Automatica 81 (2017) 142–147

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Supervisor localization of discrete-event systems under partial observation[☆]

Renyuan Zhang[a], Kai Cai[b], W. Murray Wonham[c]

[a] School of Automation, Northwestern Polytechnical University, China

[b] Department of Electrical and Information Engineering, Osaka City University, Japan

[c] Systems Control Group, Department of Electrical and Computer Engineering, University of Toronto, Canada

ARTICLE INFO

Article history:
Received 12 February 2016
Received in revised form
3 November 2016
Accepted 24 February 2017

Keywords:
Discrete-event systems
Supervisory control
Supervisor localization
Partial observation

ABSTRACT

Recently we developed *supervisor localization*, a top-down approach to distributed control of discrete-event systems. Its essence is the allocation of monolithic (global) control action among the local control strategies of individual agents. In this paper, we extend supervisor localization by considering partial observation; namely not all events are observable. Specifically, we employ the recently proposed concept of *relative observability* to compute a partial-observation monolithic supervisor, and then design a suitable localization procedure to decompose the supervisor into a set of local controllers. In the resulting local controllers, only observable events can cause state change. We finally illustrate our result by a Transfer Line example.

1. Introduction

In Cai and Wonham (2010a,b, 2015, 2016) and Zhang, Cai, Gan, Wang, and Wonham (2013) we developed a top-down approach, called *supervisor localization*, to the distributed control of multi-agent discrete-event systems (DES). This approach first synthesizes a monolithic supervisor (or a heterarchical array of modular supervisors) assuming that *all* events can be observed, and then decomposes the supervisor into a set of local controllers for the component agents. Localization creates a purely distributed control architecture in which each agent is controlled by its own local controller; this is particularly suitable for applications consisting of many autonomous components, e.g. multi-robot systems. Moreover, localization can significantly improve the comprehensibility of control logic, because the resulting local controllers typically have many fewer states than their parent supervisor. The assumption of full event observation, however, may be too strong in practice, since there often lack enough sensors to observe every event.

In this paper and its conference precursor (Zhang & Cai, 2016a), we extend supervisor localization to address the issue of partial observation. Our approach is as follows. We first synthesize a partial-observation monolithic supervisor using the concept of *relative observability* in Cai, Zhang, and Wonham (2015, 2016). Relative observability is generally stronger than observability (Cieslak, Desclaux, Fawaz, & Varaiya, 1988; Lin & Wonham, 1988), weaker than normality (Cieslak et al., 1988; Lin & Wonham, 1988), and the supremal relatively observable (and controllable) sublanguage of a given language exists. The supremal sublanguage may be effectively computed (Cai et al., 2015), and then implemented by a partial-observation (feasible and nonblocking) supervisor (Wonham, 2016, Chapter 6). We then suitably extend the localization procedure in Cai and Wonham (2010a) to decompose the supervisor into local controllers for individual agents, and moreover prove that the derived local controlled behavior is equivalent to the monolithic one.

The main contributions of this work are as follows. First, we propose the combination of supervisor localization (Cai & Wonham, 2010a) with relative observability (Cai et al., 2015), which leads to a systematic, computationally effective approach to distributed control of DES under partial observation. In particular, local controllers with only observable state transitions are automatically synthesized, and the collective local controlled behavior is guaranteed to be the same as the global nonblocking behavior.

☆ This work was supported in part by the National Nature Science Foundation of China, Grant no. 61403308; JSPS KAKENHI Grant no. JP16K18122; the Natural Sciences and Engineering Research Council, Canada, Grant no. 7399. The material in this paper was partially presented at the 2016 American Control Conference, July 6–8, 2016, Boston, MA, USA. This paper was recommended for publication in revised form by Associate Editor Christoforos Hadjicostis under the direction of Editor Christos G. Cassandras.

E-mail addresses: ryzhang@nwpu.edu.cn (R. Zhang), kai.cai@eng.osaka-cu.ac.jp (K. Cai), wonham@ece.utoronto.ca (W.M. Wonham).

http://dx.doi.org/10.1016/j.automatica.2017.03.018
0005-1098/

Second, we identify the linguistic essence of partial-observation localization by developing the following key mappings and concepts (as extensions to their full-observation counterparts). The mappings include E_α, D_α, M and T (see definitions in Section 3.2) which capture the control and marking information of the partial-observation supervisor. Based on these mappings, the new concepts are introduced, including partial-observation control covers and local controllers. In particular, a partial-observation control cover is defined on the state set of the partial-observation supervisor; roughly speaking, the latter corresponds to the powerset of the full-observation supervisor's state set. Moreover, a partial-observation local controller contains only observable state transitions, and uses control functions to determine the existence of selfloops of unobservable controllable events.

Our proposed localization procedure can in principle be used to construct local controllers from a partial-observation supervisor computed by any synthesis method. In particular, the algorithms in Takai and Ushio (2003), and Yin and Lafortune (2016b) compute a nonblocking (maximally) observable sublanguage that is generally incomparable with the supremal relatively observable sublanguage. The reason that we adopt relative observability is first of all that its generator-based computation of the supremal sublanguage is better suited for applying our localization algorithm; by contrast (Yin & Lafortune, 2016b) uses a different transition structure called "bipartite transition system". Another important reason is that the computation of relative observability has been implemented and tested on a set of benchmark examples. This enables us to study distributed control under partial observation of more realistic systems; by contrast, the examples reported in Takai and Ushio (2003) and Yin and Lafortune (2016b) are limited to academic ones.

The paper is organized as follows. Section 2 reviews the supervisory control problem of DES under partial observation and formulates the partial-observation supervisor localization problem. Section 3 develops the partial-observation localization procedure, and Section 4 illustrates the procedure by a Transfer Line example. Finally Section 5 states our conclusions.

