A. BERGERON

**A unified approach to control problems in
discrete event processes**

Informatique théorique et applications, tome 27, n° 6 (1993), p. 555-573.

<http://www.numdam.org/item?id=ITA_1993__27_6_555_0>

© AFCET, 1993, tous droits réservés.

L'accès aux archives de la revue « Informatique théorique et applications » implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/legal.php>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques
<http://www.numdam.org/>

A UNIFIED APPROACH TO CONTROL PROBLEMS IN DISCRETE EVENT PROCESSES (*)

by A. BERGERON ⁽¹⁾

Communicated by A. ARNOLD

Abstract. – *Control theory of discrete event processes has been developed in recent years by Ramadge and Wonham. Their framework is based on formal language theory and control theoretic ideas such as controllability and observability. In this paper, we focus on the automata modeling the processes and the specifications, rather than on the sets of possible or acceptable sequences. By adopting this slightly different point of view, we show that the main results of control theory can be obtained in a simple, uniform way.*

Centralized control problem, observation problems and distributed control problems are stated as particular cases of what we call design problems. With the help of two basic constructions on automata, we show that necessary and sufficient conditions for the existence of solutions are all particular cases of a general lemma on the solutions of design problems. All the results are effective, in the sense that they yield immediate algorithms for the design of controllers.

Résumé. – *Un processus discret est un système qui change abruptement d'état selon l'occurrence de suites d'événements. La théorie du contrôle de tels processus a été développée dans les dernières années par Ramadge et Wonham. Fondée à la fois sur la théorie des langages formels et sur des idées provenant de la théorie du contrôle, cette théorie explore l'existence et la construction de solutions à des problèmes de contrôle centralisé – ou réparti – dans un contexte d'observation partielle.*

Dans cet article, nous adoptons un point de vue légèrement différent en centrant l'attention sur les automates qui modélisent les processus et leur spécification. Ce nouveau point de vue permet d'obtenir l'ensemble des résultats de la théorie d'une manière simple et uniforme. Nous exposons d'abord un cadre général qui permet d'exprimer les problèmes de contrôle ou d'observation avec un formalisme identique. Nous montrons ensuite que deux constructions de base sur les automates permettent d'obtenir des conditions nécessaires et suffisantes pour la solution de ces problèmes. Tous les résultats obtenus sont constructifs, au sens où les énoncés fournissent immédiatement des algorithmes pour la construction de contrôleurs ou d'observateurs lorsqu'ils existent.

This work was partially supported by grants from BNR Ltd, FCAR of Québec, and NSERC of Canada.

(*) Received June 1992, accepted April 1993.

⁽¹⁾ Département de mathématiques et d'informatique, Université du Québec à Montréal, C. P. 8888 Succ 'A', Montréal, Canada, H3C 3P8.

1. INTRODUCTION

A *discrete event process* is a system that changes state according to the occurrence of a sequence of events. The focus is on the order and the nature of the events affecting the system rather than their time of occurrence or their duration. The central problems can be described as:

- 1) Given a discrete process and a specification of its acceptable behaviors, is it possible to construct a new process, called a *controller*, that will implement the specification by prohibiting or allowing certain events?
- 2) How is the solution affected if the controller does not have complete information about the occurrence of events in the original process?
- 3) When is it possible to distribute the controlling task among several controllers?

Various solutions to these design problems have been found in the recent years within the framework developed by Ramadge and Wonham. Necessary and sufficient conditions for the existence of solutions are expressed in terms of formal language theory and control theoretic ideas such as *controllability* and *observability* [3; 4], [6], [5]. Controllability and observability are defined, in these works, as properties of the acceptable sequences of events (in the specification) with respect to possible sequences of events (in the original process).

In this paper, we focus on the automata modeling the processes and the specifications rather than on the sets of possible or acceptable sequences. By adopting this slightly different point of view, we will show that the main results of control theory can be obtained in a simple, uniform way. This approach also provides an interesting bridge to the theory of transition systems and concurrent processes developed around the Arnold-Nivat model [1], [2].

Section 2 reviews the basic tools used in modelling discrete processes with deterministic automata. In Section 3, we present a general setting for design problems as solutions of equations of the form $\mathbf{Z} = \mathbf{D} \times \mathbf{X}$ where $\mathbf{Z}$ is a specification of a process $\mathbf{D}$ and $\mathbf{X}$ belongs to a restricted class of automata. We apply these tools in Section 4 and 5 to obtain constructive solutions of control and observation problems. Section 6 discusses control under partial observation and distributed control.

2. MODELING DISCRETE EVENT PROCESS WITH AUTOMATA

Let Σ be a finite set whose elements are called *events*. The set of all finite sequences of events is denoted by Σ^* , and the empty sequence is denoted by λ . The concatenation of two sequence x and y is denoted by xy .

Let S be a set whose elements are called *states*. An *automaton* V on the set Σ of events is given by an arbitrary partial function — multiplicatively denoted by a dot “ $\cdot$ ” — called a *transition function*:

$$\cdot: S \times \Sigma \rightarrow S$$

which, when defined, associates to a pair (s, σ) the state $s \cdot \sigma$. Among the states S , we distinguish an *initial state* i and of a subset $M \subseteq S$ of *marked states*.

