Distributed Diagnosis by using a Condensed Local Representation of the Global Diagnoses with Minimal Cardinality

Jonas Biteus *
Dept. of Electrical Engineering
Linköpings universitet, Sweden.

Erik Frisk
Dept. of Electrical Engineering
Linköpings universitet, Sweden

Mattias Nyberg
Power-train Division
Scania AB, Sweden

Abstract

The set of diagnoses is commonly calculated in consistency based diagnosis, where a diagnosis includes a set of faulty components. In some applications, the search for diagnoses is reduced to the set of diagnoses with minimal cardinality. In distributed systems, local diagnoses are calculated in each agent, and global diagnoses are calculated for the complete system. The key contribution in the present paper is an algorithm that synchronizes the local diagnoses in each agent such that these represent the global diagnoses with minimal cardinality. The resulting diagnoses only include faulty components used by the specific agent, and are therefore a condensed local representation of the global diagnoses with minimal cardinality.

1 Introduction

This paper considers distributed systems that consist of a set of agents, where an agent is a more or less independent software entity, connected to each other via some network [Hayes, 1999; Weiss, 1999]. The diagnoses can, in distributed systems, be divided into two different types, global diagnoses that are diagnoses for the complete distributed system and local diagnoses that are diagnoses for a single agent [Roos *et al.*, 2002].

It is an advantage to have the set of minimal global diagnoses in each agent. However, an agent has only an interest in knowing the fault status of the components used by that agent since the other components does not affect the specific agent. Consider for example a global diagnosis that consists of a set of components that have been found to be faulty. An agent that uses some components in its operation is interested in knowing if any of these components are included in the global diagnosis. The agent however does not have an interest in the fault status of the rest of the components in the global diagnosis. In some applications, the calculation of diagnoses is focused on to some smaller set of diagnoses, for example the most probable diagnoses [de Kleer, 1991] or the diagnoses with minimal cardinality [de Kleer, 1990].

*Corresponding author. E-mail: biteus@isy.liu.se. Address: Vehicular Systems, Electrical Engineering, Linköpings universitet, SE-581 83 Linköping, Sweden. Phone: +46 13 281994.

The key contribution in the present paper is an algorithm that synchronizes the minimal local diagnoses in one agent with the minimal local diagnoses in the other agents, such that the result is a set of diagnoses with minimal cardinality. Each resulting diagnosis is a subset of some global diagnoses, and only components used by the specific agent is included the resulting diagnosis. Since only the components used by the specific agent are included, both the size and the number of the resulting diagnoses with minimal cardinality are reduced compared to the set of global diagnoses with minimal cardinality. The resulting diagnoses with minimal cardinality are therefore a condensed local representation of the global diagnoses with minimal cardinality, and are here denoted condensed diagnoses with minimal cardinality. By reducing the size and the number, the algorithm requires a low computational load, low memory usage, and low network load. The algorithm is distributed such that it can handle both changes in the number of agents and the exchange of single agents. The algorithm is described in Section 5, using the framework for distributed diagnosis presented in Section 3–4.

Our work has been inspired by diagnosis in distributed embedded systems used in automotive vehicles. These systems typically consist of precomputed diagnostic tests that are evaluated in different agents, which in the automotive industry correspond to electronic control units (ECUs). Sets of conflicts are generated when the diagnostic tests are evaluated in the ECUs, and the ECUs then compute sets of minimal local diagnoses. These embedded distributed systems typically consist of ECUs with both limited processing power and limited RAM memory. The algorithm presented here is therefore constructed such that it requires low processing power and low memory usage. In these systems, it should be possible to exchange, add, or remove ECUs without having to do any changes to the diagnostic software. The algorithm presented in this paper is therefore constructed such that it can handle such changes. Requirements on diagnostic systems used in automotive vehicles are discussed in Section 2.

1.1 Related Work

Most research, such as [Reiter, 1987], has been aimed at the centralized diagnosis problem. These methods can also be used for distributed systems by letting a central diagnostic agent collect conflicts from the system and then calculate the minimal global diagnoses. It is not always suitable to use a dedicated central diagnostic agent due to for example limited computing resources in each agent, robustness against agent disconnection, and the possibility to add new agents to the network. It therefore exist algorithms, see for example [Provan, 2002], which compute the minimal global diagnoses in a cooperation between the agents. These algorithms aim at the complete set of global diagnoses, while the method presented here aims at the set of global diagnoses with minimal cardinality.

There also exist algorithms where the agents update the sets of local diagnoses such that these are consistent with the global diagnoses [Roos *et al.*, 2003]. The method does not guarantee that a combination of the agents' local minimal diagnoses is also a global minimal diagnosis. However, for every global minimal diagnosis, there is a combination of local minimal diagnoses. The updated sets of local diagnoses represent the global diagnoses without actually computing the complete set of global diagnoses. The method presented in this paper updates the local diagnoses such that these represent the global diagnoses with minimal cardinality.

In [Biteus *et al.*, 2005], a method was presented that calculated the global diagnoses with minimal cardinality by transmitting the minimal local diagnoses from one agent to another, which adds its minimal local diagnoses, then transmit the result to the next agent, etc. Even though the method is efficient, it might not give a sufficient distributed diagnostic system since it requires a lot of cooperation between the agents. In [Biteus *et al.*, 2005], the computational burden when calculating the global diagnoses with minimal cardinality was reduced by partitioning the system into two or more sub-systems, whose minimal local diagnoses did not share components with each other. The partition approach has similarities with the tree reduction technique used in [Wotawa, 2001] and can also be used for the algorithm presented here.

Related to this work is also the EU funded project Multi-Agents-based Diagnostic Data Acquisition and Management in Complex Systems (MAGIC) which develops an architecture useful for distributed diagnosis [Köppen-Seliger *et al.*, 2003]. The project discusses protocols for network communication, control algorithms, and other aspects of the integration of diagnosis in distributed systems.

2 Requirements on Diagnostic Systems used in the Automotive Industry

To better understand the industrial demands on diagnosis, the distributed diagnostic systems used in Scania AB heavy-duty trucks have been analyzed. These systems consist of ECUs connected to each other via a controller area network (CAN). The software embedded in the ECUs is primary used for control and monitoring.