2. Preliminaries and problem formulation

2.1. Supervisory control of DES under partial observation

A DES plant is given by a generator

$$\mathbf{G} = (Q, \Sigma, \delta, q_0, Q_m) \tag{1}$$

where Q is the finite state set; $q_0 \in Q$ is the initial state; $Q_m \subseteq Q$ is the subset of marker states; Σ is the finite event set; $\delta : Q \times \Sigma \to Q$ is the (partial) state transition function. In the usual way, δ is extended to $\delta : Q \times \Sigma^* \to Q$, and we write $\delta(q, s)!$ to mean that $\delta(q, s)$ is defined. Let Σ^* be the set of all finite strings, including the empty string ϵ. The *closed behavior* of $\mathbf{G}$ is the language $L(\mathbf{G}) = \{s \in \Sigma^* | \delta(q_0, s)!\}$ and the *marked behavior* is $L_m(\mathbf{G}) = \{s \in L(\mathbf{G}) | \delta(q_0, s) \in Q_m\} \subseteq L(\mathbf{G})$. A string s_1 is a *prefix* of a string s, written $s_1 \le s$, if there exists s_2 such that $s_1 s_2 = s$. The *(prefix) closure* of $L_m(\mathbf{G})$ is $\overline{L_m(\mathbf{G})} := \{s_1 \in \Sigma^* | (\exists s \in L_m(\mathbf{G}))\ s_1 \le s\}$. In this paper, we assume that $\overline{L_m(\mathbf{G})} = L(\mathbf{G})$; namely, $\mathbf{G}$ is *nonblocking*.

For supervisory control, the event set Σ is partitioned into Σ_c, the subset of controllable events and Σ_{uc}, the subset of uncontrollable events (i.e. $\Sigma = \Sigma_c \dot\cup \Sigma_{uc}$). For partial observation, Σ is partitioned into Σ_o, the subset of observable events, and Σ_{uo}, the subset of unobservable events (i.e. $\Sigma = \Sigma_o \dot\cup \Sigma_{uo}$). Bring in the *natural projection* $P : \Sigma^* \to \Sigma_o^*$ defined by: (i) $P(\epsilon) = \epsilon$; (ii) $P(\sigma) = \sigma$ if $\sigma \in \Sigma_o$; (iii) $P(\sigma) = \epsilon$ if $\sigma \notin \Sigma_o$; (iv) $P(s\sigma) = P(s)P(\sigma)$, for all $s \in \Sigma^*$ and $\sigma \in \Sigma$. As usual, P is extended to $P : Pwr(\Sigma^*) \to Pwr(\Sigma_o^*)$, where $Pwr(\cdot)$ denotes powerset. Write $P^{-1} : Pwr(\Sigma_o^*) \to Pwr(\Sigma^*)$ for the *inverse-image function* of P.

A *supervisory control* for $\mathbf{G}$ is any map $V : L(\mathbf{G}) \to \Gamma$, where $\Gamma := \{\gamma \subseteq \Sigma | \gamma \supseteq \Sigma_{uc}\}$. Then the closed-loop system is $V/\mathbf{G}$, with closed behavior $L(V/\mathbf{G})$ and marked behavior $L_m(V/\mathbf{G})$ (Wonham, 2016). Under partial observation $P : \Sigma^* \to \Sigma_o^*$, we say that V is *feasible* if

$$(\forall s, s' \in L(\mathbf{G}))\quad P(s) = P(s') \Rightarrow V(s) = V(s'), \tag{2}$$

and V is *nonblocking* if $\overline{L_m(V/\mathbf{G})} = L(V/\mathbf{G})$.

It is well-known (Lin & Wonham, 1988) that under partial observation, a feasible and nonblocking supervisory control V exists which synthesizes a (nonempty) sublanguage $K \subseteq L_m(\mathbf{G})$ if and only if K is both controllable and observable (Wonham, 2016). When K is not observable, however, there generally does not exist the supremal observable (and controllable) sublanguage of K. Recently in Cai et al. (2015), a new concept of *relative observability* is proposed, which is stronger than observability but permits the existence of the supremal relatively observable sublanguage.

Formally, a sublanguage $K \subseteq L_m(\mathbf{G})$ is *controllable* (Wonham, 2016) if $\overline{K}\Sigma_{uc} \cap L(\mathbf{G}) \subseteq \overline{K}$. Let $C \subseteq L_m(\mathbf{G})$. A sublanguage $K \subseteq C$ is *relatively observable* with respect to C (or C-observable) if for every pair of strings $s, s' \in \Sigma^*$ that are lookalike under P, i.e. $P(s) = P(s')$, the following two conditions hold (Cai et al., 2015):

$$\text{(i)}\ (\forall \sigma \in \Sigma)\quad s\sigma \in \overline{K}, s' \in \overline{C}, s'\sigma \in L(\mathbf{G}) \Rightarrow s'\sigma \in \overline{K} \tag{3}$$

$$\text{(ii)}\ s \in K, s' \in \overline{C} \cap L_m(\mathbf{G}) \Rightarrow s' \in K. \tag{4}$$

For $E \subseteq L_m(\mathbf{G})$ write $\mathcal{CO}(E)$ for the family of controllable and C-observable sublanguages of E. Then $\mathcal{CO}(E)$ has a unique supremal element $\sup \mathcal{CO}(E)$ which may be effectively computed (Cai et al., 2015).

2.2. Formulation of partial-observation localization problem

Let the plant $\mathbf{G}$ be comprised of N (>1) component agents

$$\mathbf{G}_k = (Q_k, \Sigma_k, \delta_k, q_{0,k}, Q_{m,k}), \quad k = 1, \ldots, N.$$

Then $\mathbf{G}$ is the synchronous product (Wonham, 2016) of $\mathbf{G}_k$ (k in the integer range $\{1, \ldots, N\}$, denoted as $[1, N]$), i.e. $\mathbf{G} = \|_{k \in [1,N]} \mathbf{G}_k$. Here Σ_k need not be pairwise disjoint. These agents are implicitly coupled through a specification language $E \subseteq \Sigma^*$ that imposes a constraint on the global behavior of $\mathbf{G}$ (E may itself be the synchronous product of multiple component specifications). For the plant $\mathbf{G}$ and the imposed specification E, let the generator $\mathbf{SUP} = (X, \Sigma, \xi, x_0, X_m)$ be such that

$$L_m(\mathbf{SUP}) := \sup \mathcal{CO}(E \cap L_m(\mathbf{G})) \tag{5}$$

and $L(\mathbf{SUP}) = \overline{L_m(\mathbf{SUP})}$ (i.e. $\mathbf{SUP}$ is nonblocking). We call $\mathbf{SUP}$ the *controllable and observable behavior.*[1] To rule out the trivial case, we assume that $L_m(\mathbf{SUP}) \neq \emptyset$.

