

Distributed Diagnosis Using a Condensed Representation of Diagnoses With Application to an Automotive Vehicle

Jonas Biteus, Erik Frisk and Mattias Nyberg

Linköping University Post Print

N.B.: When citing this work, cite the original article.

©2011 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Jonas Biteus, Erik Frisk and Mattias Nyberg, Distributed Diagnosis Using a Condensed Representation of Diagnoses With Application to an Automotive Vehicle, 2011, IEEE transactions on systems, man and cybernetics. Part A. Systems and humans, (41), 6, 1262-1267.

<http://dx.doi.org/10.1109/TSMCA.2011.2147311>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-72015>

Distributed Diagnosis using a Condensed Representation of Diagnoses with application to an Automotive Vehicle

Jonas Biteus*, Erik Frisk, and Mattias Nyberg

Abstract—In fault detection and isolation, diagnostic test results are commonly used to compute a set of diagnoses, where each diagnosis points at a set of components which might behave abnormally. In distributed systems consisting of multiple control units, the test results in each unit can be used to compute local diagnoses while all test results in the complete system give the global diagnoses. It is an advantage for both repair and for fault tolerant control to have access to the global diagnoses in each unit since these diagnoses represent all test results in all units. However, when the diagnoses, for example, are to be used to repair a unit, only the components that are used by the unit are of interest. The reason for this is that it is only these components that can have caused the abnormal behavior. However, the global diagnoses might include components from the complete system, and therefore often include components that are superfluous for the unit. Motivated by this observation, a new type of diagnosis is proposed, the condensed diagnosis. Each unit has a unique set of condensed diagnoses which represents the global diagnoses. The benefit of the condensed diagnoses is that they only include components used by the unit while still representing the global diagnoses. The proposed method is applied to an automotive vehicle and the results from the application study show the benefit of using condensed diagnoses compared to global diagnoses.

Index Terms—Fault diagnosis, Fault location, Road vehicles, Distributed algorithms

I. INTRODUCTION

Fault diagnosis is becoming more common in many industrial applications, and one of the most widespread approaches for diagnosis is the *consistency based diagnosis* approach developed within the AI field [1]. In this approach, a *minimal diagnosis* is a minimal set of *components* whose abnormal behavior is a possible explanation to why the system is faulty. Such minimal diagnoses are commonly used for both repair and fault tolerant control.

This paper considers fault diagnosis for *distributed systems* that consist of a set of *agents* [2], [3], [4]. In such distributed systems, the *global diagnoses* are diagnoses for the complete distributed system and can be computed from all *diagnostic test* results in all agents. Further, *local diagnoses* can be computed from the diagnostic test results via one agent. The

diagnoses can, for example, be used when the agent is to be repaired, by guiding the repair technician to the faulty components.

A drawback with minimal local diagnoses is that they might not be minimal considering the complete distributed system. Therefore, the global minimal diagnoses are preferred. However, a drawback when using minimal global diagnoses in an agent is that they may include many components that are not used and therefore of no interest for the agent. When performing, for example, repair of the agent, all components not used by the agent are superfluous since only the components used by the agent could have caused the abnormal behavior of the agent. Besides being superfluous, the inclusion of these components leads to an unnecessarily high usage of memory and processing power, and thereby leads to unnecessary costs.

Motivated by these observations, the contribution of the paper is a new type of diagnosis, denoted as *condensed diagnosis*, which represent the set of minimal global diagnoses. Each agent has a unique set of minimal condensed diagnoses where all components not used by the agent have been removed. To ensure that the condensed diagnoses represent the global diagnoses, a scalar variable is included in each condensed diagnosis. The value of the variable equals the number of removed components. Both the memory needed to store and the processing power needed to compute the sets of minimal condensed diagnoses are significantly reduced compared to the case where the set of minimal global diagnoses is used. However, in contrast to the minimal local diagnoses, which are also significantly smaller than the minimal global diagnoses, the minimal condensed diagnoses still represent the set of global diagnoses. To compute the condensed diagnoses, we develop an algorithm which efficiently computes a set of minimal condensed diagnoses in each agent. The benefits of the condensed diagnoses are illustrated for the distributed system in an automotive vehicle.

As indicated by the application study, our work has been inspired by diagnostic systems used in automotive vehicles [5], [6] that typically contain ECUs with both limited processing power and limited RAM memory. Another important characteristic of these distributed systems is that agents can be disconnected, replaced, or added to the system without notification. To gain a scalable [7] distributed system, the algorithm has been designed such that, for each agent, it computes the set of minimal condensed diagnoses in a distributed cooperation between the agents.

