Toward Distributed Diagnosis of Complex Physical Systems

Gautam Biswas Sherif Abdelwahed Xenofon Koutsoukos Jyoti Gandhe Eric Manders
Department of Electrical Engineering and Computer Science,
and the Institute for Software Integrated Systems, Vanderbilt University.
Email: gautam.biswas@vanderbilt.edu

Abstract

The size and complexity of present day systems motivates the need for developing distributed fault diagnosis algorithms. This paper uses the TRANSCEND approach to fault diagnosis to develop a methodology that partitions the set of possible fault candidates in a physical system to independent sets of faults given a set of measurements. Separate diagnosers that do not interact with each other can be constructed for each independent fault set while maintaining complete diagnosability of the system. The implication of this approach is that a computationally expensive diagnosis task is decomposed into a set of computationally simpler tasks that can be performed independently.

1 Introduction

In the model based approach to fault diagnosis [5], an important entity is the system model that provides the basis for analyzing observed discrepancies in system behavior for isolating faulty components in the system. For large systems, building accurate models becomes is difficult, and analyzing faulty behaviors is computationally expensive. This motivates the need to develop methodologies for decomposing the diagnosis task into sets of smaller problems that reduce the overall computational complexity of the online fault isolation. This may be achieved by decomposing the system into subsystems, where each subsystem has a diagnosis module, and results of the modules are composed to derive the system-level diagnosis.

Previous work on distributed diagnosis has focused on systems with discrete behaviors, such as a set of interconnected processors [3, 7]. An approach for decentralized diagnosis has been presented in [4] where the local diagnosers communicate with a coordination process that assembles a global diagnosis, but these methods are subject to robustness and scalability issues. Distributed diagnosis approaches where local diagnosers communicate directly with esvh other are presented in [8, 14]. A distributed diagnosis method that does not require coordination between local diagnosers has been proposed in [2], but the structure of the local models makes the local diagnosis and communication complicated.

In this paper, we extend previous work and develop a distributed diagnosis solution for continuous systems. In continuous systems, the energy flow between components results in fault effects propagating to all parts of the system as time progresses. Therefore, the traditional notion of independence among subsystems does not apply to the fault isolation task. Our approach is to partition the set of faults into subsets so that we can construct independent diagnosers for each subset. Two diagnosers are independent if they do not have to share information in establishing unique diagnosis results that are globally valid. We establish this by ensuring that the two diagnosers do not require the same set of measurements to achieve complete diagnosability. Complete diagnosability is the ability to uniquely isolate every fault candidate in the system given a set of measurements. Although such system decomposition will not always be possible, our approach can be the first step for designing distributed diagnosers.

Intuitively, fewer measurements imply the number of independent fault sets that can be constructed will be smaller, and the number of faults in each subset will be larger. Assuming the complexity of each diagnoser is directly linked to the number of faults it has to isolate, fewer measurements will produce more complex diagnosers. In the other extreme, by choosing a large number of relevant measurements, one may establish a one-to-one relationship between every fault and a corresponding measurement. This implies that we can construct a diagnoser for each fault. In practice, this will not be feasible because not all system variables are accessible, and the cost of placing sensors in particular locations may be prohibitive. Our proposed solution assumes we have a sufficient measurement set that makes the system completely diagnosable. We apply the notion of independent fault sets to develop the distributed diagnosis scheme as an extension to TRANSCEND , a qualitative model-based fault isolation scheme [11].

