

A Model of Component Consistency in Distributed Diagnosis

R. Su and W.M. Wonham¹

Abstract. In distributed diagnosis, each local diagnoser first generates a local estimate then communicates with other local diagnosers to reach suitably defined consistency. In this paper we propose two types of consistency and provide efficient computational procedures to achieve them.

1 Introduction

So far there is a large body of research on fault diagnosis for discrete-event systems, e.g. finite-state automaton approaches [10] [12], and Petri net approaches [1]. These can be roughly classified as follows: (1) centralized approaches like [10] [16]; (2) decentralized [4] [12]; and (3) distributed [11] [6]. It is well known that centralized approaches may lead to high space complexity in order to store the final diagnoser. Meanwhile, weak scalability and robustness are also features of centralized approaches. Decentralized approaches usually lead to relatively small local diagnosers and improved scalability and robustness. Nevertheless, designing a decentralized diagnoser still requires existence of a centralized target system model at the intermediate stage, which again forces a designer to face the space complexity issue. As for distributed approaches, some methods like [11] [9] [1] adopt the same design procedure as decentralized approaches except that, instead of requiring a global clock during communication as in a decentralized approach, those distributed approaches utilize concurrent communication governed by a set of local clocks. In general, they still have to face the high space complexity introduced by the intermediate centralized system model during the diagnoser design stage.

Recently years have seen “pure” distributed approaches in the sense that: (1) the system model is distributed; (2) diagnosis is locally generated; (3) constraints imposed by interaction among local components are satisfied by achieving a suitable consistency property. Approaches like [3] [6] [13] belong to this category. By introducing inter-component consistency, these approaches can avoid the potentially enormous intermediate system model. So space complexity in the design stage and during online diagnosis can be brought under control. To capture the interaction among local components, there are roughly two types of consistency in the literature: one is *global consistency*, which was more or less described in, e.g. [2] [7] [8] [6] [3]; the other one is *local consistency*, as explained in [13]. By using purely local computation and achieving appropriate inter-diagnoser consistency, we can greatly reduce space complexity of online diagnosis in many practical systems and

substantially improve the scalability and robustness of diagnosis. This paper proposes a unified model to describe the different types of consistency and presents efficient computational algorithms to achieve them.

The paper is organized as follows. In Section 2 we present two types of consistency. In Section 3, we introduce the concept of supremal global support and propose an efficient computational procedure to achieve it. Then in Section 4, we introduce the concept of supremal local support and propose a computational procedure to achieve it. We introduce hierarchical computation in Section 5 and draw conclusions in Section 6.

2 Global and Local Consistencies

Let I be an index set. For each $i \in I$ let Σ_i be an event set. As usual [14], we use “ \parallel ” to represent synchronous product. All natural projections in this paper obey the following rules of notation: given an index set J with a predefined collection of event sets $\{\Sigma_j | j \in J\}$, we define

$$\Sigma_J := \cup_{j \in J} \Sigma_j$$

Given two event sets Σ_i and Σ_j , we use $P_{i,j}$ to mean the natural projection

$$P_{i,j} : \Sigma_i^* \rightarrow (\Sigma_i \cap \Sigma_j)^*$$

and $P_{J,i}$ to mean the natural projection

$$P_{J,i} : \Sigma_J^* \rightarrow (\Sigma_J \cap \Sigma_i)^*$$

If we have another index set K with a predefined set $\{\Sigma_k | k \in K\}$, then we use $P_{J,K}$ to mean the natural projection

$$P_{J,K} : \Sigma_J^* \rightarrow (\Sigma_J \cap \Sigma_K)^*$$

Let $\mathcal{L} = \{L_i \subseteq \Sigma_i^* | i \in I\}$ be a set of languages. Now we introduce two types of consistency.

Definition 2.1 $\{L_i | i \in I\}$ is *globally consistent* if for each $i \in I$, $L_i = P_{I,i}(\parallel_{j \in I} L_j)$. \square

Here $\parallel_{j \in I} L_j$ denotes synchronous product of all local components of I . It is well known that synchronous product is one way to combine local transitional behavior to form a global (or composite) transitional behavior. If we treat each local language L_i as a local estimate in diagnosis, then Def. 2.1 says that when a system reaches global consistency, knowing all other local estimates will not improve the current local estimate L_i .

¹ Dept. ECE, University of Toronto, 10 King’s College Road, Toronto, ON M5S 3G4, Canada email: surong.wonham@control.utoronto.ca, telephone: 416-9784124, fax: 416-9780804

Definition 2.2 $\{L_i | i \in I\}$ is *locally consistent* if for each $i, j \in I$, $P_{i,j}(L_i) = P_{j,i}(L_j)$. \square

Again, if we treat each local language L_i as a local estimate in diagnosis, then Def. 2.2 says that when a system reaches local consistency, knowing the local estimate of any local component that is a neighbor of component i will not improve the current local estimate L_i .

From above description we can see that the two types of consistency provide ways to determine whether or not all local components have reached an “agreement”. On that basis we can decide when communication among local components should be terminated. The relation between global and local consistencies is described as follows.

