



Centrum voor Wiskunde en Informatica
REPORT*RAPPORT*

Supervisory control for nondeterministic systems

A. Overkamp

Department of Operations Research, Statistics, and System Theory

BS-R9411 1994

Supervisory Control for Nondeterministic Systems

Ard Overkamp

CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

Abstract

In this paper we present the first results of our attempt to set up a supervisory theory for nondeterministic discrete event systems. The supervisory control problem for nondeterministic systems is specified and motivated. A necessary and sufficient condition for the existence of a supervisor is derived. Also an algorithm is presented that will automatically generate a supervisor. The results are extended to deal with uncontrollable events.

AMS Subject Classification (1991): 68Q75, 93B25

Keywords & Phrases: Nondeterministic Discrete Event Systems, Failure Semantics, Supervisory Control.

1. INTRODUCTION

Up till now discrete event systems have been mostly modelled as deterministic systems. The current state and the next event uniquely determine the next state. Deterministic systems are effectively described by the language they can generate. Conditions under which a supervisor exists and the behavior of the controlled system are also stated in a language context. But in case of partial observation (and partial specification) systems are more appropriately modelled as nondeterministic systems.

It is not sufficient to describe nondeterministic systems only by the language they generate. The blocking properties of the systems are also important [5]. Consider the following example. Given two systems, A and B (Fig. 1 a, b). A is nondeterministic and B is its deterministic equivalent. The languages that the systems can generate are the same, but when connected to system C , (Fig. 1 c) $A||C$ can deadlock but $B||C$ can not (Fig. 1 d, e).




Figure 1: Illustration of the difference between a nondeterministic and a deterministic system.

The problem that we are trying to solve in this paper is the following. Given a nondeterministic system and a nondeterministic specification, find a supervisor such that the controlled system behaves according to the specification. Here 'behaves like' means that the controlled system may only do what the specification allows, and that it may only block an event if the specification can also block that event. A new aspect in this setup is that the specification may also be nondeterministic. This gives the possibility to specify that some nondeterminism is still allowed in the controlled system.

In Section 3 an alternative for language semantics will be introduced: failure semantics [1, 4]. It is a commonly used semantics in computer science. We will use it to state necessary and sufficient conditions for the existence of a supervisor such that the controlled system behaves according to the

specification. An algorithm to automatically generate a supervisor will be presented. In Section 5 these results will be extended to deal with uncontrollable events.

2. NOTATION AND DEFINITIONS

DEFINITION 2.1 (FINITE STATE MACHINE) A (nondeterministic) finite state machine (FSM) is a four tuple (Q, Σ, δ, Q_0) where,

Q is the state space,

Σ is the set of event labels,

$\delta : Q \times \Sigma \rightarrow 2^Q$ is the transition function, and

$Q_0 \subseteq Q$ is the nonempty set of initial states.

The domain of the transition function can be extended in a natural way to strings and sets of states. Let $A = (Q_a, \Sigma, \delta_a, A_0)$ be a FSM, $\mathcal{A} \subseteq Q_a$, $q_a \in Q_a$, $\sigma \in \Sigma$, and $s \in \Sigma^*$. We define

- $\delta_a(q_a, \varepsilon) = q_a$,
- $\delta_a(q_a, s\sigma) = \delta_a(\delta_a(q_a, s), \sigma)$,
- $\delta_a(\mathcal{A}, s) = \bigcup_{q \in \mathcal{A}} \delta_a(q, s)$.

For notational convenience we will use the state machine itself as an argument instead of the initial states. $\delta(A, s) = \delta_a(A_0, s)$.

For a deterministic finite state machine (DFSM) $\delta(q, s)$ is always a singleton set. In this case we will write $q' = \delta(q, s)$ instead of $\{q'\} = \delta(q, s)$. Also, we will simply write the initial state instead of a singleton set containing the initial state.

A useful function is the out-function, which describes the outgoing arcs of a state. $\text{out}(q) = \{\sigma \in \Sigma \mid \delta(q, \sigma) \neq \emptyset\}$.