2 Transcend Architecture

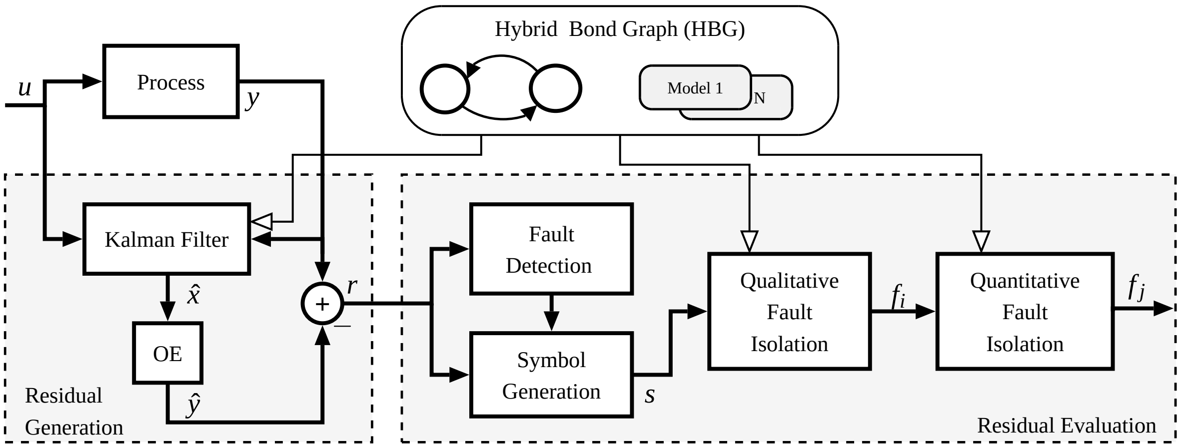

Figure 1: Architecture for TRANSCEND .

The TRANSCEND architecture, shown in Fig. 1 combines a qualitative and quantitative, model-based approach for isolation and identification of abrupt faults in process components. System models are constructed using bond graphs [6] that capture dynamic interactions between components of a physical system using a topological description. The faults correspond to component parameter changes in the bond graph. Furthermore, TRANSCEND focuses on faults that are modeled as discrete and persistent changes in component parameters, referred to as abrupt faults.

Definition 2.1 An *abrupt fault* is a change in a component parameter value that occurs at a much faster rate than the nominal dynamics of the system. ■

The occurrence of an abrupt fault results in transient behavior in the system. Fault isolation in TRANSCEND is based on the analysis of the fault transient dynamics using a qualitative framework. Specifically, the magnitude and slope of the transient residual, derived from measurements, are mapped onto $(+, 0, -)$ symbols (after energy-based filtering [9]) for qualitative matching against fault signatures.

We illustrate fault isolation in TRANSCEND , and the new algorithms developed in this paper, using a hypothetical physical system that consists of six coupled fluid tanks. We model this system using the bond graph modeling paradigm. Fig. 2 shows a single tank fluid system with a single in- and out-flow, with its corresponding bond graph. The tank is modeled as a capacity, C, the inflow as a source of flow, f, and the pipe as a resistance, R. The 0−junction represents a lossless common pressure (effort) point, where energy transfer takes place among components.

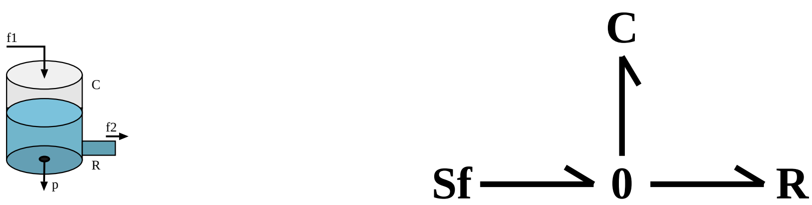

Figure 2: One-tank fluid system and bond graph.

In the multi-tank system, the tanks are connected with each other by pipes, with a source of flow into the first tank and a pipe draining each of the tanks. All the pipes are modeled by linear resistances and tanks are modeled as capacities, making this a sixth order system. Fig. 2 show this system with its corresponding bond graph. Pipe R_i drains tank C_i and pipe R_{ij} connects tanks C_i and C_j. The set possible faults contains all component parameters in the model, i.e. $F = \{C_1, \ldots, C_6, R_1, \ldots, R_6\ R_{12}, \ldots, R_{56}\}$. The set of possible measurements consists of the pressures in the tanks and flows through the pipes, i.e. $M = \{e1, e6, e11, e16, e21, e26, f2, f7, f12, f17, f22, f27, f4, f9, f14, f19, f24, f27\}$.

