Distributed knowledge for communication in decentralized discrete-event systems*

S.L. Ricker
CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

K. Rudie
Department of Electrical
and Computer Engineering
Queen's University
Kingston, Ontario K7L 3N6
Canada

E-Mail: S.L.Ricker@cwi.nl, rudie@ee.queensu.ca

Abstract

An extension to a formal model for reasoning about knowledge and communication in discrete-event systems is presented. The model is based on a modal logic where propositions describe the knowledge that agents in the system possess. Decentralized discrete-event control problems, where agents must communicate in order to effect control, are examined. Previously the identification of a state where agents should communicate was expressed solely in terms of properties of a formal language. The updated knowledge model presented here yields equivalent propositional logic expressions so that agents can determine where to communicate based on whether certain propositions are true at certain states. This amounts to agents making control decisions based on their "knowledge" of the system. The key relevant logic proposition corresponds to a check on whether a group of agents, if they pooled their knowledge, would possess sufficient information about the legality of a given event sequence. The new formulation is applied to a previously-solved example.

1. Introduction

We have proposed a knowledge model for analyzing decentralized discrete-event control problems [5, 6]. Specifically, we are interested in studying control problems where decentralized supervisors must communicate with other supervisors to achieve the desired system behavior. Other strategies for incorporating communication into a decentralized discrete-event systems (DES) framework have been proposed [1, 7, 9, 10]. Previously we used the structure of sequences in the plant language to identify states in the plant where communication could occur and eventually allow supervisors to make the correct control decision. In this paper we further examine our knowledge model (which we summarize in section 2) and update a previous hypothesis [5] concerning the nature of knowledge that the group of supervisors requires to solve such a control problem.

In [5] we speculated about the knowledge of each supervisor after pooling its information. This notion, called *distributed observability*, expressed whether a supervisor would know an event should be disabled if it had access to all the information of the other agents. We described an example where the pooled information would not be sufficient for a supervisor to make the correct control decision.

In this earlier work, we focussed on the knowledge a supervisor had about the impending control decision. Here we propose that instead, we consider whether agents know whether the system is progressing along a legal path or along an illegal path. We illustrate the extended knowledge model using an example from [6].

2. Background

The framework for decentralized discrete-event control that we assume here is taken from [4]. A discrete-event system is considered to be the generator of a formal language. A change in the system state is precipitated by the occurrence of events that describe the behavior of the system. The essence of the control problem is to construct controllers or *supervisors* that recognize some specified subset of sequences of the original system by either disabling or enabling events.

The plant is modeled by an automaton $G = (Q^G, \Sigma, \delta^G, q_0^G)$. In addition, the subset of specified or *legal* behavior of the plant is also described as an automaton $E = (Q^E, \Sigma, \delta^E, q_0^E)$. The language $L(G)$ is the set of sequences generated by G and $L(E)$ is a

*This work is partially supported by an ERCIM fellowship for the first author and NSERC Grant OGP0138887 for the second author.

subset of $L(G)$ that describes desired behavior. We assume that E is a subautomaton of G. We can also define a subset of $\Sigma_c \subseteq \Sigma$ that denotes *controllable* events. The set of controllable events for a decentralized supervisor i is denoted $\Sigma_{i,c}$ (where we assume $i \in \{1, \ldots, n\}$ if there are n supervisors).

When not all events in the system can be seen by a supervisor, we use a *canonical projection* P to describe a supervisor's view of the sequences in $L(G)$. We assume that a supervisor sees some set of events $\Sigma_o \subseteq \Sigma$ and the projection operator is a mapping from Σ^* to Σ_o^*. That is, the events $\sigma \in (\Sigma \setminus \Sigma_o)$ that are not visible to a supervisor "disappear" from its view of a sequence $s \in L(G)$. A system is said to be *observable* [3] if a supervisor that could see all the observable events makes the same control decision for all sequences that look alike via the projection operator. The view of the plant for a decentralized supervisor i is described by the canonical projection $P_i : \Sigma^* \rightarrow \Sigma_{i,o}^*$, where $\Sigma_{i,o}$ is the set of events observable to supervisor i.

When a supervisor has a partial view of a sequence, we want to be able to describe the states in the plant the supervisor considers possible. We repeat our notion from [6] of a supervisor's *local view* of a plant state.

DEFINITION **1** *The* **local view** ℓ_i *of a state* $\ell \in Q^G$ *reached via sequence* t *(i.e.,*$\exists t \in \Sigma^*$ *where* $\delta^G(t, q_0^G) = \ell$*) is the set of all the states in the plant that supervisor* i *considers the plant could be in upon seeing* $P_i(t)$*:* $\ell_i := \{q^G \mid q^G \in Q^G \wedge \exists u \in P_i^{-1}(P_i(t))$ *such that* $\delta^G(u, q_0^G) = q^G\}$.