The set of traces, or language, generated by a FSM is defined as $L(A) = \{s \in \Sigma^* \mid \delta(A, s) \neq \emptyset\}$. A language is prefix closed if $sv \in L \Rightarrow s \in L$. All languages generated by a FSM are prefix closed.

Control will not be enforced by a control map as in the Ramadge-Wonham approach, but by synchronization on common events. The controlled system (i.e. the synchronous composition of the plant and the supervisor) can only execute those events that both the supervisor and the plant can execute.

DEFINITION 2.2 (SYNCHRONOUS COMPOSITION) Let $A = (Q_a, \Sigma, \delta_a, A_0)$ and $B = (Q_b, \Sigma, \delta_b, B_0)$. The synchronous composition of A and B is $A \parallel B = (Q_a \times Q_b, \Sigma, \delta_{ab}, A_0 \times B_0)$ where $\delta_{ab}((q_a, q_b), \sigma) = \delta_a(q_a, \sigma) \times \delta_b(q_b, \sigma)$. Note that for $Q_a \subseteq Q_a, Q_b \subseteq Q_b : \emptyset \times Q_b = Q_a \times \emptyset = \emptyset$.

The next lemma is a direct result from the definition.

LEMMA 2.3 For FSMs A, B , and $s \in \Sigma^*$;

1. $(q_a, q_b) \in \delta(A \parallel B, s) \Leftrightarrow q_a \in \delta(A, s) \wedge q_b \in \delta(B, s)$.
2. For $(q_a, q_b) \in \delta(A \parallel B, s)$ $\text{out}((q_a, q_b)) = \text{out}(q_a) \cap \text{out}(q_b)$.
3. $L(A \parallel B) = L(A) \cap L(B)$.

3. FAILURE SEMANTICS

In the introduction we already argued that the language alone does not provide enough information for nondeterministic systems. The blocking properties of the systems are also important. In computer science failure semantics is introduced to deal with such problems. In failure semantics a system is not only described by the traces it can generate but also by the set of events that it can block after a trace (fail to execute). We will not elaborate on it here. The reader is referred to [1, 4] for more information. We will only define what we need and show that it fits our interests.

First, let us define what we mean by blocking. In the specification it may be stated that after certain traces the system is allowed to block, but after other traces it may not. So we have to define nonblockingness with respect to an already executed trace.

DEFINITION 3.1 (BLOCKING) Let A be a FSM and $s \in L(A)$; A is nonblocking after s if $\forall q_a \in \delta(A, s) \text{ out}(q_a) \neq \emptyset$. A FSM A is nonblocking if it is nonblocking after all $s \in L(A)$.

As stated in the introduction, the controlled system may only do what the specification allows and it may only block an event if the specification can also block that event. Because the specification is nondeterministic it is possible that sometimes an event is blocked in one branch of a nondeterministic choice and allowed in another branch. So, the controlled system has the choice to either block or execute that event. Because of this freedom the system does not have to be equivalent to the specification. It only has to reduce (implement) it [2, 3].

DEFINITION 3.2 (REDUCTION) Let A, B be FSMs; A reduces B ($A \sqsubseteq B$) if

- i) $L(A) \subseteq L(B)$, and
- ii) $\forall s \in L(A), \forall q_a \in \delta(A, s), \exists q_b \in \delta(B, s) \text{ s.t. } \text{out}(q_b) \subseteq \text{out}(q_a)$.

Here, point i states that system A may only do what system B allows, and point ii states that A may only block what B can also block.

The reduction relation guarantees that the controlled system does not block in an environment if the specification does not block in the same environment. This is formally stated in Theorem 3.3.

THEOREM 3.3 Let A, B be FSMs such that $L(A) \subseteq L(B)$. For all FSM C and $s \in \Sigma^*$, $B \parallel C$ nonblocking after s implies $A \parallel C$ nonblocking after s , if and only if $A \sqsubseteq B$.