Now let $\alpha \in \Sigma_c$ be an arbitrary controllable event, which may or may not be observable. We say that a generator

$$\mathbf{LOC}_\alpha = (Y_\alpha, \Sigma_\alpha, \eta_\alpha, y_{0,\alpha}, Y_{m,\alpha}), \quad \Sigma_\alpha \subseteq \Sigma_o \cup \{\alpha\}$$

is a *partial-observation local controller* for α if (i) $\mathbf{LOC}_\alpha$ enables/disables the event α (and only α) consistently with $\mathbf{SUP}$, and (ii) if α is unobservable, then α-transitions are selfloops in $\mathbf{LOC}_\alpha$, i.e.

$$(\forall y \in Y_\alpha)\quad \eta_\alpha(y, \alpha)! \Rightarrow \eta_\alpha(y, \alpha) = y.$$

[1] Note that **SUP**, defined over the entire event set Σ, is *not* a representation of a partial-observation supervisor. The latter can only have observable events as state transitions, according to the definition in Section 3.1.

Condition (i) means that for all $s \in \Sigma^*$ there holds

$$P_\alpha(s)\alpha \in L(\mathbf{LOC}_\alpha),\ s\alpha \in L(\mathbf{G}),\ s \in L(\mathbf{SUP}) \Leftrightarrow s\alpha \in L(\mathbf{SUP}) \qquad (6)$$

where $P_\alpha : \Sigma^* \to \Sigma_\alpha^*$ is the natural projection. Condition (ii) requires that only observable events may cause a state change in $\mathbf{LOC}_\alpha$, i.e.

$$(\forall y, y' \in Y_\alpha, \forall \sigma \in \Sigma_\alpha) \quad y' = \eta_\alpha(y, \sigma)!,\ y' \neq y \Rightarrow \sigma \in \Sigma_o. \qquad (7)$$

This requirement is a distinguishing feature of a partial-observation local controller as compared to its full-observation counterpart in Cai and Wonham (2010a).

Note that the event set Σ_α of $\mathbf{LOC}_\alpha$ in general satisfies

$$\{\alpha\} \subseteq \Sigma_\alpha \subseteq \Sigma_o \cup \{\alpha\};$$

in typical cases, both subset containments are strict. The events in $\Sigma_\alpha \setminus \{\alpha\}$ may be viewed as communication events that are critical to achieve synchronization with other partial-observation local controllers (for other controllable events). The event set Σ_α is not fixed *a priori*, but will be determined as part of the localization result presented in the next section.

We now formulate the *Partial-Observation Supervisor Localization Problem*:

Construct a set of partial-observation local controllers $\{\mathbf{LOC}_\alpha \mid \alpha \in \Sigma_c\}$ such that the collective controlled behavior of these local controllers is equivalent to the controllable and observable behavior $\mathbf{SUP}$ in (5) with respect to $\mathbf{G}$, i.e.

$$L(\mathbf{G}) \cap \left(\bigcap_{\alpha\in\Sigma_c} P_\alpha^{-1} L(\mathbf{LOC}_\alpha)\right) = L(\mathbf{SUP})$$

$$L_m(\mathbf{G}) \cap \left(\bigcap_{\alpha\in\Sigma_c} P_\alpha^{-1} L_m(\mathbf{LOC}_\alpha)\right) = L_m(\mathbf{SUP}).$$

Having obtained a set of partial-observation local controllers, one for each controllable event, we allocate each controller to the agent(s) owning the corresponding controllable event. Thereby we build for a multi-agent DES a nonblocking distributed control architecture under partial observation.

3. Partial-observation localization procedure

3.1. Uncertainty set

Let $\mathbf{G} = (Q, \Sigma, \delta, q_0, Q_m)$ be the plant, $\Sigma_o \subseteq \Sigma$ the subset of observable events, and $P : \Sigma^* \to \Sigma_o^*$ the corresponding natural projection. Also let $\mathbf{SUP} = (X, \Sigma, \xi, x_0, X_m)$ be the controllable and observable behavior (as defined in (5)).

Under partial observation, when a string $s \in L(\mathbf{SUP})$ occurs, what is observed is $P(s)$; namely, the events in Σ_{uo} $(=\Sigma \setminus \Sigma_o)$ are erased. Hence two different strings s and s' may look alike, i.e. $P(s) = P(s')$. For $s \in L(\mathbf{SUP})$, let $U(s)$ be the subset of states that may be reached by some string s' that looks like s, i.e.

$$U(s) = \{x \in X \mid (\exists s' \in \Sigma^*)\, P(s) = P(s'),\ x = \xi(x_0, s')\}.$$

It is always true that the state $\xi(x_0, s) \in U(s)$. We call $U(s)$ the *uncertainty set* of the state $\xi(x_0, s)$ associated with string s. Let

$$\mathcal{U}(X) := \{U(s) \subseteq X \mid s \in L(\mathbf{SUP})\} \qquad (8)$$

i.e. $\mathcal{U}(X)$ is the set of uncertainty sets of all states (associated with strings in $L(\mathbf{SUP})$) in X. The size of $\mathcal{U}(X)$ is $|\mathcal{U}(X)| \leq 2^{|X|}$ in general.

The transition function associated with $\mathcal{U}(X)$ is $\hat{\xi} : \mathcal{U}(X) \times \Sigma_o \to \mathcal{U}(X)$ given by

$$\hat{\xi}(U, \sigma) = \bigcup \{\xi(x, u_1\sigma u_2) \mid x \in U, u_1, u_2 \in \Sigma_{uo}^*\}. \qquad (9)$$

If there exist $u_1, u_2 \in \Sigma_{uo}^*$ such that $\xi(x, u_1\sigma u_2)!$, then $\hat{\xi}(U, \sigma)$ is defined, denoted as $\hat{\xi}(U, \sigma)!$. With $\mathcal{U}(X)$ and $\hat{\xi}$, define the *partial-observation monolithic supervisor* (Wonham, 2016)

$$\mathbf{SUPO} = (\mathcal{U}(X), \Sigma_o, \hat{\xi}, U_0, \mathcal{U}_m) \qquad (10)$$

where $U_0 = U(\epsilon)$ and $\mathcal{U}_m = \{U \in \mathcal{U}(X) \mid U \cap X_m \neq \emptyset\}$. It is known (Wonham, 2016) that $L(\mathbf{SUPO}) = P(L(\mathbf{SUP}))$ and $L_m(\mathbf{SUPO}) = P(L_m(\mathbf{SUP}))$.