Thus if supervisor i cannot determine if t or t' has occurred in the plant (i.e., $P_i(t) = P_i(t')$) and if $\delta^G(t, q_0^G) = q$ while $\delta^G(t', q_0^G) = q'$, the local view of supervisor i at state q will contain q and q'.

We will also want to construct a "monitoring" automaton to capture the overall view of the plant and each supervisor's view of the plant: $A = (Q^A, \Sigma, \delta^A, q_0^A)$, where a state $q^A \in Q^A$ is an (n+1)-tuple of the form $(q^G, (q^G)_1, \ldots, (q^G)_n)$ and $(q^G)_i$ is supervisor i's local view of state q^G; and the initial state is $q_0^A = (q_0^G, (q_0^G)_1, \ldots, (q_0^G)_n)$. Note that $\delta^A(\sigma, q^A) = q'^A$ if there is a transition of σ in G from the first entry in q^A to the first entry in q'^A. If $\sigma \in \Sigma_{i,o}$, the $(i+1)^{th}$ entry of q'^A is supervisor i's local view of the first entry in q'^A, otherwise the $(i+1)^{th}$ entry of q'^A is $(q^A)_i$. If there are two different sequences that lead to the same plant state (and if there is at least one supervisor that can distinguish them), there will be two states in A that have the same first entry. By the way in which A is constructed, we have $L(A) = L(G)$.

We introduced Theorem 1 in [6] to describe the conditions under which we could identify a place where one supervisor communicates to another so that the latter supervisor can make the correct control decision. We assumed that supervisor i could not make the correct control decision about one of its controllable events σ because there are two sequences t and t' in $L(G)$ that supervisor i could not tell apart (i.e., $P_i(t) = P_i(t')$) and $t'\sigma$ is in $L(E)$ but $t\sigma$ is not. The theorem states that as long as E is observable with respect to G, P, we can find a state q where supervisor j sends its local view q_j to supervisor i so that supervisor i will be able to distinguish t from t' and thus make the correct control decision for σ.

We developed our knowledge model for decentralized discrete-event control by adapting a model for distributed systems[2] where multiple agents reason about their knowledge of the world. The model assumes that when agents in the distributed system do not have complete knowledge of the world, they consider that several worlds are possible. In general, the set of agents is denoted by G and without loss of generality we assume here that $\mathsf{G} = \{1, 2\}$. When recasting decentralized DES into a knowledge model, an agent is considered equivalent to a supervisor. The set of agents that control event σ is denoted $\mathsf{G}_\sigma := \{i \mid \sigma \in \Sigma_{i,c}\}$. A world $w = (w_e, w_1, w_2)$ is a triple that describes the plant state (denoted w_e) and each agent's set of possible worlds (called a *local state* and denoted w_i for agent i) that it cannot distinguish from the actual plant state. Note that the local state of an agent is defined in the knowledge model as a supervisor's local view of a plant state.

A run r is a (possibly) infinite sequence of worlds. A *point* of the system consists of a run r and an index m (where m ranges over the natural numbers), denoted (r, m). At a point (r, m) the system is in some world w—the m^{th} entry along run r. We will alternately denote the world at a point (r, m) as $r(m) = (r_e(m), r_1(m), r_2(m))$. When a run can be unambiguously identified we will use a state name to identify the run. That is, if there is only one run r that reaches world w we use w to identify the run.

We can express our knowledge model as a labeled graph or *Kripke structure*. In particular, nodes are worlds and worlds that look alike to agent i are joined by an edge with a label "i". Figure 1(i) shows a plant where $\mathsf{G} = \{1, 2\}$ and agent 1 sees event a_1 and agent 2 sees event b_2. The corresponding Kripke structure for this plant is shown in figure 1(ii). The Kripke structure of this system contains five possible worlds. For example, there are three worlds in figure 1(ii)

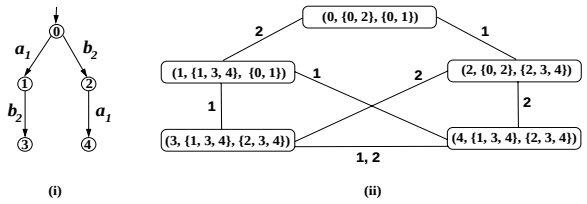

Figure 1: (i) A plant G; (ii) Kripke structure for G.

where agent 2 has a local state of $\{2,3,4\}$. This means that agent 2 cannot distinguish plant states 2, 3 or 4 (for the plant shown in figure 1(i)). This is indicated in the Kripke structure by the edge label of "2" joining of all worlds w where $w_2 = \{2,3,4\}$. Although omitted from this diagram, each world should also have a self-loop with an edge label of "1,2", since a world cannot be distinguished from itself.

