New Results for Joint Diagnosability of Self-observed Distributed Discrete Event Systems

Lina Ye, Philippe Dague

New Results for Joint Diagnosability of Self-observed Distributed Discrete Event Systems

Lina YE [1], **and Philippe DAGUE** [1]

[1] *LRI, UMR8623, Univ. Paris-Sud & CNRS, Orsay, F-91405*
lina.ye@lri.fr
philippe.dague@lri.fr

ABSTRACT

Diagnosability is an important property that determines at design stage how accurate any diagnosis algorithm can be on a partially observable system. Most existing approaches assumed that each observable event in the system is globally observed. Considering the cases where there is no global information, a recent work has proposed a new framework to check diagnosability in a system where each component can only observe its own observable events to keep the internal structure private in terms of observations. However, the authors implicitly assume that the local paths in each component can be exhaustively enumerated, which is not true in a general case where there are embedded cycles. In this paper, we get some new results about diagnosability in such a system, i.e., what we call joint diagnosability in a self-observed distributed system. First we prove its undecidability with unobservable communication events by reducing the Post's Correspondence Problem (PCP) to an observation problem, inspired from an existing work. Then we propose an algorithm to check a sufficient but not necessary condition of joint diagnosability. Finally we briefly discuss about the decidable case where communication events are all observable.

1 INTRODUCTION

Over the latest decades, with more performance requirements imposed on the complex systems, they are subject to more errors. However, it is not realistic to detect faults manually for such complex systems. Automated diagnosis mechanisms are therefore required to monitor large distributed applications. Generally speaking, diagnosis reasoning aims at detecting possible faults that can explain the observations. The possibility to achieve such a diagnosis reasoning depends on the diagnosability of the system. Diagnosability is an important property that determines at design stage how accurate any diagnosis algorithm can be on a partially observable system. The diagnosability analysis problem has received considerable attention in the literature. In this paper, the systems we discuss about are Discrete Event Systems (DES).

Some existing works analyzed diagnosability in a centralized way ((Sampath *et al.*, 1995), (Jiang *et al.*, 2001), etc.), i.e., a monolithic model of a given system is hypothesized, which is a very powerful information for diagnosability analysis. But such a global model is not always available and moreover the centralized approaches lead to combinatorial explosion of the search space. This is why very recently distributed approaches for diagnosability began to be investigated ((Pencolé, 2004), (Schumann and Pencolé, 2007), etc.), relying on local objects. However, all these approaches assume that each observable event in the system can be observed by all components, i.e., globally observed. However, there are some cases where it is not possible to assume the presence of global information. For example, networked control systems are characterized by the fact that multiple distributed components possess their own part of available information instead of a global knowledge. Then (Ye and Dague, 2010) has proposed a new framework to check diagnosability in a system where each component can only observe its own observable events to keep the internal structure private in terms of observations. However, the authors implicitly assume that the local paths in each component can be exhaustively enumerated, which is not true in a general case where there are embedded cycles. In this paper, we generalize this work to get some new results about the diagnosability of what we call self-observed distributed systems, i.e., systems where locally observable events can only be observed by their own component.

We make several contributions in this paper. The first one is to extend diagnosability of globally observed systems to what we call joint diagnosability of self-observed systems and then to prove its undecidability, in the general case where communication events are unobservable, by reducing the Post's Correspondence Problem (PCP) to a joint observability problem, which is inspired from the undecidability result of joint observability (Tripakis, 2001). The second one is to propose an algorithm for testing a sufficient

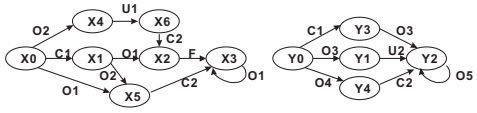

Figure 1: A system with two components G_1 (left) and G_2 (right).

condition of joint diagnosability with unobservability of communication events, where we first obtain pairs of local trajectories in the faulty component, such that for each pair only one trajectory contains the fault but both trajectories have the same local observations, and then check their global consistency into two phases. We provide the proof that it is a sufficient condition and point out why it is not necessary. The third one is to discuss about the decidable case where communication events are observable.

2 PRELIMINARIES

In this section, we model self-observed distributed DES and then recall joint diagnosability features (Ye and Dague, 2010).

We consider a self-observed distributed DES composed of a set of components $\{G_1, G_2,..., G_n\}$ that communicate by communication events. Each component can only observe its own observable events and thus keeps its internal structure private in terms of observations. Such a system is modeled by a set of finite state machines (FSM), each one representing the local model of one component. The local model of a component G_i is a FSM, denoted by $G_i = (Q_i, \Sigma_i, \delta_i, q_i^0)$, where Q_i is the set of states; Σ_i is the set of events; $\delta_i \subseteq Q_i \times \Sigma_i \times Q_i$ is the set of transitions; and q_i^0 is the initial state. The set of events Σ_i is partitioned into four subsets: Σ_{i_o}, the set of locally observable events that can be observed only by their own component G_i; Σ_{i_u}, the set of unobservable normal events; Σ_{i_f}, the set of unobservable fault events; and Σ_{i_c}, the set of communication events shared by at least one other component, which are the only shared events between components. Figure 1 depicts a self-observed distributed system with two components: G_1 (left) and G_2 (right), where the events Oi denote locally observable events, the event F denotes an unobservable fault event, the events Ui denote unobservable normal events and the events Ci denote communication events.