Proposition 2.1 If \mathcal{L} is globally consistent then it must be locally consistent. \square

The reverse statement need not always hold. For example, let

$$\begin{aligned} \Sigma_1 &:= \{\alpha_1, \alpha_2, \beta_1, \beta_2\} & L_1 &:= \{\alpha_1\beta_1, \alpha_2\beta_2\} \\ \Sigma_2 &:= \{\beta_1, \beta_2, \gamma_1, \gamma_2\} & L_2 &:= \{\beta_1\gamma_2, \beta_2\gamma_1\} \\ \Sigma_3 &:= \{\alpha_1, \alpha_2, \gamma_1, \gamma_2\} & L_3 &:= \{\alpha_1\gamma_1, \alpha_2\gamma_2\} \end{aligned}$$

Then it is not difficult to verify that

$$\begin{aligned} P_{1,2}(L_1) &= P_{2,1}(L_2) = \{\beta_1, \beta_2\} \\ P_{1,3}(L_1) &= P_{3,1}(L_3) = \{\alpha_1, \alpha_2\} \\ P_{2,3}(L_2) &= P_{3,2}(L_3) = \{\gamma_1, \gamma_2\} \end{aligned}$$

So $\{L_1, L_2, L_3\}$ is locally consistent. But it is not globally consistent because $L_1 || L_2 || L_3 = \emptyset$.

It is interesting to know under what condition(s) global consistency is equivalent to local consistency. To this end, we need more technical preparation.

Given $\{\Sigma_i | i \in I\}$ we define a graph $\text{Gr} = \langle \text{Ver}, \text{Edg} \rangle$ where $\text{Ver} = I$ and $\text{Edg} \subseteq \text{Ver} \times \text{Ver}$ such that

$$(\forall i, j \in \text{Ver}) (i, j) \in \text{Edg} \iff i \neq j \ \& \ \Sigma_i \cap \Sigma_j \neq \emptyset$$

Clearly, if $(i, j) \in \text{Edg}$ then so is (j, i) . Thus Gr is an undirected graph. Figure 1 depicts such a graph.

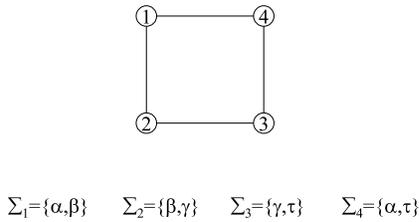


Figure 1. Example of Graph Gr

Proposition 2.2 Given $\mathcal{L} = \{L_i \subseteq \Sigma_i^* | i \in I\}$, if the corresponding graph Gr is a tree as commonly defined in graph theory [5], then \mathcal{L} is globally consistent if and only if it is locally consistent. \square

Proposition 2.1 indicates that local consistency is a necessary condition for global consistency, and Proposition 2.2 indicates that this necessary condition becomes sufficient if the graph Gr is a tree.

3 Supremal Global Support

Given $\mathcal{L} = \{L_i \subseteq \Sigma_i^* | i \in I\}$, a set $\{E_i \subseteq L_i | i \in I\}$ is a *global support* of \mathcal{L} if it is globally consistent, namely

$$(\forall i \in I) E_i = P_{I,i}(\prod_{j \in I} E_j)$$

Let $\Delta(\mathcal{L})$ be the set of all global supports of \mathcal{L} . Clearly $\Delta(\mathcal{L})$ is not empty because it contains $\{E_i = \emptyset | i \in I\}$. We define a binary relation “ \leq ” among elements in $\prod_{i \in I} \text{Pwr}(\Sigma_i^*)$, where $\text{Pwr}(\Sigma_i^*)$ means the power set of Σ_i^* .

$$\{E_i | i \in I\} \leq \{\tilde{E}_i | i \in I\} \iff (\forall i \in I) E_i \subseteq \tilde{E}_i$$

$\mathcal{E} \in \Delta(\mathcal{L})$ is called the *supremal global support* of $\Delta(\mathcal{L})$, written $\text{Sup}\Delta(\mathcal{L})$, if $(\forall \mathcal{E}' \in \Delta(\mathcal{L})) \mathcal{E}' \leq \mathcal{E}$.

Proposition 3.1 $\text{Sup}\Delta(\mathcal{L}) = \{P_{I,i}(\prod_{j \in I} L_j) | i \in I\}$. \square

It is interesting in practice, e.g. for conservative diagnosis, to know how to compute such a supremal global support. Suppose $I = \{1, 2, \dots, n\}$ and we want to compute $P_{I,n}(\prod_{j \in I} L_j)$. We first adopt an ordering on nodes in I which makes n the last node in the ordering. For simplicity, suppose the ordering is $[1, 2, \dots, n]$. The ordering is an important factor for the space complexity of the following computational procedure. Due to limited space, we leave the ordering issue to a full version of this paper, where a heuristic ordering algorithm is given. Now we present our computational procedure, where the notation used below follows the notational rules described in Section 2.1.