We want our agents to reason about their view of the world. Facts about the world are described in terms of *propositions*. The facts in our knowledge model are the events in Σ. Each event is translated into two propositions: one corresponding to whether or not the event is defined at a particular state of the plant, and the other indicating whether the event is defined in the legal automaton. The following definition summarizes some concepts from [5].

DEFINITION **2** *(i) The proposition σ_G is "event σ can occur" and σ_E is "event σ is legal". (ii) A proposition σ_G is* **true** *at point (r,m) if the event σ happens at the plant state described by $r_e(m)$. A proposition σ_G is* **false** *(denoted $\neg\sigma_G$) at point (r,m) if the event σ is not defined at the plant state described by $r_e(m)$. (iii) A proposition σ_E is* **true** *at point (r,m) if the event σ happens at the plant state $r_e(m)$ and is part of the legal behavior of the plant. A proposition σ_E is* **false** *(denoted $\neg\sigma_E$) at point (r,m) if either $\sigma_G =$* **false** *or if the event σ is part of the illegal behavior of the plant at $r_e(m)$.*

We construct an interpreted system [2]—which we denote $\mathcal{I}(G,E)$ —that allows us to speak about compound statements in our knowledge model[1]. Where G and E are understood, we will drop these arguments. The set of worlds for $\mathcal{I}$ is simply the set of states Q^A from the monitoring automaton A. Two points (r,m) and (r',m') are indistinguishable to agent i if they have the same local state. Formally this is denoted $(r,m) \sim_i (r',m')$ and means that $r_i(m) = r'_i(m')$. In addition to standard propositional logic, the model includes the modal operator K_i, $i = 1,2$, where $K_i\phi$ means "agent i knows ϕ", and the modal operator D_{G}, where $D_{\mathsf{G}}\phi$ means "among the set of agents G there is distributed knowledge of ϕ". An agent i has knowledge of fact ϕ at point (r,m) if that fact has the same truth value at all points that agent i cannot distinguish from (r,m). There is distributed knowledge of fact ϕ at point (r,m) iff ϕ is true at all points (r',m') such that for each $i \in \mathsf{G}$ it is the case that $r_i(m) = r'_i(m')$.

Using this knowledge model, a condition called Kripke-observability (defined in section 3 and equivalent to co-observability[8]) is necessary and sufficient for solving—without communication—the decentralized discrete-event control problem.

3. Extending the knowledge model

We want to use the result of Theorem 1 from [6] to lead us to an expression of the knowledge that exists in the system when communication between supervisors allows them to make the correct control decisions.

We first translate the notion of *observability* [3] in terms of our knowledge model.

DEFINITION **3** *An interpreted system $\mathcal{I}$ is* **observable** *if for all $(r,m),(r',m') \in \mathcal{I}$, for all $i \in \mathsf{G}$, if $(r,m) \sim_i (r',m')$ there do not exist propositions σ_G, σ_E such that $(\mathcal{I},r,m) \models \sigma_G \wedge \sigma_E$ and $(\mathcal{I},r',m') \models \sigma_G \wedge \neg\sigma_E$.*

To emulate the knowledge of an omniscient centralized agent (i.e., one that sees everything that each agent sees) we consider only those possible worlds that *all* agents find indistinguishable. That is, a system is observable if one agent could see everything that the other sees, there should not be two possible worlds that the agent finds indistinguishable where different control decisions for event σ are required.

Definition 3 provides a condition necessary for centralized control. In [5] we discussed a condition for decentralized control called *Kripe-observability*. If the system is not Kripke-observable then there is some point (r',m') where an illegal event σ is about to occur but none of the agents capable of controlling σ knows that σ should be disabled. The lack of knowledge arises from the existence of another point (r,m) where σ is also about to occur but is a legal event and this same set of agents cannot distinguish (r,m) from (r',m'). More formally, there exists $(r',m') \in \mathcal{I}$ and a pair of primitive propositions (σ_G,σ_E) where $(\mathcal{I},r',m') \not\models \neg\sigma_G \vee \sigma_E$, and for all $i \in \mathsf{G}_\sigma$, $(\mathcal{I},r',m') \not\models K_i\neg\sigma_E$. That is, $(\mathcal{I},r',m') \models \sigma_G \wedge \neg\sigma_E$ and for all $i \in \mathsf{G}_\sigma$ there exists $(r,m) \in \mathcal{I}$ such that $(r',m') \sim_i (r,m)$ and

[1] In [6] we used the notation $\mathcal{I}^{DES'}$ because we were considering the construction of several different interpreted systems.

