A Spectrum of Diagnosis Approaches

Alban Grastien[1,2]
[1]Optimisation Research Group, NICTA, Australia
[1]Artificial Intelligence Group, Australian National University, Australia
e-mail: first.last@nicta.com.au

Abstract

This article argues that, regardless of the modeling framework, most diagnosis approaches can be classified in three categories or a combination of those categories. The purpose of this article is therefore to improve the understanding of diagnostic techniques used across the different modeling frameworks.

1 Introduction

Several aspects of diagnosis have been studied by the research community: problem modeling, i.e., how to represent the system, the observations, and the problem that is to be solved; problem solving, i.e., how to find the solution to the diagnosis problem; and diagnosis properties such as diagnosability, minimisation of sensors, etc. This paper concentrates on the second issue.

Much work has been dedicated to solving diagnosis problems because it is hard in general. In many situations, naive, brute-force, approaches are simply not practical, and even sophisticated methods fail to scale to real-world problems.

This paper presents a classification of diagnosis approaches in three categories: system state tracking, fault mode characterisations, and consistency checking. The objective here is to draw a coherent picture of the approaches that have been developed independently in different sub-communities primarily defined by the system modeling framework they rely on. For instance diagnosis of discrete event systems has long ignored the 'consistency checking' approach mostly for historical reasons (the seminal work been developed by Control theorists).

The document is divided as follows. Section 2 gives a generic definition of diagnosis. Section 3 presents the three approaches to diagnosis. Section 4 digresses on how diagnosability can influence the diagnostic approach. Finally Section 5 illustrates how algorithms have been developed for the three approaches, regardless of the modeling framework.

2 A Diagnosis Definition

The definition of diagnosis presented here and used later in the paper is very generic and practical implementations may be quite different from the definitions.

Model-based diagnosis uses a *model* that represents the possible set of system behaviours. Let us write $\mathbb{B}$ the set of behaviours allowed by the model and $\beta \in \mathbb{B}$ a single behaviour. The model also indicates what observations a given behaviour will produce. The *observation* (assumed deterministic) of behaviour β is $\sigma = obs(\beta)$ and the set of possible observations is $\mathbb{O} \ni \sigma$. Finally, diagnosis is about detecting and identifying the fault mode that the system is running under. The *fault mode* of behaviour β is $\delta = mode(\beta)$ and the set of modes is $\mathbb{M} \ni \delta$, which is generally the power set of a set of faults.

A diagnosis problem is defined as a model and an observation. The (consistency-based) diagnosis is defined as the set of modes that are logically consistent with the model and the observation, or a proper subset of this set, i.e., the set of modes exemplified by a behaviour consistent with the observation.

$$\Delta(\langle \mathbb{B}, obs, mode, \sigma \rangle) = \{\delta \in \mathbb{M} \mid \exists \beta \in \mathbb{B}.\ obs(\beta) = \sigma \wedge mode(\beta) = \delta\}. \tag{1}$$

3 Three Approaches

3.1 System State Tracking

The first approach to solve diagnosis is to have a literal interpretation of Equation 1. In this approach, the diagnosis is computed in two successive phases:

1. The set of system behaviours consistent with the observations is computed. We should call these behaviours *explanations* (although they should not be mistaken with the homonymous term from abductive diagnosis):
$$Expl = \{\beta \in \mathbb{B} \mid obs(\beta) = \sigma\}.$$
2. The diagnosis is then extracted from the set of explanations. It is generally assumed that this task is rather easy since it simply consists in ignoring the non-diagnosis-related information (internal variables) of the explanations:
$$\Delta = \{\delta \in \mathbb{M} \mid \exists \beta \in \mathbb{B}.\ mode(\beta) = \delta\}.$$

One clear advantage of this approach is that it is quite interactive. Because the set of explanations is computed, it is possible to request more details on-the-fly (for instance, return a full behaviour of the system that supports the diagnosis) or to refine the diagnosis by adding more information (a typical application being the sequential diagnosis or the incremental diagnosis for dynamic systems). Practical implementations often keep only the relevant part of the explanations: for dynamic systems, only the final belief state.