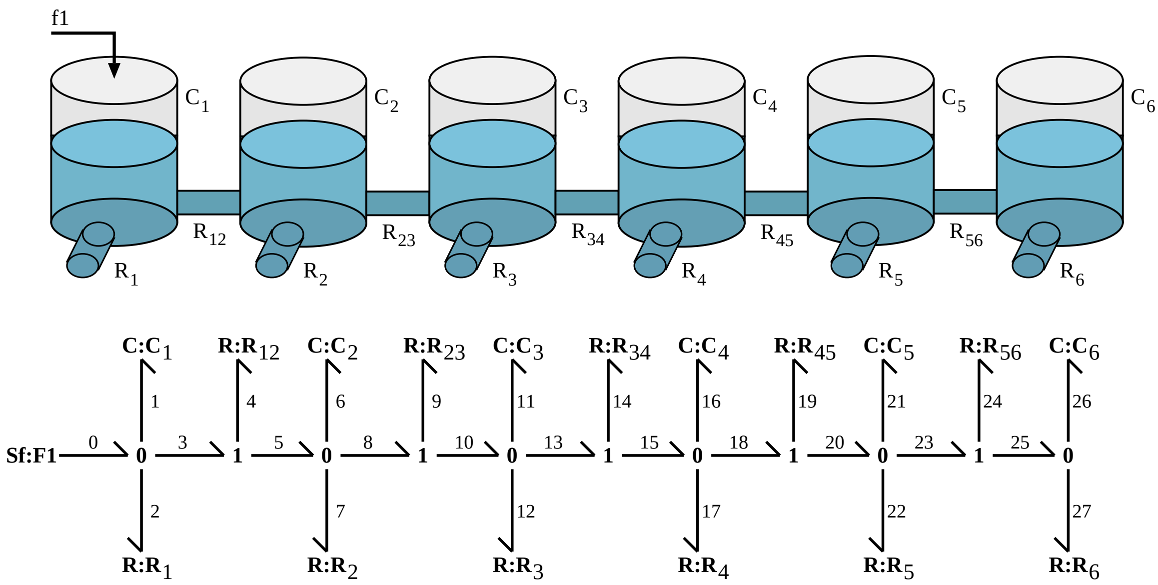

Figure 3: Six-tank system diagram and corresponding bond graph model.

Qualitative Fault Isolation

Fault isolation is developed as a graph based algorithm that operates on the Temporal Causal Graph (TCG), that is derived automatically from the bond graph [11].

Definition 2.2 A *Temporal Causal Graph* (TCG) is a directed graph $< V, L, D >$, where $V = E \cup F$ (E is a set of efforts and F is a set of flows of a bond graph), L is the label set $\{=, 1, -1, p, p^{-1}, pdt, p^{-1}dt\}$ (p is a parameter name of the physical system model). The dt specifier indicates a temporal edge relation, which implies that a vertex affects the derivative of its successor vertex across the temporal edge, and $D \subseteq V \times L \times V$ is a set of edges. ■

Thus the TCG represents the causality of physical effects in the system, and retains the dynamics of the bond graph model. The TCG in effect specifies the *signal flow graph*, albeit in a form where each edge relation contains at most one component parameter value. Fig. 2 shows the TCG for the six-tank system.

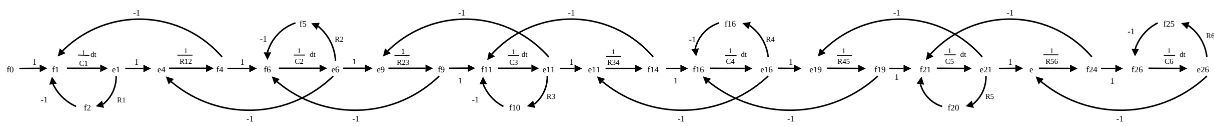

Figure 4: Temporal causal graph for the six-tank system.

TRANSCEND exploits the fact that system dynamics can be expressed in a qualitative form using the TCG, and follows a hypothesize-and-test approach to diagnosis. The key aspect of the approach is the notion of the *fault signature*, which captures the predicted transient behavior around the point of fault occurrence.