* Corresponding author.

1. Manuscript received December 22, 2006.

2. Acknowledgment: This work was funded by the Swedish Foundation for Strategic Research and Scania CV AB.

3. J. Biteus, e-mail jonas.biteus@scania.com, and M. Nyberg, e-mail mattias.nyberg@scania.com, are with the Systems architecture division at heavy duty truck manufacturer Scania CV in Sweden. E. Frisk, e-mail frisk@isy.liu.se, is with the Electrical Engineering department at Linköping universitet in Sweden.

Related Work

In contrast to the distributed algorithm presented in this paper, most research on fault isolation, such as [8], [9], has been aimed at the centralized diagnosis problem. These centralized methods can also be used for distributed systems by letting a single diagnostic agent collect all diagnostic test results from all agents and then compute the minimal global diagnoses. However, as noted above, it is not always suitable to use such a dedicated single diagnostic agent. For distributed systems, there exist algorithms, for example those presented in [10], [11], which are aimed at computing the set of minimal global diagnoses in a distributed cooperation between the agents. The methods in these previous papers differ from the method presented here in that they aim at computing the complete set of minimal global diagnoses. There also exist algorithms where the agents update the sets of signals transmitted between the agents, such that the set of minimal local diagnoses in each agent is consistent with all signals available on the network, see for example [12]. However, these algorithms assume that the set of components is defined such that each component has an output and an input signal, an assumption which is not made in this paper.

II. CONSISTENCY BASED DIAGNOSIS

A system consists of a set of components \mathcal{C} which are to be supervised for abnormal behavior. A component is an object that can be diagnosed, such as sensors, actuators, cables, and pipes, and it can be in different behavioral modes. Here, only the abnormal and the normal modes are considered, where the abnormal mode does not have a model.

Within consistency based diagnosis, a diagnosis D is a set of components whose abnormal behaviors is a consistent explanation to the diagnostic test results, see for example [1]. Further, a minimal diagnosis is a minimal such explanation. If a diagnostic test responds, i.e. if it detects that any component it supervises is behaving abnormally, then a conflict π is generated, where π is a set consisting of the supervised components. Similarly to minimal diagnoses, a minimal conflict is a minimal such set of supervised abnormal components. The minimal diagnoses can be computed based on the set of minimal conflicts.

III. DISTRIBUTED DIAGNOSIS

This section will first present a framework for distributed systems consisting of agents, components, and signals. The framework will then be used to define condensed diagnosis.

A. Relation Between Local and Global Diagnoses

A distributed system consists of a set of agents $\mathcal{A} = \{A_1, \dots, A_n\}$ where each agent includes a diagnostic system. In such distributed systems, a local diagnosis is a diagnosis that is determined by the set of conflicts in one agent, while a global diagnosis is determined by all conflicts in all agents. The conflicts are typically generated from responded diagnostic tests. Here, a set of minimal local diagnoses in agent A_i is denoted by \mathbb{D}^{A_i} and a set of minimal global diagnoses is denoted by \mathcal{D} .

The set of global diagnoses can be computed from the sets of minimal conflicts in all agents. However, the minimal global diagnoses can also be computed by extracting the minimal diagnoses in the set resulting from a merge of all sets of minimal local diagnoses.

This idea of computing the minimal global diagnoses from the sets of minimal local diagnoses will be used in the algorithm which computes the set of minimal condensed diagnoses in each agent.

B. Signals and Components in Distributed Systems

A condensed diagnosis in one agent should only include the components that the agent uses. To be able to decide which components that an agent uses, the components are here partitioned into *private components* $\mathcal{P} \subseteq \mathcal{C}$ and *common components* $\mathcal{G} \subseteq \mathcal{C}$. A private component is only used by one agent, while a common component is used by two or more agents. The set of private components is therefore further partitioned into different sets belonging to different agents, where the set $\mathcal{P}^{A_i} \subseteq \mathcal{P}$ is used by agent A_i .

An agent in a distributed system could, in addition to the components, also use signals available over the network. A signal is typically a value from a sensor, to an actuator, or some computed value. Here \mathcal{S} is the set of signals, $\mathcal{S}_{in}^{A_i} \subseteq \mathcal{S}$ is the subset of signals which are used by agent A_i and $\mathcal{S}_{out}^{A_i} \subseteq \mathcal{S}$ is the subset of signals which are outputs from agent A_i .