The main issue of this approach is that it requires some heavy on-line processing. Computing *Expl* can be quite demanding and even the second phase can be hard (in particular if *Expl* is exponentially large compared to the original model). Most of the effort in this approach has therefore been spent on finding representations of *Expl* that i) are compact and that ii) allow for easy extraction of the diagnostic information. This approach also sometimes requires a strong-fault model, i.e., a model that can predict the behaviour of the system in faulty situations. Another issue is that it can be hard to distribute the computation (for instance, using multi-core processors).

3.2 Fault Mode Characterisation

The second approach to solve diagnosis consists in computing and using distinctive features of observations of fault modes. For instance, certain alarms may be explicit enough that the diagnosis can be derived without finding a precise explanation that supports this diagnosis; a perhaps extreme example would be an 100% reliable alarm "component A broken", which would be sufficient to produce the diagnosis "component A broken" without need to explain all other observations.

One way to formalise this approach would be to define a set of pair $\langle \mathbb{O}_i, \mathbb{M}_i \rangle$ with the following semantics:

$$\sigma \in \mathbb{O}_i \Rightarrow \Delta \cap \mathbb{M}_i = \emptyset.$$

For instance, if two sensors are supposed to return the same value in nominal condition, then $\mathbb{O}_i$ would represent the set of observations where the two sensors return different values and $\mathbb{M}_i$ would be (a subset of) the set that does not contain the nominal mode. It is assumed that $\mathbb{O}_i$ can be represented compactly and the inclusion $\sigma \in \mathbb{O}_i$ can be tested efficiently. Notice that the model $\mathbb{B}$ is missing in this definition, which does not mean that the approach is not model-based since i) the characterisations can be generated from a model and ii) the characterisations can be seen as models.

This approach is very appealing because it is generally quite efficient (most of the work is done off-line). Furthermore the characterisations are often readable and can be validated by an expert or even hand-written by an expert.

There are several issues with this approach. First, the number of characterisations is often very large, for instance if a unique characterisation is defined for each fault mode, and the number of fault mode is exponential (power set of a set of faults). Worst though, there are often many characterisations per fault mode (or per fault), especially in highly reconfigurable systems or systems that can be used, or react differently, in a wide range of situations.

A second line of research for this approach is the automatic generation of the characterisations.

On a side note, state-based observations can be seen as particular cases of characterisation (each state is characterised by what observations can be made in this state). Such characterisation can be trivial (problem inputs [Cordier and Largouët, 2001]) or not [Blackhall and Kan John, 2008; Daigle *et al.*, 2010; Bayoudh *et al.*, 2009].

3.3 Consistency Checking

The third approach to diagnosis is based on consistency checks. Each check is defined as a set of assumptions on the system behaviour. An oracle (for instance, a theorem prover) verifies whether these assumptions are consistent with the model and the observations. The checks are generated dynamically and the diagnosis is deduced from the result of the checks.

Formally, a check is a subset of system modes $\mathbb{M}_i \subseteq \mathbb{M}$ and the result of the check is i) either a mode $\delta \in \mathbb{M}_i \cap \Delta$ if one such mode exists or ii) a superset $\mathbb{M}'_i \supseteq \mathbb{M}_i$ such that $\mathbb{M}'_i \cap \Delta = \emptyset$ ($\mathbb{M}'_i$ is called a *conflict*; in the worst-case $\mathbb{M}'_i = \mathbb{M}_i$).

This approach is very appealing in that it can adapt quite efficiently to the problem at hand. For instance, assume that i) the system behaviour is nominal and ii) we are looking only for the minimal diagnosis. Then the diagnosis can be found with only one consistency check. The other two approaches on the other hand could require quite expensive (pre-)computation.

There are several issues with this approach however: it is unable to generate all the diagnoses in general, as this would require an exponential number of checks; it is often limited to (possibly all) minimal diagnoses. Another problem is that the computation time is very unpredictable in general. An issue that is being addressed is the solving of the checks, especially for dynamic systems (reduces to model-checking).

3.4 Discussion

The three categories of algorithms provided above bear some conceptual similarities but behave quite differently. State tracking and fault mode characterisation can be similar in that the characterisations are often computed by computing all possible results a state tracking could provide (typically, the Sampath diagnoser for DES [Sampath *et al.*, 1995]). On the other hand, a characterisation can be computed by asserting assumptions (as in consistency checking) and deriving the logical consequences.