We denote the synchronized FSM of components $G_1, ..., G_n$ by $\|(G_1, ..., G_n)$, where the synchronized events are the shared events between components and any one of them always occurs simultaneously in all components that define it. The state space of the synchronized FSM is the Cartesian product of the state spaces of components. The global model of the entire system is implicitly defined as the synchronized FSM of all components based on their shared events, i.e., communication events. However, the global model will not be calculated in this paper since in a self-observed distributed system, the global occurrence order of observable events is not accessible. In the following, we call subsystem of G the synchronization of a subset of components of G, i.e., $\|(G_{s_1}, ..., G_{s_m})$, where $\{s_1, ...s_m\} \subseteq \{1, ...n\}$. Note that one component or the entire system can also be considered as a subsystem.

Given the system model $G = (Q, \Sigma, \delta, q^0)$, the set of words produced by the FSM G is a prefix-closed language $L(G)$ that describes the normal and faulty behaviors of the system. Formally, $L(G) = \{s \in \Sigma^* | \exists q \in Q, (q^0, s, q) \in \delta\}$, where the transition δ has been extended from events to words. In the following, we call a word of $L(G)$ also a **trajectory** in the system G and a sequence $q_0\sigma_0q_1\sigma_1...$ a **path** in G, where $\sigma_0\sigma_1...$ is a trajectory and for all i, we have $(q_i, \sigma_i, q_{i+1}) \in \delta$. Given $s \in L(G)$, we denote the post-language of $L(G)$ after s by $L(G)/s$ and denote the projection of s to observable events of G (resp. G_i) by $P(s)$ (resp. $P_i(s)$). For example, if $s = O1.U2.O3^*$, then we have $P(s) = O1.O3^*$, where Oi denotes an observable event. We adopt the assumption described in (Pencolé, 2004), i.e., the projection of the global language on each local model is observable live, in particular there is no unobservable cycle in any component. For the sake of simplicity, our approach is illustrated by dealing with only one fault, which can be extended to the case with multiple faults.

Next we rephrase for our context the definition of reconstructibility introduced in (Cori and Métivier, 1985).

Definition 1 *(Reconstructibility). Given a system G that is composed of several subsystems, i.e., $G = \|(G_{s_1}, ..., G_{s_m})$, a set of trajectories in these subsystems is said to be reconstructible with respect to G if it is obtained by projection on this set of subsystems of a trajectory ρ in G.*

If there is no common communication event between two subsystems, then any trajectory in one subsystem and any one in the other subsystem are reconstructible.

For the sake of consistency, now we rename what is called cooperative diagnosability in (Ye and Dague, 2010) as joint diagnosability. We denote a trajectory ending with the fault f by s^f.

Definition 2 *(Joint diagnosability). A fault f is jointly diagnosable in a self-observed distributed system G with components $\{G_1, ...G_n\}$, iff*

$$\exists k \in N, \forall s^f \in L(G), \forall t \in L(G)/s^f, (\forall i \in \{1, ..., n\}, |P_i(t)| \geq k) \Rightarrow (\forall p \in L(G) (\forall i \in \{1, ..., n\}, P_i(p) = P_i(s^f.t)) \Rightarrow f \in p).$$

Joint diagnosability of a fault f means that for each trajectory s^f in G, for each t that is an extension of s^f with enough locally observable events in all components, every trajectory p in G that is equivalent to $s^f.t$ for local observations in each component should contain in it f. In a self-observed system, we call a pair of trajectories p and $p\prime$ satisfying the three following conditions a (global) **indeterminate pair**: 1) p contains f and $p\prime$ does not; 2) p has arbitrarily long local observations in all components after the occurrence of f; 3) $\forall i \in \{1, ..., n\}$, $P_i(p) = P_i(p\prime)$. Here arbitrarily long local observations can be considered as infinite local observations. Now we have the following theorem (Ye and Dague, 2010).

Theorem 1 *Given a self-observed distributed system G, a fault f is jointly diagnosable in G iff there is no (global) indeterminate pair in G.*

3 UNDECIDABLE CASE

To discuss about joint diagnosability, we consider two cases: 1) communication events are unobservable; 2) communication events are observable. In this and the next sections, we consider the first general case, i.e., unobservability of communication events.

From theorem 1, we know that checking joint diagnosability boils down to check the existence of (global) indeterminate pairs that witness non joint diagnosability. Inspired from (Tripakis, 2001), we discuss first about whether it is decidable or not. To be self-contained, we rephrase joint observability (Tripakis, 2001), where the global model of a system is assumed.

Definition 3 ***(Joint observability)*** *Given the global model G of a system and the sets of observable events $\Sigma_{i_o} \subseteq \Sigma, i = 1, ..., k$, then a fault f is jointly observable in $L(G)$ w.r.t. these sets if the following condition holds, where the projection of s to Σ_{i_o} is denoted by $P_{\Sigma_{i_o}}(s)$:*

$$\forall s \in L(G), s\prime \in L(G), f \in s, f \notin s\prime,$$
$$\exists i = 1, ...k, P_{\Sigma_{i_o}}(s) \neq P_{\Sigma_{i_o}}(s\prime)$$

Joint observability of a fault in a system means that there are not two system behaviors such that only one of them contains the fault but their projections on each set of observable events are the same. Then the undecidability of joint observability with at least two sets of observable events is proved by reducing the Post's Correspondence Problem (PCP) to an observation problem. Now we briefly describe the outline of the proof with two sets (Tripakis, 2001).

1)PCP: given a finite alphabet Σ, two sets of words $v_1, v_2, ..., v_k$ and $z_1, z_2, ..., z_k$ over Σ, then a solution to PCP is a sequence of indices $(i_m)_{1 \leq m \leq n}$ with $n \geq 1$ and $1 \leq i_m \leq k$ for all m such that $v_{i_1} v_{i_2} ... v_{i_n} = z_{i_1} z_{i_2} ... z_{i_n}$.