It will be useful to consider the extension of a transition function to the set Σ^* of all sequences of events. Every transition function can be uniquely extended to a partial function:

$$\cdot: S \times \Sigma^* \rightarrow S$$

that has the following properties:

- (i) $s \cdot \lambda = s$.
- (ii) $(s \cdot x) \cdot \sigma = s \cdot (x \sigma)$ whenever the left hand side is defined.

The language $L(V)$ *recognized* by the automaton V is the subset of Σ^* defined by:

$$L(V) = \{x \mid i \cdot x \in M\}.$$

Finally, we will denote by $\bar{V}$ the automaton obtained by marking all the states of an automaton V . We have:

$$L(\bar{V}) = \{x \mid i \cdot x \text{ is defined}\}.$$

An automaton is often described by a *directed graph* whose vertices are labeled by the states S and whose arrows are labeled by the symbols σ of Σ . For each defined transition $s \cdot \sigma = s'$, we have the corresponding arrow:

$$\begin{array}{c} \sigma \\ s \rightarrow s' \end{array}$$

If x is a sequence of events, then $x \in L(\bar{V})$ or x is defined in V if there is a path labeled by x beginning at the initial state i . If the state reached by the path is marked, then $x \in L(V)$.

DEFINITION 2.1: *Non-blocking automata.*

(i) If L is a subset of Σ^* then $Pref_L = \{x \in \Sigma^* \mid xy \in L\}$ is the set of *prefixes* of sequences in L . Since y can be the empty sequence, we always have $L \subseteq Pref_L$.

(ii) An automaton V is said to be *non-blocking* if

$$L(\bar{V}) = Pref_L(V). \quad \blacksquare$$

Consider, for example, the following automaton on $\Sigma = \{a, b\}$:

where only r is marked. This automaton is blocking since the sequence ab belongs to $L(\bar{V})$ but is not a prefix of any sequence in $L(V)$.

In the following sections, we will be mainly concerned with the design and analysis of special classes of automata. For that purpose, we need ways to compare them, and the usual equivalences defined in automata theory (mostly based on the language recognized by an automaton) are not powerful enough. In design problems, we will need the non-blocking property to be preserved under equivalence.

DEFINITION 2.2: *Partial order between automata.*

The relation $V \leq W$ holds whenever

- (i) $L(V) \subseteq L(W)$ and
- (ii) $L(\bar{V}) \subseteq L(\bar{W})$.

When $V \leq W$ and $W \leq V$, we simply write $V = W$. $\blacksquare$

An immediate consequence of this definition is:

PROPOSITION 2.3: *If $V = W$, then V is non-blocking if and only if W is non-blocking.* $\blacksquare$

Given two automata V and W , with states S_V and S_W , initial states i_V and i_W , and marked states M_V and M_W . The *product* $V \times W$ has states $S = S_V \times S_W$ with the transition function:

$$\cdot : S \times \Sigma^* \rightarrow S \text{ where } (v, w) \cdot x = (v \cdot x, w \cdot x) \text{ whenever both are defined.}$$

The initial state of the product is $i = (i_V, i_W)$ and the marked states

$$M = \{(v, w) \mid v \in M_V \text{ and } w \in M_W\}.$$

We can think of the two processes as simultaneously functioning: they synchronize on the occurrence of every event. Of course, in many examples, one of the process will 'generate' the events and the other will 'react' to them, but the modelling is the same in both cases. The following theorem expresses a basic relation between products and inequalities:

THEOREM 2.4: $V \leq W$ if and only if $V = W \times V$.

Proof: If $V \leq W$ then,

$$L(V) = L(V) \cap L(W) = L(V \times W)$$

since a state is marked in the product if and only if it is marked by both automata. We also have

$$L(\bar{V}) = L(\bar{V}) \cap L(\bar{W}) = L(\bar{V} \times \bar{W}) = L(\overline{V \times W})$$

since those last automata are identical. On the other hand, if $V = W \times V$, we have easily

$$L(V) = L(V) \cap L(W) \subseteq L(W) \text{ and}$$

$$L(\bar{V}) = L(\bar{V}) \cap L(\bar{W}) \subseteq L(\bar{W}). \quad \blacksquare$$

Further properties of these operations are easily verified:

PROPOSITION 2.5: Let Y, Z, V and W be automata. Then

- (i) If $Y \leq V$ and $Z \leq W$ then $Y \times Z \leq V \times W$.
- (ii) $Z \leq V \times W$ if and only if $Z \leq V$ and $Z \leq W$. $\blacksquare$

3. A GENERAL FRAME FOR DESIGN PROBLEMS

Let D be an 'existing' discrete process and suppose we want to restrict, observe, or modify the behavior of this process. Such design problems can be solved by constructing one or more new processes that will function simultaneously with D , and that will be able to prevent, authorize, observe or report the occurrence of events in D .