Fault mode characterisation relies on significant off-line preprocessing but expects little on-line operations (which is why it has been pushed by the Control community). On the other hand, system tracking and consistency checking require an important computation at runtime but expect that the specificity of the diagnostic problem, i.e., the observations of the behaviour to explain, will allow for swift resolution.

In this sense, any off-line preprocessing aimed at altering the model (for instance, model abstractions) to simplify the diagnosis problem has a fault mode characterisation flavor.

4 The Role of Diagnosability

Diagnosability is the property that the faults can be identified when they occur on the system. This property makes the system easier to control and efforts should be made at design stage to ensure diagnosability.

There is another benefit from diagnosability that was recognised in the last decade, which is that diagnosability can make the diagnosis problem easier to solve.

Essentially, diagnosability tells us that there is a (precise) solution to the problem, and it is therefore worth searching of it.

Rather than performing diagnosis with a set of fault modes $\mathbb{M}$, one might actually want to solve k independent diagnosis subproblems with different fault modes $\mathbb{M}_i$ such that $\mathbb{M} \subseteq \mathbb{M}_1 \times \ldots \times \mathbb{M}_k$. The diagnosis is then the set of modes accepted by each subdiagnosis:

$$\Delta = \mathbb{M} \cap (\Delta_i \times \ldots \times \Delta_k) .$$

This decomposition improves performance if the time necessary to solve the diagnosis problem with $\mathbb{M}$ is bigger than that of solving all the subproblems. This generally requires i) that $|\mathbb{M}| \gg |\mathbb{M}_1| + \ldots + |\mathbb{M}_k|$ and ii) that abstractions can be applied for each subproblem.

The potential issue here is a decrease in precision. Consider a diagnosis with two faults a and b and four fault modes: $\mathbb{M} = \{\{\neg a, \neg b\}, \{a, \neg b\}, \{\neg a, b\}, \{a, b\}\}$. This problem can be decomposed in two subproblems each dedicated to one specific fault: the first subproblem has modes $\{\neg a\}$, $\{a\}$ and the second one has modes $\{\neg b\}$, $\{b\}$. If the diagnoses are $\{\neg a, b\}$ and $\{a, \neg b\}$ (exactly one fault occurred), then the diagnoses will be $\{\neg a\}$, $\{a\}$ for the first subproblem and $\{\neg b\}$, $\{b\}$ for the second one, which will translate to $\Delta = \mathbb{M}$, hence the precision loss.

It has been recognised that, if each fault is diagnosable, then the decomposition presented above is precise [Cordier *et al.*, 2006], and this result has been extended to other decompositions [Grastien and Torta, 2011].

Decompositions as presented above have been widely used. Characterisation methods can be seen as implementations of such decompositions: the set $\mathbb{M}_i$ presented in Subsection 3.2 bears many similarities with decompositions. The abstraction is implicit in this type of approach since the characterisations will generally refer to a specific subset of the sensors.

However characterisations-based approaches are not the only ones that can benefit from decomposition. Since each subproblem can be solved in parallel, tracking-based approaches and consistency-based approaches can benefit from it (consistency-based approaches can be seen as a dynamic implementation of this scheme with binary sets $\mathbb{M}_i$). Notice however that when abstraction is applied (which is one reason to perform decomposition, at least for the tracking approaches), this can reduce the ability to produce global explanations that support a diagnosis, since the diagnosis is not computed globally.

Diagnosability can also be used to parallelise computation in consistency checking approaches. For instance, if fault f is diagnosable, one can start two searches in parallel, one looking for an explanation including fault f and the other one excluding f; the solution will be found as soon as one search terminates (either by finding an explanation, or by proving that no such explanation exists).

5 Approaches within Modeling Frameworks

The purpose of this section is to illustrate how those different algorithms have been implemented in different modeling frameworks.

5.1 Circuits

The Theory of diagnosis from first principles from Reiter [1987] is an implementation of the consistency checking approach. Starting from the best diagnosis (no fault), consistency of a fault mode is tested and certified or nullified together with a larger set of modes through a conflict. Under the assumption that the model is weakly-faulty (the effect of the faults on the model is not modeled), this algorithm even returns all the diagnoses.