Definition 2.3 A fault signature of order N is an N-tuple consisting of the predicted magnitude and 1^{st} through N^{th} order time-derivative effects of a residual signal in response to a fault, expressed as qualitative values: below normal ($-$), normal (0), and above normal (+). Typically N is chosen to be the order of the system. ■

During hypothesis generation the algorithm identifies the set of component parameters that may explain the observed deviation, together with a hypothesized direction of change. During hypothesis refinement, new symbolic measurement variables are matched against the fault signatures for each possible candidate. When the signature is no longer consistent with the observed behavior in the system, the candidate is dropped. A systematic analysis of the qualitative diagnosis approach employed in TRANSCEND , established the discriminatory power of qualitative fault signatures:

Lemma 2.1 In a purely qualitative framework, only the following characteristics of a signal can be used to discriminate among faults:

1. if there is an abrupt change, the direction of abrupt change plus the direction of change immediately following the abrupt change. This implies that there are four distinct fault signatures. i.e. $(+,+)$, $(+,-)$, $(-,+)$ and $(-,-)$.

2. If there is no abrupt change, then only the first direction of change in the measured signal provides discriminatory information. Therefore, this case has two distinct fault signatures. i.e., $(0..+)$ and $(0..-)$. ■

The above lemma is employed to determine when the qualitative fault isolation scheme would no longer be able to discriminate among fault candidates [10]. The results were also used in measurement selection algorithms [12] to find the minimum number of measurements that establish complete diagnosability in the faults of interest. These ideas are exploited to establish independent fault sets among a set of fault hypotheses.

3 Designing a Distributed Diagnosis System

It is clear when one deals with continuous dynamic systems, the system topology and the relations between variables are more complex. The state space representation for continuous systems clearly indicates that system variable dependencies are expressed as continuous functions of time. However, decomposition of a system based on the physical connections between components and the temporal properties of event propagation cannot easily be achieved in the state space representation. In this work, we exploit properties of the bond graph model of continuous systems, and the TCG models derived from these bond graphs to establish independence among sets of faults given a set of measurements, and then develop non-interacting diagnosers for the independent subsets. This extends our previous work in the area of measurement selection algorithms [12]. The rest of this section establishes the framework for partitioning the fault set into independent subsets, and establishes an effective method for finding good set partitions.

3.1 Complete Diagnosability and Measurement Selection

A good fault isolation system must have the capability to uniquely isolate all single faults of interest in a physical system. Given that we have an adequate system model, the ability to diagnose faults depends on the measurements used for diagnostic analysis.

Definition 3.1 Given a set of faults $F = \{f_1, f_2, \ldots, f_l\}$ and a set of measurements $M = \{m_1, m_2, \ldots, m_n\}$, a fault isolation scheme achieves *complete diagnosability* if it can uniquely isolate all possible single faults f_i, for $i \in [1 \ldots l]$ given M. ■

In more detail, a fault, f_i is diagnosable if there is at least one distinguishing fault signature between f_i and any other fault in the system. A system is diagnosable if all its faults are diagnosable. In the TRANSCEND structure this translates to

$$(\forall i, j \in [1, l], i \neq j)(\exists m_k \in M) \;\; \mathsf{FS}(f_i, m_k) \neq \mathsf{FS}(f_j, m_k)$$

where $\mathsf{FS}(f_i, m_j)$ is the observed fault signature for measurement m_k given that fault f_i occurs, and the inequality of fault signatures is defined in terms of the discriminatory power of measurements given by Lemma 2.1.

Complete diagnosability is based on the number of measurements available to the fault isolation algorithm. For example, consider the six-tank system in Fig. 2. Let us assume that the fault parameters of interest the tank capacities, C1 and C2, and the outlet pipe resistances, R1 and R2. If the pressures at the bottom of the two tanks, i.e., the effort variables e1 and e6 in the bond graph model, are the measured variables, then by inspecting the fault signatures in Table 1, we see that we can uniquely isolate all possible faults. For the sake of discussion, if e6 is the only measured variable, then the faults C1− (decrease in tank 1 capacitance) and R2+ (partial block in outlet pipe 2 resistance) can no longer be differentiated from each because their signatures for e6 are the same i.e. $\{0 + - \ldots\}$, and, therefore, the system is not completely diagnosable. On the other hand, fault C2− is still uniquely diagnosable from the rest of the faults, because its fault signature for measurement e6 is $\{+ - \ldots\}$ which is different from signatures of all other faults for e6.