Computational Procedure for Global Consistency: (CPGC)

1. Initially set $T_0 := \Sigma_1$ and $W_0 := \Sigma_1^*$.
2. For $k = 1, 2, \dots, n-1$,
 - Set $J_k := \{1, 2, \dots, k\}$ and $T_k := \Sigma_{J_k} \cap \Sigma_{I-J_k}$.
 - $W_k := P_{T_{k-1} \cup \Sigma_k, T_k}(W_{k-1} || L_k)$
3. $E_n := W_{n-1} || L_n$. \square

Proposition 3.2 Let E_n be as above. Then $E_n = P_{I,n}(\prod_{j \in I} L_j)$. \square

To demonstrate CPGC, consider the following example. The system consists of 18 local components, namely $I = \{1, 2, \dots, 17, 18\}$, whose transition structures are depicted in the left side of Figure 2. The right picture shows the system’s network structure. Suppose we

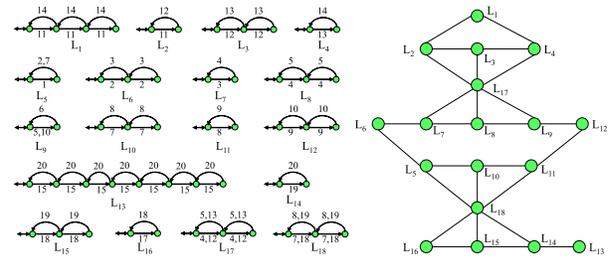


Figure 2. Plant \mathcal{L} (Left) and Graph Gr (Right)

want to compute $E_{18} = P_{I,18}(\prod_{k \in I} L_k)$. We select an arbitrary ordering, say $[1, 2, 3, \dots, 17, 18]$. The computational result of CPGC is shown as follows.

Table 3.1

k	L_k	$W_{k-1} L_k$	$W_k = P_{T_{k-1} \cup \Sigma_k, T_k}(W_{k-1} L_k)$	Null Events
1	[4,6]	[4,6]	[4,6]	\emptyset
2	[2,2]	[8,13]	[3,8]	11
3	[3,4]	[13,30]	[13,30]	\emptyset
4	[2,2]	[6,9]	[3,4]	14
5	[2,3]	[6,17]	[6,17]	\emptyset
6	[3,4]	[18,60]	[3,10]	12
7	[2,2]	[6,20]	[3,10]	3
8	[3,4]	[9,33]	[9,33]	\emptyset
9	[2,3]	[18,78]	[9,42]	6
10	[3,4]	[27,135]	[27,135]	\emptyset
11	[2,2]	[54,279]	[54,279]	\emptyset
12	[3,4]	[162,756]	[27,108]	9,10
13	[8,14]	[216,1242]	[27,135]	15
14	[2,2]	[54,270]	[27,135]	20
15	[3,4]	[81,432]	[81,432]	\emptyset
16	[2,2]	[162,891]	[81,432]	17
17	[3,8]	[54,252]	[9,24]	4,5,12,13
18	[3,8]	[6,12]	-	-

In Table 3.1, the column of “Null Events” represents the set of events that are projected out by $P_{T_{k-1} \cup \Sigma_k, T_k}$ for each $k : 1 \leq k \leq 17$. The entry $[a_1, a_2]$ means that the resultant finite-state automaton has a_1 states and a_2 transitions. From the table we see that our objective E_{18} is realized by an automaton with 6 states and 12 transitions. All results are obtained by using the software CTCT [15]. Columns under the titles $W_{k-1} || L_k$ and W_k describe the size of all intermediate computational results before we obtain E_{18} . The maximum size of intermediate results is 216 states with 1242 transitions. The brute force computation completely fails in this example because the size of the result $||_{k \in I} L_k$ is too big to be handled by our software CTCT, which exhausts memory before it finishes the computation. By that time the computation has already reached 10^6 states and 8×10^6 transitions. If we compare [216, 1242] with $[10^6, 8 \times 10^6]$, then clearly CPGC is much more efficient than the brute-force method. The correctness of E_{18} is guaranteed by Prop. 3.1. In order to obtain $\text{Sup}\Delta(\mathcal{L})$, we need to repeat CPGC for each local component $i \in I$. At this point, we can see that CPGC may entail high time complexity if the system consists of many local components. We will see later that hierarchical computation may come to the rescue.

4 Supremal Local Support

4.1 Concept of Supremal Local Support

In the spirit of section 3, we say that a set $\{E_i \subseteq L_i | i \in I\}$ is a *local support* of \mathcal{L} if it is locally consistent, namely

$$(\forall i, j \in I) P_{i,j}(E_i) = P_{j,i}(E_j)$$

Let $\Gamma(\mathcal{L})$ be the set of all local supports of \mathcal{L} . Clearly $\Gamma(\mathcal{L})$ is not empty because it contains the support $\{E_i = \emptyset | i \in I\}$. Then we have the following result.