In order to state the problem formally, we first suppose that the process D is modeled by an automaton such that any undefined transition for a pair (s, σ) corresponds to the fact that the event σ is *impossible* in state s . All the *possible* sequences of events are taken into account in the model. If the behavior of D is to be restricted or observed, we need a specification of what are the acceptable behaviors or what are the properties to observe.

Such a specification will again be modeled by an automaton Z , and we assume that $Z \leq D$. This restriction is natural since if there was a sequence of events defined in Z but not in D , it would correspond to an impossible sequence by our assumption on D .

Finally, in order to solve design problems, we want to construct automata that will be able to perform specialized tasks. For a given task (such as controlling, observing, reporting, etc.) the identification of these specialized automata will correspond to a particular class $\mathcal{A}$ of automata.

In this framework, we can state the *general design problem* as:

Let $\mathcal{A}$ be a class of automata. Given a process D , and a specification Z such that $Z \leq D$, does there exist an automaton $X \in \mathcal{A}$ such that

$$Z = D \times X?$$

If the class $\mathcal{A}$ is unrestricted, Theorem 2.4 gives immediately the trivial solution $X = Z$. By restricting the solutions to have certain additional properties on their states or transition functions, we will solve many variants of control and analysis problems. These solutions will be constructive in the sense that the existence theorems will give explicit construction for X , and are based on the following result:

LEMMA 3.1: *Let $\mathcal{A}$ be a class of automata and suppose that there exists a mapping A which associates to every automata Z an automaton $A_Z \in \mathcal{A}$, such that $Z \leq A_Z$ and for each $X \in \mathcal{A}$,*

$$Z \leq X \Rightarrow A_Z \leq X.$$

Then, if $Z \leq D$, there exists an automaton $X \in \mathcal{A}$ such that:

$$Z = D \times X \quad \text{if and only if } Z = D \times A_Z.$$

Proof: The *if* part is immediate since we assumed that $A_Z \in \mathcal{A}$. On the converse, the fact that $Z \leq A_Z$ and $Z \leq D$, gives us:

$$Z \leq D \times A_Z.$$

Suppose now that there exists an automaton $X \in \mathcal{A}$ such that $Z = D \times X$. Then $Z \leq X$ so, we have $A_Z \leq X$. Multiplying by D on both sides we get:

$$D \times A_Z \leq D \times X = Z. \quad \blacksquare$$

Thus, if a suitable mapping (or construction) A is provided for the class $\mathcal{A}$, Lemma 3.1 gives elegant necessary and sufficient conditions for the design

problem with respect to the class $\mathcal{A}$. In the next sections, we will discuss two basic constructions corresponding respectively to the control and observation problems.

4. THE CENTRALIZED CONTROL PROBLEM

Suppose that $\mathbf{D}$ is a process on the set of events Σ . Among those events, some of them may be *controllable*, that is there exists physical means to prevent their occurrence. For example, a communication process can suspend message sending until it receives a given signal. Given such a subset Σ_c of controllable events (and its complement the *uncontrollable* events Σ_{uc}), the centralized control problem can be stated informally as designing an automaton which changes state according to the sequence of events generated by $\mathbf{D}$, and prevent or authorize controllable events according to its state.

If an automaton is to perform such a controlling task, its only 'impossible' events must be exactly those it prevented, thus we have the following definition:

DEFINITION 4.1: *Control Automata*

An automaton $\mathbf{C}$ is a *control automaton* with respect to Σ_c if, for all states s ,

$$\sigma \in \Sigma_{uc} \text{ implies } s \cdot \sigma \text{ is defined}$$

The class of all control automata (with respect to Σ_c) will be noted $\mathcal{A}_c$.

If $\mathbf{C}$ is a control automaton, we can define the following control strategy for $\mathbf{C}$: in state s , prevent the occurrence of any event such that $s \cdot \sigma$ is undefined, and allow all others. It is clear, from the definition of control automata, that such a strategy is always realizable.

The *centralized control problem* can be formulated as:

Let $\mathcal{A}_c$ be the class of control automata with respect to Σ_c . Given a process $\mathbf{D}$, and a specification $\mathbf{Z}$ such that $\mathbf{Z} \leq \mathbf{D}$, does there exist an automaton $\mathbf{X} \in \mathcal{A}_c$ such that

$$\mathbf{Z} = \mathbf{D} \times \mathbf{X}?$$

Consider now the following construction $\mathbf{C}_Z$ on an automaton $\mathbf{Z}$: we add a new unmarked state u to $\mathbf{Z}$ and whenever $s \cdot \sigma$ is undefined—including the new state u —and σ is uncontrollable, we set $s \cdot \sigma = u$ in $\mathbf{C}_Z$. With this

construction we have:

THEOREM 4.2: *The mapping C_Z has the following properties:*

- (i) $C_Z \in \mathcal{A}_c$.
- (ii) $Z \leq C_Z$.
- (iii) For each $X \in \mathcal{A}_c$, $Z \leq X \Rightarrow C_Z \leq X$.

Proof: (i) This is immediate from the construction of C_Z .

(ii) Every path in Z is also a path in C_Z .