2.1 An Example of a Distributed System

One configuration of the distributed system in Scania's heavy-duty vehicles is shown in Figure 1. The system includes three separate CAN buses, the red, the yellow, and the green. Each of the ECUs is connected to sensors and actuators, and both sensor values and control signals can be shared with the other ECUs over the network. One example of an

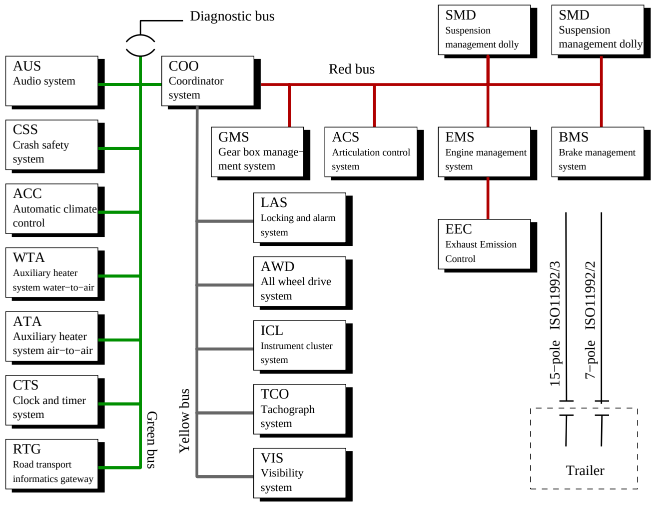

Figure 1: The distributed system in current Scania heavy-duty vehicles.

ECU is the engine management system, which is connected to sensors and actuators related to the engine. There can be up to about 30 ECUs in the system, depending on the type of the truck, and roughly between 4 and 110 components are diagnosed by each ECU.

The ECUs' CPUs have typically a clocking speed of 8 to 64 MHz, and a RAM memory capacity of about 4 to 150 kB. CAN buses can typically transfer 100 to 500 kbit/s. As these numbers indicate there is not much computational, memory, nor network capacity available, especially when considering that the ECUs should be used for both control and diagnosis.

2.2 Overall Requirements on Distributed Systems

A distributed system that can present the same information to users, as if it were a centralized system, can be denoted transparent [Tanenbaum and van Steen, 2002]. Considering fault diagnosis, one interpretation of transparency is that the minimal diagnoses presented by the distributed diagnostic system should be the same as those presented by a centralized diagnostic system, meaning that the minimal global diagnoses should be presented, not only the minimal local diagnoses. Another interpretation of transparency is that, even though one ECU fails to deliver its minimal local diagnoses the remaining system should still be able to deliver the minimal global diagnoses. This means that the diagnostic processes should be distributed among the ECUs, or if a centralized diagnostic ECU is used, backup ECUs should exist.

If it is possible to increase or decrease the size of the system, without changes in the software, the system can be said to be scalable [Tanenbaum and van Steen, 2002]. Considering a truck, it should be possible to attach new parts including new ECUs to the network without having to change the software in the ECUs.

If it is possible to exchange an ECU to for example a new version without having to change the software in the other ECUs, the system can be said to be interoperable [Tanenbaum and van Steen, 2002]. This is especially important in automotive systems where it frequently occurs that parts are replaced by parts from other manufacturers.

2.3 Requirement Conclusions

An algorithm for distributed diagnosis used in automotive vehicles, should require a limited processing power, memory usage, and network load. The algorithm should further result in a transparent, scalable, and interoperable system.

The algorithm presented in this paper synchronizes the minimal local diagnoses, which results in a transparent, scalable, and interoperable system. The condensed diagnoses with minimal cardinality do not include components that are not of interest and this reduces computational load, memory usage, and network load.

3 Consistency Based Diagnosis

A system consists of a set of components $\mathcal{C}$, where a component is something that can be diagnosed. This not only includes components directly connected to the agents, such as sensors and actuators, but also includes components shared between the agents, e.g. cables and pipes.

To reduce the complexity of the diagnostic system, it is sometimes preferable to only consider the abnormal AB and the not abnormal $\neg AB$ mode, where the AB mode does not have a model. This means that the minimal diagnosis hypothesis is fulfilled [de Kleer *et al.*, 1992], and therefore the notation in for example GDE will be employed.

A diagnosis is a set of components $D \subseteq \mathcal{C}$, such that the components' abnormal behaviors, the remaining components' normal behaviors, the system description, and the observations are consistent. Since the minimal diagnosis hypothesis is fulfilled and D is a diagnosis, all supersets of D are also diagnoses. Further, a diagnosis D' is a minimal diagnosis if there is no proper subset $D \subset D'$ where D is a diagnosis [de Kleer *et al.*, 1992].

An evaluation of a diagnostic test results in a conflict if some components, checked by the test, have been found to behave abnormal. A conflict is a set of components $\pi \subseteq \mathcal{C}$, such that the components' normal behaviors, the system description, and the observations are inconsistent. A set $D \subseteq \mathcal{C}$ is a diagnosis if and only if it has a nonempty intersection with every conflict in a set of conflicts. A consequence of this is that the set of minimal diagnoses is exactly determined by the set of minimal conflicts [de Kleer *et al.*, 1992].

In some cases, it is computationally intractable to calculate the complete set of minimal diagnoses. To reduce the computational cost, the search can be focused on the diagnoses with minimal cardinality, as described in for example [de Kleer, 1990]. Let $\mathbb{D}$ be a set of diagnoses, then the set of minimal cardinality diagnoses is the set $\mathbb{D}_{mc} = \{D \in \mathbb{D} : |D| = \min_{D\in\mathbb{D}} |D|\}$.

4 Distributed Diagnosis

This section will present the framework for distributed systems that will be used to describe how condensed diagnoses with minimal cardinality can be calculated.

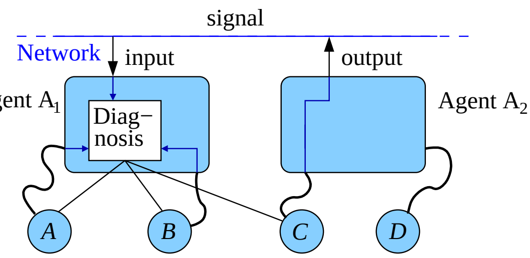

Figure 2: Agents, network, components, and diagnosis.

4.1 Outputs, Inputs, and Components in Distributed Systems

A distributed system consists of a set of agents $\mathcal{A}$. A local diagnosis is determined by the conflicts in a single agent, while a global diagnosis is determined by all agents' conflicts.