Proposition 4.1 $(\exists \mathcal{E} \in \Gamma(\mathcal{L})) (\forall \mathcal{E}' \in \Gamma(\mathcal{L})) \mathcal{E}' \leq \mathcal{E}$ \square

We call \mathcal{E} in Prop. 4.1 the *supremal local support* of $\Gamma(\mathcal{L})$, written $\text{Sup}\Gamma(\mathcal{L})$. By Prop. 2.1 we can derive that $\text{Sup}\Delta(\mathcal{L}) \leq \text{Sup}\Gamma(\mathcal{L})$. In distributed diagnosis, $\text{Sup}\Delta(\mathcal{L})$ leads to a tighter estimate than $\text{Sup}\Gamma(\mathcal{L})$ does. But we have two reasons to be interested in $\text{Sup}\Gamma(\mathcal{L})$. First, mathematically it imposes a much more difficult problem than $\text{Sup}\Delta(\mathcal{L})$ does. Prop. 3.1 defines a concrete map from the initial set \mathcal{L} to $\text{Sup}\Delta(\mathcal{L})$. But we do not have a similar map between \mathcal{L} and $\text{Sup}\Gamma(\mathcal{L})$. Second, from a computational point of view, we will see that we can use a rather *scalable* computational procedure to compute $\text{Sup}\Gamma(\mathcal{L})$, as long as the procedure finitely terminates. But to find a similar scalable procedure for $\text{Sup}\Delta(\mathcal{L})$ is still an open question. In distributed diagnosis, by computing $\text{Sup}\Gamma(\mathcal{L})$ we can find a tradeoff between quality of diagnosis and scalability and robustness of the diagnoser. Next, we introduce a computational procedure for the supremal local support.

4.2 Computing Supremal Local Support

In [13] the authors proposed a scalable computational algorithm that can be used here to achieve the supremal local support defined above. A similar algorithm appeared later in [3] in the form of a Petri net computation. In this paper we revise our algorithm by imposing a binary relation on the nodes of the target network.

Given \mathcal{L} , let $\text{Gr} = (\text{Ver}, \text{Edg})$ be the associated graph. Suppose we pick one node of Ver as the *root* of Gr . Without loss of generality, suppose the root is node 1. From graph theory [5], a *simple path* between two different nodes $i, j \in \text{Ver}$ is a sequence of edges

$$(v_0, v_1), (v_1, v_2), \dots, (v_{k-1}, v_k) \in \text{Edg}$$

such that

$$v_0 = i \ \& \ v_k = j \ \& \ (\forall l, r : 0 \leq l, r \leq k) \ l \neq r \Rightarrow v_l \neq v_r$$

The *length* of such a simple path is k . We define the distance between the root node 1 and a node $v \in \text{Ver}$, written $d(1, v)$, as the *length* of the shortest *path* between v and node 1. We assume that **the graph Gr is connected**. Let $\phi : \text{Ver} \rightarrow \{1, \dots, |\text{Ver}|\}$ be a one-to-one map such that

$$(\forall (i, j) \in \text{Edg}) \ d(1, i) < d(1, j) \Rightarrow \phi(i) < \phi(j)$$

In general there may exist more than one choice for ϕ . Fix one choice of ϕ and define the following binary relation among vertices:

$$(\forall i, j \in \text{Ver}) \ i \downarrow j \iff (i, j) \in \text{Edg} \ \& \ \phi(i) < \phi(j)$$

We call i a *father node* of j and the binary relation \downarrow the *father-son* relation. With the proposed father-son relation \downarrow , we can partition Ver of Gr as follows.

Proposition 4.2 Let $\text{Gr} = (\text{Ver}, \text{Edg})$ be a graph as described above, where $\text{Ver} = I$. Suppose a father-son relation \downarrow among nodes in Ver is given. Then there exists a partition $\{V_0, V_1, \dots, V_k\}$ on Ver such that for each set V_i ($0 \leq i \leq k$) the following two conditions hold,

1. $(\forall j \in V_i) \{r \in I | r \downarrow j\} \subseteq \cup_{m=0}^{i-1} V_m$
2. $(\forall j \in V_i) \{r \in I | j \downarrow r\} \subseteq \cup_{m=i+1}^k V_m$

where $\cup_{m=0}^{-1} V_m := \emptyset$ and $\cup_{m=k+1}^k V_m := \emptyset$. \square

With such a partition we present the following computational procedure to compute the supremal local support $\text{Sup}\Gamma(\mathcal{L})$.

Computational Procedure for Local Consistency: (CPLC)

1. (Initialization) $(\forall i \in I) E_i^0 := L_i$
2. (Centripetal Propagation) At odd round $n \in 2\mathbb{N} + 1$, starting from V_k and ending at V_0 , for each node $i \in V_j$ ($k \geq j \geq 0$) compute

$$E_i^n := \begin{cases} E_i^{n-1} & \text{if } (\nexists r \in I) i \downarrow r \\ E_i^{n-1} \Big| \Big| (||_{r: i \downarrow r} P_{r,i}(E_r^{n-1})) & \text{otherwise} \end{cases}$$

3. (Centrifugal Propagation) At even round $n \in 2\mathbb{N}^+$, starting from V_0 and ending at V_k , for each node $i \in V_j$ ($0 \leq j \leq k$) compute

$$E_i^n := \begin{cases} E_i^{n-1} & \text{if } (\nexists r \in I) r \downarrow i \\ E_i^{n-1} \Big| \Big| (||_{r: r \downarrow i} P_{r,i}(E_r^{n-1})) & \text{otherwise} \end{cases}$$

4. **Termination Condition:** $(\exists n \in \mathbb{N}^+)(\forall i \in I) E_i^n = E_i^{n+1}$ \square

Proposition 4.2 guarantees that in step 2 before we compute E_i^n all those E_r^n with $i \downarrow r$ have been computed already. Similarly in step 3 before we compute E_i^n all those E_r^n with $r \downarrow i$ have been computed. Therefore CPLC is well defined. The main feature of CPLC is its *scalability*, in the sense that when we add a new node to the graph or remove one from the graph, only a few neighboring nodes need to update their local communication protocol, which is usually done by simply changing their list of father nodes and son nodes. As a comparison, in CPGC when we add a new node or remove one, we need to redesign the ordering among nodes. So CPGC is not scalable.