(iii) Suppose that $Z \leq X$, we first have $L(C_Z) = L(Z) \subseteq L(X)$ since any path in C_Z which goes in a marked state does not go through the new state u – there are no outgoing arrows from u .

Now, let $x \in L(\bar{C}_Z)$, if $i \cdot x \neq u$ then $x \in L(\bar{Z}) \subseteq L(\bar{X})$. Otherwise, if $i \cdot x = u$, x can be decomposed as $x = x' \sigma y$ such that $x' \in L(\bar{Z})$ and σy is a sequence of uncontrollable events. But $x' \in L(\bar{X})$, so $x' \sigma y \in L(\bar{X})$ since X is a control automaton. ■

COROLLARY 4.3: *The centralized control problem has a solution if and only if $Z = D \times C_Z$.* ■

Notes: Other sets of necessary and sufficient conditions for the existence of a solution to the centralized control problem appear in [3]. These conditions are based on the notion of *controllable* languages. A language K is *controllable* with respect to L to if

$$x \in \text{Pref}_K \text{ and } x\sigma \in L \text{ with } \sigma \text{ uncontrollable implies } x\sigma \in \text{Pref}_K$$

and their results are established by considering the controllability of $L(Z)$ with respect to $L(\bar{D})$.

5. THE OBSERVATION PROBLEM

Suppose that D is a process on the set of events Σ . Among those events, the occurrence of some of them may be *unobservable*. For example, a communication process may not be able to tell if a given message has been received. Given a specification Z for the process and such a subset Σ_{uo} of unobservable events (and its complement the *observable* events Σ_o), the observation problem is whether or not certain properties of a system can be determined without knowledge about the occurrence of unobservable events.

In order to perform such an observation task, an automaton should have the same ‘reactions’ whether or not an unobservable event occurred. The

simplest observation automata are the ones that loop on all (defined) unobservable events:

DEFINITION 5.1: Observation Automata

An automaton $\mathbf{O}$ is an *observation automaton* with respect to Σ_o if, for all states s ,

$$s \cdot \sigma \text{ is defined and } \sigma \in \Sigma_{uo} \text{ implies } s \cdot \sigma = s.$$

The class of all observation automata (with respect to Σ_o) will be noted $\mathcal{A}_o$.

The *observation problem* can be formulated as:

Let $\mathcal{A}_o$ be the class of observation automata with respect to Σ_o . Given a process $\mathbf{D}$, and a specification $\mathbf{Z}$ such that $\mathbf{Z} \leq \mathbf{D}$, does there exist an automaton $\mathbf{X} \in \mathcal{A}_o$ such that

$$\mathbf{Z} = \mathbf{D} \times \mathbf{X}?$$

This problem will be solved in the same fashion as the control problem. We want to construct the smallest observation automata $\mathbf{O}_Z$ that contains a given automata $\mathbf{Z}$. The idea behind the construction is the following: suppose an observable event occurs in $\mathbf{Z}$ changing its state to s , in order to ignore unobservable events we have to assume that $\mathbf{Z}$ could be either in state s or in any other state reachable from s by a sequence of unobservable events. For example, the initial state of $\mathbf{O}_Z$ will be the *set* of all states reachable from the initial state of $\mathbf{Z}$ with sequences of unobservable events.

More formally, let S be the set of states of $\mathbf{Z}$, with initial state i and marked states M . Let Σ_{uo}^* be the set of unobservable sequences. The states of the automaton $\mathbf{O}_Z$ are non-empty subsets $\mathcal{P}^+(S)$ of S . The initial state is:

$$I = \{i \cdot u \mid i \cdot u \text{ is defined and } u \in \Sigma_{uo}^*\}$$

and the marked states of $\mathbf{O}_Z$ are all subsets that contain at least one marked state of $\mathbf{Z}$. The transition function is defined as:

$$\circ : \mathcal{P}^+(S) \times \Sigma \rightarrow \mathcal{P}^+(S)$$

Let $T \subseteq S$, if $s \cdot \sigma$ is defined in $\mathbf{Z}$ for at least one s in T then $T \cdot \sigma$ is defined and

$$(i) \text{ If } \sigma \in \Sigma_o, T \cdot \sigma = \{s \cdot \sigma \mid s \in T, s \cdot \sigma \text{ is defined and } u \in \Sigma_{uo}^*\}$$

$$(ii) \text{ If } \sigma \in \Sigma_{uo}, T \cdot \sigma = T$$

otherwise $T \cdot \sigma$ is undefined.

In order to prove that the construction $\mathbf{O}_Z$ satisfies the hypothesis of Lemma 3.1, we need the following results on the relations between Z and $\mathbf{O}_Z$:

LEMMA 5.2: *If $i \cdot x$ is defined in Z , then $I \circ x$ is defined in $\mathbf{O}_Z$ and $i \cdot x \in I \circ x$.*

Proof: The proof is an induction on the length of x . If $x = \lambda$, then $i \cdot x = i$ and $I \circ x = I$ are always defined and we have, by construction, that $i \in I$. Now, let $x = y \sigma$ be defined in Z , then y is defined in Z and we can suppose that $I \circ y$ is defined and $i \cdot y \in I \circ y$. Since $(i \cdot y) \cdot \sigma$ is defined we have immediately $(I \circ y) \circ \sigma$ is defined.