Proof (if part) if $A \sqsubseteq B$ then $\forall q_a \in \delta(A, s) \exists q_b \in \delta(B, s) \text{ s.t. } \text{out}(q_b) \subseteq \text{out}(q_a)$. If $B \parallel C$ is nonblocking after s then, by Lemma 2.3, $\forall q_c \in \delta(C, s) \text{ out}(q_b) \cap \text{out}(q_c) \neq \emptyset$. Then also $A \parallel C$ is nonblocking after s because $\text{out}(q_a) \cap \text{out}(q_c) \supseteq \text{out}(q_b) \cap \text{out}(q_c) \neq \emptyset$.

(only if part). We will prove that if A does not reduce B then there exists a FSM C and a string s such that $B \parallel C$ is nonblocking after s , but $A \parallel C$ can block after s . If not $A \sqsubseteq B$ then $\exists s \in L(A), \exists q_a \in \delta(A, s) \text{ s.t. } \forall q_b \in \delta(B, s) \text{ out}(q_b) \not\subseteq \text{out}(q_a)$. Let C be a DFSM such that $q_c = \delta(C, s)$ with $\text{out}(q_c) = \Sigma - \text{out}(q_a)$. Then $\forall q_b \in \delta(B, s) \text{ out}(q_b) \cap \text{out}(q_c) = \text{out}(q_b) \cap (\Sigma - \text{out}(q_a)) = \text{out}(q_b) - \text{out}(q_a) \neq \emptyset$. But $\text{out}(q_a) \cap \text{out}(q_c) = \text{out}(q_a) \cap (\Sigma - \text{out}(q_a)) = \emptyset$. \square

4. CONTROLLER SYNTHESIS

In this section we will first state under what condition there exists a supervisor such that the controlled system reduces the specification. We will call this condition reducibility. Then, an algorithm will be presented that generates such a supervisor. In the following G will denote an uncontrolled system and E a specification.

DEFINITION 4.1 (REDUCIBILITY) Let G, E be FSMs. A language K is reducible (w.r.t. G, E) if

$$\forall s \in K, \forall q_g \in \delta(G, s), \exists q_e \in \delta(E, s) \text{ s.t. } \text{out}(q_e) \subseteq \text{out}(q_g) \cap \{\sigma \in \Sigma \mid s\sigma \in K\}.$$

LEMMA 4.2 Let G, E be FSMs. Let S be a DFSM such that $L(S)$ is reducible. Then $G \parallel S \sqsubseteq E$.

Proof (Point i of the definition of reduction) If $s \in L(S) \cap L(G)$ then $\delta(G, s)$ is not empty. So, by the definition of reducibility, there exists a $q_e \in \delta(E, s)$. Hence $s \in L(E)$, and $L(G \parallel S) \subseteq L(E)$.

(Point ii of the definition of reduction). For $s \in L(G \parallel S)$ let $(q_g, q_s) \in \delta(G \parallel S, s)$, then, by Lemma 2.3, $q_g \in \delta(G, s)$ and $q_s = \delta(S, s)$. Because S is deterministic $\{\sigma \in \Sigma \mid s\sigma \in L(S)\} = \text{out}(q_s)$. Then, by reducibility of $L(S)$, we know $\exists q_e \in \delta(E, s)$ such that $\text{out}(q_e) \subseteq \text{out}(q_g) \cap \text{out}(q_s)$. Hence also, by Lemma 2.3, $\text{out}(q_e) \subseteq \text{out}((q_g, q_s))$. \square

THEOREM 4.3 Let G, E be FSMs. There exists a supervisor S such that $G \parallel S \sqsubseteq E$ if and only if there exists a nonempty, prefix closed, and reducible language K .

Proof (if part) Let S_K be a deterministic state machine generating K . The proof follows directly from Lemma 4.2.