Figure 3 depicts how the CPLC procedure works in a simple example. In this example, the father-son relation is defined as follows:

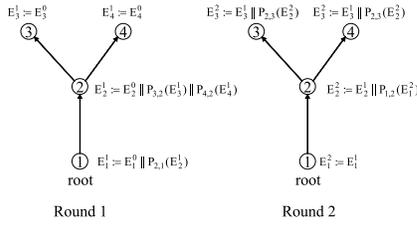


Figure 3. 2-Round CPLC

$$1 \downarrow 2 \text{ and } 2 \downarrow 3 \text{ and } 2 \downarrow 4$$

So the partition is $\{V_0 = \{1\}, V_1 = \{2\}, V_2 = \{3, 4\}\}$. We can see that during the computational process, in odd-number rounds computation starts from nodes in V_2 , then the node in V_1 and finally the node in V_0 . In even-number rounds computation goes in the opposite direction, namely starts from the node in V_0 , then the node in V_1 and finally nodes in V_2 . So the computational process is well organized.

Proposition 4.3 If CPLC terminates at $n \in \mathbb{N}^+$, then we have that $\{E_i^n | i \in I\} = \text{Sup}\Gamma(\mathcal{L})$. \square

So far we have shown in Prop. 4.3 that as long as CPLC finitely terminates, the result is the supremal local support $\text{Sup}\Gamma(\mathcal{L})$. Nevertheless, it turns out that CPLC may not always terminate. So next we briefly discuss termination issue.

4.3 Termination of CPLC

It has been pointed out in [3] that the original algorithm may not always terminate. This is also true of CPLC. It is natural to seek conditions for termination. We have the following results.

Proposition 4.4 If for each $i \in I$, L_i contains a finite number of strings, then CPLC terminates. \square

Proposition 4.5 If Gr is a tree then the termination condition holds at the end of round 2. \square

The advantage of imposing the father-son relation is demonstrated in Prop. 4.5, namely we know the maximum number of rounds that we have to wait before CPLC terminates in a tree. In another paper

on CPLC, we show that the existence of a similar upper bound of rounds in a graph called *simple net* which contains disjoint loops. Such information is very helpful for online diagnosis because we can predict when to report diagnostic result at the latest. The original algorithm proposed in [13] did not have this feature.

So far we have introduced concepts of supremal global support and supremal local support, and provided concrete computational procedures. Nevertheless, we have seen that CPLC may not always terminate and CPGC can only compute the supremal global support in a sequential way, which may result in high time complexity. So next we introduce hierarchical computation, which can partially reduce high time complexity of CPGC.

5 Hierarchical Computation

5.1 Concept of Hierarchical Computation

First, we describe how to construct a two-level hierarchy. Given the index set I , let \hat{I} be a second index set based on a partition $\{J_{\hat{k}} \neq \emptyset | \hat{k} \in \hat{I}\}$ of I . From now on we use i, j, k, \dots for elements of I , and $\hat{i}, \hat{j}, \hat{k}, \dots$ for elements of \hat{I} . Similarly, we use J, K, \dots for subsets of I and \hat{J}, \hat{K}, \dots for subsets of \hat{I} . For each $\hat{k} \in \hat{I}$, recall that by the notational rules we have

$$\Sigma_{J_{\hat{k}}} := \cup_{i \in J_{\hat{k}}} \Sigma_i \text{ and } \Sigma_{I-J_{\hat{k}}} := \cup_{j \in I-J_{\hat{k}}} \Sigma_j$$

We define a new event set $\Sigma_{\hat{k}} \subseteq \Sigma_{J_{\hat{k}}}$ such that

$$\Sigma_{J_{\hat{k}}} \cap \Sigma_{I-J_{\hat{k}}} \subseteq \Sigma_{\hat{k}} \quad (1)$$

In other words, the new event set $\Sigma_{\hat{k}}$ should contain all events that are shared by $\{L_i | i \in J_{\hat{k}}\}$ and $\{L_j | j \in I - J_{\hat{k}}\}$. In fault diagnosis, besides those shared events, $\Sigma_{\hat{k}}$ usually contains only observable events which represent observation in the system.

As an illustration, let $I = \{1, 2, 3, 4\}$ and suppose $\Sigma_1 = \{\alpha, \beta\}$, $\Sigma_2 = \{\beta, \gamma\}$, $\Sigma_3 = \{\sigma, \gamma, \delta\}$ and $\Sigma_4 = \{\alpha, \delta\}$. Suppose we partition I into two subsets $J_a = \{1, 2\}$ and $J_b = \{3, 4\}$ with $\hat{I} = \{a, b\}$. Then

$$\Sigma_{J_a} = \{\alpha, \beta, \gamma\} \text{ and } \Sigma_{J_b} = \{\alpha, \gamma, \delta, \sigma\}$$

So $\Sigma_{J_a} \cap \Sigma_{I-J_a} = \Sigma_{J_a} \cap \Sigma_{J_b} = \{\alpha, \gamma\}$. To define a new set $\Sigma_a \subseteq \Sigma_{J_a}$ such that (1) holds, we can pick $\Sigma_a = \{\alpha, \gamma\}$ or $\Sigma_a = \{\alpha, \gamma, \beta\}$.