Fault	e1	e6
C_1^-	$\{+ - + - + - +\}$	$\{0 + - + - + -\}$
R_2^+	$\{0\ \ 0 + - + - +\}$	$\{0 + - + - + -\}$
C_2^-	$\{0 + - + - + -\}$	$\{+ - + - + - +\}$
R_1^+	$\{0 + - + - + -\}$	$\{0\ \ 0 + - + - +\}$

Table 1: Fault signatures for the six-tank system example.

Definition 3.2 Given a set of faults $F = \{f_1, f_2, f_3...f_l\}$ and a set of measurements, $M = \{m_1, m_2, m_3...m_n\}$, *measurement selection* identifies the smallest set of measurements $MS \subseteq M$, that can uniquely isolate every single fault in the system, i.e., the subset of measurements makes the system completely diagnosable. ■

Depending on the initial choice of M, it may not be possible to find MS such that the F is completely diagnosable. Also, it is possible to find more than one subset MS that guarantees complete diagnosability.

The problem of measurement selection is an instance of the set cover problem [1, 13], which is known to be NP-complete. Consequently, the optimal solution for this problem (i.e., one that guarantees that MS has the least number of measurements) cannot be time efficient. However, for large systems, this algorithm can be run at design time to address the sensor placement problem and provide certain performance guarantees for the implemented diagnoser. We have developed and analyzed measurement selection algorithms as a partitioning problem [12]. Given an initial partition $P_o = F$, find a set of measurements, $MS = \{m_p | m_p \in M\}$ that refines P_o towards the set $\{\{f_1\}, \{f_2\}, \cdots, \{f_l\}\}$, i.e., each element of the final partition contains only one fault. In this work, we extend the measurement selection algorithm for deriving sets of independent faults.

3.2 Independence of faults for diagnosis

As discussed earlier, our goal is to implement the overall diagnosis system as a set of distributed noninteracting local diagnosers. To this end, we develop the notion of independence of faults within a system, and then develop an algorithm for partitioning the fault set F into subsets that are independent of each other. An independent diagnoser can then be employed for each fault subset. Two fault sets are independent for diagnosis purposes if they can be completely diagnosed with measurement sets that do not overlap.

Definition 3.3 Two sets of faults $F_1, F_2 \in F, F_1 \cap F_2 = \varnothing$ are said to be *independent* for diagnosis given measurement set M if there exists two sets $M_1, M_2 \subseteq M$ such that,

- F_1 is diagnosable for the measurement set M_1
- F_2 is diagnosable for the measurement set M_2
- $M_1 \cap M_2 = \varnothing$ ■

For example, consider the two sets of faults $F_1 = \{C1-, R1+\}$ and $F_2 = \{C2-, R2+\}$ in the six-tank system (Fig. 2). Assuming pressures e1 and e6 are the two measured variables, Table 3.1 indicates that the measurement e1 is sufficient to isolate the faults in F_1, whereas measurement e6 is sufficient to isolate the faults in F_2. Therefore, F_1 and F_2 are independent of each other for the purpose of diagnosis.

3.3 Algorithm for Partitioning the Fault Set

Since independent fault sets can be diagnosed with non-interacting diagnosers, and the complexity of designing and running a diagnosis algorithm increases as the size of the model and the number of faults to be diagnosed increase, our goal for the partitioning algorithm is to create a maximal number of independent fault subsets as the set of given measurements will allow.

To develop a systematic formulation for this problem we first define a fault signature matrix given F and M. The rows in this matrix correspond to the faults of interest in the system $\{f_i | 1 \leq i \leq l\}$, and the columns correspond to the available measurements $\{m_j | 1 \leq j \leq n\}$. This matrix is given as

$$FSM = \left[\mathsf{FS}(f_i, m_j)\right]_{l \times n},$$

where $\mathsf{FS}(f_i, m_j)$ is the fault signature for measurement m_j given fault f_i. Next, we define the notion of a distinguishing set for fault, $f_i \in F$.