More recently, SAFARI [Feldman *et al.*, 2010] computes some diagnoses and tries to improve them by removing faults that are not essential in the current mode. Reduction to SAT has been proposed [Smith *et al.*, 2005] and recently made scalable [Metodi *et al.*, 2012]. Those algorithms are based on consistency-checking.

The approach from De Kleer and Williams [1987] is very interesting in that it does not quite fall in any of the approaches presented in this paper. It consists in using an ATMS to derive facts from any assumption on the faults. Conflicts are then generated, which allow to compute the diagnoses (the hitting sets of the conflicts). It sits somewhere between state tracking and consistency checking: it derives all the facts it can (which is similar to the tracking approach) but it also dynamically generates conflicts before producing the diagnosis.

The work initiated by Provan and Darwiche sits between the tracking zone and the mode characterisation: on the one hand, the goal is to find a compact representation of the system belief state, on the other hand, a great deal of energy is performed off-line to build a structure that will allow this on-line efficiency, including by applying abstraction [Provan, 2001]. The model, the observations, and the fault modes are all propositional formula, and the diagnosis is computed by computing the conjunction of the model and the observations and projecting the result on the mode variables [Huang and Darwiche, 1996; Darwiche, 1998]. Because such an algorithm is exponential in general, and often also in practice, efforts were made to identified classes of representations for propositional formula (improved BDDs [Torta and Torasso, 2004]) that would limit the constraints on the algorithm [Darwiche and Marquis, 2002; Huang and Darwiche, 2007; Siddiqi and Huang, 2011]. Substantial work was also dedicated to finding structural properties that can be used to speed up the computation of diagnosis [Huang and Darwiche, 1996; Darwiche, 1998; Siddiqi and Huang, 2007; Siddiqi, 2011].

5.2 Discrete Event Systems

Discrete event systems (DES) are dynamic systems whose evolution is modeled by a finite set of events.

The first works in (model-based) diagnosis of DES have used characterisation and are the chronicles and the Sampath diagnoser.

A chronicle is a collection of observable events that are temporally constraints [Cordier and Dousson, 2000]. If this collection of events is recognised in the flow of observations generated by the system, what information associated with the chronicle (generally a fault) is identified. The work in this topic has concentrated on making the recognition fast [Dousson and Le Maigat, 2007] and on generating the chronicles automatically [Guerraz and Dousson, 2004; Pencolé and Subias, 2009]

The Sampath diagnoser is an extreme version of characterisation [Sampath *et al.*, 1995]. For each fault mode δ, the set $\mathbb{O}_\delta$ of the observations consistent with this mode is generated: $\mathbb{O}_\delta = \{\sigma \in \mathbb{O} \mid \exists \beta \in \mathbb{B}.\ \sigma = obs(\beta) \wedge \delta = mode(\beta)\}$. These sets are languages that can be represented by DFAs (deterministic finite automata, assuming the state space of the DES was finite) where each state is labeled with the diagnosis. Given a sequence of observations, the diagnosis algorithm only needs to follow the single path on the diagnoser associated with these observations, and extract the diagnosis at the final state. The main problem of the diagnoser is its size: exponential in the number of state in the DES and in the number of fault modes, i.e., generally double exponential in the number of state variables and in the number of faults, which makes it impractical but for tiny problems [Rintanen, 2007]. There has been some work on reducing the size of the diagnoser using abstraction (bisimulation) [Marchand and Rozé, 2002], using BDDs [Schumann *et al.*, 2004], or by decomposing the fault modes [Pencolé *et al.*, 2006]. In this last paper, the decomposition is based on building a diagnoser for each fault; there could be potential for research here in building several diagnosers on different types of decomposition (cf. [Grastien and Torta, 2011]).

The Artificial Intelligence community has mostly worked on the tracking approach [Zanella and Lamperti, 2003]. Because this approach is very expensive in general, many algorithms have been used to exploit the symmetries in the model, e.g., decentralised/distributed techniques [Pencolé and Cordier, 2005; Su and Wonham, 2005; Cordier and Grastien, 2007; Kan John and Grastien, 2008; Kan John *et al.*, 2010], symbolic tools (BDDs) [Schumann *et al.*, 2007], Petri nets [Aghasaryan *et al.*, 1998; Jiroveanu and Boël, 2005; Cabasino *et al.*, 2010].