Now let $U \in \mathcal{U}(X), x \in U$ be any state in $\mathbf{SUP}$ and $\alpha \in \Sigma_c$ be a controllable event. We say that (1) α is *enabled* at $x \in U$ if $\xi(x, \alpha)!$; (2) α is *disabled* at $x \in U$ if $\neg\xi(x, \alpha)!$ and $(\exists s \in \Sigma^*)\xi(x_0, s) = x\ \&\ \hat{\xi}(U_0, Ps) = U\ \&\ \delta(q_0, s\alpha)!$; (3) α is *not defined* at $x \in U$ if $\neg\xi(x, \alpha)!$ and $(\forall s \in \Sigma^*)\xi(x_0, s) = x\ \&\ \hat{\xi}(U_0, Ps) = U \Rightarrow \neg\delta(q_0, s\alpha)!$.

Under partial observation, the control actions after string $s \in L(\mathbf{SUP})$ depend not on the individual state $\xi(x_0, s) \in X$, but just on the uncertainty set $U(s) \in \mathcal{U}(X)$ (i.e. the state of $\mathbf{SUPO}$). Since the language $L_m(\mathbf{SUP})$ is (relatively) observable, the following is true.

Lemma 1. *Given* $\mathbf{SUP}$ *in* (5), *let* $U \in \mathcal{U}(X)$, $x \in U$, *and* $\alpha \in \Sigma_c$. *If* α *is enabled at* $x \in U$, *then for all* $x' \in U$, *either* α *is also enabled at* $x' \in U$, *or* α *is not defined at* $x' \in U$. *On the other hand, if* α *is disabled at* $x \in U$, *then for all* $x' \in U$, *either* α *is also disabled at* $x' \in U$, *or* α *is not defined at* $x' \in U$.

For a proof of Lemma 1, see Zhang and Cai (2016b).

3.2. Localization procedure

The procedure of partial-observation localization proceeds similarly to Cai and Wonham (2010a), but is based on the set $\mathcal{U}(X)$ of the uncertainty sets and its associated transition function $\hat{\xi}$, i.e. based on the partial-observation monolithic supervisor $\mathbf{SUPO}$ in (10).

First, consider the following four functions which capture the control and marking information on the uncertainty sets. Fix a controllable event $\alpha \in \Sigma_c$. Define $E_\alpha : \mathcal{U}(X) \to \{0, 1\}$ according to

$$E_\alpha(U) = \begin{cases} 1, & \text{if } (\exists x \in U)\xi(x, \alpha)!, \\ 0, & \text{otherwise.} \end{cases}$$

Thus $E_\alpha(U) = 1$ if event α is enabled at some state $x \in U$. Then by Lemma 1 at any other state $x' \in U$, α is either enabled or not defined. Also define $D_\alpha : \mathcal{U}(X) \to \{0, 1\}$ according to

$$D_\alpha(U) = \begin{cases} 1, & \text{if } (\exists x \in U)\neg\xi(x, \alpha)!\ \&\ (\exists s \in \Sigma^*) \\ & \left[\xi(x_0, s) = x\ \&\ \hat{\xi}(U_0, Ps) = U\ \&\ \delta(q_0, s\alpha)!\right], \\ 0, & \text{otherwise.} \end{cases}$$

Hence $D_\alpha(U) = 1$ if α is disabled at some state $x \in U$. Again by Lemma 1 at any other state $x' \in U$, α is either disabled or not defined.

Next, define $M : \mathcal{U}(X) \to \{0, 1\}$ according to

$$M(U) = \begin{cases} 1, & \text{if } U \in \mathcal{U}_m, \\ 0, & \text{otherwise.} \end{cases}$$

Thus $M(U) = 1$ if U is marked in $\mathbf{SUPO}$ (i.e. U contains a marker state of $\mathbf{SUP}$). Finally define $T : \mathcal{U}(X) \to \{0, 1\}$ according to

$$T(U) = \begin{cases} 1, & \text{if } (\exists s \in \Sigma^*)\xi(x_0, s) \in U\ \& \\ & \hat{\xi}(U_0, Ps) = U\ \&\ \delta(q_0, s) \in Q_m, \\ 0, & \text{otherwise.} \end{cases}$$

So $T(U) = 1$ if U contains some state that corresponds (via a string s) to a marker state of $\mathbf{G}$.

With the above four functions capturing control and marking information of the uncertainty sets in $\mathcal{U}(X)$, we define the *control consistency relation* $\mathcal{R}_\alpha \subseteq \mathcal{U}(X) \times \mathcal{U}(X)$ as follows.

Definition 1. For $U, U' \in \mathcal{U}(X)$, we say that U and U' are *control consistent* with respect to α, written $(U, U') \in \mathcal{R}_\alpha$, if

(i) $E_\alpha(U) \cdot D_\alpha(U') = 0 = E_\alpha(U') \cdot D_\alpha(U)$,
(ii) $T(U) = T(U') \Rightarrow M(U) = M(U')$.

Thus a pair of uncertainty sets (U, U') satisfies $(U, U') \in \mathcal{R}_\alpha$ if (i) event α is enabled at at least one state of U, but is not disabled at any state of U', and vice versa; (ii) U, U' both contain marker states of **SUP** (resp. both do not contain) provided that they both contain states corresponding to some marker states of **G** (resp. both do not contain). It is easily verified that $\mathcal{R}_\alpha$ is generally not transitive, thus not an equivalence relation. This fact leads to the following definition of a *partial-observation control cover*.

Definition 2. Let I be some index set, and $\mathcal{C}_\alpha = \{\mathcal{U}_i \subseteq \mathcal{U}(X) | i \in I\}$ be a cover on $\mathcal{U}(X)$. We say that $\mathcal{C}_\alpha$ is a *partial-observation control cover* with respect to α if

(i) $(\forall i \in I, \forall U, U' \in \mathcal{U}_i)\ (U, U') \in \mathcal{R}_\alpha$,
(ii) $(\forall i \in I, \forall \sigma \in \Sigma_o)\ (\exists U \in \mathcal{U}_i)\ \hat{\xi}(U, \sigma)! \Rightarrow \big[(\exists j \in I)$
$(\forall U' \in \mathcal{U}_i)\ \hat{\xi}(U', \sigma)! \Rightarrow \hat{\xi}(U', \sigma) \in \mathcal{U}_j\big]$.