Definition 3.4 The *distinguishing measurement set* for fault $f_i \in F$ is defined by the map $\mathsf{Dis} : F \rightarrow \mathcal{P}^2(M)$ where

$$\mathsf{Dis}(f_i) = \{M' \subseteq M \mid f_i \text{ is diagnosable given } M'\}$$

■

In general, $\mathsf{Dis}(f_i)$ will contain multiple measurement subsets. (We assume $\mathsf{Dis}(f_i) \neq \varnothing$). The partitioning problem is then to find a maximal size partition P of F that satisfies

$$(\forall p_i, p_j \in P) \quad \left[\bigcup_{f_i \in p_i} \mathsf{Dis}(f_i)\right] \bigcap \left[\bigcup_{f_j \in p_j} \mathsf{Dis}(f_j)\right] = \varnothing$$

It is clear that the solution to the partitioning problem is at least as complex as the measurement selection problem described in the last section. Therefore, this problem too is NP-complete, and a time constrained practical solution to this problem will require the use of heuristics. We apply knowledge of the diagnosis task in TRANSCEND to solve the partitioning problem in multiple steps, applying reasonable heuristics to reduce the complexity of the search for a good partition.

As a first step to the solution, we construct the distinguishing measurement sets, $\mathsf{Dis}(f_i)$ for all the faults in set F. The second step identifies the sets in the $\mathsf{Dis}(f_i)$'s that include *measurements with discontinuities (mwd)*. Measurement m_j for fault f_i is *mwd*, if the corresponding fault signature has a non zero magnitude value, i.e., the first element of the fault signature list, $FS(f_i, m_j) \neq 0$. The non zero magnitude value in a fault signature implies that the measurement m_j shows a discontinuous change at the time point of fault occurrence. In the tank system, for example, if fault C1− occurs, measurement e1 is *mwd* because its residual value will show an immediate positive jump. In previous work, we have established that discontinuities in measurements, if reliably detected provide quick discriminatory evidence for fault isolation [11]. Therefore, discontinuity detection typically improves the overall time for fault isolation.

We define a distinguishing measurement set with a discontinuity as an element of $\mathsf{Dis}(f_i)$ that includes an *mwd* measurement, which is more likely to distinguish the current fault. This set will be denoted `MWDs`.

Definition 3.5 $M' \in \mathsf{Dis}(f_i)$ is defined to be a measurement set with discontinuity, if $M' \cap \texttt{MWDS} \neq \varnothing$. The set of measurement sets with discontinuity in $\mathsf{Dis}(f_i)$ is denoted $\mathsf{MWD}(f_i)$. ■

To complete this step, we identify all the $M' \in \mathsf{Dis}(f_i) \mid 1 \leq i \leq l$ such that $M' \in \mathsf{MWD}(f_i)$.

Extrapolating the intuition that all measurement in `MWD` contain a lot of discriminatory information, the corresponding $\mathsf{MWD}(f_i)$ will typically include a small number of measurements. This observation leads to the next step, which is to form the *seed elements* of partition P. The seed elements are selected as follows. First, pick F_{seed}, where $F_{seed} = \{f_i \in F \mid \mathsf{MWD}(f_i) \neq \varnothing\}$. Here $f_i \in F_{seed}$ becomes a seed in P if

$$\mathsf{MWD}(f_i) \bigcap \left[\bigcup_{f_j \in F \setminus f_i} \mathsf{MWD}(f_j) \right] = \varnothing$$

Each of the chosen f_i's then becomes the seed about which the elements $p_1,\ p_2, \ldots$ of partition P are formed. The partition completion step takes the remaining faults in F that have not been assigned to an element in P one by one, and sequentially searches for the best p_i to add the fault to. There are two considerations in making the addition.

1. Avoid adding f_k to p_i's which then cause this p_i to merge with some other p_j. Two elements of the partition will merge if $\mathsf{Dis}(p_i) \bigcap \mathsf{Dis}(p_j) \neq \varnothing$.
2. Add f_k to that p_i, such that the distinguishing measurements for p_i increases by the smallest amount. The intuition here is that keeping $\mathsf{Dis}(p_i)$ small is more likely to keep elements of the partition being merged as more faults are added to it.

