicating with this state change, that the other process is now free to enter its critical section. Below is pictured the state graph for one such automaton. (At each state, both ports α and γ are activated, where the output sets for each are $\alpha:\{1,2\}$ and $\gamma:\{\#\}$. This labelling will be deleted from each node, since it is the same for all nodes of the state graph.)

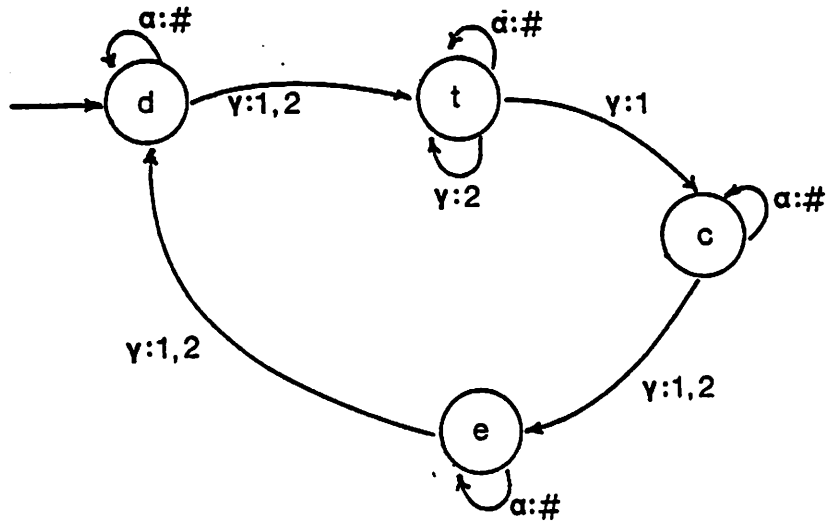


Figure 6

We now connect together the two processes, P^1 and P^2 through ports α_1 and γ_2 , and ports γ_1 and α_2 , and allow these processes to communicate in any state, even when one of them is in its critical section. (Notice that after the port connections, the network no longer has any external ports, and is thus incapable of any visible state transitions. To avoid this problem, we add an external port η to each process P^i , which functions as a signal to the environment that process P^i has completed its critical section; the output set of η is simply $\{\#\}$. Port η will only be activated in the "end" state of process P^i .)

The state graph for the specified port connection of P^1 and P^2 is pictured below. (Normally, we only depict the external communications of a network in a state diagram. However, in order to facilitate the reader's comprehension of the example, we explicitly show all state transitions, resulting from external

and internal communications. Internal port labels are enclosed in parentheses. Since all ports are activated in all states, with the exception of the "end" state of each process, then the list of activated and inactivated ports and their output sets will be dropped from the node labellings. Each node will be labelled with a pair of states, each element of the pair representing the state of one of the port automata, P^1 and P^2 , subscripted accordingly.)