If σ is observable we have immediately that $(i \cdot y) \cdot \sigma \in (I \circ y) \circ \sigma$. On the other hand, if σ is unobservable then $(I \circ y) \circ \sigma = I \circ y$ so we have to prove that $(i \cdot y) \cdot \sigma \in I \circ y$. But if an accessible state of $\mathbf{O}_Z$ contains a given state s it contains, by construction, all the states reachable from s by unobservable events, so if $i \cdot y \in I \circ y$ we have $(i \cdot y) \cdot \sigma \in I \circ y$. ■

The second lemma expresses a deeper property: each sequence defined in $\mathbf{O}_Z$ looks like some sequence in Z in the following sense. Consider the projection

$$\pi : \Sigma^* \rightarrow \Sigma^*$$

that erases unobservable events, then:

LEMMA 5.3: *If $I \circ x$ is defined in $\mathbf{O}_Z$ and $t \in I \circ x$, there exists a sequence x' defined in Z such that $i \cdot x' = t$ and $\pi(x) = \pi(x')$.*

Proof: The proof is an induction on the number of observable events of x .

Suppose that $\pi(x) = \lambda$, that is $x \in \Sigma_{uo}^*$, then $I \circ x = I = \{i \cdot u \mid i \cdot u \text{ is defined and } u \in \Sigma_{uo}^*\}$ thus if $t \in I \circ x$, there exists a sequence $u' \in \Sigma_{uo}^*$ defined in Z such that $t = i \cdot u'$ and $\pi(u') = \lambda = \pi(x)$.

Suppose now that $x = y \sigma z$ with σ being the last observable event of x , that is $\pi(x) = \pi(y) \sigma$. Let t be a state in

$$I \circ x = I \circ y \sigma = (I \circ y) \circ \sigma = \{s \cdot \sigma u \mid s \in I \circ y, s \cdot \sigma u \text{ is defined and } u \in \Sigma_{uo}^*\}$$

so there exists an s in $I \circ y$ and a u' in Σ_{uo}^* such that $s \cdot \sigma u' = t$.

By the induction hypothesis, there is a y' defined in Z such that $i \cdot y' = s$ and $\pi(y) = \pi(y')$. But the fact that $i \cdot y' = s$, and $s \cdot \sigma u'$ is defined, implies that $i \cdot y' \sigma u' = t$ is defined in Z , and $\pi(y' \sigma u') = \pi(y') \sigma = \pi(y) \sigma = \pi(x)$. ■

We can now prove:

THEOREM 5.4: *The mapping $\mathbf{O}_Z$ has the following properties:*

- (i) $\mathbf{O}_Z \in \mathcal{A}_o$.
- (ii) $Z \leq \mathbf{O}_Z$.
- (iii) For each $X \in \mathcal{A}_o$, $Z \leq X \Rightarrow \mathbf{O}_Z \leq X$.

Proof: (i) is immediate by construction, since $\mathbf{O}_Z$ loops on all unobservable events.

(ii) If $x \in L(\bar{Z})$ then $i \cdot x$ is defined in Z , so $I \circ x$ is defined in $\mathbf{O}_Z$ by Lemma 5.2. If $x \in L(Z)$ then the state $i \cdot x$ is marked, and since $i \cdot x \in I \circ x$, the state $I \circ x$ will be marked.

(iii) Suppose that $Z \leq X$ with $X \in \mathcal{A}_o$. First we will show that if x is defined in $\mathbf{O}_Z$ then $\pi(x)$ is defined in X . If x is defined in $\mathbf{O}_Z$ then, by Lemma 5.3, there exists an x' defined in Z such that $\pi(x') = \pi(x)$. But since $Z \leq X$, x' is in X . Since X loops on every unobservable event, $\pi(x')$ can be traced in X by following the same path as x' and skipping the unobservable loops.

We will prove that $L(\bar{\mathbf{O}}_Z) \subseteq L(\bar{X})$ by induction on the number of unobservable events of sequences in $L(\bar{\mathbf{O}}_Z)$. If $x \in L(\bar{\mathbf{O}}_Z)$, and x has no unobservable events, then $\pi(x) = x$ and, by the preceding remark, we have $x \in L(\bar{X})$.

Suppose now that $x = y\sigma z$ with σ being the last unobservable event of x . We will show that both $y\sigma$ and yz are defined in X . Since X loops on every unobservable σ , we will be able to conclude that $y\sigma z$ is defined in X .

If $x = y\sigma z$ is defined in $\mathbf{O}_Z$ with σ being the last unobservable event of x , then $\pi(x) = \pi(y)z$ is defined in X . By the induction hypothesis, y also is defined in X , and since $\pi(y)$ and y reach the same state in X , if $\pi(y)z$ is defined, then yz is also defined.