C. Signals Depending on Components

The normal behavior of a signal s depends on the normal behavior of all the components in a set $\text{Dep}(s) \subseteq \mathcal{C}$. If any of the components in the set $\text{Dep}(s)$ behaves abnormally then s is assumed to also behave abnormally.

To make the algorithm for the computation of the minimal condensed diagnoses more readable, some special properties of the dependencies will be assumed.

Assumption 1: a) The dependency for a signal s which is an output from agent A_i is limited such that $\text{Dep}(s) \subseteq \mathcal{P}^{A_i}$. b) The dependencies for two different signals are disjoint.

This assumption makes the algorithm more readable, but, as shown in [13], it is possible to compute the minimal condensed diagnoses even if Assumption 1 is not fulfilled. Besides making the algorithm more readable, it is also the case that Assumption 1 is often fulfilled in industrial applications by construction.

Since the diagnostic system in an agent could supervise both components and signals, a diagnosis $D \subseteq \mathcal{C}$, as defined in Section II, has to be extended to the set of components and signals. However, a diagnosis including signals can always be propagated to the set of components by replacing the signals with their dependencies. Due to this propagation, the word diagnosis will here be used both for a true diagnosis and a diagnosis including both components and signals.

IV. CONDENSED DIAGNOSES REPRESENTING GLOBAL DIAGNOSES

In this section, a minimal condensed diagnosis will be defined with respect to the set of minimal global diagnoses.

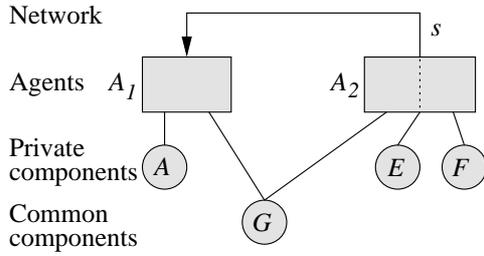


Fig. 1. An example of a system consisting of two agents, four components A , E , F , and G , and one signal s . The agents use the components connected with lines.

A. Relation Between Global and Condensed Diagnoses

A condensed diagnosis should include the components and inputs that an agent A_i use, which here are the common components \mathcal{G} , the private components for the specific agent \mathcal{P}^{A_i} , and the signals which are not outputs from the specific agent $\mathcal{S} \setminus \mathcal{S}_{out}^{A_i}$. For a global diagnosis, the unused components should be removed to gain a condensed diagnosis. However, to gain a correct cardinality of the condensed diagnosis, the information that additional components did exist in the global diagnosis has to be preserved. Here, the condensed diagnosis will therefore be represented by a tuple $\langle D, k \rangle$, where the set D includes the used components and signals, and the scalar variable k is the number of removed components. The variable k is important since it preserves the cardinality of the global diagnoses.

B. Condensed Diagnosis

A condensed diagnosis is formally defined as follows.

Definition 1 (Condensed diagnosis): Let \mathcal{D} be a set of minimal global diagnoses, where, for each diagnosis $\bar{D} \in \mathcal{D}$, $\bar{D} \subseteq \mathcal{C}$. The tuple $\langle D, k \rangle$, where $D \subseteq \mathcal{P}^{A_i} \cup \mathcal{G} \cup (\mathcal{S} \setminus \mathcal{S}_{out}^{A_i})$ and $k \in \mathbb{Z}$, is a condensed diagnosis in agent A_i if there exists $\bar{D} \in \mathcal{D}$ such that

- $|D| + k = |\bar{D}|$
- $D \cap \mathcal{P} = \bar{D} \cap \mathcal{P}^{A_i}$
- $D \cap \mathcal{G} = \bar{D} \cap \mathcal{G}$
- $D \cap \mathcal{S} = \{s \in \mathcal{S} \setminus \mathcal{S}_{out}^{A_i} : \text{Dep}(s) \cap \bar{D} \neq \emptyset\}$.

Interpretation of the different requirements for a condensed diagnosis: a) means that the cardinality of D plus k should equal the cardinality of the global diagnosis \bar{D} ; b) means that D should only include private components used by agent A_i ; c) means that all common components should be included; d) means that signals, that might be abnormal due to the dependency on some abnormal component, should be included in D .

Example 1: Consider the system shown in Fig. 1. There exists a signal s whose dependency is $\text{Dep}(s) = \{E\}$, represented by the dotted line. The sets of private components are $\mathcal{P}^{A_1} = \{A\}$ and $\mathcal{P}^{A_2} = \{E, F\}$. The set of common components is $\mathcal{G} = \{G\}$. The components connected by lines are used by the corresponding agent. Let $\{A, E, F, G\}$ be a minimal global diagnosis. A condensed diagnosis in agent A_1 is $\langle \{A, G, s\}, 1 \rangle$, where component A is included since it is

a private component in A_1 , G since it is a common component, and s since it depends on the abnormal component E . Component F is represented by $k = 1$.