With the same notational rules for natural projections, for each $\hat{k} \in \hat{I}$ we define a new language on $\Sigma_{\hat{k}}^*$ as follows:

$$L_{\hat{k}} := P_{J_{\hat{k}}, \hat{k}}(\cup_{i \in J_{\hat{k}}} L_i) \quad (2)$$

Thus we have a new set of languages $\hat{\mathcal{L}} = \{L_{\hat{k}} | \hat{k} \in \hat{I}\}$ which serves as the abstraction of \mathcal{L} . Each element of $\hat{\mathcal{L}}$ is called an *abstract local module* of \mathcal{L} . The relation between \mathcal{L} and $\hat{\mathcal{L}}$ can be described as follows:

$$\begin{aligned} P_{I, \hat{I}}(\cup_{i \in I} L_i) &= P_{I, \hat{I}}(\cup_{\hat{k} \in \hat{I}} (\cup_{i \in J_{\hat{k}}} L_i)) \\ &= \cup_{\hat{k} \in \hat{I}} P_{J_{\hat{k}}, \hat{k}}(\cup_{i \in J_{\hat{k}}} L_i) \\ &= \cup_{\hat{k} \in \hat{I}} L_{\hat{k}} \end{aligned}$$

Thus we finally have

$$P_{I, \hat{I}}(\cup_{i \in I} L_i) = \cup_{\hat{k} \in \hat{I}} L_{\hat{k}} \quad (3)$$

We can define the supremal global support of $\Delta(\hat{\mathcal{L}})$ in the same way as for $\Delta(\mathcal{L})$. Then by Prop. 3.1 we have

$$\{E_i := P_{\hat{I}, \hat{i}}(\|\hat{k} \in \hat{I} L_k\)|\hat{i} \in \hat{I}\} = \text{Sup}\Delta(\hat{\mathcal{L}})$$

From (3) we can derive the relation between $\text{Sup}\Delta(\mathcal{L}) = \{E_j | j \in I\}$ and $\text{Sup}\Delta(\hat{\mathcal{L}}) = \{E_i | i \in \hat{I}\}$ as follows:

$$(\forall \hat{i} \in \hat{I}) E_i = P_{\hat{I}, \hat{i}}(\|_{j \in I} E_j) \quad (4)$$

Let

$$\mathbb{W} := \prod_{i \in I} \text{Pwr}(\Sigma_i^*) \text{ and } \hat{\mathbb{W}} := \prod_{\hat{k} \in \hat{I}} \text{Pwr}(\Sigma_{\hat{k}}^*)$$

Everything we have done so far in this section can be summarized in Figure 4, where the maps in the diagram are defined as follows:

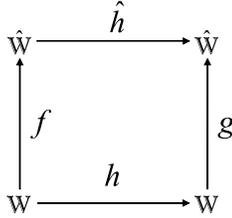


Figure 4. Two-Level Hierarchy

$$1. f : \mathbb{W} \rightarrow \hat{\mathbb{W}},$$

$$f(\{L_i | i \in I\}) := \{L_{\hat{k}} = P_{J_{\hat{k}}, \hat{k}}(\|_{j \in J_{\hat{k}}} L_j) | \hat{k} \in \hat{I}\}$$

$$2. g : \mathbb{W} \rightarrow \hat{\mathbb{W}},$$

$$g(\{E_i | i \in I\}) := \{E_{\hat{k}} = P_{\hat{I}, \hat{k}}(\|_{j \in I} E_j) | \hat{k} \in \hat{I}\}$$

$$3. h : \mathbb{W} \rightarrow \mathbb{W} : \mathcal{L} \mapsto h(\mathcal{L}) := \text{Sup}\Delta(\mathcal{L}).$$

$$4. \hat{h} : \hat{\mathbb{W}} \rightarrow \hat{\mathbb{W}} : \hat{\mathcal{L}} \mapsto \hat{h}(\hat{\mathcal{L}}) := \text{Sup}\Delta(\hat{\mathcal{L}}).$$

By expressions (2) and (4) we can see that the above diagram commutes, which tells us how to obtain the high-level structure from the one at low-level. Nevertheless, it is more interesting in practice to know how to obtain the supremal global support $\{E_i | i \in I\}$ of $\Delta(\mathcal{L})$ at the low level if the supremal global support $\{E_{\hat{k}} | \hat{k} \in \hat{I}\}$ of $\Delta(\hat{\mathcal{L}})$ at high level is given. To this end we have the following result.