A partial-observation control cover $\mathcal{C}_\alpha$ lumps the uncertainty sets $U \in \mathcal{U}(X)$ into (possibly overlapping) *cells* $\mathcal{U}_i \in \mathcal{C}_\alpha$, $i \in I$, according to (i) the uncertainty sets U that reside in the same cell $\mathcal{U}_i$ must be pairwise control consistent, and (ii) for every observable event $\sigma \in \Sigma_o$, the uncertainty set that is reached from any uncertainty set $U' \in \mathcal{U}_i$ by a one-step transition σ must be covered by the same cell $\mathcal{U}_j$. Inductively, two uncertainty sets U and U' belong to a common cell of $\mathcal{C}_\alpha$ if and only if U and U' are control consistent, and two future uncertainty sets that can be reached respectively from U and U' by a given observable string are again control consistent.

The partial-observation control cover $\mathcal{C}_\alpha$ differs from its counterpart in Cai and Wonham (2010a) in two aspects. First, $\mathcal{C}_\alpha$ is defined on $\mathcal{U}(X)$, not on X; this is due to state uncertainty caused by partial observation. Second, in condition (ii) of $\mathcal{C}_\alpha$ only observable events in Σ_o are considered, not Σ; this is to generate partial-observation local controllers whose state transitions are triggered only by observable events. We call $\mathcal{C}_\alpha$ a *partial-observation control congruence* if $\mathcal{C}_\alpha$ happens to be a partition on $\mathcal{U}(X)$, namely its cells are pairwise disjoint.

Having defined a partial-observation control cover $\mathcal{C}_\alpha$ on $\mathcal{U}(X)$, we construct a generator $\mathbf{J}_\alpha = (I, \Sigma_o, \zeta_\alpha, i_0, I_m)$ defined over Σ_o and a control function $\psi_\alpha : I \to \{0, 1\}$ as follows. Recall from (10) that $U_0 = U(\epsilon)$ and thus $x_0 \in U_0$.

(i) $i_0 \in I$ such that $U_0 \in \mathcal{U}_{i_0}$; (11)
(ii) $I_m := \{i \in I | (\exists U \in \mathcal{U}_i) X_m \cap U \neq \emptyset\}$; (12)
(iii) $\zeta_\alpha : I \times \Sigma_o \to I$ with $\zeta_\alpha(i, \sigma) = j$
if $(\exists U \in \mathcal{U}_i)\ \hat{\xi}(U, \sigma) \in \mathcal{U}_j$; (13)
(iv) $\psi_\alpha(i) = 1$ iff $(\exists U \in \mathcal{U}_i)\ E_\alpha(U) = 1$. (14)

The control function $\psi_\alpha(i) = 1$ means that event α is enabled at state i of $\mathbf{J}_\alpha$. Note that owing to cell overlapping, the choices of i_0 and ζ_α may not be unique, and consequently $\mathbf{J}_\alpha$ may not be unique. In that case we pick an arbitrary instance of $\mathbf{J}_\alpha$.

Finally we define the *partial-observation local controller* $\mathbf{LOC}_\alpha = (Y_\alpha, \Sigma_\alpha, \eta_\alpha, y_{0,\alpha}, Y_{m,\alpha})$ as follows.

(i) $Y_\alpha = I$, $y_{0,\alpha} = i_0$, and $Y_{m,\alpha} = I_m$. Thus the control function ψ_α is $\psi_\alpha : Y_\alpha \to \{0, 1\}$.

(ii) $\Sigma_\alpha = \{\alpha\} \cup \Sigma_{com,\alpha}$, where

$$\Sigma_{com,\alpha} := \{\sigma \in \Sigma_o \setminus \{\alpha\} | (\exists i, j \in I) i \neq j,\ \zeta_\alpha(i, \sigma) = j\}. \quad (15)$$

Thus $\Sigma_{com,\alpha}$ is the set of observable events that are not merely selfloops in $\mathbf{J}_\alpha$. It holds by definition that $\{\alpha\} \subseteq \Sigma_\alpha \subseteq \Sigma_o \cup \{\alpha\}$, and $\Sigma_{com,\alpha}$ contains the events of other local controllers that need to be communicated to $\mathbf{LOC}_\alpha$.

(iii) If $\alpha \in \Sigma_o$, then $\eta_\alpha := \zeta_\alpha|_{Y_\alpha \times \Sigma_\alpha} : Y_\alpha \times \Sigma_\alpha \to Y_\alpha$, i.e. η_α is the restriction of ζ_α to $Y_\alpha \times \Sigma_\alpha$. If $\alpha \in \Sigma_{uo}$, first obtain $\eta_\alpha := \zeta_\alpha|_{Y_\alpha \times \Sigma_\alpha}$ and then add α-selfloops $\eta_\alpha(y, \alpha) = y$ to those $y \in Y_\alpha$ with $\psi_\alpha(y) = 1$.

Lemma 2. *The generator* $\mathbf{LOC}_\alpha$ *is a partial-observation local controller for* α*, i.e.* (6) *and* (7) *hold.*

For a proof of Lemma 2, see Zhang and Cai (2016b).

3.3. Main result

By the same procedure as above, we construct a set of partial-observation local controllers $\mathbf{LOC}_\alpha$, one for each controllable event $\alpha \in \Sigma_c$. We shall verify that these local controllers collectively achieve the same controlled behavior as represented by **SUP** in (5).

Theorem 1. *The set of partial-observation local controllers* $\{\mathbf{LOC}_\alpha | \alpha \in \Sigma_c\}$ *is a solution to the Partial-Observation Supervisor Localization Problem, i.e.*

$$L(\mathbf{G}) \cap L(\mathbf{LOC}) = L(\mathbf{SUP}) \quad (16)$$
$$L_m(\mathbf{G}) \cap L_m(\mathbf{LOC}) = L_m(\mathbf{SUP}) \quad (17)$$

where $L(\mathbf{LOC}) = \bigcap_{\alpha \in \Sigma_c} P_\alpha^{-1} L(\mathbf{LOC}_\alpha)$ *and* $L_m(\mathbf{LOC}) = \bigcap_{\alpha \in \Sigma_c} P_\alpha^{-1} L_m(\mathbf{LOC}_\alpha)$.

Proof. First, we prove ($\subseteq$) of (16), i.e. $L(\mathbf{G}) \cap L(\mathbf{LOC}) \subseteq L(\mathbf{SUP})$, by induction on the length of strings.

For the base case, as it was assumed that $L_m(\mathbf{SUP})$ is nonempty, it follows that the languages $L(\mathbf{G})$, $L(\mathbf{LOC})$ and $L(\mathbf{SUP})$ are all nonempty, and as they are closed, the empty string ϵ belongs to each.