2) Now let $\Sigma\prime = \{a_1, ...a_k\}$ be a set of new letters, not in Σ. Then consider the language L over $\Sigma \cup \Sigma\prime \cup \{good, bad\}$, defined by the regular expression: $good(v_1 a_1 + ... + v_k a_k)^+ + bad(z_1 a_1 + ... + z_k a_k)^+$, where Σ^+ denotes the set of all finite words over Σ except ϵ. All words in L that start with $good$ constitute the normal behaviors and all those that start with bad constitute the faulty behaviors.

3) If there is a solution for the above PCP, i.e., there exist indices $i_1, ..i_n \in \{1...k\}, n \geq 1$, such that $v_{i_1} v_{i_2} ... v_{i_n} = z_{i_1} z_{i_2} ... z_{i_n}$, then the fault is not jointly observable w.r.t. observable events sets Σ and $\Sigma\prime$. The reason is that in this case, we have a pair of words $\rho = good v_{i_1} a_{i_1} v_{i_2} a_{i_2} ... v_{i_n} a_{i_n}$ and $\rho\prime = bad z_{i_1} a_{i_1} z_{i_2} a_{i_2} ... z_{i_n} a_{i_n}$ such that both $\rho, \rho\prime \in L$ with $\rho\prime$ only containing the fault, and $\rho, \rho\prime$ have the same observations both for Σ and $\Sigma\prime$, which violates joint observability.

4) On the other side, if the fault is not jointly observable, there is at least one pair of words violating joint observability, denoted by ρ and $\rho\prime$. Since only one of them is a normal behavior, say ρ, then ρ must be of the form $good v_{i_1} a_{i_1} v_{i_2} a_{i_2} ... v_{i_n} a_{i_n}$ and $\rho\prime$ must be of the form $bad z_{j_1} a_{j_1} z_{j_2} a_{j_2} ... z_{j_l} a_{j_l}$. Furthermore, we know that ρ and $\rho\prime$ have the same observations both for Σ and $\Sigma\prime$. So we have $a_{i_1} a_{i_2} ... a_{i_n} = a_{j_1} a_{j_2} ... a_{j_l}$, which means that $l = n, i_1 = j_1, i_2 = j_2, ..., i_n = j_n$ since a_{i_m} is a letter not a word. And then we also get $v_{i_1} v_{i_2} ... v_{i_n} = z_{i_1} z_{i_2} ... z_{i_n}$, which means that there does exist a solution for the above PCP.

5) Thus checking joint observability, i.e., checking the existence of such a pair violating joint observability, is equivalent to check for a solution to PCP.

Since PCP is a well known undecidable problem, then checking joint observability is also undecidable.

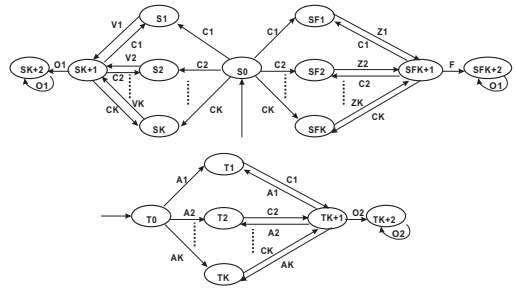

Figure 2: A system with two components G_1 (top) and G_2 (bottom) for proving undecidability of joint diagnosability.

Now, from definition 2, we know that joint diagnosability is violated iff it exists two infinite trajectories, one with the fault and the other without the fault, and both give the same infinite sequence of observations for each component (observer), i.e. it exists an indeterminate pair. So, for just two components, the example of the proof above shows that joint diagnosability checking boils down to check the existence of an infinite sequence $i_1 ... i_n ...$ such that $v_{i_1} ... v_{i_n} .... = z_{i_1} ... z_{i_n} ...$, which is actually the infinite PCP. There are two major differences between joint diagnosability in our framework and joint observability in (Tripakis, 2001). One is that the former assumes that local observers are attached to local components that are synchronized by common communication events to get a global model while the latter separates arbitrarily the observable events in the global model into several sets. The other one is that joint diagnosability consists in separating infinite trajectories while joint observability consists in separating finite ones. Thus, if any communication event is assumed to be unobservable, joint diagnosability checking boils down to infinite PCP. But this one has also been proved to be undecidable (Halava and Harju, 2006), which gives the result.

For the sake of simplicity, we give now a simpler and self-contained proof of undecidability of joint diagnosability that relies only on finite PCP.

Theorem 2 *Given a self-observed distributed system where communication events are unobservable, checking joint diagnosability of a given fault is undecidable.*

Proof :

1) To prove this theorem, without loss of generality, consider the example depicted in figure 2, where the system is composed of two components G_1 and G_2. In G_1, each one of $Vi, i \in \{1, ..., k\}$, and each one of $Zi, i \in \{1, ..., k\}$, denotes a sequence of observable events all different from $O1$, $C1, ..., Ck$ are unobservable communication events, F denotes a fault event and $O1$ is an observable event. In G_2, each one of $Ai, i \in \{1, ..., k\}$, denotes an observable event different from $O2$, $C1, ..., Ck$ are unobservable communication events and $O2$ is an observable event. Then the observations in G_1 can be described as

$Vi_1Vi_2...Vi_nO1^*$ without fault or $Zi_1Zi_2...Zi_nO1^*$ with fault, where $\forall i_j, j \in \{1,...,n\}, i_j \in \{1,...,k\}$. In G_2, the observations are $Ai_1Ai_2...Ai_nO2^*$. In this system, the occurrence of the fault can be confirmed by the observation of $O1$. Next we discuss about the case without the observation of $O1$.