The work on the consistency-based approach has mostly concentrated on the implementation of the consistency checker, either a planner [McIlraith, 1998; Sohrabi *et al.*, 2010; Haslum and Grastien, 2011], a model-checker [Cordier and Largouët, 2001], a SAT solver [Grastien *et al.*, 2007] or an integer linear programming solver [Basile *et al.*, 2009; Dotoli *et al.*, 2009]. The second level of the approach, i.e., what consistency checks should be generated is a relatively novel topic [Grastien *et al.*, 2011; 2012].

There has been some attempts at combining different approaches. Lamperti and Zanella compute the diagnosis with the state-tracking approach, but reuse the computed diagnosis to build characterisations [Lamperti and Zanella, 2004].

It should be possible to combine characterisations with other approaches: for instance, use characterisation to quickly find a diagnosis, and use more expensive approaches were this one to fail. Similarly consistency-based approaches could implement decentralised approaches used in state tracking, by trying to prove/disprove the occurrence of a fault locally, before considering more system-wide model.

5.3 Continuous and Hybrid Systems

One obvious way to diagnose continuous and hybrid systems is to discretise them and use techniques from diagnosis of DES [Bresolin and Capiluppi, 2011]

Characterisation approaches are often used for diagnosis of continuous and hybrid systems, possibly because they run very fast, once the characterisations have been found. These characterisations are known as possible conflicts [Pulido and Alonso González, 2004], analytical redundancy relations [Staroswiecki and Comtet-Varga, 2001], minimal structurally overdetermined sets [Krysander *et al.*, 2010], etc.

In essence, these *indicators* are built by infering the state of internal variables from some observations and predicting other observations from these inferences. However certain functions are difficult to invert and work has been devoted to generate these indicators nevertheless [de Flaugergues *et al.*, 2010]. The second major issue is coping with the number of indicators by generating them on-the-fly [Heintz *et al.*, 2008; Bayoudh *et al.*, 2009].

The tracking approach has also been used extensively. Multi-mode Kalman filters compute the probability distribution on the state space [Blom and Bar-Shalom, 1988] and most of the work has been dedicated to finding good approximations of these distributions [Benazera and Travé-Massuyès, 2009]. Particle filters [Arulampalam *et al.*, 2002] samples the probability distribution to keep the representation of this distribution small and the simulation easy. Similarly, HyDE [Narasimhan and Brownston, 2007] maintains a list of candidates that are discarded when their weight goes below a threshold.

The consistency checking approach has been proposed by Dague [1994] but seems to have been abandonned until recently. Ernits & Dearden [2011] use an SMT solver to solve the diagnosis checks. Both approaches are done on the steady state.

6 Conclusion

This article argued that diagnosis algorithms can be classified in three categories: tracking the state of the system, characterising the different system modes, and performing consistency checks between the model, the observations, and some assumptions.

Interestingly, the same approaches have been developed regardless of the modeling framework. Researchers working on different types of model but with the same approach to diagnose have been faced with the same issues.

Acknowledgments

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

[Aghasaryan *et al.*, 1998] A. Aghasaryan, É. Fabre, A. Benveniste, R. Boubour, and C. Jard. Fault detection and diagnosis in distributed systems: an approach by partially stochastic Petri Nets. *Journal of Discrete Event Dynamical Systems (JDEDS)*, 8(2):203–231, 1998.

[Arulampalam *et al.*, 2002] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing (TSP)*, 50(2):174–188, 2002.

[Basile *et al.*, 2009] F. Basile, P. Chiacchio, and G. De Tommasi. An efficient approach for online diagnosis of discrete event systems. *IEEE Transactions on Automatic Control (TAC)*, 54(4):748–759, 2009.

[Bayoudh *et al.*, 2009] M. Bayoudh, L. Travé-Massuyès, and X. Olive. Diagnosis of a class of non linear hybrid systems by on-line instantiation of parameterized analytic redundancy relations. In *20th International Workshop on Principles of Diagnosis (DX-09)*, pages 283–290, 2009.