On the other hand, if $x = y\sigma z$ is defined in $\mathbf{O}_Z$, then $y\sigma$ is also defined, and $s \cdot \sigma$ is defined in Z for at least one s in $I \circ y$. By Lemma 5.3, there exists a y' defined in Z such that $i \cdot y' = s$ and $\pi(y') = \pi(y)$. But since $s \cdot \sigma$ is defined in Z , and $i \cdot y' = s$, we have that $i \cdot y' \sigma$ is defined in Z , implying that $y' \sigma$ is defined in X . By the induction hypothesis, y is defined in X , and since $\pi(y') = \pi(y)$, y and y' reach the same state in X . Thus, if $y' \sigma$ is defined so is $y \sigma$.

In order to prove that $L(\mathbf{O}_Z) \subseteq L(X)$, we note that if $x \in L(\mathbf{O}_Z)$ then x is in $L(\bar{X})$ by the preceding paragraphs, and we only have to prove that x reaches a marked state of X . If $x \in L(\mathbf{O}_Z)$ then $I \circ x$ is marked in $\mathbf{O}_Z$, so there is a marked state t in Z such that $t \in I \circ x$. By Lemma 5.3, there exists a sequence x' defined in Z such that $i \cdot x' = t$ and $\pi(x) = \pi(x')$. But x' and x are

both defined in $\mathbf{X}$ and reach the same state since $\pi(x) = \pi(x')$. Since x' reaches a marked state in $\mathbf{Z}$, it reaches a marked state in $\mathbf{X}$. ■

COROLLARY 5.5: *The observation problem has a solution if and only if $\mathbf{Z} = \mathbf{D} \times \mathbf{O}_Z$.* ■

Notes: In the framework of the theory developed by Ramadge and Wonham, the observation problem is treated through the notion of *observability* [6], and Corollary 5.5 can be used as an algorithm to determine if $L(\mathbf{Z})$ is an observable language with respect to $L(\bar{\mathbf{D}})$.

Finally, we have the following propositions relating the constructions $\mathbf{O}_Z$ and $\mathbf{C}_Z$. The first one says that the ‘observation’ construction preserves the property of being a control automaton:

PROPOSITION 5.6: *Let $\mathcal{A}_c$ be the class of control automata with respect to Σ_c . If $\mathbf{C} \in \mathcal{A}_c$ and $\mathbf{O}_C$ is the minimal observation automaton with respect to Σ_o that contains $\mathbf{C}$, then $\mathbf{O}_C \in \mathcal{A}_c$.*

Proof: Let $\mathbf{T}$ be a state of $\mathbf{O}_C$ and σ an uncontrollable event. Recall that $\mathbf{T}$ is a non-empty subset of the states of $\mathbf{C}$ and $\mathbf{T} \cdot \sigma$ is defined if $\mathbf{s} \cdot \sigma$ is defined in $\mathbf{C}$ for at least one $\mathbf{s}$ in $\mathbf{T}$. If σ is uncontrollable, the fact that $\mathbf{C}$ is a control automaton implies that $\mathbf{s} \cdot \sigma$ is always defined. ■

Applying the constructions in the reverse order does not always yield an observation automaton. We must take care of the possibility that some unobservable and uncontrollable events are not defined:

PROPOSITION 5.7: *Let $\mathcal{A}_o$ be the class of control automata with respect to Σ_o . If $\mathbf{O} \in \mathcal{A}_o$ and $\mathbf{C}_O$ is the minimal control automaton with respect to Σ_c that contains $\mathbf{C}$, then $\mathbf{C}_O \in \mathcal{A}_o$ if and only if the transition function of $\mathbf{O}$ is defined for all $\sigma \in \Sigma_{uc} \cap \Sigma_{uo}$.*

Proof: If the transition function of $\mathbf{O}$ is defined for all $\sigma \in \Sigma_{uc} \cap \Sigma_{uo}$, then we do not add any unobservable transition in the construction of $\mathbf{C}_O$. So if $\mathbf{O}$ did loop on all unobservable events, $\mathbf{C}_O$ has the same property. Conversely, if there is a state $\mathbf{s}$ in $\mathbf{O}$ such that $\mathbf{s} \cdot \sigma$ is not defined with $\sigma \in \Sigma_{uc} \cap \Sigma_{uo}$, then in constructing $\mathbf{C}_O$ we will add a new transition $\mathbf{s} \cdot \sigma = \mathbf{u}$ and the resulting automaton will not be in $\mathcal{A}_o$. ■

6. CONTROL UNDER PARTIAL OBSERVATION AND DISTRIBUTED CONTROL

6.1 The Control Problem with Partial Observation

In this section, we want to characterize solutions to the problem of constructing a controller that can function under partial observation. Here we are given two subsets of Σ : Σ_c are the controllable events, and Σ_o are the observable events. We want to construct controllers that prevent or authorize events in Σ_c while getting only the information in Σ_o .