Agent A_2 is not interested in the same components as agent A_1 , therefore, in this agent, the condensed diagnosis $\langle \{E, F, G\}, 1 \rangle$ represents the global diagnosis. Here, component A has been removed and is represented by $k = 1$. \diamond

Why should a condensed diagnosis include the signals instead of propagating these signals to the components on which they depends? The reason for this is that a signal might depend on many components which are not directly used by the agent that use the signal. For repair, signals are typically available and can be checked by a technician for correctness. Since it might be more or less difficult to check all components that a signal depends on, it is an advantage to focus on the signal itself.

C. Minimal Condensed Diagnosis

Similar to minimal diagnoses, the minimal condensed diagnoses can be used to represent all condensed diagnoses.

Definition 2 (Minimal condensed diagnosis): A condensed diagnosis $t = \langle D, k \rangle$ is a *minimal condensed diagnosis* if there exists no condensed diagnosis $t' = \langle D', k' \rangle$ such that $D' \subset D$ and $|D'| + k' < |D| + k$.

The set of minimal condensed diagnoses in one agent represents the set of minimal global diagnoses.

V. AN ALGORITHM FOR THE COMPUTATION OF THE SETS OF MINIMAL CONDENSED DIAGNOSES

The outline of the algorithm is that each agent first computes its set of minimal local diagnoses, and then transmits, to all other agents, the minimal local diagnoses that include components used by other agents. After this, each agent merges the received sets with its own set of minimal local diagnoses which results in the set of minimal condensed diagnoses.

Example 2: Consider two agents A_1 and A_2 . Agent A_1 has a minimal local diagnosis $\{A, G\}$ where A is a private component and G is a common component. Agent A_2 has a minimal local diagnosis $\{E, G\}$ where E is a private component. To compute the minimal condensed diagnoses, the algorithm transmits the tuple $\langle \{G\}, 1 \rangle$ from agent A_1 to agent A_2 . This tuple includes the part of the local diagnosis in agent A_1 which is used by the other agent. This transmitted tuple is then merged with the minimal local diagnosis $\{E, G\}$ in A_2 which give the minimal condensed diagnosis $\langle \{E, G\}, 1 \rangle$. As should be, this coincides with the minimal condensed diagnosis computed from the minimal global diagnosis $\{A, E, G\}$. \diamond

The transmitting procedure is described in detail in Section V-A, and the receiving and merging process is described in Section V-B. Finally, the main algorithm is described in Section V-C. In the algorithms, D is a diagnosis, $\Sigma \subseteq \mathcal{S}$, $\Omega \subseteq \mathcal{S}_{out}^{A_i}$, $P \subseteq \mathcal{P}$, and finally $G \subseteq \mathcal{G}$.

Algorithm 1 – $\text{Transmit}(\mathbb{D}^{A_i})$. Transmit subsets of minimal local diagnoses.

Input: A set of minimal local diagnoses \mathbb{D}^{A_i} .

Output: Set of tuples TX^{A_i} representing the diagnoses in \mathbb{D}^{A_i} that should be transmitted.

```

1:  $\mathbb{D}^{TX} := \{D \in \mathbb{D}^{A_i} : D \cap (\mathcal{S} \cup \mathcal{G} \cup (\cup_{s \in \mathcal{S}} \text{Dep}(s))) \neq \emptyset\}$ 
2:  $TX^{A_i} := \emptyset$ 
3: for each  $D = P \cup G \cup \Sigma \in \mathbb{D}^{TX}$  do
4:    $\Omega := \{s \in \mathcal{S} : \text{Dep}(s) \cap P \neq \emptyset\}$ 
5:    $k := |P| - |\Omega|$ 
6:    $TX^{A_i} := TX^{A_i} \cup \{\langle G \cup \Sigma \cup \Omega, k \rangle\}$ 
7: end for
8:  $TX^{A_i} := \text{RemoveNonMinimal}(TX^{A_i})$ 
9: if  $\mathbb{D}^{A_i} \setminus \mathbb{D}^{TX} \neq \emptyset$  then
10:   $k := \min_{D \in \mathbb{D}^{A_i} \setminus \mathbb{D}^{TX}} |D|$ 
11:   $TX^{A_i} := TX^{A_i} \cup \{\langle \emptyset, k \rangle\}$ 
12: end if

```

A. Transmit Subsets of Local Diagnoses to Other Agents

Using the definition of condensed diagnosis, it can be seen that a minimal local diagnosis in one agent is of interest for the other agents if it includes components or signals used by some other agent or if a component is included that some signal depends on. The local diagnoses including such components or signals should be transmitted to the other agents. However, before the diagnoses are transmitted, the private components can be removed since these are not used by any other agent.