2) Without the observation of $O1$, the local observations are $wO1^+$ for G_1 and $Ai_1Ai_2...Ai_nO2^*$ for G_2, where $w = Vi_1Vi_2...Vi_n$ when there is no fault or $w = Zi_1Zi_2...Zi_n$ when there is a fault. Clearly, if PCP has a solution, i.e., $\exists(i_m)_{1\leq m\leq n}$ such that $Vi_1Vi_2...Vi_n = Zi_1Zi_2...Zi_n$, we have two trajectories p and $p\prime$ such that the observations of p in G_1 are $Vi_1Vi_2...Vi_nO1^+$, which is a trajectory without fault, while the observations of $p\prime$ in G_1 are $Zi_1Zi_2...Zi_nO1^+$, which is a trajectory with a fault. And both p and $p\prime$ have the same observations for G_2, i.e., $Ai_1Ai_2...Ai_nO2^*$. Thus we get that p and $p\prime$ have the same observations for both G_1 and G_2, i.e., $Vi_1Vi_2...Vi_nO1^+$=$Zi_1Zi_2...Zi_nO1^+$ for G_1 and $Ai_1Ai_2...Ai_nO2^*$ for G_2, then the fault is not jointly diagnosable.

3) On the other hand, if the fault is not jointly diagnosable, then we obtain at least one indeterminate pair, denoted by p and $p\prime$ such that the projection of p on G_1 is $Ci_1Vi_1Ci_2Vi_2...Ci_nVi_nO1^*$, on G_2 is $Ai_1Ci_1Ai_2Ci_2...Ai_nCi_nO2^*$ and that of $p\prime$ on G_1 is $Cj_1Zj_1Cj_2Zj_2...Cj_mZj_mFO1^*$ and on G_2 is $Aj_1Cj_1Aj_2Cj_2...Aj_mCj_mO2^*$. From the fact that p and $p\prime$ have the same observations for G_2, we get $Ai_1Ai_2...Ai_nO2^* = Aj_1Aj_2...Aj_mO2^*$ and thus we have $m = n$ and $i_1 = j_1, ..., i_n = j_n$. And then from the same observations of p and $p\prime$ on G_1, we get $Vi_1Vi_2...Vi_nO1^* = Zi_1Zi_2...Zi_nO1^*$, i.e., $Vi_1Vi_2...Vi_n = Zi_1Zi_2...Zi_n$, which means that there is a solution for PCP.

This proves that the existence of a solution for PCP is equivalent to the fault being not jointly diagnosable. Since PCP is an undecidable problem, then checking joint diagnosability is undecidable.

4 ALGORITHM TO TEST A SUFFICIENT CONDITION

We have proved that joint diagnosability in self-observed distributed systems with unobservable communication events is undecidable. We can nevertheless propose an algorithm to test a sufficient condition, which is still quite useful in some circumstances. The first step is to construct the local diagnoser from a given local model, which allows one to get fault information for any local trajectory. Then we show how to build the local twin plant, which allows one to obtain original information about indeterminate pairs (also called local indeterminate pairs in the following), based on the local diagnoser. The next step is to check the global consistency, i.e., to check whether the local indeterminate pairs can be extended into (global) indeterminate pairs, whose existence testifies non joint diagnosability. We give an algorithm to only test a sufficient but not necessary condition for global consistency. The proof is consistent with the undecidability result obtained in the precedent section. Actually our algorithm remains trivially applicable when the assumption of unobservability of communication events is partially relaxed, i.e.., in the most general case where some communication events are observable and others unobservable.

4.1 Original diagnosability information

From theorem 1, we know that joint diagnosability verification consists in checking the existence of indeterminate pairs in the system. In the distributed framework, we use the structure called local twin plant defined in (Jiang *et al.*, 2001) to analyze joint diagnosability. In particular, the considered fault is assumed to only occur in one component, denoted by G_f. Then the local twin plant for G_f contains original information for indeterminate pairs: actually this twin plant is a FSM that compares every pair of local trajectories to search for the pairs with the same arbitrarily long local observations, but exactly one of the two containing a fault, which precisely we call local indeterminate pairs. First, we define an operation called delay closure with respect to a subset Σ_d of Σ to preserve all information about the events in Σ_d by abstracting away irrelevant parts.

Definition 4 *(Delay Closure). Given a FSM $G = (Q, \Sigma, \delta, q^0)$, its delay closure with respect to $\Sigma_d \subseteq \Sigma$ is $\complement_{\Sigma_d}(G) = (Q, \Sigma_d, \delta_d, q^0)$ where $(q, \sigma, q') \in \delta_d$ iff $\exists s \in (\Sigma \backslash \Sigma_d)^*, (q, s\sigma, q') \in \delta$.*

We now describe how to construct the local diagnoser of a given component, based on which we build the local twin plant. Given a local model, we get a modified one by attaching fault label, denoted by $l \in \{N, F\}$, where N for normal and F for fault, to each state. In other words, before the occurrence of the fault, each state is labeled with label N and, after its occurrence, with label F.