$(\mathcal{I}, r, m) \models \sigma_G \wedge \sigma_E$. Note that it is not possible for $(r, m) = (r', m')$ because G is assumed to be a deterministic automaton. That is, if $(r, m) = (r', m')$ then one transition involving event σ at plant state $r_e(m) = r'_e(m')$ would be part of the legal behavior while a second transition of σ would not be part of the legal behavior.

We want to identify those points where Kripke-observability does not hold. That is, we seek a point where an agent does not have sufficient knowledge to make the correct control decision and, in addition, there is an event that must be disabled.

DEFINITION **4** *A point* (r', m') *is* $\boldsymbol{K_i}$**-bad** *with respect to event* σ *for agent* $i \in \mathsf{G}_\sigma$ *if* $(\mathcal{I}, r', m') \not\models K_i \neg\sigma_E$ *and* $(\mathcal{I}, r', m') \models \sigma_G \wedge \neg\sigma_E$.

In addition, we want to identify those points that are indistinguishable from a K_i-bad point for agent i where event σ must be allowed to happen.

DEFINITION **5** *A point* (r, m) *is* $\boldsymbol{K_i}$**-good** *with respect to event* σ *and point* (r', m') *for agent* $i \in \mathsf{G}_\sigma$ *if* (r', m') *is a* K_i*-bad point with respect to* σ *for agent* $i \in \mathsf{G}_\sigma$, $(r', m') \sim_i (r, m)$ *and* $(\mathcal{I}, r, m) \models \sigma_G \wedge \sigma_E$.

Note that there could be more than one K_i-good point associated with a K_i-bad point (r', m').

We want to define new propositions $\text{good}(r, m)$ and $\text{bad}(r', m')$ with respect to all the K_i-bad points (r', m') and their associated K_i-good points. Using the Kripke structure for $\mathcal{I}$ (constructed from the states of A) we can identify runs r and r' that contain (r, m) and (r', m'). A point (r, m'') that occurs along run r for $m'' < m$ is deemed a "good" point with respect to (r, m). A point (r', m'') that occurs along run r' for $m'' < m$ is deemed a "bad" point with respect to (r', m').

DEFINITION **6** (*i*) *The proposition* $\text{good}(r, m)$ *is "this is a good point with respect to point* (r, m)*" and* $\text{bad}(r', m')$ *is "this is a bad point with respect to point* (r', m')*". (ii) A proposition* $\text{bad}(r', m')$ *is* **true** *at point* (r'', m'') *if* $r'' = r'$, $m'' \leq m'$, *and* (r', m') *is a* K_i*-bad point; otherwise* $\text{bad}(r', m')$ *is* **false**. (*iii*) *A proposition* $\text{good}(r, m)$ *is* **true** *at point* (r'', m'') *if* $r'' = r$, $m'' \leq m$ *and* (r, m) *is a* K_i*-good point associated with a* K_i*-bad point* (r', m')*; otherwise* $\text{bad}(r', m')$ *is* **false**.

In the formulation of Theorem 1 in [6], we assumed that at least one of the agents could not distinguish legal sequences t and t', where one lead to an illegal sequence $t'\sigma$ and the other lead to a legal sequence $t\sigma$. We sought a state q along either t or t' where the intersection of the agents' local views of q (with respect to sequences t and t') yielded a set of states that did not contain both a state that lay on the path of the legal sequence $t\sigma$ and a state that lay on the path of the illegal sequence $t'\sigma$. When we translate this in terms of the new propositions in definition 6, it seems as if we want a knowledge condition that indicates that there is distributed knowledge that the system is either along a bad path to point (r', m') (i.e., $D_\mathsf{G}\ \text{bad}(r', m')$) or that the system is along a good path to point (r, m) (i.e., $D_\mathsf{G}\ \text{good}(r, m)$). The problem with this approach is that it does not take into account the fact that there could be other paths—neither "good" nor "bad"—that an agent cannot distinguish from the others. Thus we rephrase the knowledge requirement so that there is distributed knowledge that the system is *not* along a good path or *not* along a bad path.