The minimal local diagnoses that are not transmitted to the other agents can be represented by a tuple including the empty diagnosis and the variable k whose value is the minimal cardinality of any of the not transmitted minimal local diagnoses, i.e. a tuple $\langle \emptyset, k \rangle$. The agents receiving this tuple will then be aware of that there exist one or more non-transmitted minimal local diagnoses with cardinality k or higher.

Algorithm 1 performs the steps described above. Row 1 decides which minimal local diagnoses that include components or signals used by other agents, resulting in the set \mathbb{D}^{TX} . Rows 2–7 construct a tuple for each local diagnosis $D = P \cup G \cup \Sigma$ in the set \mathbb{D}^{TX} , where the set P is the private components, G is the common components, and Σ is the signals included in the local diagnosis. Each tuple in the set TX^{A_i} includes a set consisting of G and Σ , which are taken directly from the local diagnosis, and the set Ω which consists of the signals which depend on any of the removed private components in P . Each tuple also includes the variable k , which equals the number of removed private components in P minus the number of added signals in Ω . In row 8, the function $\text{RemoveNonMinimal}(\cdot)$ removes all non-minimal tuples, where a minimal tuple is defined in the same way as a minimal condensed diagnosis, see Definition 2. Finally, rows 9–12 add a tuple representing the not transmitted minimal local diagnoses.

Example 3: Consider the system shown in Fig. 2, where the sets of private components are $\mathcal{P}^{A_1} = \{A, B, C\}$, $\mathcal{P}^{A_2} = \{D\}$, and $\mathcal{P}^{A_3} = \{E, F\}$, while the set of common compo-

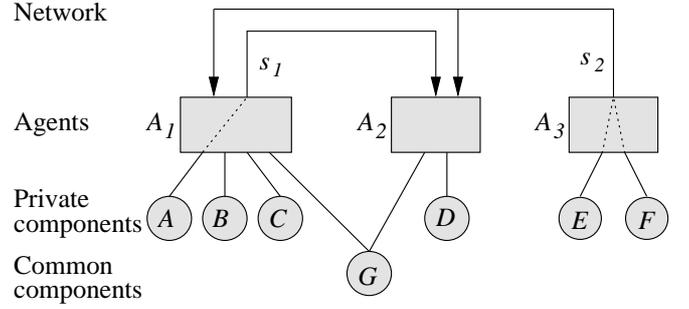


Fig. 2. An example of a system consisting of three agents, seven components, and two signals. The components connected with lines are used by the corresponding agents.

nents is $\mathcal{G} = \{G\}$. There exist two signals with dependencies $\text{Dep}(s_1) = \{A\}$ and $\text{Dep}(s_2) = \{E, F\}$. Assume that the following set of minimal local diagnoses has been computed in agent A_1

$$\mathbb{D}^{A_1} = \{\{B, s_2\}, \{A, B\}, \{G, B\}, \{C\}\}.$$

Using Algorithm 1, the private components are removed and the set of tuples to be transmitted is thereby

$$TX^{A_1} = \{\langle \{s_2\}, 1 \rangle, \langle \{s_1\}, 1 \rangle, \langle \{G\}, 1 \rangle, \langle \emptyset, 1 \rangle\}$$

where the tuple $\langle \emptyset, 1 \rangle$ represents the non-transmitted diagnosis $\{C\}$. The set TX^{A_1} will represent the set of local diagnoses in A_1 when the agents A_2 and A_3 are computing their unique set of minimal condensed diagnoses. \diamond

B. Receive and Merge the Transmitted sets of Tuples

The second step when calculating the sets of minimal condensed diagnoses is for each agent to receive the sets of tuples transmitted from the other agents, transform them into an appropriate form, and then compute the minimal condensed diagnoses. If a received tuple includes a signal s that is an output from the receiving agent then the receiver is interested in knowing which private components that could have caused the abnormal behavior of the signal. The signal is therefore replaced with the components which it depends on. After all such signals have been replaced by their dependencies, the minimal local diagnoses in the receiving agent and the sets of received tuples are merged which results in the set of minimal condensed diagnoses.