Definition 5 *(Local diagnoser). Given a local model G_i, its local diagnoser D_i is obtained by operating the delay closure with respect to the set of communication events and observable events on the modified model: $D_i = \complement_{\Sigma_{i_o} \cup \Sigma_{i_c}}(G_i^m)$, where G_i^m is the modified version of G_i.*

Based on the local diagnoser, the corresponding local twin plant is obtained by synchronizing the local diagnoser with itself based on the locally observable events, allowing one to obtain all pairs of local trajectories with the same observations to search for local indeterminate pairs. To simplify this synchronization, the two identical local diagnosers, denoted by D_i^l for the left instance and D_i^r for the right instance, can be reduced as follows: D_i^l is obtained by retaining only paths with at least one fault cycle and D_i^r is obtained by retaining only paths with at least one non-fault cycle. This reduction keeps all necessary original diagnosability information since what we are interested in here are only local indeterminate pairs, i.e., pairs of local trajectories with only one containing the fault and both of them with the same observations. However, this reduction is only applicable for the construction of the local twin plant of the faulty component G_f; for other components, the local twin plant is obtained by synchronizing the non reduced left instance and the non reduced right instance since there is no fault information. Considering that this synchronization is based

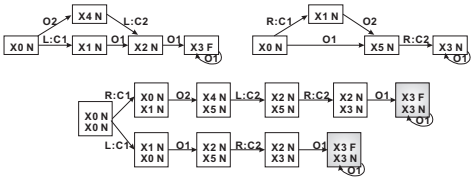

Figure 3: Two reduced instances of the diagnoser for G_1 (top) and part of the corresponding local twin plant (bottom).

on the set of locally observable events Σ_{i_o}, the non-synchronized events are distinguished between the two instances by the prefix L or R: in D_i^l (D_i^r), each communication event $c \in \Sigma_{i_c}$ from D_i is renamed by $L:c$ ($R:c$). The names of all locally observable events are left unchanged.

Definition 6 *(Local twin plant). Given a local diagnoser D_i for the component G_i, the corresponding local twin plant is a FSM, denoted by $T_i = D_i^l \| D_i^r$, where the synchronized events are locally observable events in G_i.*

Each state of a local twin plant is a pair of local diagnoser states providing two possible diagnoses with the same local observations. Given a twin plant state $((q^l, l^l)(q^r, l^r))$, where (q^l, l^l) and (q^r, l^r) are two diagnoser states and each one contains both a system state and a corresponding pattern state, if the considered fault $f \in l^l \cup l^r$ but $f \notin l^l \cap l^r$, which means that the occurrence of f is not certain up to this state, then this state is called an ambiguous state with respect to the fault f. An ambiguous state cycle is a cycle containing only ambiguous states. In a local twin plant, if a path contains an ambiguous state cycle with at least one locally observable event, then it is called a **local indeterminate path**, which corresponds to a local indeterminate pair. Note that local indeterminate paths contain original diagnosability information and can be obtained only in the local twin plant of the component G_f. If a local indeterminate pair can be extended into a global indeterminate pair, then we say that its corresponding local indeterminate path is globally consistent. Figure 3 shows the left and right instances of the local diagnoser for the faulty component G_1 of Figure 1 (top) as well as a part of the corresponding local twin plant (bottom). Clearly, in the local twin plant, we have local indeterminate paths since they have ambiguous state cycles with observable events.

4.2 Global consistency checking

Joint diagnosability verification consists in checking the existence of globally consistent local indeterminate paths, whose existence proves non joint diagnosability. To check the global consistency of local indeterminate paths in a subsystem, we consider now two important issues: 1) the global consistency of the corresponding left trajectories of the local indeterminate paths in the local twin plant of the subsystem, shortly called left consistency checking in the following; 2) the global consistency of the corresponding right trajectories of the local indeterminate paths in the local twin plant of the subsystem, shortly called right consistency checking.

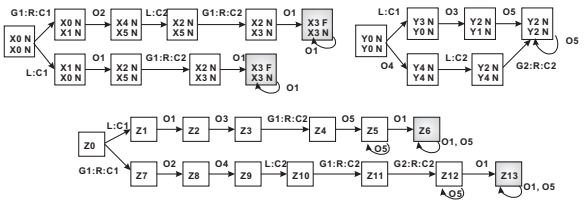

Figure 4: Part of the renamed local twin plants for G_1 and G_2 (top) and part of the left consistent plant T_f^l (bottom).

Definition 7 *(Left (Right) consistent plant). Given a subsystem G_S composed of components $G_{i_1}, ..., G_{i_m}$ and their corresponding local twin plants $T_{i_1}, ..., T_{i_m}$, to obtain a left (right) consistent plant with respect to the subsystem G_S, denoted by T_f^l (T_f^r), we perform the following two steps:*

- *Distinguish right (left) communication events between local twin plants by renaming them with the prefix of component ID. For example, $R{:}C2$ ($L{:}C2$) in the local twin plant of G_2 is renamed as $G_2{:}R{:}C2$ ($G_2{:}L{:}C2$).*
- *Synchronize the renamed local twin plants with the synchronized events being the common left (right) communication events, which works because observable events do not intersect between components and non-synchronized right (left) communication events are distinguished by the prefix of component ID.*

In the left (right) consistent plant with respect to a subsystem G_S, each path p corresponds to a set of paths $p_{i_1}, ..., p_{i_m}$ in the local twin plants of all components in G_S such that the set of left (right) trajectories of $p_{i_1}, ..., p_{i_m}$ are reconstructible with respect to G_S. For our example, the bottom part of figure 4 shows a part of the left consistent plant T_f^l, which is obtained by synchronizing the renamed local twin plant of G_1 and that of G_2 (top part of figure 4) based on the common left communication events.