Proposition 5.1 Let $\mathcal{L} = \{L_i | i \in I\} \in \mathbb{W}$, $\hat{\mathcal{L}} = \{L_{\hat{k}} | \hat{k} \in \hat{I}\} \in \hat{\mathbb{W}}$. Let f be defined as in Figure 4. Suppose $f(\mathcal{L}) = \hat{\mathcal{L}}$. If we have $\{E_i | i \in I\} = \text{Sup}\Delta(\mathcal{L})$ and $\{E_{\hat{k}} | \hat{k} \in \hat{I}\} = \text{Sup}\Delta(\hat{\mathcal{L}})$, then we get that $(\forall \hat{k} \in \hat{I})(\forall i \in J_{\hat{k}}) E_i = P_{J_{\hat{k}}, i}(\|_{j \in J_{\hat{k}}} L_j) | E_{\hat{k}}$. \square

By Prop. 3.1 we have

$$E_i = P_{I, i}(\|_{j \in I} L_j)$$

Since computation is over I , if the size of I is large, then the time complexity may be a problem. Nevertheless, Proposition 5.1 tells us that if we can achieve the supremal global support $E_{\hat{k}}$ at the high level, where $\hat{k} \in \hat{I}$ and $i \in J_{\hat{k}}$, then we only need the module $\{L_j | j \in J_{\hat{k}}\}$ plus $\{E_{\hat{k}}\}$ to achieve E_i . Notice that $J_{\hat{k}}$ is usually much smaller than I in a large scale system, and furthermore, concurrent

computation can be applied on different local modules which are disjoint from each other. Both factors contribute to low time complexity of hierarchical computation. This is the main advantage gained from hierarchical computation. Due to the nature of abstraction, achieving the supremal global support in the abstract model at the high level is relatively easy. Prop. 5.1 has crucial application in multi-resolution diagnosis. It is not difficult to see that the current theory can be extended to a hierarchy with more levels such that, between every two adjacent levels we have a commutative diagram as Figure 4 and Proposition 5.1 holds. To demonstrate the hierarchical computational procedure, consider the following example.

5.2 Example of Hierarchical Computation

Suppose we have four languages L_1, \dots, L_4 which are depicted in Figure 5 with the corresponding graph Gr. The index set $I = \{1, 2, 3, 4\}$. Suppose we partition I into two sets $J_a = \{1, 2\}$ and

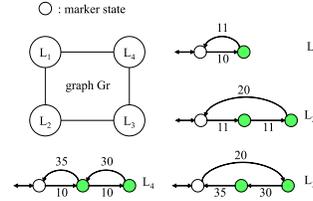


Figure 5. Plant Model and Network Graph

$J_b = \{3, 4\}$, where $\hat{I} = \{a, b\}$. Suppose we assign new event sets at the high level based on expression (1) as follows:

$$\Sigma_a := \{10, 20\} \text{ and } \Sigma_b := \{10, 20, 35\}$$

Then we can compute abstract local modules based on expression (2) as follows:

$$L_a := P_{J_a, a}(L_1 || L_2) \text{ and } L_b := P_{J_b, b}(L_3 || L_4)$$

where the results are shown in Figure 6. After that we can compute

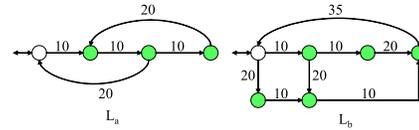


Figure 6. Construction of $\hat{\mathcal{L}} = \{L_a, L_b\}$

the supremal global support at the high level according to

$$E_a := P_{\hat{I}, a}(L_a || L_b) \text{ and } E_b := P_{\hat{I}, b}(L_a || L_b)$$

where the results are depicted in Figure 7. After we get E_a and E_b we can compute the supremal global support at the bottom level. Suppose we want to compute E_1 and E_3 , which are shown as follows:

$$E_1 = P_{J_a, 1}(L_1 || L_2 | E_a) \text{ and } E_3 = P_{J_b, 3}(L_3 || L_4 | E_b)$$

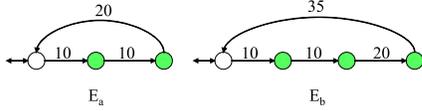


Figure 7. $\text{Sup}\Delta(\hat{\mathcal{L}}) = \{E_a, E_b\}$

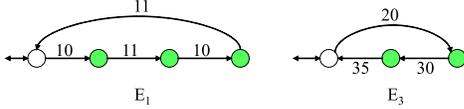


Figure 8. Results of E_1 and E_3

The results are depicted in Figure 8. Notice that E_1 and E_3 can be computed simultaneously because they belong to different modules. The results are verified to be isomorphic to the ones obtained by the brute-force method, namely

$$E_1 := P_{I,1}(L_1 || \dots || L_4) \text{ and } E_3 := P_{I,3}(L_1 || \dots || L_4)$$

The computation demonstrated above is realized by the software CTCT [15].