Algorithm 2 performs the steps described above. Rows 1–9 transform each set of transmitted tuples, such as TX^{A_j} transmitted by agent A_j , into sets of received tuples, such as RX^{A_j} . For each tuple $\langle G \cup \Sigma \cup \Omega, k \rangle$ in the set TX^{A_j} , row 4 computes a set $\bar{\Omega}$ which consists of the signals in Σ which are outputs from the receiving agent A_i , and the set $\bar{\Sigma}$ which are the signals in Σ except those in $\bar{\Omega}$. The signals in $\bar{\Omega}$ are then replaced by a set consisting of one component from the dependency of each signal, resulting in a set of components \bar{P} . The received tuple $\langle \bar{P} \cup G \cup \bar{\Sigma} \cup \Omega, k \rangle$ is thereafter constructed and stored in the set of tuples RX^{A_j} . The computation of each set \bar{P} is done in row 5 using the function $\text{CrossProduct}(M)$ which computes the cross product for the collection of sets M .

Algorithm 2 – ReceiveCondense($\{TX^{A_1}, \dots, TX^{A_n}\}$). Compute the minimal condensed diagnoses in agent A_i .

Input: For each agent A_j except A_i , a received set TX^{A_j} resulting from the evaluation of Transmit(\cdot). The set of minimal local diagnoses \mathbb{D}^{A_i} .

Output: The set of minimal condensed diagnoses $\mathbb{D}_c^{A_i}$.

```

1: for each  $A_j$  except  $A_i$  do
2:    $RX^{A_j} := \emptyset$ 
3:   for each  $\langle G \cup \Sigma \cup \Omega, k \rangle \in TX^{A_j}$  do
4:      $\bar{\Omega} := \Sigma \cap S_{out}^{A_i}$  and  $\bar{\Sigma} := \Sigma \setminus \bar{\Omega}$ 
5:     for each  $\bar{P} \in \text{CrossProduct}(\cup_{s \in \bar{\Omega}} \{\text{Dep}(s)\})$  do
6:        $RX^{A_j} := RX^{A_j} \cup \{\langle \bar{P} \cup G \cup \bar{\Sigma} \cup \Omega, k \rangle\}$ 
7:     end for
8:   end for
9: end for
10:  $RX^{A_i} := \{\langle D, 0 \rangle : D \in \mathbb{D}^{A_i}\}$ 
11:  $\mathbb{D}_c^{A_i} := \text{CrossProductTuple}(\{RX^{A_1}, \dots, RX^{A_n}\})$ 
12:  $\mathbb{D}_c^{A_i} := \text{RemoveNonMinimal}(\mathbb{D}_c^{A_i})$ 

```

Row 10 transforms the minimal local diagnoses in agent A_i to the same format as the received tuples. In rows 11–12 in the algorithm, the sets $RX^{A_1}, \dots, RX^{A_n}$ are merged and the non-minimal condensed diagnoses are removed. The function $\text{CrossProductTuple}(M)$ for a set M of sets of tuples performs a normal cross product with respect to the sets included in the tuples and a summation of the variables k . The function $\text{RemoveNonMinimal}(\cdot)$ removes all non-minimal condensed diagnoses.

C. Main Algorithm for the Computation of the Sets of Minimal Condensed Diagnoses

Based on the algorithms presented above, Algorithm 3 can be used to compute the set of minimal condensed diagnoses in each agent. For each agent, the algorithm computes the set of minimal local diagnoses \mathbb{D}^{A_i} in row 1 using a minimal hitting set algorithm. The algorithm then evaluate rows 2–5 which include calls to Algorithm 1, where each call has a set of minimal local diagnoses as input and gives a set of transmitted diagnoses TX^{A_i} as output. Finally, rows 6–9 are evaluated which includes calls to Algorithm 2, where each call has the set of transmitted sets $TX^{A_1}, \dots, TX^{A_n}$ except TX^{A_i} as input. After that Algorithm 3 has been evaluated, each agent has a unique set of minimal condensed diagnoses $\mathbb{D}_c^{A_i}$.

VI. APPLICATION STUDY OF THE PROPOSED ALGORITHM TO AN AUTOMOTIVE VEHICLE

The development of the proposed algorithm was motivated by efficiency compared to computing the complete set of minimal global diagnoses via a central agent. To verify these motivations, the algorithm has in this section been applied to an automotive vehicle and compared to a centralized implementation which computes the minimal global diagnoses. For a fuller description of the example, see [13].

Algorithm 3 – Main. Computation of the set of minimal condensed diagnoses in each agent.

Input: The set of conflicts Π^{A_i} in each agent A_i .