Previously (in [5]) we speculated that the distributed knowledge required to determine places to share or pool information involved the propositions from definition 2. We update our definition here:

DEFINITION **7** *A system has* **distributed observability** *if for all* $\sigma \in \Sigma$, *for all points* (r', m') *in* $\mathcal{I}$ *that are* K_i*-bad with respect to* σ *and their associated* K_i*-good points* (r, m), *there exists a point* (r'', m'') *in* $\mathcal{I}$ *where either* $r'' = r'$ *and* $m'' \leq m'$ *or* $r'' = r$ *and* $m'' \leq m$ *such that*

$$(\mathcal{I}, r'', m'') \models D_\mathsf{G}(\neg \mathit{good}(r, m)) \vee D_\mathsf{G}(\neg \mathit{bad}(r', m')), \quad (1)$$

and therefore either $(\mathcal{I}, r', m'') \models D_\mathsf{G}(\neg \mathit{good}(r, m))$ *where* $m'' \leq m'$, *or* $(\mathcal{I}, r, m'') \models D_\mathsf{G}(\neg \mathit{bad}(r', m'))$ *where* $m'' \leq m$.

That is, there is distributed knowledge of whether the system is not along a good path or not along a bad path. This means that somewhere before a control decision regarding event σ must be made by some agent $i \in \mathsf{G}_\sigma$, there is a point where, if all the agents pooled their knowledge, agent i would no longer confuse the path to (r', m') with the path to (r, m).

The restatement of Theorem 1 in [6] within the knowledge framework is as follows:

THEOREM **1** *Suppose that* $\mathcal{I}$ *is observable and not Kripke-observable. Then* $\mathcal{I}$ *has distributed observability.*

Proof: Let (r', m') be a K_i-bad point with respect to σ for agent $i \in \mathsf{G}_\sigma$. Let (r, m) be a K_i-good point with respect to σ and (r', m'). From definition 4 we have

$$(\mathcal{I}, r', m') \models \sigma_G \wedge \neg\sigma_E. \quad (2)$$

Similarly, from definition 5 we have

$$(\mathcal{I}, r, m) \models \sigma_G \wedge \sigma_E. \quad (3)$$

If $(r', m') \sim_i (r, m)$ then, by definition, (r', m') and (r, m) have the same local state (i.e., local view) according to agent i. This means that in A we can find a path t' along run r' that leads to $r'(m')$ and a path t along run r that leads to $r(m)$ such that agent i cannot distinguish between these paths, i.e., $P_i(t) = P_i(t')$.

Since G generates the same language as A, we also know that $\delta^G(t, q_0^G)$ is defined and that $\delta^G(t', q_0^G)$ is defined. In particular, $\delta^G(t, q_0^G) = r_e(m)$ and $\delta^G(t', q_0^G) = r'_e(m')$.

By (2) and from above we have $t'\sigma \notin L(E)$. Similarly, by (3) and above we have $t\sigma \in L(E)$.

We can now apply Theorem 1 from [6] and find a state $\ell \in Q^G$ along either t or t' such that the intersection of ℓ_i and ℓ_j (for $i, j \in \{1, 2\}, i \neq j$) does not contain both a state along t and a state along t'. Therefore we can also find $w \in Q^A$ such that $w = (\ell, \ell_1, \ell_2)$ where w lies along run r or run r'.

Let $w = r(m'')$ where $m'' \leq m$, i.e., w lies along run r before (r, m). From definition 6, we have

$$(\mathcal{I}, r, m'') \models \text{good}(r, m). \quad (4)$$

Claim 1: $(\mathcal{I}, r, m'') \models D_{\mathsf{G}} \neg \text{bad}(r', m')$

Suppose that $(\mathcal{I}, r, m'') \not\models D_{\mathsf{G}} \neg \text{bad}(r', m')$. Then by (4) there exists $(r''', m''') \in \mathcal{I}$, $(r, m'') \neq (r''', m''')$ such that $(r, m'') \sim_i (r''', m''')$ for $i = 1, 2$ and $(\mathcal{I}, r''', m''') \models \text{bad}(r', m')$.

Therefore, by definition 6, $r''' = r'$ and $m''' \leq m'$.

Let the world at (r', m''') be denoted by $w' = (\ell', \ell'_1, \ell'_2)$. Since $(r', m''') \sim_i (r, m'')$ for $i = 1, 2$ then

$$\begin{aligned} \ell_1 &= \ell'_1 \\ \ell_2 &= \ell'_2. \end{aligned} \quad (5)$$

Since ℓ' is along run r' that leads to $r'(m')$ and $m''' \leq m'$, ℓ' lies along t'. Further, since $\ell \in \ell_1 \cap \ell_2$ (because a local view ℓ_i of a state ℓ always contains ℓ), this means that $\ell_1 \cap \ell_2$ contains a state along t (since $m'' \leq m$ and t is a path on run r that leads to $r(m)$) and that $\ell'_1 \cap \ell'_2$ contains ℓ', a state along t'.