The *control problem with partial observation* is the following:

Let $\mathcal{A}_{oc} = \mathcal{A}_o \cap \mathcal{A}_c$ be the class of observation automata with respect to Σ_o and control automata with respect to Σ_c . Given a process $\mathbf{D}$, and a specification $\mathbf{Z}$ such that $\mathbf{Z} \leq \mathbf{D}$, does there exist an automaton $\mathbf{X} \in \mathcal{A}_{oc}$ such that

$$\mathbf{Z} = \mathbf{D} \times \mathbf{X}?$$

We will solve this problem by combining the results of Theorem 4.2 on the control problem, and Theorem 5.4 on the observation problem. We define, for each automaton $\mathbf{Z}$, the automaton

$$\mathbf{OC}_Z = \mathbf{O}_{C_Z}.$$

THEOREM 6.1: *The mapping $\mathbf{OC}_Z$ has the following properties:*

- (i) $\mathbf{OC}_Z \in \mathcal{A}_{oc}$.
- (ii) $\mathbf{Z} \leq \mathbf{OC}_Z$.
- (iii) For each $\mathbf{X} \in \mathcal{A}_{oc}$, $\mathbf{Z} \leq \mathbf{X} \Rightarrow \mathbf{OC}_Z \leq \mathbf{X}$.

Proof: (i) follows immediately from Proposition 5.6.

(ii) Holds, since $\mathbf{Z} \leq \mathbf{C}_Z$ by Theorem 4.2, and $\mathbf{C}_Z \leq \mathbf{O}_{C_Z}$ by Theorem 5.4.

(iii) Let $\mathbf{X} \in \mathcal{A}_{oc}$. Then if $\mathbf{Z} \leq \mathbf{X}$, we have $\mathbf{C}_Z \leq \mathbf{X}$ by Theorem 4.2, since $\mathbf{X} \in \mathcal{A}_c$. So $\mathbf{O}_{C_Z} \leq \mathbf{X}$ by Theorem 5.4, since $\mathbf{X} \in \mathcal{A}_o$. Thus $\mathbf{OC}_Z \leq \mathbf{X}$. ■

COROLLARY 6.2: *The control problem with partial observation has a solution if and only if*

$$\mathbf{Z} = \mathbf{D} \times \mathbf{OC}_Z. \quad \blacksquare$$

6.2 Distributed Control

We now have the tools to solve the problem of distributing the control among several controllers that can control or observe different sets of events. The problem, called the distributed control problem with global specification,

can be stated as:

Let $\mathcal{A}_{oc_i}$ be a finite family of classes of observation and control automata with respect to Σ_{o_i} and Σ_{c_i} . Given a process $\mathbf{D}$, and a specification $\mathbf{Z}$ such that $\mathbf{Z} \leq \mathbf{D}$, do there exist automata $\mathbf{X}_i \in \mathcal{A}_{oc_i}$ such that

$$\mathbf{Z} = \mathbf{D} \times \prod \mathbf{X}_i?$$

To solve this problem, we simply apply to $\mathbf{Z}$ the construction of the preceding section for each class $\mathcal{A}_{oc_i}$, yielding a family of mappings $\mathbf{OC}_{Z_i}$.

THEOREM 6.3: *The distributed control problem with global specification has a solution if and only if*

$$\mathbf{Z} = \mathbf{D} \times \prod \mathbf{OC}_{Z_i}.$$

Proof: Since $\mathbf{Z} \leq \mathbf{D}$, and $\mathbf{Z} \leq \mathbf{OC}_{Z_i}$ for each i , we have that

$$\mathbf{Z} \leq \mathbf{D} \times \prod \mathbf{OC}_{Z_i}.$$

On the other hand, if a solution exists, we have that $\mathbf{Z} \leq \mathbf{X}_i$ for each i , thus $\mathbf{OC}_{Z_i} \leq \mathbf{X}_i$. Multiplying those together, along with $\mathbf{D}$, we get

$$\mathbf{D} \times \prod \mathbf{OC}_{Z_i} \leq \mathbf{D} \times \prod \mathbf{X}_i = \mathbf{Z}. \quad \blacksquare$$

Notes: Necessary and sufficient conditions for the solution of the distributed control problem with global specification appear in [5], where they are expressed in terms of language properties.

7. AN EXAMPLE

Suppose three rooms are connected with doors as indicated in following diagram:

Each door, labelled A, B, C or D can allow objects to move in either direction. Observable and/or controllable events will be, for example, of the form A_{in} or C_{out} meaning that an object just passed through a door in a given direction (fixing arbitrarily the “in” direction as left to right). We are

to construct two controllers, one at Site 1 and one at Site 2, that should prevent two objects from being in the same room simultaneously. The controller at Site 1 can observe events at doors A, B and C, and the controller at Site 2 events at doors A, B and D. Six events are controllable, i. e. all except A_{out} and B_{out} . Is it possible to distribute these controllable events among the two controllers such that they can enforce the constraint, while allowing maximum freedom to objects?

The following automaton represents the specification Z . Each state displays explicitly the position of objects and we can take the "empty" state as initial. (For clarity, we labeled only the "in" arrows since all events are reversible.)

Figure 1. — The specification Z .

We can easily construct the observation automata O_{z_1} and O_{z_2} (Figure 2) giving the partial view of the specification Z from Site 1 and Site 2 respectively.