[Benazera and Travé-Massuyès, 2009] E. Benazera and L. Travé-Massuyès. Set-theoretic estimation of hybrid system configurations. *IEEE Transactions on Systems, Man, and Cybernetics (TSMC)*, 39(5):1277–1290, 2009.

[Blackhall and Kan John, 2008] L. Blackhall and P. Kan John. Model-based diagnosis of hybrid dynamical networks for fault tolerant control. In *19th International Workshop on Principles of Diagnosis (DX-08)*, pages 239–246, 2008.

[Blom and Bar-Shalom, 1988] H. Blom and Y. Bar-Shalom. The interacting multiple model algorithm for systems with Markovian switching coefficients. *IEEE Transactions on Automatic Control (TAC)*, 33(8):780–783, 1988.

[Bresolin and Capiluppi, 2011] D. Bresolin and M. Capiluppi. A game-theoretic approach to fault diagnosis of hybrid systems. In *Second International Symposium on Games, Automata, Logics and Formal Verification (GandALF-11)*, pages 237–249, 2011.

[Cabasino *et al.*, 2010] M. P. Cabasino, A. Giua, and C. Seatzu. Fault detection for discrete event systems using Petri nets with unobservable transitions. *Automatica*, 46(9):1531–1539, 2010.

[Cordier and Dousson, 2000] M.-O. Cordier and Ch. Dousson. Alarm driven monitoring based on chronicles. In *Fourth IFAC Symposium on Fault Detection, Supervision and Safety of Technical Processes (SafeProcess-00)*, pages 286–291, 2000.

[Cordier and Grastien, 2007] M.-O. Cordier and A. Grastien. Exploiting independence in a decentralised and incremental approach of diagnosis. In *20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, 2007.

[Cordier and Largouët, 2001] M.-O. Cordier and Ch. Largouët. Using model-checking techniques for diagnosing discrete-event systems. In *12th International Workshop on Principles of Diagnosis (DX-01)*, pages 39–46, 2001.

[Cordier *et al.*, 2006] M.-O. Cordier, L. Travé-Massuyès, and X. Pucel. Comparing diagnosability in continuous and discrete-event systems. In *17th International Workshop on Principles of Diagnosis (DX-06)*, pages 55–60, 2006.

[Dague, 1994] Ph. Dague. Model-based diagnosis of analog electronic circuits. *Annals of Mathematics and Artificial Intelligence (AMAI)*, 11:439–492, 1994.

[Daigle *et al.*, 2010] M. Daigle, I. Roychoudhury, G. Biswas, and X. Koutsoukos. An event-based approach to distributed diagnosis of continuous systems. In *21st International Workshop on Principles of Diagnosis (DX-10)*, pages 15–22, 2010.

[Darwiche and Marquis, 2002] A. Darwiche and P. Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research (JAIR)*, 17:229–264, 2002.

[Darwiche, 1998] A. Darwiche. Model-based diagnosis using structured system descriptions. *Journal of Artificial Intelligence Research (JAIR)*, 8:165–222, 1998.

[de Flaugergues *et al.*, 2010] V. de Flaugergues, V. Cocquempot, M. Bayart, and M. Pengov. On non-invertibilities for structural analysis. In *21st International Workshop on Principles of Diagnosis (DX-10)*, pages 23–30, 2010.

[de Kleer and Williams, 1987] J. de Kleer and B. Williams. Diagnosing multiple faults. *Artificial Intelligence (AIJ)*, 32:97–130, 1987.

[Dotoli *et al.*, 2009] M. Dotoli, M. P. Fanti, A. M. Mangini, and W. Ukovich. On-line detection in discrete event systems by Petri nets and integer linear programming. *Automatica*, 45:2665–2672, 2009.

[Dousson and Le Maigat, 2007] Chr. Dousson and P. Le Maigat. Chronicle recognition improvement using temporal focusing and hierarchization. In *20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 324–329, 2007.

[Ernits and Dearden, 2011] J. Ernits and R. Dearden. Towards diagnosis modulo theories. In *22nd International Workshop on Principles of Diagnosis (DX-11)*, pages 249–256, 2011.