From (5), $\ell_1 \cap \ell_2 = \ell'_1 \cap \ell'_2$ and therefore $\ell_1 \cap \ell_2$ contains ℓ', a state along t'. This leads to a contradiction. □ *Claim* 1

We can use a similar argument to show that if w lies along a path to $r'(m')$, then $(\mathcal{I}, r', m'') \models D_{\mathsf{G}} \neg \text{good}(r, m)$. □ THEOREM 1

One of the difficulties we encountered in [6] regarding the identification of a communication state is illustrated in the left-hand side of figure 2. Suppose that legal sequence $t\sigma$ passes through states q, u and x while illegal sequence $t'\sigma$ passes through q', u and x'. Let agent i be the only agent that controls event σ and further suppose that agent i cannot distinguish t from t' (i.e., $P_i(t) = P_i(t')$). Suppose that the result of applying Theorem 1 from [6] to this example is that the communication state is q. It might be, though, that although agent i would be able to distinguish the paths of t from t' at q, this information would be lost at state u, where the two paths intersect again. That is, the knowledge that an agent had acquired earlier does not persist as the system evolves.

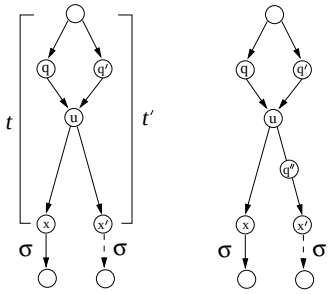

Figure 2: Is the system is along legal sequence $t'\sigma$ or illegal sequence $t\sigma$?

Our previous solution to the scenario in the left-hand side of figure 2 was to "split" state u. Thus we assumed that the only state common to t and t' was the initial state. Yet it might be possible to find a state such as q'' in the right-hand side of figure 2 where $q''_i \cap q''_j$ does not contain both a state along t and along t'. Subsequently, agent i retains the knowledge it acquired at q''_i and is able to make the correct control decision about σ.

We hypothesize that one way to avoid splitting states is to find a state where once an agent knows that the system is not along a good path (equivalently, not along a bad path) this knowledge persists until the correct control decision must be made. Thus we seek a state w along t (respectively, along t') satisfying the hypothesis of Theorem 1 such that (a) distributed observability holds at w in $\mathcal{I}$ and (b) once information pooling occurs at w and we update the states of $\mathcal{I}$ accordingly to reflect the effect of pooling at w, distributed knowledge as described in (1) holds at every successor state of the revised w along t (respectively, states along t') in the updated $\mathcal{I}$.

In the next section, instead of explicitly referencing a point by its run r and an index m, we will alternately use the notation for worlds w to refer to the point. We extend this abbreviation to the predi-

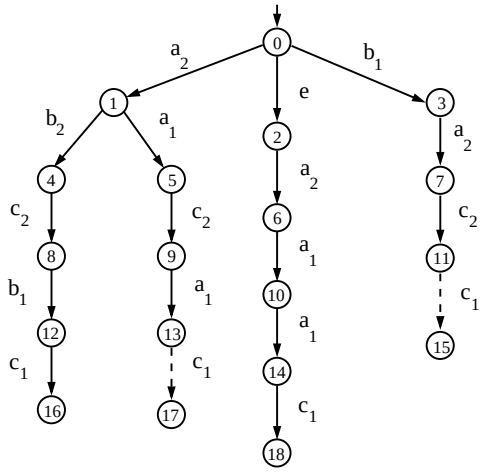

Figure 3: Combined plant G and legal automaton E.

cates "good" and "bad" from definition 6, and refer to good(w) and bad(w). In addition we refer to K_i-good (-bad) worlds.

4. Example and Discussion

Figure 3 shows a plant and legal automaton where supervisor 1 can see and control events a_1, b_1 and c_1 while supervisor 2 sees and controls events a_2, b_2 and c_2. Because it sees only part of the defined plant behavior, supervisor 1 cannot make the correct control decision to disable c_1 at state 13 because it cannot distinguish state 14 (where c_1 *should* be allowed to occur) from state 13. A similar situation occurs at states 11 and 12.