For the inductive step, suppose that $s \in L(\mathbf{G}) \cap L(\mathbf{LOC})$ implies $s \in L(\mathbf{SUP})$, and $s\alpha \in L(\mathbf{G}) \cap L(\mathbf{LOC})$ for an arbitrary event $\alpha \in \Sigma$; we must show that $s\alpha \in L(\mathbf{SUP})$. If $\alpha \in \Sigma_u$, then $s\alpha \in L(\mathbf{SUP})$ because $L_m(\mathbf{SUP})$ is controllable. Otherwise, we have $\alpha \in \Sigma_c$ and there exists a partial-observation local controller $\mathbf{LOC}_\alpha$ for α. It follows from $s\alpha \in L(\mathbf{LOC})$ that $s\alpha \in P_\alpha^{-1} L(\mathbf{LOC}_\alpha)$ and $s \in P_\alpha^{-1} L(\mathbf{LOC}_\alpha)$. So $P_\alpha(s\alpha) \in L(\mathbf{LOC}_\alpha)$ and $P_\alpha(s) \in L(\mathbf{LOC}_\alpha)$, namely, $\eta_\alpha(y_{0,\alpha}, P_\alpha(s\alpha))!$ and $\eta_\alpha(y_{0,\alpha}, P_\alpha(s))!$. Let $y := \eta_\alpha(y_{0,\alpha}, P_\alpha(s))$; then $\eta_\alpha(y, \alpha)!$ (because $\alpha \in \Sigma_\alpha$). Since α may be observable or unobservable, we consider the following two cases.

Case (1) $\alpha \in \Sigma_{uo}$. It follows from the construction (iii) of $\mathbf{LOC}_\alpha$ that $\eta_\alpha(y, \alpha)!$ implies that for the state $i \in I$ of the generator $\mathbf{J}_\alpha$ corresponding to y (i.e. $i = \zeta_\alpha(i_0, P(s))$), there holds $\psi_\alpha(i) = 1$. By the definition of ψ_α in (14), there exists an uncertainty set $U \in \mathcal{U}_i$ such that $E_\alpha(U) = 1$. Let $U' = \hat{\xi}(U_0, Ps)$; by (13) and $i = \zeta_\alpha(i_0, Ps)$, $U' \in \mathcal{U}_i$. According to (9), $\xi(x_0, s) \in U'$. Since U and U' belong to the same cell $\mathcal{U}_i$, by the definition of partial-observation control cover they must be control consistent, i.e. $(U, U') \in \mathcal{R}_\alpha$. Thus $E_\alpha(U) \cdot D_\alpha(U') = 0$, which implies $D_\alpha(U') = 0$. The latter means that for all states $x \in U'$, either (i) $\xi(x, \alpha)!$ or (ii) for all $t \in \Sigma^*$ with $\xi(x_0, t) = x$ and $\hat{\xi}(x_0, Pt) = U'$, $\delta(q_0, t\alpha)$ is not defined. Note that (ii) is impossible because for $\xi(x_0, s) \in U'$, $s\alpha \in L(\mathbf{G})$. Thus by (i), $\xi(\xi(x_0, s), \alpha)!$, and therefore $s\alpha \in L(\mathbf{SUP})$.

Case (2) $\alpha \in \Sigma_o$. In this case, for the state $i \in I$ of the generator $\mathbf{J}_\alpha$ corresponding to y (i.e. $i = \zeta_\alpha(i_0, P(s))$), there holds $\zeta_\alpha(i, \alpha)!$. By the definition of ζ_α in (13), there exists an uncertainty set $U \in \mathcal{U}_i$ such that $\hat{\xi}(U, \alpha)!$, i.e. $E_\alpha(U) = 1$. The rest of the proof is identical to Case (1) above, and we conclude that $s\alpha \in L(\mathbf{SUP})$ in this case as well.

The ($\supseteq$) direction of (16), as well as Eq. (17), can be established similarly to Cai and Wonham (2010a), and we refer to Zhang and Cai (2016b) for a proof. □

Remark 1. The developed localization procedure in the preceding section may be applied without change to decompose a partial-observation supervisor with other properties (e.g. diagnosability Sampath, Lafortune, & Teneketzis, 1998; Yin & Lafortune, 2016a and opacity Dubreil, Darondeau, & Marchand, 2010). As asserted by Theorem 1, any properties enforced by **SUP** will be preserved and collectively achieved by the derived partial-observation local controllers.

Remark 2. As in Cai and Wonham (2010a,b, 2016), our proposed partial-observation supervisor localization can be applied to deal with large-scale systems, by combining it with an efficient heterarchical supervisory synthesis approach. We refer to Zhang and Cai (2016b) for the details of its application to the distributed control of a group of automated guided vehicles serving a manufacturing workcell under partial observation.

3.4. Localization algorithm

In the following, we adapt the supervisor localization algorithm in Cai and Wonham (2010a) to compute the partial-observation local controllers.

Let $\mathbf{SUP} = (X, \Sigma, \xi, x_0, X_m)$ be the controllable and observable behavior (as in (5)), with controllable Σ_c and observable Σ_o. Fix $\alpha \in \Sigma_c$. The algorithm in Cai and Wonham (2010a) would construct a control cover on X. Here instead, owing to partial observation, we first find the set $\mathcal{U}(X)$ of all uncertainty sets and label it as $\mathcal{U}(X) = \{U_0, U_1, \ldots, U_{n-1}\}$. Also we calculate the transition function $\hat{\xi} : \mathcal{U}(X) \times \Sigma_o^* \to \mathcal{U}(X)$. These steps are done by constructing the partial-observation monolithic supervisor **SUPO** as in (10) (Wonham, 2016).

Next, we apply the localization algorithm in Cai and Wonham (2010a) to construct a partial-observation control cover $\mathcal{C}_\alpha$ on $\mathcal{U}(X)$. Initially $\mathcal{C}_\alpha$ is set to be the singleton partition on $\mathcal{U}(X)$, i.e.

$$\mathcal{C}_\alpha = \{\{U_0\}, \{U_1\}, \ldots, \{U_{n-1}\}\}.$$

Write $\mathcal{U}_i$, $\mathcal{U}_j$ for two cells in $\mathcal{C}_\alpha$. Then the algorithm 'merges' $\mathcal{U}_i$, $\mathcal{U}_j$ into one cell if for every uncertainty set $U_i \in \mathcal{U}_i$ and every $U_j \in \mathcal{U}_j$, U_i and U_j, as well as their corresponding future uncertainty sets reachable by identical strings, are control consistent in terms of $\mathcal{R}_\alpha$. The algorithm loops until all uncertainty sets in $\mathcal{U}(X)$ are checked for control consistency. We call this algorithm the *partial-observation localization algorithm*.