Here, the complete set of components are partitioned into private components $\mathcal{P} \subseteq \mathcal{C}$ and common components $\mathcal{G} \subseteq \mathcal{C}$, where $P \cap \mathcal{G} = \emptyset$. 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 partitioned into different sets belonging to different agents such that for two such sets, $\mathcal{P}^{A_i} \cap \mathcal{P}^{A_j} = \emptyset$, where A_i, $A_j \in \mathcal{A}$. The set X^A is the subset of X that is used in agent A.

In addition to components, an agent in a distributed system should also be able to diagnose inputs from other agents. The outputs are values from sensors, to actuators, or from calculations, which are made available to the other agents over the network. The complete set of signals $\mathcal{S}$, a set of inputs $\mathcal{IN} \subseteq \mathcal{S}$, and a set of outputs $\mathcal{OUT} \subseteq \mathcal{S}$ is here used. Each output $\sigma \in \mathcal{OUT}$ is connected to a subset of inputs $\Gamma \subseteq \mathcal{IN}$.

Example 1: Figure 2 shows a typical layout of agents and components. The system consists of two agents, a network, and four sensor components, A to D. The sensors A and B are physically connected to agent A_1, while the sensors C and D are connected to A_2. A diagnosis in agent A_1 could for example include the components A, B, and C, connected with dashed lines. Agent A_1 diagnoses indirectly sensor C through the signal transmitted over the network. ⋄

An output might depend on other components, such as sensors, and the information about this relationship is stored in the output's assumptions.

Definition 1 (Assumption) *Let s be a signal which is an output from agent A or an input that is connected to the output from agent A. Let the set $\text{ass}(s) \subseteq \mathcal{P}^A \cup \mathcal{G} \cup \mathcal{IN}^A$. If s is abnormal if and only if some non-empty subset $C \subseteq \text{ass}(s)$ is abnormal, then $\text{ass}(s)$ is the set of assumptions for s.*

Each output depends on components and other inputs. This dependency can be propagated to a set consisting only of components.

Definition 2 (Dependency) *Let $s \in \mathcal{S}$ be a signal, then the dependency for s is*

$$\text{dep}(s) = \text{ass}(s) \cap \mathcal{C} \cup \bigcup_{t\in\text{ass}(s)\cap\mathcal{S}} \text{dep}(t).$$

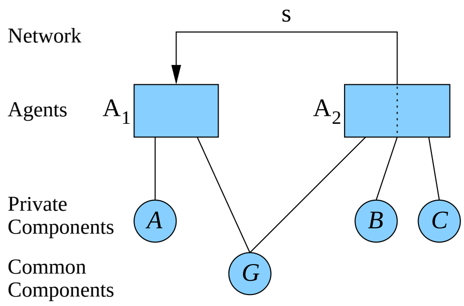

Figure 3: An example of a two agent system.

Since the $\mathrm{dep}(\cdot)$ function is defined implicit, the possibility of loops has to be considered in an implementation.

Example 2: Continuation of Example 1. The assumption of the output is $\mathrm{ass}(\mathrm{output}) = \{C\}$. The dependency equals the assumption since the assumption does not include any signal. $\diamond$

4.2 Diagnoses & Conflicts on Components & Inputs

An agent should state diagnoses that include both components and inputs. The diagnoses in Section 3 cannot be used directly since these only include components. Instead, a component-input diagnosis is defined on the set $\mathcal{C} \cup \mathcal{IN}$.

Definition 3 (Component-input diagnosis) *A set $D = C \cup \Gamma \subseteq \Theta$, $C \subseteq \mathcal{C}$, $\Gamma \subseteq \mathcal{IN}$, is a component-input diagnosis if the set $C \cup \hat{C}$, where $\forall i \in \Gamma : \hat{C} \cap \mathrm{dep}(i) \neq \emptyset$, is a diagnosis.*

A component-input diagnosis will simply be denoted a diagnosis when no misunderstanding is imminent. As with diagnoses, conflicts can also be defined upon $\mathcal{C} \cup \mathcal{IN}$.

4.3 Condensed Diagnoses Representing Global Diagnoses

It was mentioned in the introduction that the condensed diagnoses with minimal cardinality should be calculated in each agent.

Definition 4 (Condensed diagnosis) *Let $\mathcal{D}$ be a set of global diagnoses. The tuple $\langle D, k\rangle$, where $D \subseteq P^A \cup G \cup \mathcal{IN}^A$ and $k \in \mathbb{Z}$, is a condensed diagnosis in agent A if $\exists \bar{D} \in \mathcal{D}$, such that*

$$|D| + k = |\bar{D}|, D \cap \mathcal{P} = \bar{D} \cap \mathcal{P}^A, D \cap \mathcal{G} = \bar{D} \cap \mathcal{G},$$
$$D \cap \mathcal{IN} = \{i : \mathrm{dep}(i) \cap \bar{D} \backslash D \neq \emptyset, i \in \mathcal{IN}\}.$$

The condensed diagnosis $\langle D, x\rangle$ in agent A, is a tuple where the set D represents the subset of some global diagnoses, diagnosis $\bar{D}$ in the definition, including components used by agent A. Variable k represents the components not included in D but included in $\bar{D}$.

Interpretation of the different requirements for a condensed diagnosis: $|D| + k = |\bar{D}|$ means that the cardinality plus k should equal that of the global diagnosis with minimal cardinality; $D \cap \mathcal{P} = \bar{D} \cap \mathcal{P}^A$ means that D should only include private components used by agent A; $D \cap \mathcal{G} = \bar{D} \cap \mathcal{G}$ means that the common components should be included; $D \cap \mathcal{IN} = \{i : \mathrm{dep}(i) \cap \bar{D} \neq \emptyset, i \in \mathcal{IN}\}$ means that inputs, that might be faulty due to its dependency on some faulty components, should be included in D.

Example 3: Consider the system shown in Figure 3. There exist a signal s whose dependency $\mathrm{dep}(s) = \{B\}$, represented by the dotted line. The sets of private components are $\mathcal{P}^{A_1} = \{A\}$ and $\mathcal{P}^{A_2} = \{B, C\}$. The set of common components is $\mathcal{G} = \{G\}$. Let $\{A, B, C, G\}$ be a global diagnosis.

A condensed diagnosis in agent A_1 is $\langle\{A, G, s_1\}, 1\rangle$. 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 faulty component B. Component C is represented by $k = 1$ since it does not affect agent A_1. $\diamond$

The *minimal cardinality condensed diagnoses* is the set of condensed diagnoses where each condensed diagnosis is a subset of some minimal cardinality global diagnoses , i.e. in Definition 4, the set of global diagnoses $\mathcal{D}$ is the set of minimal cardinality global diagnoses. The objective of the algorithm described in the next section is to calculate the sets of minimal cardinality condensed diagnoses in each agent.