For the example shown in figure 3, we can construct the Kripke structure and determine that Kripke-observability does not hold at two worlds: $v' = (11, \{3, 7, 11, 12\}, \{9, 11, 13, 15, 17\})$ and $\widehat{v'} = (13, \{13, 14\}, \{9, 11, 13, 15, 17\})$ for $\mathsf{G}_\sigma = \{1\}$. At both v' and $\widehat{v'}$ it is the case that $c_{1_G} \wedge \neg c_{1_E}$ holds. Agent 1 cannot distinguish v' and $\widehat{v'}$ from $v = (12, \{3, 7, 11, 12\}, \{8, 12, 16\})$ and $\hat{v} = (14, \{13, 14\}, \{1, 5, 6, 7, 10, 14, 18\})$, respectively. At both worlds v and $\hat{v}$ it is the case that $c_{1_G} \wedge c_{1_E}$ holds. Thus at v' and $\widehat{v'}$ agent 1 does not have enough knowledge to make the correct control decision about event c_1.

According to the result of Theorem 1 in [6], we would choose world $w = (1, \{0, 1, 2, 4, 6, 8\}, \{1, 5, 6, 7, 10, 14, 18\})$ as the communication state for v and v'. Similarly, $\hat{w} = (9, \{5, 9, 10\}, \{9, 11, 13, 15, 17\})$ would be the communication state for $\hat{v}$ and $\widehat{v'}$. Is there distributed observability in this system? We identify two K_1-bad worlds for agent 1—both with respect to event c_1—v' and $\widehat{v'}$. The associated K_1-good worlds are, respectively, v and $\hat{v}$. The system $\mathcal{I}$ for the example in figure 3 has distributed observability if we can find the following: (i) a world w that occurs either along a path to v or v' such that there is distributed knowledge of either $\neg$good(v) or $\neg$bad(v'); and (ii) a world $\hat{w}$ that occurs either along a path to $\hat{v}$ or $\widehat{v'}$ such that there is distributed knowledge of either $\neg$good($\hat{v}$) or $\neg$bad($\widehat{v'}$). We will describe only the effect of communication at w for K_1-bad world $v' = (11, \{3, 7, 11, 12\}, \{9, 11, 13, 15, 17\})$ and its associated K_1-good world $v = (12, \{3, 7, 11, 12\}, \{8, 12, 16\})$.

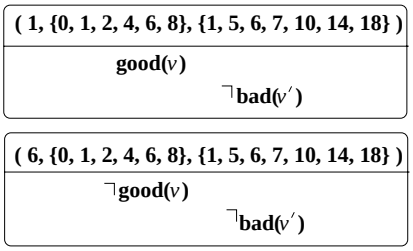

Figure 4: Checking distributed knowledge of $\neg$bad(v').

The only world that is indistinguishable from w according to both agents is $(6, \{0, 1, 2, 4, 6, 8\}, \{1, 5, 6, 7, 10, 14, 18\})$. Figure 4 shows these two worlds and the truth assignments to the propositions good(v) and bad(v'). Since $\neg$bad(v') holds at both these worlds, we can say that there is distributed knowledge of $\neg$bad(v') at w. That is, if agent 1 received a communication of agent 2's local state then agent 1 would know that the system is not along a bad path.

When agent 2 is at its local state $w_2 = \{1, 5, 6, 7, 10, 14, 18\}$ it sends this local state to agent 1. If agent 1 was at its local state of $\{0, 1, 2, 4, 6, 8\}$ and received a communication of w_2 from agent 2, agent 1 would update its local state to $\{1, 6\}$, the intersection of the two local states. When agent 1 updates its local state, this affects the view it has of the rest of the system. We consider the following procedure for describing the propagation of information from agent j to agent i along the path from a communication state w to the K_i-good world v (equivalently, to a K_i-bad world v'). In particular, note that w_j does not change.

Procedure 1

1. w_j is sent to agent i; at w, w_i is updated to $w_i \cap w_j$. We will denote the updated local view as w_i^{com}.

2. For each world w'' such that w'' lies between w and v (and $w'' \neq w$) along legal sequence $t\sigma$:

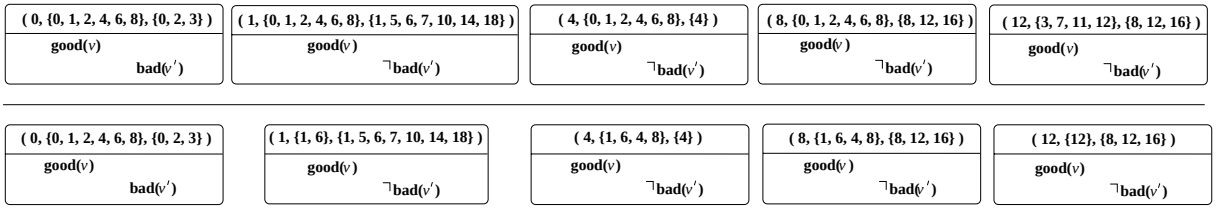

Figure 5: Worlds that occur along the path to $v = (11, \{3, 7, 11, 12\}, \{9, 11, 13, 15, 17\})$ before communication (above line) and after communication (below line).