4.3 Algorithm

Algorithm 1 presents the procedure to verify a sufficient condition of joint diagnosability. As shown in the pseudo-code, algorithm 1 performs as follows. Given the input as the set of component models, the fault F that may occur in the component G_f, we initialize the parameters as empty, i.e., G_S^l, the subsystem for the left consistency checking and G_S^r, the subsystem for the right consistency checking. The procedure of the algorithm can be separated into two parts: left consistency checking (line 3-12) and right consistency checking (line 13-24).

Left consistency checking begins with the local twin plant construction of G_f, the subsystem G_S^l being now G_f (line 3-4). As long as both the left consistent plant T_f^l with respect to the current left subsystem G_S^l and $DirectCC(G, G_S^l)$ are not empty (line 5), where $DirectCC(G, G_S^l)$ is the set of directly connected components to the subsystem G_S^l (a directly connected component being one sharing at least one common communication event with the subsystem)

Algorithm 1 Algorithm to check a sufficient condition of joint diagnosability in the general case

1: INPUT: the system model $G = (G_1, ..., G_n)$; the fault F and the faulty component G_f
2: Initializations: $G_S^l \leftarrow \emptyset$ (subsystem for left consistency checking); $G_S^r \leftarrow \emptyset$ (subsystem for right consistency checking)
3: $T_f^l \leftarrow ConstructLTP(G_f)$
4: $G_S^l \leftarrow G_f$
5: **while** $T_f^l \neq \emptyset$ and $DirectCC(G, G_S^l) \neq \emptyset$ **do**
6: $G_i \leftarrow SelectDirectCC(G, G_S^l)$
7: $T_i \leftarrow ConstructLTP(G_i)$
8: $T_f^l \leftarrow T_f^l \| T_i$
9: $G_S^l \leftarrow Add(G_S^l, G_i)$
10: $T_f^l \leftarrow RetainConsisPaths(T_f^l)$
11: **if** $T_f^l = \emptyset$ **then**
12: return "F is jointly diagnosable in G"
13: **else**
14: $T_f^r \leftarrow AbstractRight(G_f, T_f^l)$
15: $G_S^r \leftarrow G_f$
16: **while** $T_f^r \neq \emptyset$ and $G_S^l \neq G_S^r$ **do**
17: $G_i \leftarrow SelectDirectCC(G_S^l, G_S^r)$
18: $T_f^r \leftarrow T_f^r \| AbstractRight(G_i, T_f^l)$
19: $G_S^r \leftarrow Add(G_S^r, G_i)$
20: $T_f^r \leftarrow RetainConsisPaths(T_f^r)$
21: **if** $T_f^r = \emptyset$ **then**
22: return "F is jointly diagnosable in G"
23: **else**
24: return "Joint diagnosability cannot be determined"

the algorithm repeatedly performs the following steps to further check left consistency in an extended subsystem:

1. Select one directly connected component G_i to the subsystem G_S^l and construct its local twin plant T_i (line 6-7).
2. Synchronize T_f^l with T_i to obtain left consistent plant for this extended subsystem based on the set of common left communication events (line 8). To do this, the set of non-synchronized right communication events are distinguished by the prefix of component ID.
3. Update the subsystem G_S^l by adding G_i and reduce the newly obtained T_f^l by retaining only paths with ambiguous state cycles containing observable events for all components in G_S^l (line 9-10).

If the left consistent plant T_f^l is empty, then there is no local indeterminate path that corresponds to a set of paths in the local twin plants of all components in the subsystem such that their left trajectories are reconstructible, which implies the non existence of a globally consistent local indeterminate path. In this case joint diagnosability information is returned (line 11-12). Otherwise, if T_f^l is not empty (line 13), then we proceed to check right consistency of the corresponding paths in T_f^l that have been already verified to be left consistent in the whole system.
Right consistency checking begins with the function $AbstractRight(G_f, T_f^l)$ (line 14), which performs delay closure with respect to the set of right communication events and observable events of G_f. In this way, what we obtain does not contain left communication events. Then the subsystem G_S^r is assigned as G_f (line 15). As long as the right consistent plant T_f^r for the current right subsystem G_S^r is not empty and $G_S^l \neq G_S^r$ (line 16), we repeatedly perform the following steps to check right consistency in an extended subsystem (note that since left consistency checking does explore all connected components, during right consistency checking we only need to consider the subsystem G_S^l instead of the whole system):

1. Select a directly connected component G_i to G_S^r from G_S^l (line 17).
2. Perform the function $AbstractRight(G_i,\ T_f^l)$, which has already been described above, and then synchronize with T_f^r based on the set of common right communication events (line 18). To do this, we rename the right communication events by removing the prefix of component ID, e.g., G_i:R:$C2$ renamed as R:$C2$.
3. Update the subsystem G_S^r by adding G_i and reduce the newly obtained T_f^r by retaining only paths with ambiguous state cycles containing observable events for all components in G_S^r (line 19-20).

If the right consistent plant T_f^r is empty, then there is no local indeterminate path that corresponds to a set of paths in the local twin plants such that their left trajectories and right trajectories are reconstructible respectively, i.e., there is no globally consistent local indeterminate path. In this case, the algorithm returns joint diagnosability information (line 21-22). Otherwise, if T_f^r is not empty, actually we cannot determine whether the fault is jointly diagnosable or not. Then the algorithm returns the information about the indetermination of joint diagnosability (line 23-24). In other words, empty left consistent plant T_f^l or empty right consistent plant T_f^r is a sufficient condition but not a necessary condition of joint diagnosability.