5 Algorithm for Calculating the Minimal Cardinality Condensed Diagnoses

The main idea of the algorithm presented in this Section is that each agent first find and then transmits the subset of minimal local diagnoses that other agents might be interested of. Each agent then receives the transmitted diagnoses and merges these with its own set of minimal local diagnoses resulting in the minimal cardinality condensed diagnoses.

The transmitting part is described in Section 5.2, the receiving and merging is described in Section 5.3, and finally the main algorithm is described in Section 5.4.

In the algorithms, D is some diagnosis, $\Gamma \subseteq \mathcal{IN}$, $\Omega \subseteq \mathcal{OUT}$, $P \subseteq \mathcal{P}$, and $G \subseteq \mathcal{G}$.

5.1 Outputs Dependent on Inputs

The algorithm, as written, requires that $\mathrm{ass}(s) \subseteq \mathcal{C}$, i.e. an output's value does not depend on any signal. This can be fulfilled for a general system, where $\mathrm{ass}(s) \subseteq \mathcal{C} \cup \mathcal{IN}$, by replacing the assumptions with the corresponding dependencies, such that $\mathrm{ass}(s) := \mathrm{dep}(s) \subseteq \mathcal{C}$.

5.2 Transmit Diagnoses

The first step is to find and transmit the subset of minimal local diagnoses that is of interest to the other agents.

A minimal local diagnosis should be transmitted if it includes common components, inputs, or components that outputs depends on, since these might affect other agents. For each diagnosis that should be transmitted, the private components can be removed since these do not directly affect other agents. If some outputs depend on any of the removed private components these outputs are instead added to the diagnosis. The minimal local diagnoses that are not transmitted on the network can be represented by a variable $n \in \mathbb{N}$, which is the minimal cardinality of the non-transmitted local diagnoses. The agents receiving the diagnoses will then be aware that

Algorithm 1 Transmit diagnoses $\texttt{transmit}(A, m)$.

Require: A set of minimal local diagnoses $\mathbb{D}^A$, limit m.
Ensure: Set TX broadcasted on the network.
1: $\mathbb{D}^{TX} := \{D \in \mathbb{D}^A : |D| \leq m\}$
2: $\mathbb{D}^{TX} := \{D \in \mathbb{D}^{TX} : D \cap (\mathcal{IN} \cup \mathcal{G} \cup (\cup_{\sigma \in \mathcal{OUT}^A} \mathrm{ass}(\sigma))) \neq \emptyset\}$
3: $n := \min_{D \in \mathbb{D}^A \setminus \mathbb{D}^{TX}} |D|$
4: $TX := \{\langle \bar{D}, k\rangle : D \in \mathbb{D}^{TX}, D = P \cup G \cup \Gamma, \bar{D} = G \cup \Gamma \cup \Omega, \Omega = \{\sigma \in \mathcal{OUT}^A : \mathrm{ass}(\sigma) \cap P \neq \emptyset\}, k = |P| - |\Omega|\}$
5: **if** $|\mathbb{D}^{TX}| \neq |\mathbb{D}^A|$ **then**
6: $\quad TX := TX \cup \{\langle \emptyset, n\rangle\}$
7: **end if**
8: Broadcast TX on the network.

there exist one or more non-transmitted local diagnoses with cardinality n.

Algorithm 1 performs the steps described above. Since the minimal cardinality condensed diagnoses are searched for, the algorithm accepts a maximum limit m on the cardinality of the local diagnoses to be transmitted. Row 2 decides which diagnoses that should be transmitted. Row 4 constructs a tuple including the diagnosis without private components and a variable $k \in \mathbb{N}$ representing the removed private components.

Example 4: Consider the system shown in Figure 4, where objects connected to the agents with solid lines are included in some minimal local diagnosis, while those connected with dashed lines are not included. The sets of private components are $\mathcal{P}^{A_1} = \{A\}$, $\mathcal{P}^{A_2} = \{B, C, D\}$, and $\mathcal{P}^{A_3} = \{E\}$. The set of common components is $\mathcal{G} = \{G\}$. There exist two signals with assumptions $\mathrm{ass}(s_1) = \{B\}$ and $\mathrm{ass}(s_2) = \{E\}$.

Assume that the following set of minimal local diagnoses has been calculated in agent A_2

$$\mathbb{D}^{A_2} = \{\{C, s_2\}, \{B, C\}, \{G, C\}, \{D\}\}.$$

Using Algorithm 1 with $m \geq 2$, the set

$$\mathbb{D}^{TX} = \{\{C, s_2\}, \{B, C\}, \{G, C\}\}$$

is first calculated. The variable $n = 1$, i.e. the cardinality of $\{D\}$. The transmitted set of tuples is

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

where the private components have been removed. This set will represent $\mathbb{D}^{A_2}$ in agents A_1 and A_3. $\diamond$

5.3 Receive and Merge Diagnoses

The second step is to receive the transmitted sets, transform them into an appropriate form, and then calculate the minimal cardinality condensed diagnoses.

If a received diagnosis include a signal s that are an output from the receiving agent then the receiver know which components that s depend on, and s is therefore replaced with the assumption $\mathrm{ass}(s)$. After the replacement, the minimal local diagnoses and the received diagnoses can be merged to form a set of condensed diagnoses. If a condensed diagnosis includes several inputs and if several of these inputs depend

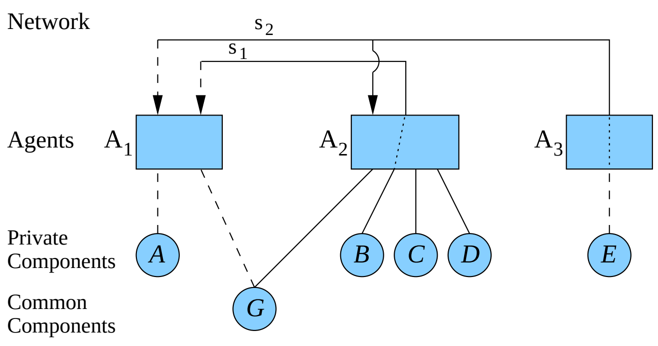

Figure 4: An example of a three agent system.

on the same component then the cardinality of the condensed diagnosis will not be correct. Consider for example two signals s_1 and s_2 depending on component A. The cardinality of $\langle\{s_1, s_2\}, 0\rangle$ is two while the cardinality of the corresponding global diagnosis $\{A\}$ is one. The condensed diagnosis should be $\langle\{s_1, s_2\}, -1\rangle$ where the minus one a compensation.