(only if part). Let $K = L(S||G)$. Then K is nonempty and prefix closed. We will prove that K is reducible. $\forall s \in K$, $s \in L(S||G)$, so $\delta(S||G, s) \neq \emptyset$. Then, by lemma 2.3, $\delta(G, s) \neq \emptyset$ and $\delta(S, s) \neq \emptyset$. Thus $\forall s \in K$, $\forall q_g \in \delta(G, s) \exists q_s$ s.t. $(q_g, q_s) \in \delta(G||S, s)$. By the definition of reduction $\exists q_e \in \delta(E, s)$ s.t. $\text{out}(q_e) \subseteq \text{out}((q_g, q_s))$. By lemma 2.3 $\text{out}((q_g, q_s)) \subseteq \text{out}(q_g)$. And by construction of K $\text{out}((q_g, q_s)) \subseteq \{\sigma \in \Sigma | s\sigma \in K\}$. So, $\text{out}(q_e) \subseteq \text{out}(q_g) \cap \{\sigma \in \Sigma | s\sigma \in K\}$. \square

ALGORITHM 4.4 The following algorithm will construct a supervisor (if it exists) such that the controlled system reduces the specification.

1. Generate a deterministic state machine $R^0 = (Q_r^0, \Sigma, \delta_r^0, R_0^0)$, where

$$Q_r^0 = 2^{Q_s} \times 2^{Q_e},$$

$$R_0^0 = (G_0, E_0),$$

$$\delta_r^0((\mathcal{G}, \mathcal{E}), \sigma) = \begin{cases} (\delta_g(\mathcal{G}, \sigma), \delta_e(\mathcal{E}, \sigma)), & \text{if } \delta_g(\mathcal{G}, \sigma) \neq \emptyset \wedge \delta_e(\mathcal{E}, \sigma) \neq \emptyset, \\ \text{empty}, & \text{otherwise.} \end{cases}$$

2. Construct R^{i+1} from R^i by removing all reachable states $q_r^i = (\mathcal{G}, \mathcal{E})^i$ (including its in- and out-going transitions) from the state space of R^i that do *not* satisfy the following condition,

$$\forall q_g \in \mathcal{G} \exists q_e \in \mathcal{E} \text{ s.t. } \text{out}(q_e) \subseteq \text{out}(q_g) \cap \text{out}(q_r^i).$$

3. Repeat the previous step until all reachable states satisfy the condition or until there are no more reachable states left. Let R be the last R^i .

THEOREM 4.5 *Let G and E be FSMs. Then Algorithm 4.4 produces a DFSM R in a finite number of steps. If the state space of R is nonempty then $G||R \sqsubseteq E$. If the initial state is removed from R then no supervisor exists such that the controlled system reduces the specification.*

Proof (The algorithm stops in a finite number of steps) If in one step of the algorithm no state is removed, then the algorithm stops because all states satisfy the condition. If in every step at least one state is removed from the state space then the algorithm will eventually halt because the state space is finite.

(The algorithm returns a correct solution). We will prove that when the algorithm finds a solution then $L(R)$ is reducible. And thus, by Lemma 4.2, $G||R \sqsubseteq E$. Let $s \in L(R)$ and $q_r = (\mathcal{G}, \mathcal{E}) = \delta(R, s)$. Then $\mathcal{G} = \delta(G, s)$, $\mathcal{E} = \delta(E, s)$, and because R is deterministic $\text{out}(q_r) = \{\sigma \in \Sigma | s\sigma \in L(R)\}$. Then, by step 3 of the algorithm, $\forall q_g \in \delta(G, s) \exists q_e \in \delta(E, s)$ s.t. $\text{out}(q_e) \subseteq \text{out}(q_g) \cap \text{out}(q_r) = \text{out}(q_g) \cap \{\sigma \in \Sigma | s\sigma \in L(R)\}$.

(The algorithm finds a solution if one exists). Assume there exists a nonempty, prefix closed, and reducible language K . First, we will prove that $K' = K \cap L(G)$ is also nonempty, prefix closed and reducible. Then we will prove that $K' \subseteq L(R)$. Thus the algorithm will return a nonempty solution.