6 Conclusions

In this paper we introduced two main concepts: the supremal global support and the supremal local support. In distributed diagnosis each type of supremal support can be used to model consistency among local diagnoses. Since global consistency is stronger than local consistency, applying the supremal global support in diagnosis usually results in a tighter estimate about fault status of a target system than applying the supremal local support does. Nevertheless, we can see that the computational procedure CPLC for achieving the supremal local support is much more scalable than the computational procedure CPGC for achieving the supremal global support, as long as CPLC terminates. At this point, the supremal local support, which only considers neighborhood information of each local component, offers a tradeoff between diagnosis quality, in terms of the number of fault candidates contained in each local diagnosis, and scalability and robustness of the distributed diagnoser. To the best of our knowledge, CPLC first appeared in [13] then in [3] a similar algorithm was given in the form of a Petri net computation. In this paper we improved it further by imposing the father-son relation \downarrow among nodes. With this improved CPLC, we can tell in the worst case, how many rounds a user has to wait before termination, if indeed it terminates. Nevertheless, we admit that with a general network graph, the decidability of the termination of CPLC is still unsolved. As for CPGC, from the example we can see that it is superior to the brute-force method as far as space complexity is concerned. Nevertheless, since it is a purely sequential computational procedure, if a system consists of many local components, CPGC may take a long time to compute the supremal global support. To reduce time complexity of CPGC, we introduced hierarchical computation, which achieves the supremal global support at the high level first, then descends to the low level to achieve the supremal global support at the low level. CPGC

can be used at each single level during hierarchical computation. By Prop. 5.1 we can see that when computation comes to the low level, concurrent computation can be done in each disjoint local abstract module. This is another application of divide-and-conquer. In online diagnosis, this may substantially reduce time complexity. Nevertheless, we admit that we have to pay the price of using extra memory to store the high-level abstract model. A rigorous space complexity analysis remains to be carried out. But a brief informal analysis will be given in the full version paper. Our experience shows that, if we consider both static memory to store the hierarchical distributed model and dynamic memory to store the intermediate computational results, the space complexity of the hierarchical method is no worse, and in many cases much better, than the non-hierarchical method.

REFERENCES

- [1] A. Aghasaryaiu, E. Fabre, A. Benveniste, R. Boubour, and C. Jard, 'A petri net approach to fault detection and diagnosis in distributed systems', in *Proc. 1997 IEEE Conference on Decision and Control (CDC'97)*, pp. 720–725, December 1997.
- [2] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanella, 'Diagnosis of large active systems', *Artificial Intelligence*, **110**(1), 135–183, May 1999.
- [3] A. Benveniste, S. Haar, E. Fabre, and C. Jard, 'Distributed monitoring of concurrent and asynchronous systems', in *Proc. ATPN-Workshop on Discrete Event Systems Control*, pp. 97–142, Eindhoven, The Netherlands, June 2003.
- [4] R. Debouk, S. Lafortune, and D. Teneketzis, 'Coordinated decentralized protocols for failure diagnosis of discrete event systems', *Discrete Event Dynamic Systems: Theory and Applications*, **10**(1/2), 33–86, January 2000.
- [5] Reinhard Diestel, *Graph Theory*, Springer-Verlag, New York, 2nd edn., 2000.
- [6] E. Fabre, A. Benveniste, and C. Jard, 'Distributed diagnosis for large discrete event dynamic systems', in *Proc. 15th IFAC World Congress*, Barcelona, Spain, July 2002.
- [7] Gianfranco Lamperti, Marina Zanella, and Paolo Pogliano, 'Diagnosis of active systems by automata-based reasoning techniques', *Artificial Intelligence*, **12**(3), 217–237, May 2000.
- [8] Yannick Pencol, Marie-Odile Cordier, and Laurence Roz, 'Incremental decentralized diagnosis approach for the supervision of a telecommunication network', in *Proc. 2002 IEEE Conference on Decision and Control (CDC'02)*, pp. 435–440, Las Vegas, USA, December 2002.
- [9] Gregory Provan, 'A model-based diagnosis framework for distributed systems', in *Proc. of International Workshop on Principles of Diagnosis (DX'02)*, pp. 16–25, Austria, May 2002.
- [10] M. Sampath, S. Lafortune, and D. Teneketzis, 'Active diagnosis of discrete-event systems', *IEEE Transactions on Automatic Control*, **40**, 908–929, July 1998.
- [11] R. Sengupta, 'Diagnosis and communication in distributed systems', in *Proc. International Workshop on Discrete Event Systems (WODES'98)*, pp. 144–151, IEE, London, August 1998.
- [12] R. Su and W.M. Wonham, 'Decentralized fault diagnosis for discrete-event systems', in *Proc. 2000 CISS*, pp. TP1:1–6, Princeton, New Jersey, March 2000.
- [13] R. Su, W.M. Wonham, J. Kurien, and X. Koutsoukos, 'Distributed diagnosis for qualitative systems', in *Proc. 6th International Workshop on Discrete Event Systems (WODES'02)*, pp. 169–174, Zaragoza, Spain, October 2002.
- [14] W. M. Wonham, *Notes on Control of Discrete-Event Systems: ECE 1636F/1637S 2003-2004*, Systems Control Group, Dept. ECE, University of Toronto. URL: www.control.utoronto.ca/DES.html, 2003.
- [15] W.M. Wonham, *Software Package for Supervisory Control*, Systems Control Group, Dept. of ECE, University of Toronto. URL: www.control.utoronto.ca/DES.html, 2003.
- [16] S. Hashtrudi Zad, R.H. Kwong, and W.M. Wonham, 'Fault diagnosis in discrete-event systems', in *Proc. 1998 IEEE Conference on Decision and Control (CDC'98)*, pp. 3769–3774, Tampa, Florida, USA, December 1998.