Output: The set of minimal condensed diagnoses $\mathbb{D}_c^{A_i}$ in each agent A_i .

```

1: for each agent  $A_i$  do  $\mathbb{D}^{A_i} := \text{MinimalDiagnoses}(\Pi^{A_i})$  in  $A_i$ 
2: for each Agent  $A_i$  do
3:   compute  $TX^{A_i} := \text{Transmit}(\mathbb{D}^{A_i})$  in agent  $A_i$ 
4:   broadcast  $TX^{A_i}$  on the network
5: end for
6: for each Agent  $A_i$  do
7:   receive  $TX^{A_j}$  in  $A_i$  from all agents  $A_j$  except  $A_i$ 
8:   compute  $\mathbb{D}_c^{A_i} := \text{ReceiveCondense}(\{TX^{A_1}, \dots, TX^{A_n}\} \setminus TX^{A_i})$  in  $A_i$ 
9: end for

```

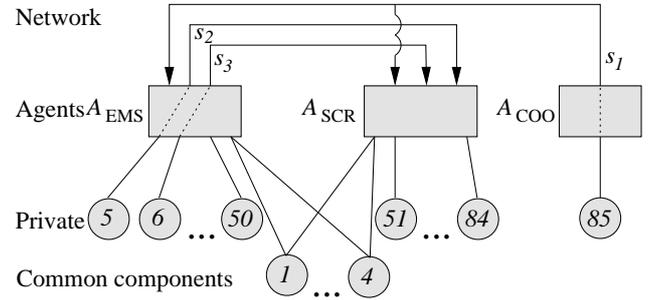


Fig. 3. Schematic overview of the distributed system in the application.

A. Test Cases Used in the Application Study

The algorithm proposed in this paper has been applied to a part of the diagnostic system used in a series production heavy duty vehicle. The part includes the *engine management system* (EMS), the *selective catalytic reduction system* (SCR), and the part of the *coordinator system* (COO) affecting the EMS and the SCR.

There are 85 components and three signals, see Fig. 3 for a schematic overview. The diagnostic systems in the A_{EMS} and the A_{SCR} consist of 57 and 55 diagnostic tests, respectively.

B. Evaluation

The proposed algorithm is compared to a centralized algorithm based on the method for computing minimal diagnoses from minimal conflicts presented in [14]. In the centralized algorithm, all conflicts from all agents are transmitted to A_{COO} where the set of minimal global diagnoses is computed.

To compare the algorithms, the number of abnormal components in the complete system has been varied and a test suite of simulations has been created by randomly making some components behave abnormally. The mean of the maximum number of needed operations in any agent can be seen in Fig. 4(a). It can be seen that for the test cases with more than one abnormal component, the mean of the maximum number of needed operations with the proposed algorithm is lower than the centralized algorithm, and it can also be seen that the growth rate is significantly lower.

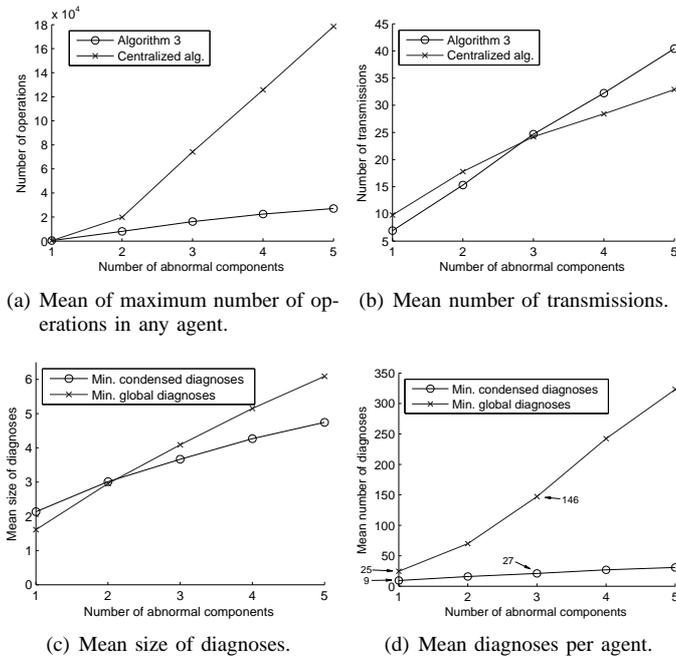


Fig. 4. Results from the evaluation of the automotive application.