Algorithm 2 performs the steps described above. The algorithm transforms the received sets of tuples by replacing inputs that are outputs from the current agent, with the components in the outputs' assumptions, row 1–3. The function $MHS(M)$ calculates the minimal hitting set for the collection of sets M. For example, the minimal hitting set $MHS(\{\{A, B\}, \{B, C\}\}) = \{\{A, C\}, \{B\}\}$. To be able to calculate the compensation discussed above, the set of inputs is partitioned into the set $\tilde{\Gamma}^j$, which is outputs that was added to the diagnosis when A_j transmitted the tuple, and the set $\bar{\Gamma}^j$, which is the inputs to agent A_j without the outputs from A_i. In row 5 the received sets of tuples are merged. In row 6, the condensed diagnoses are compensated for signals depending on the same components. Finally, in row 7, the condensed diagnoses, which do not have minimal cardinality, is removed.

Algorithm 2 Calculate the minimal cardinality condensed diagnoses $\texttt{condense}(A_i)$.

Require: Received sets TX^{A_j} from all agents $A_{j \neq i}$ as a result of evaluating $\texttt{transmit}(A_j, m)$. Set of minimal local diagnoses $\mathbb{D}^{A_i}$.
Ensure: Set of minimal cardinality condensed diagnoses $\mathbb{D}_s^{A_i}$.
1: **for all** $j \neq i$ **do**
2: $\quad RX^{A_j} := \{\langle \bar{P} \cup G \cup \bar{G} \cup \bar{\Gamma} \cup \tilde{\Gamma}, k\rangle : \langle G \cup \Gamma \cup \Omega, k\rangle \in TX^{A_j}, \tilde{\Gamma} = \Omega, \bar{\Omega} = \Gamma \cap \mathcal{OUT}^{A_i}, \bar{\Gamma} = \Gamma \setminus \bar{\Omega}\}, H \in MHS(\cup_{\sigma \in \bar{\Omega}} \{\mathrm{ass}(\sigma)\}), \bar{P} = H \cap \mathcal{P}, \bar{G} = H \cap \mathcal{G}\}$
3: **end for**
4: $RX^{A_i} := \{\langle D, 0\rangle : D \in \mathbb{D}^{A_i}\}$
5: $\mathbb{D}_s^{A_i} := \{\langle H, k\rangle : \langle D_j, k_j\rangle \in RX^{A_j}, H = \cup_j D_j, k = \sum_j k_j\}$.
6: $\mathbb{D}_s^{A_i} := \{\langle D, k + \texttt{comp}(D)\rangle : \langle D, k\rangle \in \mathbb{D}_s^{A_i}\}$
7: $\mathbb{D}_s^{A_i} := \{\langle \bar{D}, k\rangle \in \mathbb{D}_s^{A_i}, k = \min_{\langle \bar{D}, k\rangle \in \mathbb{D}_s^{A_i}} k\}$

The value of $k + \texttt{comp}(D)$ should be the difference between the cardinality of D and a corresponding minimal car-

Algorithm 3 Function $\texttt{comp}(D)$.

Input: Diagnosis D.
Output: Variable k.

1: Each diagnosis is constructed such that
$D = (P^i \cup G^i \cup \Gamma^i) \cup_{j \neq i} (\bar{P}^j \cup G^j \cup \bar{G}^j \cup \bar{\Gamma}^j \cup \tilde{\Gamma}^j)$
2: $Z^j := (\Gamma^i \cup_{k \neq i} \bar{\Gamma}^k) \cap \mathcal{OUT}^{A_j}$
3: $k^j_{21} := \min_{H \in MHS(\cup_{\sigma \in Z^j} \{\text{ass}(\sigma)\})} |H \cap \mathcal{P}| - |Z^j|$
4: $k^j_{22} := \min_{H \in MHS(\cup_{\sigma \in \tilde{\Gamma}^j \cap Z^j} \{\text{ass}(\sigma)\}) \cap \mathcal{P}} |H| - |\tilde{\Gamma}^j \cap Z^j|$
5: $k_2 := \sum_j (k^j_{21} + k^j_{22})$
6: $Z_G := G^i \cup_{j \neq i} (G^j \cup \bar{G}^j)$
7: $k_3 := \min_{H \in MHS(\cup_{\sigma \in \cup_{j \neq i} Z^j} \{\text{ass}(\sigma)\}) \cap \mathcal{G}} |H| - |Z_G \cap H|$
8: $k := k_2 + k_3$

dinality global diagnosis. The variable k is calculated in Algorithm 1, while the function $\texttt{comp}(D)$ is given by Algorithm 3. In the algorithm, k_2 is the compensation for signals depending on the same private components, while k_3 is the compensation for signals depending on the same common components. The following example will illustrate $\texttt{comp}(D)$.

Example 5: Consider an agent A_2 with outputs s_1 and s_2, private components $\mathcal{P}^{A_2} = \{A\}$, and assumptions $\text{ass}(s_1) = \{A\}$ and $\text{ass}(s_2) = \{A\}$.

Let a minimal local diagnosis in agent A_1 be $D = \{s_1, s_2\}$ and let A_2 have an empty set of minimal local diagnoses. The minimal cardinality condensed diagnosis in A_2, before evaluating $\texttt{comp}(D)$, is in this case $\langle \{s_1, s_2\}, 0 \rangle$. A minimal cardinality global diagnosis is $\{A\}$, and the cardinality of the minimal cardinality condensed diagnosis is therefore not correct, $|\{s_1, s_2\}| + 0 \neq |\{A\}|$. Using $\texttt{comp}(D)$, row 1 give that $\Gamma^2 = \{s_1, s_2\}$. Row 3 give that $k^2_{21} = |\{A\}| - |\{s_1, s_2\}| = -1$. The result is the minimal cardinality condensed diagnosis $\langle \{s_1, s_2\}, -1 \rangle$, which has correct cardinality.