(a) if $w_i'' = w_i, w_i'' \longleftarrow w_i^{com} \cup \{q \mid \exists\, u \in (\Sigma \setminus \Sigma_{i,o})^*, \exists\, y \in w_i^{com}$ and $\delta^G(u, y) = q\}$;

(b) else $w_i'' \longleftarrow w_i'' \setminus \{x \in Q^G \mid$ for all $y \in w_i^{com}$, $\nexists u \in \Sigma^*$ where $\delta^G(u, y) = x\}$

The top of figure 5 contains the possible worlds that occur along legal sequence $t\sigma$. The bottom of the same figure shows the effect that receiving communication from agent 2 has on the local states of agent 1. At the point at which the agents in this example pool information, agent 1 knows that the system is either at state 1 or state 6. This corresponds to agent 2 having seen a_2. If agent 2 had seen $a_2 b_2$ and communicated that information to agent 1, then agent 1 would have known the system was at state 4. It is possible that after receiving communication from agent 2 subsequent system behavior consists of events that are unobservable to agent 1. So even though there is a point in the system evolution where agent 1 is "certain" that the system is either at state 1 or state 6, our definition of a local view still takes into account the occurrence of unobservable events. Thus when using Procedure 1 with $w_1^{com} = \{1, 6\}$ and worlds (4, $\{0, 1, 2, 4, 6, 8\}$, $\{4\}$) and (8, $\{0, 1, 2, 4, 6, 8\}$, $\{8, 12, 16\}$), the updated local view for both these worlds is $\{1, 6, 4, 8\}$. Note that in a Kripke structure, if a local state for agent i contains n elements then there would be n worlds with such a local state for agent i. This is the case in the Kripke structure of figure 1(ii) where the local state for agent 2 being $\{2, 3, 4\}$ corresponded to three worlds being joined by edges labeled with a "2". Yet when the complete Kripke structure for the plant in figure 3 is updated according to Procedure 1, only *two* possible worlds with the local state of $\{1, 6, 4, 8\}$ are indistinguishable to agent 1.

An interesting aspect of the updated local views for agent 1 in figure 5 is that when the "successor" states of communication state w are updated using Procedure 1, distributed knowledge of $\neg \mathrm{bad}(v')$ does in fact hold at all the updated versions of these "successor" states. We know from [6] that communication at w does give agent 1 knowledge to make the correct control decision. We are currently examining combining Theorem 1—as presented here—with the notion that distributed observability could persist as the system evolves. This could allow us to avoid the splitting activity noted in figure 2 and identify a range of possible worlds where communication from an agent will still lead to a control solution.

References

[1] G. Barrett and S. Lafortune. On the synthesis of communicating controllers with decentralized information structures for DES. In *Proc. 37th Conf. Decision Contr.*, pages 3281–3286, 1998.

[2] J.Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *J. ACM*, 37(3):549–587, 1990.

[3] F. Lin and W.M. Wonham. On observability of discrete-event systems. *Info. Sci.*, 44:173–198, 1988.

[4] P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete event processes. *SIAM J. Contr. Optim.*, 25(1):206–230, 1987.

[5] S.L. Ricker and K. Rudie. Know means no: Incorporating knowledge into decentralized discrete-event control. In *Proc. Am. Contr. Conf.*, pages 2348–2353, 1997.

[6] S.L. Ricker and K. Rudie. Incorporating communication and knowledge into decentralized discrete-event systems. In *Proc. 38th Conf. Decision Contr.*, pages 1326–1332, 1999.

[7] K. Rudie, S. Lafortune, and F. Lin. Minimal communication in a distributed discrete-event control system. In *Proc. Am. Contr. Conf.*, pages 1965–1970, 1999.

[8] K. Rudie and W.M. Wonham. Think globally, act locally: Decentralized supervisory control. *IEEE Trans. Automat. Contr.*, 37(11):1692–1708, 1992.

[9] J.H. van Schuppen. Decentralized supervisory control with information structures. In *Proc. Int. Workshop on Discrete Event Systems*, pages 36–41, 1998.

[10] K.C. Wong and J.H. van Schuppen. Decentralized supervisory control of discrete-event systems with communication. In *Proc. Int. Workshop on Discrete Event Systems*, pages 284–289, 1996.