Because $\varepsilon \in K$ and $\varepsilon \in L(G)$, $K \cap L(G) \neq \emptyset$. Because K and $L(G)$ are prefix closed, $sv \in K \cap L(G) \Rightarrow s \in K \wedge s \in L(G) \Rightarrow s \in K \cap L(G)$. So K' is prefix closed. By reducibility of K , $\forall s \in K \cap L(G)$, $\forall q_g \in \delta(G, s)$, $\exists q_e \in \delta(E, s)$ s.t. $\text{out}(q_e) \subseteq \text{out}(q_g) \cap \{\sigma \in \Sigma | s\sigma \in K\} = \text{out}(q_g) \cap \{\sigma \in \Sigma | s\sigma \in K \cap L(G)\}$. So K' is also reducible.

Now, we will prove by induction on the number of steps of the algorithm that $K' \subseteq L(R^i)$ for all i . So, also $K' \subseteq L(R)$.

Initial step: $\forall s \in K', s \in L(G)$. Then, by reducibility, $\delta(E, s) \neq \emptyset$. So $s \in L(E)$, and $K' \subseteq L(G) \cap L(E) = L(R^0)$.

Inductive hypothesis: $K' \subseteq L(R^i)$.

Let $s \in K'$ and $q_r^i = \delta(R^i, s)$. By the hypothesis, $s\sigma \in K' \Rightarrow s\sigma \in L(R^i) \Rightarrow \delta(R^i, s\sigma) = \delta_{R^i}(q_r^i, \sigma) \neq \emptyset \Rightarrow \sigma \in \text{out}(q_r^i)$. So, $\{\sigma \in \Sigma \mid s\sigma \in K'\} \subseteq \text{out}(q_r^i)$. By reducibility, $\forall q_g \in \mathcal{G}, \exists q_e \in \mathcal{E}$ s.t. $\text{out}(q_e) \subseteq \text{out}(q_g) \cap \{\sigma \in \Sigma \mid s\sigma \in K'\} \subseteq \text{out}(q_g) \cap \text{out}(q_r^i)$. So, q_r^i will not be removed from R^i . Hence, $s \in L(R^{i+1})$ and $K' \subseteq L(R^{i+1})$. \square

From the last part of the proof it can be deduced that the algorithm generates the least restrictive supervisor. (The supremal element with respect to language inclusion).

5. UNCONTROLLABLE EVENTS

Sometimes a system can generate events that can not be blocked by a supervisor (e.g. machine breakdown). Ramadge and Wonham showed that in the presence of these uncontrollable events we need the condition of controllability to guarantee the existence of a supervisor. We will show that for nondeterministic systems the same condition is needed.

Recall from [6] the definition of controllability. Let G be a FSM, $\Sigma_u \subseteq \Sigma$. A prefix closed language K is controllable (w.r.t G, Σ_u) if $K\Sigma_u \cap L(G) \subseteq K$. Note that this is equivalent to: $\forall s \in K, \forall q_g \in \delta(G, s) \text{ out}(q_g) \cap \Sigma_u \subseteq \{\sigma \in \Sigma \mid s\sigma \in K\}$.

Ramadge and Wonham called a supervisor that always accepts an uncontrollable event complete. We have to adapt the definition of completeness to deal with nondeterministic systems and control by synchronization.

DEFINITION 5.1 (COMPLETENESS) A supervisor S is complete (w.r.t a FSM G) if

$$\forall s \in L(S||G), \forall q_s \in \delta(S, s), \forall q_g \in \delta(G, s), \text{out}(q_g) \cap \Sigma_u \subseteq \text{out}(q_s) .$$

THEOREM 5.2 *Let G and E be FSMs. There exists a supervisor S such that $G||S \sqsubseteq E$, and S complete w.r.t G if and only if there exists a nonempty, prefix closed, reducible and controllable language K .*

Proof (if part) Let S_K be a DFSM generating K . Then, by lemma 4.2, $G||S_K \sqsubseteq E$. $\forall s \in L(S_K||G)$, let $q_s = \delta(S_K, s)$. Then, by controllability of $L(S_K)$, $\forall q_g \in \delta(G, s) \text{ out}(q_g) \cap \Sigma_u \subseteq \{\sigma \in \Sigma \mid s\sigma \in L(S_K)\} = \text{out}(q_s)$. so S_K is complete.