Assume now that agent A_2 has the minimal local diagnosis $\{A\}$, and agent A_1 has an empty set of minimal local diagnoses. Agent A_2 transmit the set $\{\langle \{s_1, s_2\}, -1 \rangle\}$ which is the set of minimal cardinality condensed diagnosis in A_1. Using $\texttt{comp}(D)$, the set $\tilde{\Gamma}^2 = \{s_1, s_2\}$ and the result is that $k_2 = 0$. In the second example, the compensation was done in the transmitting agent, while in the first example, the compensation had to be done in the receiving agent. $\diamond$

How computationally difficult is the $\texttt{comp}(D)$ function? Consider the special case where $\text{ass}(\sigma) \subseteq \mathcal{P}^{A_i}$ and $\text{ass}(\sigma_k) \cap \text{ass}(\sigma_l) = \emptyset$, which means that a signal only depends on private components and that no two signals depends on the same component. In this simplified case $\texttt{comp}(D) = 0$. If $\text{ass}(\sigma) \subseteq \mathcal{P}$, i.e. a signal only depends on private components, then the variable $k_3 = 0$ while k_2 might be some non-zero value. The more connections that exist between signals dependencies, and the more common components that exist, the more computationally complex will $\texttt{comp}(D)$ be.

The following lemma shows how the minimal cardinality condensed diagnoses could be calculated.

Lemma 1 *Let $\mathcal{D}^{mc}$ be the set of minimal cardinality global diagnoses, and let $\texttt{transmit}(A_j, \infty)$ be evaluated for all agents $A_j \in \mathcal{A}$. Let $\mathbb{D}^{A_i}_s$ be the result after evaluating Algorithm 2 in agent A_i. Then $\mathbb{D}^{A_i}_s$ is the set of minimal cardinality condensed diagnoses.*

Proof The complete proof is given in [Biteus *et al.*, 2006]. Outline of the proof: A diagnosis $\langle D^e, k \rangle \in \mathbb{D}^{A_i}_s$ is by construction a condensed diagnosis if $|D^e| + k = |D^g|$, where e referees to condensed and g to global. First show the cardinality of $D^g \in \mathcal{D}^{mc}$, second show that the cardinality of D^e in $\langle D^e, k \rangle \in \mathbb{D}^{A_i}_s$, and finally show that the cardinalities are equal. It is found that $|D^e| + X = |D^g|$ where $X = \sum_{j \neq i}(k^j_1 + k^j_2) + k_3$, and k^j_1 is identified as k_1 in Algorithm 1 in agent A_j. The variables k^j_2 and k_3 is identified as the corresponding variables in Algorithm 3. $\square$

Example 6: Continuation of Example 4. Assume that agent A_2 has received the following sets of tuples from agent A_1 and A_3,

$$TX^{A_1} = \{\langle \{s_1\}, 0 \rangle, \langle \emptyset, 1 \rangle\} \qquad TX^{A_3} = \{\langle \{s_2\}, 0 \rangle\}$$

Using Algorithm 2, the sets

$$RX^{A_1} = \{\langle \{B\}, 0 \rangle, \langle \emptyset, 1 \rangle\} \qquad RX^{A_3} = \{\langle \{s_2\}, 0 \rangle\}$$

is calculated. Assume that agent A_2 has the following set of minimal local diagnoses

$$\mathbb{D}^{A_2} = \{\{C, s_2\}, \{B, C\}, \{G, C\}, \{D\}\}$$

which is transformed with the algorithm to the set

$$RX^{A_2} = \{\langle \{C, s_2\}, 0 \rangle, \langle \{B, C\}, 0 \rangle, \langle \{G, C\}, 0 \rangle, \langle \{D\}, 0 \rangle\}.$$

The RX sets are merged, resulting in the set

$$\mathbb{D}^{A_1}_s \quad = \quad \{\langle \{B, C, s_2\}, 0 \rangle, \langle \{s_2, C\}, 1 \rangle, \langle \{s_2, D\}, 1 \rangle\},$$

which is the set of minimal cardinality condensed diagnoses. $\diamond$

Algorithm 4 Main algorithm.

Require: Set of minimal local diagnoses $\mathbb{D}^A$ in all agents.
Result: Set of minimal cardinality condensed diagnoses.

1: Decide with voting: $m_1 = \max_{A \in \mathcal{A}} \min_{D \in \mathbb{D}^A} |D|$
2: $\forall A \in \mathcal{A}$ evaluate $\texttt{transmit}(A, m_1)$
3: $\forall A \in \mathcal{A}$ evaluate $\texttt{condense}(A)$
4: Decide with voting: $m_2 = \max_{A \in \mathcal{A}} \min_{D \in \mathbb{D}^A_s} |D|$
5: $\forall A \in \mathcal{A}$ evaluate $\texttt{transmit}(A, m_2)$
6: $\forall A \in \mathcal{A}$ evaluate $\texttt{condense}(A)$

5.4 Main Algorithm

Lemma 1 shows that the sets of minimal cardinality condensed diagnoses can be calculated with Algorithm 1 to 3. However, it is possible to reduce the computational burden by using the cardinality limit m in Algorithm 1.

Algorithm 4 first calculates a lower bound m_1 on the cardinality of the minimal cardinality global diagnoses. The minimal cardinality condensed diagnoses are then calculated with this lower bound as input to $\texttt{transmit}(\cdot)$. The algorithm then computes an upper bound on the cardinality of the global

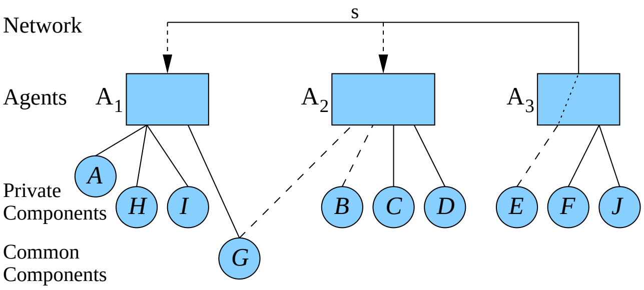

Figure 5: A system including relatively few signals compared to the number of components.

diagnoses with minimal cardinality m_2. Since m_2 is the cardinality of a global diagnosis, it is known that the cardinality of a minimal cardinality global diagnosis is less than or equal to m_2. The local diagnoses with cardinality greater than m_2 can therefore not be part of a minimal cardinality global diagnosis and can therefore be ignored. The result is described in Theorem 1. The result after evaluating Algorithm 4 is that all agents have a set of minimal cardinality condensed diagnoses $\mathbb{D}_s^A$.

The reason for using Algorithm 4 is that it is, in most cases, more efficient than using the algorithm as described in Lemma 1.

Even though Algorithm 4 is written as two separated parts, the result of part one should in an implementation be used when calculating part two. The correctness of the algorithm is shown in Theorem 1.