If these two automata function together, it is quite easy to convince oneself that the product has only eight accessible states and there is a natural correspondence with the eight original states of the specification. Indeed, states of the product can be thought of as non-empty intersections of states of each of the observers.

If the number of objects is unbounded, the whole process has an infinite number of states. But since we are interested only in enforcing the specification, we can model the whole process D by adding deadlock states (the "alarm" states) to the specification, and directing to them possible events that violate the constraint. In Figure 3, the round states represent these states.

Figure 2. — Observation automata O_{z_1} for Site 1 and O_{z_2} for Site 2.

Figure 3. — The process D.

As for control strategies, suppose first that we choose the obvious distribution of controllable events:

Site 1 : A_{in}, C_{in}, C_{out}
Site 2 : B_{in}, D_{in}, D_{out}

Controllable events distribution (1)

Figure 4 shows parts of the automata OC_{Z_1} and OC_{Z_2} that result from this choice (with the state u resulting from the 'control' construction). These fragments show explicitly that the sequence $A_{in} B_{in} C_{in} D_{in}$ is defined in both automata, and since it is also defined in D , it will be defined in the product $D \times OC_{Z_1} \times OC_{Z_2}$. But the sequence $A_{in} B_{in} C_{in} D_{in}$ is not in the specification Z , so there is no solution with this choice of strategy.

Figure 4. — Partial automata OC_{Z_1} and OC_{Z_2} for distribution (1).

An alternate strategy is the following:

Site 1 : A_{in}, C_{out}, D_{in} Site 2 : B_{in}, D_{out}, C_{in}
--

Controllable events distribution (2)

This distribution, curiously, assigns control of events that are not observable from the site. For example, Site 1 will control event D_{in} , which is unobservable from the site. For this choice, a solution exists. It can be obtained simply from the two observers of Figure 2, by disabling, in a given state, all the events that are undefined but controllable from the site. We obtain the two controllers of Figure 5.

In order to prove that this is a solution, we have to show that the product of these two automata with the process D is contained in the specification Z . It suffices to verify that all events that could lead to an alarm state are disabled by one of the controllers. There are just 6 states to check. For

Figure 5. — Control strategy for Site 1 and Site 2 (disabled events in italics).

example, if the two controllers are in states:

The events A_{in} and B_{in} must be disabled (as one can see in Figure 3). But the first controller disables events A_{in} and C_{out} , and the second disables events B_{in} and D_{out} . The verification is just as elementary for the 5 other states that could lead to an alarm state.

8. CONCLUDING REMARKS

In this paper, we obtained new sets of necessary and sufficient conditions for the solution of problems in the theory of control of discrete event systems. Lemma 3.1 provides a unifying setting for design problems, which we applied successfully to classical control and observation problems. It also suggests that other classes of automata could be considered in order to solve other kinds of design problems.

Further work includes investigation of the *range* problem. In [5] the following problem is solved. With the notations of section 6.2, given $\mathbf{Z} \leq \mathbf{Z}' \leq \mathbf{D}$ where $\mathbf{Z}$ and $\mathbf{Z}'$ are respectively a minimal and a maximal specification, does

there exists controllers $X_i \in \mathcal{A}_{oc_i}$ such that

$$Z \leq \overline{D \times \prod X_i} \leq Z'$$

A necessary and sufficient condition for the existence of a solution is easily seen to be

$$\overline{D \times \prod OC_{z_i}} \leq Z'.$$

Furthermore, if a solution exists, then the product $D \times \prod X_i$ where $X_i = \overline{OC_{z_i}}$ is always non-blocking (simply because anything defined in this product must be marked by Z' , thus by D).

A more difficult problem consists in asking for the existence of controllers such that

$$Z \leq D \times \prod X_i \leq Z'?$$

Again, a necessary and sufficient condition is that $D \times \prod OC_{z_i} \leq Z'$, but there is no simple way to ensure that the product is non-blocking.

We also intend to explore complexity issues, especially as regards the observation construction, which can grow exponentially. These issues are crucial in the commutation and telecommunication problems around which we developed this model.

REFERENCES

1. A. ARNOLD, Systèmes de transitions finis et sémantique des processus communicants, *T.S.I.*, 9,, n° 3, 1990, pp. 193-216.
2. A. ARNOLD and M. NIVAT, *Comportements de processus*, In Colloque AFCET, Les mathématiques de l'informatique, 1982, pp. 35-68.
3. P. RAMADGE and W. WONHAM, Supervisory Control of a Class of Discrete Event Processes, *SIAM J. Control and Optimization*, 25, No. 1, 1987, pp. 206-230.
4. P. RAMADGE and W. WONHAM, The Control of Discrete Event Systems, *Proceedings of the IEEE*, 77, No. 1, 1989, pp. 81-98.
5. K. RUDIE and W. WONHAM, Think Globally, Act Locally: Decentralized Supervisory Control, *IEEE Transactions on Automatic Control*, 37, No. 11, 1992, pp. 1692-1708.
6. F. LIN and M. WONHAM, On Observability of Discrete-Event Systems, *Information Science*, 44, 1988, pp. 173-198.