Theorem 3 *In algorithm 1, if the left consistent plant T_f^l or the right consistent plant T_f^r is empty, then the fault is jointly diagnosable, but the reverse is not true.*

Proof :
($\Rightarrow$) Suppose that T_f^l or T_f^r is empty and that the fault is not jointly diagnosable. From non joint diagnosability, it follows that there exists at least one globally consistent local indeterminate path. Since global consistency of a local indeterminate path implies both left consistency and right consistency, from algorithm 1 we know that, after left and right consistency checking, this local indeterminate path must correspond to a

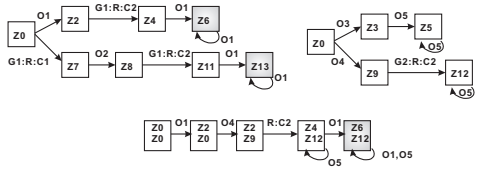

Figure 5: FSM after delay closure on the left consistent plant (figure 4) for G_1 and G_2 (top) and part of the right consistent plant (bottom).

path both in T_f^l and in T_f^r. Thus neither T_f^l nor T_f^r is empty, which contradicts the assumption.

($\Leftarrow$) Now we explain why non emptiness of both T_f^l and T_f^r does not necessarily imply that the fault is not jointly diagnosable. Suppose that T_f^l is not empty and that it contains two paths, denoted by $\rho 1$ and $\rho 2$, corresponding to two local indeterminate paths. $\rho 1$ corresponds to a set of paths $\rho_i^1, 1 \leq i \leq n$ in the local twin plants of all components and $\rho 2$ corresponds to a set of paths $\rho_i^2, 1 \leq i \leq n$ in all local twin plants. Now suppose that the right trajectories of the set of paths $\rho_i^1, 1 \leq i \leq n$ are not reconstructible and the same for that of the set of paths $\rho_i^2, 1 \leq i \leq n$. It ollows that the two local indeterminate paths cannot be extended into global indeterminate pairs and thus are not globally consistent. Then we further suppose that the right trajectories of the set of paths $\rho_1^1, ..., \rho_{n-1}^1, \rho_n^2$ are reconstructible or the same for the set of paths $\rho_1^2, ..., \rho_{n-1}^2, \rho_n^1$. In this case, from algorithm 1, it follows that finally the right consistent plant T_f^r is not empty. Now both T_f^l and T_f^r are not empty but there is no globally consistent local indeterminate paths, i.e., the fault is jointly diagnosable.

Now illustrate on our example the fact that the condition is not necessary. The top part of figure 5 shows the results of performing delay closure with respect to the set of right communication events and locally observable events both for G_1 and G_2 on the left consistent plant depicted in the bottom part of figure 4. Then, to check right consistency, we rename again the right communication events by removing the component ID such that they can be synchronized. Finally the bottom part of figure 5 shows a part of the right consistent plant, which is not empty. Now both left and right consistent plants are not empty, but this does not imply the existence of global indeterminate pairs that witness non joint diagnosability. Actually the part of the left consistent plant depicted here corresponds to two local indeterminate pairs in G_1 with their corresponding left consistent pairs in G_2, i.e., one local indeterminate pair is $((C1.O1.F.O1^*), (O1.C2.O1^*))$ in G_1 with its left consistent pair $((C1.O3.O5^*), (O3.U2.O5^*))$ in G_2 and the other local indeterminate pair is $((O2.U1.C2.F.O1^*), (C1.O2.C2.O1^*))$ in G_1 with its left consistent pair $((O4.C2.O5^*), (O4.C2.O5^*))$ in G_2. While the right consistent plant shown here corresponds to one local indeterminate pair in G_1, which is $((C1.O1.F.O1^*), (O1.C2.O1^*))$, with its corresponding right consistent pair in G_2, i.e., $((O4.C2.O5^*), (O4.C2.O5^*))$. Thus we can see that the same local indeterminate pair does not correspond to the same consistent pair in G_2 in the left consistent plant and in the right consistent plant, which means that this local indeterminate pair cannot be extended into a global indeterminate pair. So our algorithm gives indeterminate information for joint diagnosable systems that satisfy the following condition: for any set of paths, i.e., one path in the local twin plant of each component and one in the plant for faulty component being a local indeterminate path, that are left consistent and right consistent respectively, then we have that their corresponding local trajectories in the components cannot constitute an indeterminate pair through synchronization.

5 DECIDABLE CASE

We have proved the undecidability of joint diagnosability when communication events are unobservable. If we assume their observability, then this problem becomes decidable. The reason is that, when any communication event is observable, then in the local twin plant of each component, we obtain all pairs of local trajectories with the same observations, including the same observable communication events. In other words, each path in the local twin plant corresponds to a pair of local trajectories with the same sequence of communication events. It follows that to check global consistency of local indeterminate paths, what was a separate checking for left and right consistency becomes only one checking. While in algorithm 1, the checking into two separate phases is the reason why it gives only a sufficient but not necessary condition for joint diagnosability. Actually, the observability of communication events makes joint diagnosability equivalent to classical diagnosability since only one checking for global consistency implies the same global occurrence order of observations for global indeterminate pairs.