Similar to Cai and Wonham (2010a), the algorithm terminates in a finite number of steps and results in a partial-observation control *congruence* $\mathcal{C}_\alpha$ (i.e. with pairwise disjoint cells). The complexity of the algorithm is $O(n^4)$; since the size n of $\mathcal{U}(X)$ is $n \leq 2^{|X|}$ in general, the algorithm is exponential in $|X|$.

4. Case study: Transfer Line example

In this section, we illustrate the above partial-observation localization algorithm by a Transfer Line system **TL**, as displayed in Fig. 1. **TL** consists of two machines **M1**, **M2** followed by a test unit **TU**; these agents are linked by two buffers (Buffer1, Buffer2) with capacities of three slots and one slot, respectively. We model the synchronous product of **M1**, **M2**, and **TU** as the plant to be controlled; the specification is to protect the two buffers against overflow and underflow.

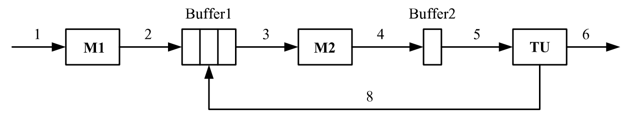

Fig. 1. Transfer Line: system configuration, with the set of controllable events $\Sigma_c = \{1, 3, 5\}$.

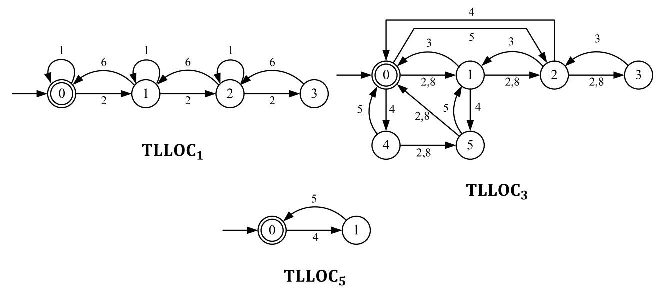

Fig. 2. Transfer Line: local controllers with full observation. Notation: a circle with → denotes the initial state, and a double circle denotes a marker state; this notation will be used in Fig. 3 as well.

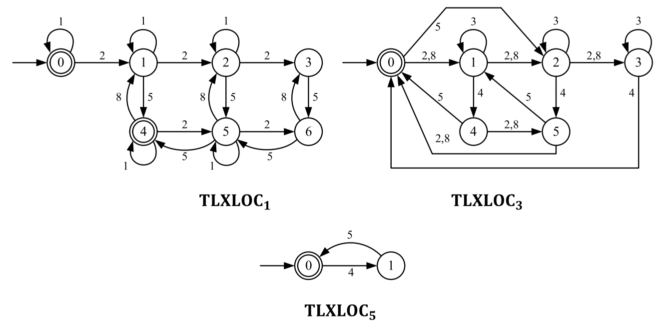

Fig. 3. Transfer Line: local controllers under partial observation with $\Sigma_o = \{1, 2, 4, 5, 8\}$.

For comparison purpose, we first present the local controllers under full observation. By Cai and Wonham (2010a), these controllers are as displayed in Fig. 2, and their control logic is as follows.

$\mathbf{TLLOC}_1$ for agent **M1** ensures that no more than three workpieces can be processed in the material-feedback loop. This is realized by counting the occurrences of event 2 (input a workpiece into the loop) and event 6 (output a workpiece from the loop).

$\mathbf{TLLOC}_3$ for agent **M2** guarantees no overflow or underflow of the two buffers. This is realized by counting events 2, 8 (input a workpiece to Buffer1), 3 (output a workpiece from Buffer1), 4 (input a workpiece to Buffer2), and 5 (output a workpiece from Buffer2).

$\mathbf{TLLOC}_5$ for agent **TU** guarantees no overflow or underflow of Buffer2. This is realized by counting event 4 (input a workpiece into Buffer2) and event 5 (output a workpiece from Buffer2).

Now consider partial observation when $\Sigma_o = \{1, 2, 4, 5, 8\}$ (i.e. events 3 and 6 are unobservable). We first compute as in (5) the controllable and observable behavior **SUP** which has 39 states. Then we apply the localization algorithm to obtain the partial-observation local controllers. The results are displayed in Fig. 3. It is verified that the collective controlled behavior of these controllers is equivalent to **SUP**.

The control logic of $\mathbf{TLXLOC}_1$ for agent **M1** is again to ensure that no more than three workpieces can be processed in the loop. But since event 6 is unobservable, the events 5 and 8 instead must be counted so as to infer the occurrences of 6: if 5 followed by 8 is observed, then 6 did not occur, but if 5 is observed and 8 is not observed, 6 may have occurred. As can be seen in Fig. 3, event 6

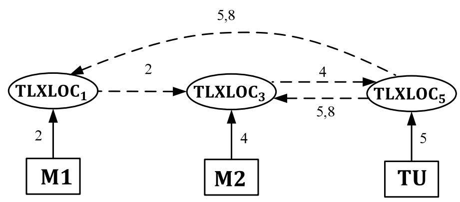

Fig. 4. Transfer Line: communication diagram of partial-observation local controllers.

being unobservable increased the structural complexity of the local controller (as compared to its counterpart in Fig. 2).

The control logic of **TLXLOC**$_3$ for agent **M2** is again to prevent overflow and underflow of the two buffers. But since event 3 is unobservable, instead the occurrences of event 4 must be observed to infer the decrease of content in Buffer1, and at the same time the increase of content in Buffer2. Also note that since the unobservable controllable event 3 is enabled at states 0, 1, 2, 3, we have selfloops of event 3 at those states. The state size of **TLXLOC**$_3$ is the same as its counterpart in Fig. 2.

TLXLOC$_5$ for agent **TU** is identical to the one in the full-observation case.

Finally, we allocate each local controller to the agent owning the corresponding controllable event; according to the transition diagrams of the local controllers, we obtain a communication diagram, as displayed in Fig. 4. A local controller either directly observes an event generated by the agent owning it, as denoted by the solid lines in Fig. 4, or imports an event by communication from other local controllers, as denoted by the dashed lines.