[Feldman *et al.*, 2010] A. Feldman, G. Provan, and A. van Gemund. Approximate model-based diagnosis using greedy stochastic search. *Journal of Artificial Intelligence Research (JAIR)*, 38:371–413, 2010.

[Grastien and Torta, 2011] A. Grastien and G. Torta. Reformulation for the diagnosis of discrete-event systems. In *Ninth Symposium on Abstraction, Reformulation and Approximation (SARA-11)*, pages 42–49, 2011.

[Grastien *et al.*, 2007] A. Grastien, Anbulagan, J. Rintanen, and E. Kelareva. Diagnosis of discrete-event systems using satisfiability algorithms. In *22nd Conference on Artificial Intelligence (AAAI-07)*, pages 305–310, 2007.

[Grastien *et al.*, 2011] A. Grastien, P. Haslum, and S. Thiébaux. Exhaustive diagnosis of discrete event systems through exploration of the hypothesis space. In *22nd International Workshop on Principles of Diagnosis (DX-11)*, pages 60–67, 2011.

[Grastien *et al.*, 2012] A. Grastien, P. Haslum, and S. Thiébaux. Conflict-based diagnosis of discrete event systems: theory and practice. In *13th International Conference on the Principles of Knowledge Representation and Reasoning (KR-12)*, 2012.

[Guerraz and Dousson, 2004] B. Guerraz and Ch. Dousson. Chronicles construction starting from the fault model of the system to diagnose. In *15th International Workshop on Principles of Diagnosis (DX-04)*, pages 51–56, 2004.

[Haslum and Grastien, 2011] P. Haslum and A. Grastien. Diagnosis as planning: two case studies. In *Fifth Scheduling and Planning Applications Workshop (SPARK-11)*, pages 37–44, 2011.

[Heintz *et al.*, 2008] F. Heintz, M. Krysander, J. Roll, and E. Frisk. FlexDx: a reconfigurable diagnosis framework. In *19th International Workshop on Principles of Diagnosis (DX-08)*, pages 79–86, 2008.

[Huang and Darwiche, 1996] C. Huang and A. Darwiche. Inference in belief networks: a procedural guide. *International Journal of Approximate Reasoning (IJAR)*, 15(3):225–263, 1996.

[Huang and Darwiche, 2007] J. Huang and A. Darwiche. The language of search. *Journal of Artificial Intelligence Research (JAIR)*, 29:191–219, 2007.

[Jiroveanu and Boël, 2005] G. Jiroveanu and R. Boël. Petri net model-based distributed diagnosis for large interacting systems. In *16th International Workshop on Principles of Diagnosis (DX-05)*, pages 25–30, 2005.

[Kan John and Grastien, 2008] P. Kan John and A. Grastien. Local consistency and junction tree for diagnosis of discrete-event systems. In *18th European Conference on Artificial Intelligence (ECAI-08)*, 2008.

[Kan John *et al.*, 2010] P. Kan John, A. Grastien, and Y. Pencolé. Synthesis of a distributed and accurate diagnoser. In *21st International Workshop on Principles of Diagnosis (DX-10)*, pages 209–216, 2010.

[Krysander *et al.*, 2010] M. Krysander, J. Åslund, and E. Frisk. A structural algorithm for finding testable submodels and multiple fault isolability analysis. In *21st International Workshop on Principles of Diagnosis (DX-10)*, pages 79–86, 2010.

[Lamperti and Zanella, 2004] G. Lamperti and M. Zanella. A bridged diagnostic method for the monitoring of polymorphic discrete-event systems. *IEEE Transactions on Systems, Man, and Cybernetics (TSMC)*, 34(5):2222–2244, 2004.

[Marchand and Rozé, 2002] H. Marchand and L. Rozé. Diagnostic de pannes sur des systèmes à événements discrets : une approche à base de modèles symboliques. In *13th Congrès francophone AFRIF-AFIA (RFIA-02)*, pages 191–200, 2002.

[McIlraith, 1998] Sh. McIlraith. Explanatory diagnosis: conjecturing actions to explain observations. In *Sixth International Conference on the Principles of Knowledge Representation and Reasoning (KR-98)*, pages 167–177, 1998.