(only if part). Take $K = L(S||G)$, then by the proof of Theorem 4.3 (only if part) K is reducible. $\forall s \in L(S||G), \forall q_s \in \delta(S, s), \forall q_g \in \delta(G, s) \text{ out}(q_g) \cap \Sigma_u \subseteq \text{out}(q_s) \Rightarrow \text{out}(q_g) \cap \Sigma_u \subseteq \text{out}(q_s) \cap \text{out}(q_g) = \text{out}((q_g, q_s)) \subseteq \{\sigma \in \Sigma \mid s\sigma \in L(S)\}$. So, $L(S||G)$ is controllable. \square

ALGORITHM 5.3 This algorithm constructs a complete supervisor (if it exists) such that the controlled system reduces the specification. The algorithm is the same as Algorithm 4.4, except that the following is added to step 2.

2. ...

Also, remove those states $q_r^i = (G, \mathcal{E})^i$ that do not satisfy the following condition,

$$\forall q_g \in \mathcal{G} \text{ out}(q_g) \cap \Sigma_u \subseteq \text{out}(q_r^i) .$$

THEOREM 5.4 *Let G and E be FSMs. Then Algorithm 5.3 produces a DFSM R in a finite number of steps. If the state space of R is nonempty then $G||R \sqsubseteq E$ and R is complete. If the initial state is removed from R then no complete supervisor exists such that the controlled system reduces the specification.*

Proof The proof goes along the same lines as the proof of Theorem 4.5. The following has to be added to the different steps of the proof.

(The algorithm returns a correct solution). $\forall s \in L(R||G)$, let $q_r = (G, \mathcal{E}) = \delta(R, s)$. Then, by step 3 of the algorithm, $\forall q_g \in \mathcal{G} = \delta(G, s)$, $\text{out}(q_g) \cap \Sigma_u \subseteq \text{out}(q_r)$. Hence R is complete.

(The algorithm finds a solution if one exists). After the sentence that starts with 'By reducibility': Also, by controllability, $\forall s \in K', \forall q_g \in \mathcal{G}$, $\text{out}(q_g) \cap \Sigma_u \subseteq \{\sigma \in \Sigma | s\sigma \in K'\} \subseteq \text{out}(q_r^i)$. So, q_r^i will not be removed from $R^i \dots$ \square

6. CONCLUSIONS

In case of systems that are partially observed or partially specified one has to realize that the behaviour of a system depends on the nondeterministic properties of that system. This paper has been an attempt to set up a supervisory theory for nondeterministic systems. A condition (reducibility) is found for the existence of a supervisor such that the controlled system behaves like the specification. An algorithm is described which synthesizes a deterministic least restrictive supervisor. These results are extended to deal with uncontrollable events.

What remains to be done is to analyze the consequences of these results for systems with partial specification and partial observation.

REFERENCES

1. J.C.M Baeten and W.P. Weijland. *Process algebra*. Cambridge University Press, Cambridge, 1990.
2. E. Brinksma, G. Scollo and C. Steenbergen. LOTOS specifications, their implementations, and their tests. *Proc. of IFIP Workshop 'Protocol Specification, Testing and Verification VI'*, pages 349–360, 1987.
3. S.D. Brooks, C.A.R. Hoare, and A.W. Roscoe. A theory for communicating sequential processes. *Journal of the ACM*, 31:560–599, 1984.
4. R.J. van Gladbeek. The linear time - branching time spectrum. *Proc. CONCUR '90*, Amsterdam, Lecture Notes in Computer Science 458, pages 278–297, 1990.
5. M. Heymann. Concurrency and discrete event control. *IEEE Control Systems Magazine*, 10:103–112, 1990.
6. P.J.G. Ramadge and W.M. Wonham. The control of discrete event systems. *Proc. of the IEEE*, 77:81–98, 1989.