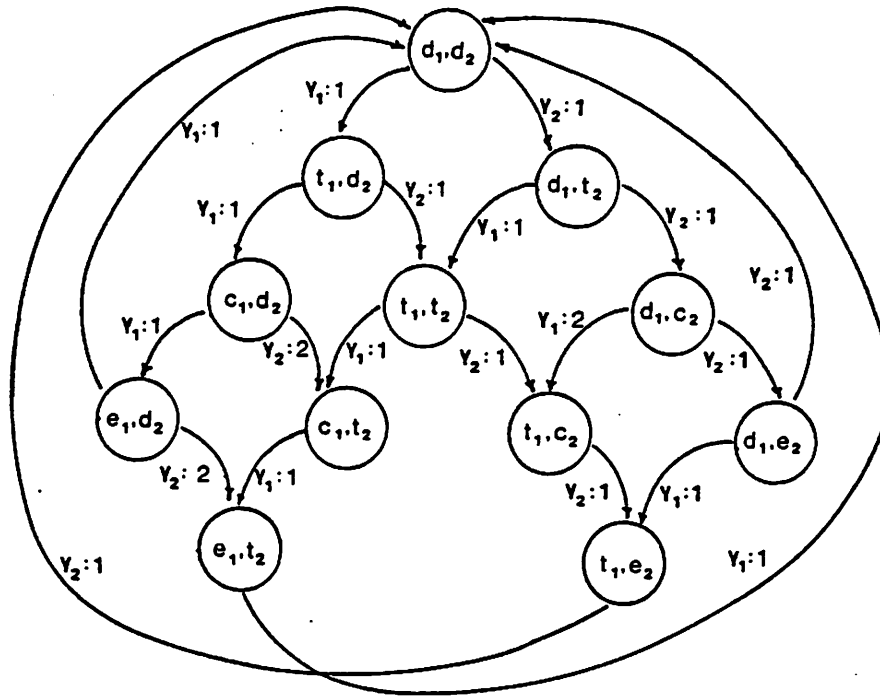


Figure 7

We notice that we do indeed have mutual exclusion from the critical section. However, in every state, there is the potential for lockout, since one of the processes may continually send its status to the other process, thus remaining in the same state while permitting the other process to advance. Since at

each state, the probability of receiving an input at α or at γ is $1/2$, and since each of the processes executes forever, then the probability of a lockout goes to zero.

This concludes our discussion of process interaction.

Conclusions -

In approaching concurrent computation from the automata-theoretic perspective, we have attempted to present a relatively simple and straightforward model of communicating processes, in the form of port automata. We have observed that our intuitive conceptions of process interconnection and decomposition are indeed formalizable in this framework. We have found that port automata lend themselves naturally to a state graph representation, a visual aid in understanding process interaction. Most importantly, we observe that, using an automata-theoretic approach to concurrency, we may successfully define an algebraic notion of process within the category of sets, without having to resort to the category of domains and the added structure therein.

References -

- [Arbib and Manes 75] Arbib, M. A. and Manes, E. G. Fuzzy machines in a category. Bulletin of the Australian Mathematical Society, 13, No. 2 (1975), 169-210.
- [Arbib and Manes 80] Arbib, M. A. and Manes, E. G. The greatest fixpoint approach to data types. Proceedings of the 3rd Workshop Meeting on "Categorical and Algebraic Methods in Computer Science and System Theory", November 3-7, 1980, Dortmund, West Germany.
- [Dijkstra 68a] Dijkstra, E. W. The structure of "THE" - multiprogramming system. Comm. ACM, 11, No. 5 (May), 341-346.
- [Dijkstra 68b] Dijkstra, E. W. Co-operating sequential processes. In F. Genuys (ed.), Programming Languages, Academic Press, New York, 43-112.
- [Brinch Hansen 78] Brinch Hansen, P. Distributed processes: a concurrent programming concept. Comm. ACM, 21, No. 11 (Nov), 934-941.
- [Campbell and Habermann 74] Campbell, R. H. and Habermann, A. N. The specification of process synchronization by path expressions. In E. Gelenbe and C. Kaiser (ed.), Operating Systems, International Symposium, Rocquencourt, 1974, LNCS, Vol. 16, Springer-Verlag, Berlin, 1974, 89-102.
- [Hewitt 77] Hewitt, C. Viewing control structures as patterns of passing messages. Artificial Intelligence, 8, 3, June 1977, 323-364.
- [Hoare 78] Hoare, C. A. R. Communicating sequential processes. Comm. ACM, 21, No. 8 (Aug), 666-677.
- [Ladner 80] Ladner, R. E. The complexity of problems in systems of communicating sequential processes. Journal of Computer and System Sciences, 21, 179-194.
- [Milne and Milner 79] Milne, G. and Milner, R. Concurrent processes and their syntax. J. ACM, 26, No. 2 (Apr), 302-321.
- [Milner 80] Milner, R. A calculus of communicating systems. In G. Goos and J. Hartmanis (ed.), LNCS, Vol. 92, Springer-Verlag, Berlin, 1980, 33-40.

[Petri 73] Petri, C. A. Concepts of net theory. In Mathematical Foundations of Concepts of Computer Science, High Tatras, Sept. 1973, 137-146.

[Plotkin 76] Plotkin, G. D. A powerdomain construction. SIAM J. Comput., 5, No. 3 (Sept 1976), 452-487.

[Smyth 76] Smyth, M. B. Power domains. Journal of Computer and System Sciences, 16, 23-36.