5. Conclusions

We have developed partial-observation supervisor localization to solve the distributed control of multi-agent DES under partial observation. This approach first employs relative observability to compute a partial-observation monolithic supervisor, and then decomposes the supervisor into a set of local controllers whose state changes are caused only by observable events. A Transfer Line example is presented for illustration. In future research we shall extend the partial-observation localization procedure to study distributed control of timed DES.

References

Cai, K., & Wonham, W. M. (2010a). Supervisor localization: a top-down approach to distributed control of discrete-event systems. *IEEE Transactions on Automatic Control*, *55*(3), 605–618.

Cai, K., & Wonham, W. M. (2010b). Supervisor localization for large discrete-event systems: case study production cell. *International Journal of Advanced Manufacturing Technology*, *50*(9–12), 1189–1202.

Cai, K., & Wonham, W. M. (2015). New results on supervisor localization, with case studies. *Discrete Event Dynamic Systems*, *25*(1–2), 203–226.

Cai, K., & Wonham, W. M. (2016). *Supervisor Localization: A Top-Down Approach to Distributed Control of Discrete-Event Systems. Lecture Notes in Control and Information Sciences, vol. 459*. Springer.

Cai, K., Zhang, R., & Wonham, W. M. (2015). Relative observability of discrete-event systems and its supremal sublanguages. *IEEE Transactions on Automatic Control*, *60*(3), 659–670.

Cai, K., Zhang, R., & Wonham, W. M. (2016). Relative observability and coobservability of timed discrete-event systems. *IEEE Transactions on Automatic Control*, *61*(11), 3382–3395.

Cieslak, R., Desclaux, C., Fawaz, A. S., & Varaiya, P. (1988). Supervisory control of discrete-event processes with partial observations. *IEEE Transactions on Automatic Control*, *33*(3), 249–260.

Dubreil, J., Darondeau, P., & Marchand, H. (2010). Supervisor control for opacity. *IEEE Transactions on Automatic Control*, *55*(5), 1089–1100.

Lin, F., & Wonham, W. M. (1988). On observability of discrete-event systems. *Information Sciences*, *44*(3), 173–198.

Sampath, M., Lafortune, S., & Teneketzis, D. (1998). Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, *43*(7), 908–929.

Takai, S., & Ushio, T. (2003). Effective computation of Lm(G)-closed, controllable, and observable sublanguage arising in supervisory control. *Systems & Control Letters*, *49*(3), 191–200.

Wonham, W.M. (2016). Supervisory control of discrete-event systems. Systems control group, ECE Dept, Univ. Toronto, Toronto, ON, Canada, July. Available at http://www.control.utoronto.ca/DES.

Yin, X., & Lafortune, S. (2016a). A uniform approach for synthesizing property-enforcing supervisors for partially-observed discrete-event systems. *IEEE Transactions on Automatic Control*, *61*(8), 2140–2154.

Yin, X., & Lafortune, S. (2016b). Synthesis of maximally permissive supervisors for partially-observed discrete-event systems. *IEEE Transactions on Automatic Control*, *61*(5), 1239–1254.

Zhang, R., & Cai, K. (2016a). On supervisor localization based distributed control of discrete-event systems under partial observation. In *Proc. 2016 American control conference*, Boston, (pp. 764–769).

Zhang, R., & Cai, K. (2016b). *Supervisor localization of discrete-event systems under partial observation. Technical Report*. Available at http://arxiv.org/abs/1509.05498.

Zhang, R., Cai, K., Gan, Y., Wang, Z., & Wonham, W. M. (2013). Supervision localization of timed discrete-event systems. *Automatica*, *49*(9), 2786–2794.

Renyuan Zhang received the B.Eng. degree and the Ph.D. degree in Electrical Engineering from Xi'an Jiaotong University, Xi'an, China, in 2007 and 2013 respectively. He is currently a lecturer at School of Automation, Northwestern Polytechnical University, Xi'an, China. He was an international visiting student in the University of Toronto 2011–2012. His research interests include distributed control of discrete-event systems and hybrid control of urban traffic systems.

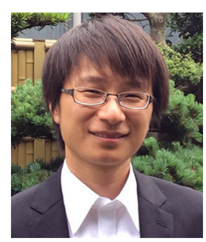

Kai Cai received the B.Eng. degree in Electrical Engineering from Zhejiang University, Hangzhou, China, in 2006; the M.A.Sc. degree in Electrical and Computer Engineering from the University of Toronto, Toronto, ON, Canada, in 2008; and the Ph.D. degree in Systems Science from the Tokyo Institute of Technology, Tokyo, Japan, in 2011. He is currently an Associate Professor at Osaka City University. Previously, he was an Assistant Professor at the University of Tokyo (2013–2014), and a postdoctoral Fellow at the University of Toronto (2011–2013).

Dr. Cai's research interests include distributed control of discrete-event systems and cooperative control of networked multi-agent systems. He is the co-author (with W.M. Wonham) of Supervisor Localization (Springer 2016). He received the Best Paper Award of SCIE in 2013, the Best Student Paper Award of the IEEE Multi-Conference on Systems and Control and the Young Author's Award of SCIE in 2010.

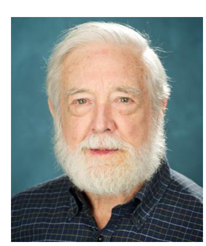

W. Murray Wonham received the B.Eng. degree in engineering physics from McGill University in 1956, and the Ph.D. in control engineering from the University of Cambridge (UK) in 1961.

From 1961 to 1969, he was associated with several US research groups in control. Since 1970, he has been a faculty member in Systems Control, with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada. His research interests have included stochastic control and filtering, geometric multivariable control, and discrete-event systems.

He is the author of Linear Multivariable Control: A Geometric Approach (3rd edition) (New York: Springer-Verlag, 1985) and co-author (with C. Ma) of Nonblocking Supervisory Control of State Tree Structures (New York: Springer-Verlag, 2005), and co-author (with K. Cai) of Supervisor Localization (New York: Springer-Verlag, 2016).

Wonham is a Fellow of the Royal Society of Canada and a Foreign Member of the (US) National Academy of Engineering. In 1987, he received the IEEE Control Systems Science and Engineering Award and in 1990 was Brouwer Medalist of the Netherlands Mathematical Society. In 1996 he was appointed University Professor in the University of Toronto, and in 2000 University Professor Emeritus.