Algorithm 2 presents the procedure to check joint diagnosability when communication events are assumed to be observable. Taking the system model and the faulty component as input, the parameter, i.e., the current subsystem G_S, is initialized as empty. Then the algorithm begins with the construction of the local twin plant of the faulty component G_f, the current subsystem becoming G_f (line 3-4). Here we emphasize that, due to the observability of communication events, then the local twin plant should be constructed by synchronizing two instances based on the set of observable events and the set of communication events, i.e., here communication events do not need to be distinguished by changing their names. As long as both T_f and $DirectCC(G, G_S)$ are not empty (line 5), then the following steps are repeatedly performed:

- Select one component G_i directly connected to the current subsystem G_S and then construct its local twin plant T_i, which is obtained by first operating delay closure with respect to the set of communication events and observable events and then by synchronizing the two instances based on all events, i.e., the set of communication events and observable events. (line 6-7)
- Synchronize the current local twin plant T_f and T_i based on the common communication events of G_S and G_i. (line 8)
- Update the current subsystem by adding this selected component and keep only the paths in the

newly obtained FSM that contain ambiguous state cycles with observations for all involved components. (line 9-10)

During this procedure, if the local twin plant T_f for the current subsystem happens to be empty, which means that there is no path that contains ambiguous state cycle with observations for all concerned components, thus there is no local critical path that is globally consistent and the algorithm returns joint diagnosability information (line 11-12). Otherwise, if at the end the final FSM is not empty, it is returned by the algorithm as non joint diagnosability information (line 13-14) as any path in it corresponds to a globally consistent local indeterminate path. The reason is that if the communication events are observable, then any path in a local twin plant corresponds to a pair of local trajectories with the same observations, including the same communication events. So with the assumption of observability of communication events, joint diagnosability checking becomes decidable, whose verification is provided by the algorithm 2.

Algorithm 2 Algorithm for checking joint diagnosability with observability of communication events

1: INPUT: the system model $G = (G_1, ..., G_n)$; the fault F and the faulty component G_f
2: Initializations: $G_S \leftarrow \emptyset$ (subsystem considered for current checking)
3: $T_f \leftarrow ConstructLTP(G_f)$
4: $G_S \leftarrow G_f$
5: **while** $T_f \neq \emptyset$ and $DirectCC(G, G_S) \neq \emptyset$ **do**
6: $G_i \leftarrow SelectDirectCC(G, G_S)$
7: $T_i \leftarrow ConstructLTP(G_i)$
8: $T_f \leftarrow T_f \| T_i$
9: $G_S \leftarrow Add(G_S, G_i)$
10: $T_f \leftarrow RetainConsisPaths(T_f)$
11: **if** $T_f = \emptyset$ **then**
12: return "F is jointly diagnosable in G"
13: **else**
14: return T_f

6 CONCLUSION

In this paper, we consider self-observed distributed systems such that observable events can only be observed by their own components. Clearly, we do not need the monolithic model of the system, thus the distributed and private (w.r.t. observation) nature of real systems is taken into account. Then we prove the undecidability of joint diagnosability checking when communication events are unobservable, before proposing an algorithm to test a sufficient condition of joint diagnosability in this case. To check the non existence of indeterminate pairs in the system, we start from local indeterminate paths in the local twin plant and then we check both in sequence left consistency and right consistency. Note that, due to the observation-privacy, the global occurrence order of observable events between different components is not known, which is taken into account through constructing left and right consistent plants separately. At the opposite, in the approaches for DES with globally observable events, twin plant is constructed by incrementally synchronizing local twin plants via both left and right communication events at the same time, which means that in their case, knowledge of the global occurrence order is required. Similar to the distributed algorithms for classical diagnosability ((Pencolé, 2004), (Schumann and Pencolé, 2007), etc.), our algorithm has to construct some part of a global structure, but much less than in the centralized approach (in particular the components not involved in the algorithm have their model completely not disclosed) and this is normally unavoidable for off-line diagnosability analysis. Then we discuss the decidable case where communication events are observable by giving an algorithm that checks joint diagnosability. However, the decidable case in (Ye and Dague, 2010) is not the same as that in this paper. The former also deals with unobservable communication events, which becomes decidable because of the assumption of exhaustive enumeration about local paths. While here the decidable case is about observability of communication events. We see that there is a gap between these two cases as the unobservable case is undecidable and the observable case is decidable. Next interesting work is to investigate where is the frontier between these two cases, i.e., to study the decidability of joint diagnosability for partial observability of communication events.

REFERENCES

(Cori and Métivier, 1985) R. Cori and Y. Métivier. Recognizable subsets of some partially abelian monoids. *Theoretical Computer Science*, 35:179–189, 1985.

(Halava and Harju, 2006) V. Halava and T. Harju. Undecidability of infinite post correspondence problem for instances of size 9. *Theoretical Informatics and Applications*, 40(4):551–557, 2006.

(Jiang *et al.*, 2001) S. Jiang, Z. Huang, V. Chandra, and R. Kumar. A polynomial time algorithm for diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 46(8):1318–1321, 2001.

(Pencolé, 2004) Y. Pencolé. Diagnosability analysis of distributed discrete event systems. In *Proceedings of 16th European Conference on Artificial Intelligence ECAI-04*, pages 43–47, Valencia, Spain, 2004.

(Sampath *et al.*, 1995) M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 40(9):1555–1575, 1995.

(Schumann and Pencolé, 2007) A. Schumann and Y. Pencolé. Scalable diagnosability checking of event-driven systems. In *Proceedings of 20th International Joint Conference on Artificial Intelligence IJCAI-07*, pages 575–580, Hyderabad, India, 2007.

(Tripakis, 2001) S. Tripakis. Undecidable problems of decentralized observation and control. In *40th IEEE Conference on Decision and Control*, Orlando, Florida, 2001.

(Ye and Dague, 2010) L. Ye and P. Dague. Diagnosability analysis of discrete event systems with autonomous components. In *Proceedings of 19th European Conference on Artificial Intelligence ECAI-10*, pages 105–110, Lisbon, Portugal, 2010.