Theorem 1 *Same assumptions as in Lemma 1, but let $\mathbb{D}_s^{A_i}$ be the result after evaluating Algorithm 4 in agent A_i. Then the set $\mathbb{D}_s^{A_i}$ is the set of minimal cardinality condensed diagnoses.*

Proof Follows from Lemma 1. □

6 Example using the Algorithms

Two examples will be studied in this section.

Example 7: Consider the system shown in Figure 5. It includes three agents with the sets of private components $\mathcal{P}^{A_1} = \{A, H, I\}$, $\mathcal{P}^{A_2} = \{B, C, D\}$, $\mathcal{P}^{A_1} = \{E, F, J\}$, the set of common components $\mathcal{G} = \{G\}$, and the set of signals $\{s\}$.

The following sets of conflicts have been detected, $\Pi^{A_1} = \{\{A, H, G\}, \{A, I, G\}\}$, $\Pi^{A_2} = \{\{C\}, \{D\}\}$, and $\Pi^{A_3} = \{\{F, J\}\}$. The sets of minimal local diagnoses are calculated from the conflicts resulting in the sets

$$\mathbb{D}^{A_1} = \{\{A\}, \{G\}, \{H, I\}\} \qquad \mathbb{D}^{A_2} = \{\{C, D\}\}$$
$$\mathbb{D}^{A_3} = \{\{F\}, \{J\}\}.$$

The following sets of tuples are transmitted to agent A_1 from A_2 and A_3.

$$TX^{A_2} = \{\langle\emptyset, 2\rangle\} \qquad TX^{A_3} = \{\langle\emptyset, 1\rangle\}.$$

The received sets in agent A_1 are

$$RX^{A_1} = \{\langle\{H, I\}, 0\rangle, \langle\{A\}, 0\rangle, \langle\{G\}, 0\rangle\}$$
$$RX^{A_2} = \{\langle\emptyset, 2\rangle\} \qquad RX^{A_3} = \{\langle\emptyset, 1\rangle\}.$$

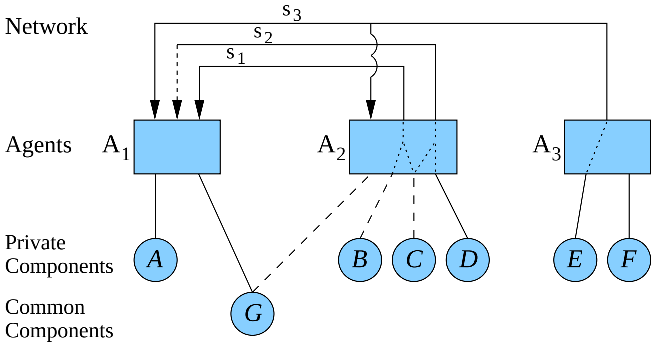

Figure 6: A system including relatively many signals compared to the number of components.

In agent A_1, the resulting set of minimal cardinality condensed diagnoses is $\mathbb{D}_s^{A_1} = \{\langle\{A\}, 3\rangle, \langle\{G\}, 3\rangle\}$. To verify this set, the set of minimal cardinality global diagnoses is calculated

$$\mathcal{D}^{mc} = \{\{A, C, D, F\}, \{A, C, D, J\}, \{G, C, D, F\}, \{G, C, D, J\}\}.$$

Is the minimal cardinality condensed diagnoses correct? Consider the condensed diagnosis $D = \langle\{A\}, 3\rangle$, and the minimal cardinality global diagnosis $\bar{D} = \{A, C, D, F\}$. Using Definition 4 to verify that D is a condensed diagnosis, $|D| + 0 = |\bar{D}|$, $D \cap \mathcal{P} = \bar{D} \cap \mathcal{P}^A = \{A\}$, and $D \cap \mathcal{G} = \bar{D} \cap \mathcal{G} = \emptyset$. Further

$$D \cap \mathcal{IN} = \{i : \text{dep}(i) \cap \bar{D} \backslash D \neq \emptyset, i \in \mathcal{IN}\} = \emptyset.$$

This shows that D is a condensed diagnosis, and since $\bar{D}$ is a minimal cardinality global diagnosis, D is a minimal cardinality condensed diagnosis.

The condensed diagnosis $\langle\{A\}, 3\rangle$ represents the first and the second minimal cardinality global diagnoses, while the condensed diagnosis $\langle\{G\}, 3\rangle$ represents the other. ⋄

In the example above, the minimal cardinality condensed diagnoses was a condensed and efficient representation of the minimal cardinality global diagnoses. The next example will be used to exemplify when the minimal cardinality condensed diagnoses is not a condensed and efficient representation.

Example 8: Consider the system shown in Figure 6. The system includes three agents with the sets of private components $\mathcal{P}^{A_1} = \{A\}$, $\mathcal{P}^{A_2} = \{B, C, D\}$, $\mathcal{P}^{A_1} = \{E, F\}$, set of common components $\mathcal{G} = \{G\}$, and the set of signals $\{s_1, s_2, s_3\}$. The assumptions are $\text{ass}(s_1) = \{B, C\}$, $\text{ass}(s_2) = \{C, D\}$, $\text{ass}(s_3) = E$.

The following sets of conflicts have been detected, $\Pi^{A_1} = \{\{s_1, A, G\}, \{s_3, A, G\}\}$, $\Pi^{A_2} = \{\{s_3, D\}\}$, and $\Pi^{A_3} = \{\{E, F\}\}$. The minimal local diagnoses is calculated from the object conflicts resulting in the sets

$$\mathbb{D}^{A_1} = \{\{s_1, s_3\}, \{A\}, \{G\}\} \qquad \mathbb{D}^{A_2} = \{\{s_3\}, \{D\}\}$$
$$\mathbb{D}^{A_3} = \{\{E\}, \{F\}\}.$$

The following sets of tuples are transmitted to agent A_1 from A_2 and A_3.

$$TX^{A_2} = \{\langle\{s_3\}, 0\rangle, \langle\{s_2\}, 0\rangle\} \qquad TX^{A_3} = \{\langle\{s_3\}, 0\rangle, \langle\emptyset, 1\rangle\}$$

The received sets in agent A_1 are

$$RX^{A_1} = \{\langle\{s_1, s_3\}, 0\rangle, \langle\{A\}, 0\rangle, \langle\{G\}, 0\rangle\}$$
$$RX^{A_2} = \{\langle\{s_3\}, 0\rangle, \langle\{s_2\}, 0\rangle\}$$
$$RX^{A_3} = \{\langle\{s_3\}, 0\rangle, \langle\emptyset, 1\rangle\}$$

Resulting set of minimal cardinality condensed diagnoses

$$\mathbb{D}_s^{A_1} = \{\langle\{s_1, s_3\}, 0\rangle, \langle\{s_1, s_2, s_3\}, -1\rangle, \langle\{A, s_3\}, 0\rangle, \langle\{G, s_3\}, 0\rangle, \}$$

where the -1 in the second condensed diagnosis means that the true cardinality is one less than the cardinality for the set $\{s_1, s_2, s_3\}$.