The numbers of needed transmissions for the two algorithms are shown in Fig. 4(b), where it can be seen that the network load for the proposed algorithm is somewhat higher. Fig. 4(c) shows the mean size of a diagnosis, and it can be seen that the memory usage for the proposed algorithm is somewhat lower. In Fig. 4(d), the mean number of minimal diagnoses per agent is shown. It can be seen that the number of minimal condensed diagnoses is significantly reduced compared to the number of minimal global diagnoses.

The memory usage needed to store the set of diagnoses is approximately the mean number of diagnoses multiplied by the mean size of the diagnoses. Primarily due to the reduction in number of diagnoses, the memory usage to store the sets of minimal condensed diagnoses is reduced compared to when the set of minimal global diagnoses is stored.

C. Recommendations Based on the Automotive Application Study

When choosing a distributed diagnostic solution, network load, processor load, memory usage, network load determinism, and possible number of abnormal components should be considered as a minimum. If the network load must behave deterministically, then it is best to transmit all test results continuously to all agents since the number of test results does not change. If the number of possible abnormal components is low, then a centralized algorithm should be used since this will give a sufficiently low processor and network load. If the number of possible abnormal components can not be guaranteed to be low and processor load and memory usage should be kept low, then the algorithm proposed in this paper could be used.

All in all, the choice of distributed diagnostic solution is a difficult one where many different aspects have to be taken into account.

VII. CONCLUSIONS

The objective, when designing the algorithm proposed in the paper, was to obtain a diagnostic system that used low processing power with low memory usage, and which made a compact representation of the minimal global diagnoses available in each agent. The proposed algorithm creates a unique set of minimal condensed diagnoses for each agent, where each set of minimal condensed diagnosis only includes components used by the agent. An application study on the distributed system in a heavy duty vehicle has shown the benefits of the minimal condensed diagnoses compared to the minimal global diagnoses. Further, by using some extra network load, both the computational load and the memory usage are significantly reduced compared to a centralized algorithm which computes the set of minimal global diagnoses.

REFERENCES

- [1] J. de Kleer and J. Kurien, "Fundamentals of model-based diagnosis," in *Proceedings of IFAC Safeprocess'03*, Washington, U.S.A., 2003.
- [2] G. Weiss, Ed., *Multiagent systems : a modern approach to distributed artificial intelligence*. Cambridge, Mass., U.S.A: MIT Press, 1999.
- [3] S. Takai and R. Kumar, "Decentralized diagnosis for nonfailures of discrete event systems using inference-based ambiguity management," *Systems, Man and Cybernetics, Part A, IEEE Transactions on*, vol. 40, no. 2, p. 406, 2010.
- [4] O. Nasri, H. Shraim, P. Dague, O. Heron, and M. Cartron, "Model-based decentralized embedded diagnosis inside vehicles: Application to smart distance keeping function," in *Control and Fault-Tolerant Systems (SysTol), 2010 Conference on*, 2010, p. 17.
- [5] G. Leen and D. Heffernan, "Expanding automotive electronic systems," *Computer*, vol. 35, no. 1, pp. 88–93, Jan 2002.
- [6] J. Gertler, *Fault Detection and Diagnosis in Engineering Systems*. Marcel Dekker, 1998.
- [7] A. S. Tanenbaum and M. van Steen, *Distributed Systems*. Prentice Hall, 2002.
- [8] R. Reiter, "A theory of diagnosis from first principles," *Artificial Intelligence*, vol. 32, no. 1, pp. 57–95, Apr 1987.
- [9] M. Nyberg, "A generalized minimal hitting-set algorithm to handle diagnosis with behavioral modes," *Systems, Man and Cybernetics, Part A, IEEE Transactions on*, vol. 41, no. 1, pp. 137–148, 2011.
- [10] G. Provan, "A model-based diagnosis framework for distributed systems," in *13th International Workshop on Principles of Diagnosis*, Semmering, Austria, May 2002.
- [11] J. Kurien, X. Koutsoukos, and F. Zhao, "Distributed diagnosis of networked, embedded systems," in *Proceedings of 13th International Workshop on Principles of Diagnosis*, Semmering, Austria, May 2002.
- [12] N. Roos, A. t. Teijje, and C. Witteveen, "Reaching diagnostic agreement in multi-agent diagnosis," in *First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2004)*, New York, U.S.A., July 2004, pp. 1254–1255.
- [13] J. Biteus, "Fault isolation in distributed embedded systems," Ph.D. dissertation, Linköpings universitet, April 2007.
- [14] J. de Kleer and B. Williams, "Diagnosing multiple faults," *Artificial Intelligence*, vol. 32, Apr 1987.