[Metodi *et al.*, 2012] A. Metodi, R. Stern, M. Kalech, and M. Codish. Compiling model-based diagnosis to Boolean satisfaction. In *26th Conference on Artificial Intelligence (AAAI-12)*, 2012.

[Narasimhan and Brownston, 2007] S. Narasimhan and L. Brownston. HyDE - a general framework for stochastic and hybrid model-based diagnosis. In *18th International Workshop on Principles of Diagnosis (DX-07)*, pages 162–169, 2007.

[Pencolé and Cordier, 2005] Y. Pencolé and M.-O. Cordier. A formal framework for the decentralised diagnosis of large scale discrete event systems and its application to telecommunication networks. *Artificial Intelligence (AIJ)*, 164(1–2):121–170, 2005.

[Pencolé and Subias, 2009] Y. Pencolé and A. Subias. A chronicle-based diagnosability approach for discrete timed-event systems: application to web-services. *Journal of Universal Computer Science (jucs)*, 15(17):3246–3272, 2009.

[Pencolé *et al.*, 2006] Y. Pencolé, D. Kamenetsky, and A. Schumann. Towards low-cost fault diagnosis in large component-based systems. In *Sixth IFAC Symposium on Fault Detection, Supervision and Safety of Technical Processes (SafeProcess-06)*, 2006.

[Provan, 2001] G. Provan. Hierarchical model-based diagnosis. In *12th International Workshop on Principles of Diagnosis (DX-01)*, pages 167–174, 2001.

[Pulido and Alonso González, 2004] B. Pulido and C. Alonso González. Possible conflicts: a compilation technique for consistency-based diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics (TSMC)*, 34(5):2192–2206, 2004.

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

[Rintanen, 2007] J. Rintanen. Diagnosers and diagnosability of succinct transition systems. In *20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 538–544, 2007.

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

[Schumann *et al.*, 2004] A. Schumann, Y. Pencolé, and S. Thiébaux. Symbolic models for diagnosing discrete-event systems. In *16th European Conference on Artificial Intelligence (ECAI-04)*, pages 1085–1086, 2004.

[Schumann *et al.*, 2007] A. Schumann, Y. Pencolé, and S. Thiébaux. A spectrum of symbolic on-line diagnosis approaches. In *22nd Conference on Artificial Intelligence (AAAI-07)*, 2007.

[Siddiqi and Huang, 2007] S. Siddiqi and J. Huang. Hierarchical diagnosis of multiple faults. In *20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 581–586, 2007.

[Siddiqi and Huang, 2011] S. Siddiqi and J. Huang. Sequential diagnosis by abstraction. *Journal of Artificial Intelligence Research (JAIR)*, 41:329–365, 2011.

[Siddiqi, 2011] S. Siddiqi. Computing minimum-cardinality diagnoses by model relaxation. In *22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, 2011.

[Smith *et al.*, 2005] A. Smith, A. Veneris, M. F. Ali, and A. Viglas. Fault diagnosis and logic debugging using Boolean satisfiability. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 24(10):1606–1621, 2005.

[Sohrabi *et al.*, 2010] Sh. Sohrabi, J. Baier, and Sh. McIlraith. Diagnosis as planning revisited. In *12th International Conference on the Principles of Knowledge Representation and Reasoning (KR-10)*, pages 26–36, 2010.

[Staroswiecki and Comtet-Varga, 2001] M. Staroswiecki and G. Comtet-Varga. Analytical redundancy relations for fault detection and isolation in algebraic dynamic systems. *Automatica*, 37:687–699, 2001.

[Su and Wonham, 2005] R. Su and W. Wonham. Global and local consistencies in distributed fault diagnosis for discrete-event systems. *IEEE Transactions on Automatic Control (TAC)*, 50(12):1923–1935, 2005.

[Torta and Torasso, 2004] G. Torta and P. Torasso. The role of OBDDs in controlling the complexity of model based diagnosis. In *15th International Workshop on Principles of Diagnosis (DX-04)*, pages 9–14, 2004.

[Zanella and Lamperti, 2003] M. Zanella and G. Lamperti. *Diagnosis of active systems*. Kluwer Academic Publishers, 2003.