To be able to verify the correctness of the minimal cardinality condensed diagnoses, the minimal cardinality global diagnoses are calculated. The set of conflicts, not including signals, is $\Pi = \{\{B, C, A, G\}, \{E, A, G\}, \{E, D\}, \{E, F\}\}$ and the set of minimal cardinality global diagnoses is

$$\mathcal{D}^{mc} = \{\{B, E\}, \{A, E\}, \{C, E\}, \{G, E\}\}.$$

Is the minimal cardinality condensed diagnoses correct? Consider the condensed diagnosis $D = \langle\{s_1, s_3\}, 0\rangle$, and the minimal cardinality global diagnosis $\bar{D} = \{B, E\}$. Using Definition 4 to verify that D is a condensed diagnosis, $|D|+0 = |\bar{D}|$, $D\cap\mathcal{P} = \bar{D}\cap\mathcal{P}^A = \emptyset$, and $D\cap\mathcal{G} = \bar{D}\cap\mathcal{G} = \emptyset$. Further

$$D \cap \mathcal{IN} = \{i : \mathrm{dep}(i) \cap \bar{D} \backslash D \neq \emptyset, i \in \mathcal{IN}\} = \{s_1, s_3\}$$

since $\mathrm{ass}(s_1) \cap \bar{D} \neq \emptyset$ and $\mathrm{ass}(s_3) \cap \bar{D} \neq \emptyset$. This shows that D is a condensed diagnosis, and since $\bar{D}$ is a minimal cardinality global diagnosis, D is a minimal cardinality condensed diagnosis. ⋄

As can be seen in the above example, there where quite some calculations that had to be performed compared to the calculation of the minimal cardinality global diagnoses. If a system has a high degree of components used by several agents, the minimal cardinality condensed diagnoses will include relatively many components. The reduction of size and the number of diagnoses will in this case be limited, and the efficiency of the algorithm reduced, as was seen in the example.

Considering automotive systems, the ECUs typically have a large number of private components compared to both the number of inputs and the number of common components. It is therefore applicable to use the algorithm for these systems.

7 Conclusions

The objective when designing the algorithm described in Section 5, was to gain a diagnostic algorithm that used low processing power, low memory usage, low network load, and resulted in a transparent, scalable, and interoperable distributed system, see Section 2.

An algorithm has been presented in Section 5, that synchronizes the minimal local diagnoses in a distributed system. The result is a set of minimal cardinality condensed diagnosis, where each minimal cardinality condensed diagnosis is a subset of some minimal cardinality global diagnoses, see Theorem 1 and Lemma 1. The minimal cardinality condensed diagnoses only include components that are used by the specific agent and is therefore a local condensed representation of the minimal cardinality global diagnoses.

A diagnostic system, using the algorithm presented in this paper, is transparent since the loss of an agent would only mean that this agent would not transmit its minimal local diagnoses on the network. It is scalable since it is directly possible to add new ECUs to the network. The algorithm requires a low processing load and memory usage since unwanted private components have been removed from the condensed diagnoses with minimal cardinality.

References

[Biteus *et al.*, 2005] J. Biteus, M. Jensen, and M. Nyberg. Distributed diagnosis for embedded systems in automotive vehicles. In *Proceedings of* IFAC *World Congress'06*, Prague, Czech Republic, 2005.

[Biteus *et al.*, 2006] Jonas Biteus, Erik Frisk, and Mattias Nyberg. Condensed representation of global diagnoses with minimal cardinality in local diagnoses – extended version. Technical report, Dept. of Electrical Engineering, Linköpings Universitet, Sweden, 2006. To be published.

[de Kleer *et al.*, 1992] J. de Kleer, A. K. Mackworth, and R. Reiter. Characterizing diagnoses and systems. *Artificial Intelligence*, 56, 1992.

[de Kleer, 1990] J. de Kleer. Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45:197–222, 1990.

[de Kleer, 1991] J. de Kleer. Focusing on probable diagnoses. In *Proceedings of 9th National Conf. on Artificial Intelligence*, Anaheim, U.S.A., 1991.

[Hayes, 1999] C. C. Hayes. Agents in a nutshell-a very brief introduction. *Knowledge and Data Engineering*, IEEE *Transactions on*, 11(1):127–132, Jan/Feb 1999.

[Köppen-Seliger *et al.*, 2003] B. Köppen-Seliger, T. Marcu, M. Capobianco, S. Gentil, M. Albert, and S. Latzel. MAGIC: An integrated approach for diagnostic data management and operator support. In *Proceedings of* IFAC *Safeprocess'03*, Washington, U.S.A., 2003.

[Provan, 2002] G. Provan. A model-based diagnosis framework for distributed systems. In *13th International Workshop on Principles of Diagnosis*, Semmering, Austria, May 2002.

[Reiter, 1987] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, Apr 1987.

[Roos *et al.*, 2002] N. Roos, A. ten Teije, A. Bos, and C. Witteveen. An analysis of multi-agent diagnosis. In *Proceedings of the Conference on Autonomous Agents and Mult-Agent Systems*, Bologna, Italy, Jul 2002.

[Roos *et al.*, 2003] N. Roos, A. ten Teije, and C. Witteveen. A protocol for multi-agent diagnosis with spatially distributed knowledge. In *Proceedings of 2nd Conference on Autonomous Agents and Mult-Agent Systems*, Australia, Jul 2003.

[Tanenbaum and van Steen, 2002] Andrew S. Tanenbaum and Maarten van Steen. *Distributed Systems*. Prentice Hall, 2002.

[Weiss, 1999] G. Weiss, editor. *Multiagent systems : a modern approach to distributed artificial intelligence*. MIT Press, Cambridge, Mass., U.S.A, 1999.

[Wotawa, 2001] F. Wotawa. A variant of Reiter's hitting-set algorithm. *Information Processing Letters*, 79, 2001.