Note that the above described algorithm does a sequential (as opposed to exhaustive) search for forming partition seeds and adding faults to partition elements, therefore, it is generally suboptimal.

4 Example: A Six-tank System

We demonstrate the algorithm on a hypothetical connected six-tank system. The six tank examples are run for different fault and measurement sets to demonstrate the effect of different measurement sets on the partitions formed.

We present two examples for the six-tank system, varying both the fault and measurement sets. In the first experiment, we measure tank pressures, i.e. $M = \{e1, e6, e11, e16, e21, e26\}$, and start with a fault set F that includes tank capacities (C_i's and outlet pipe resistances (R_j's). The algorithm picks $\{C_1\}, \{C_2\}, \{C_3\}, \{C_4\}, \{C_5\}$, and $\{C_6\}$ as the seeds for the partition P. This is because each tank pressure is a mwd for the corresponding tank capacity fault. Note that the R_j's show no discontinuity in the measurements. The algorithm then adds the R_j's sequentially. Each R_j matches up with its corresponding C_j to generate the final partition: $\{\{C_1, R_1\}, \{C_2, R_2\}, \{C_3, R_3\}, \{C_4, R_4\}, \{C_5, R_5\}, \{C_6, R_6\}\}$. The implication is that each pressure measurement can uniquely identify the corresponding tank capacity and outlet pipe resistance fault.

For the second experiment, we choose the same fault set as above, but change the measurement set to $M = \{f4, f7, f14, f17, f24, f27\}$, i.e., we assume that the outflows from each tank are the only measured values. In this situation, the fault set breaks down into three independent fault sets. The initial partition formed is $\{\{C_1\}, \{C_2\}, \{C_3\}, \{C_4\}, \{C_5\}, \{C_6\}\}$. After adding R_1 the set becomes $\{\{C_1, R_1, C_2\}, \{C_3\}, \{C_5\}\}$. When R_1 is added the distinguishing set for $\{C_1, R_1\}$ requires two measurements, $f4$ and $f7$, and this causes the two elements $\{C_1\}$ and $\{C_2\}$ to merge forming the partition set $P = \{\{C_1, R_1, C_2\}, \{C_3\}, \{C_5\}\}$. R_2 requires no additional measurements for discrimination and merges into the first element of P. This pattern continues, and one obtains the final partition with independent fault sets: $\{\{C_1, R_1, C_2, R_2\}, \{C_3, R_3, C_4, R_4\}, \{C_5, C_6, R_5, R_6\}\}$.

5 Conclusions

This paper has developed a methodology for distributed diagnosis of complex continuous systems by partitioning the given fault set into sets of independent faults for which non-interacting diagnosers can be constructed. By simplifying the models required for each diagnoser, and the number of faults and measurements that it has to deal with, we significantly reduce the computational complexity of the overall diagnosis task. In this work, we exploit the fault signatures derived from the TCG model of the physical process that capture the transient dynamics in qualitative form, to derive independence among fault sets. Like the measurement selection problem, the algorithm for generation independent partitions of faults is NP-complete, so we derive a suboptimal algorithm that uses domain-specific heuristics (measurements with discontinuities) and a sequential search process to develop a time-constrained algorithm. The algorithm produces reasonable results when applied to a cascaded six-tank example.

In future work, we will extend this algorithm to situations where the partition sets are not completely independent, i.e., elements of the partition have overlapping measurement sets. The challenge is to extend our algorithm to find minimal overlapping sets. Alternatively, if prior fault probabilities are available the set covering that minimizes the communication requirements between diagnosers must be found. Then it is still possible to design interacting diagnosers that are computationally efficient for online applications.

Acknowledgements

This work was supported in part through grants from the NASA-IS program (Contract number: NAS2-37143), DARPA SEC program (Contract number: F30602-96-2-0227), and the NASA-ALS program (Contract number: NCC 9-159). We acknowledge the help provided by Gabor Karsai and Gyula Simon and Vanderbilt in the development of the FACT architecture.

References

[1] A.V. Aho, J.E. Hopcroft, and J.D. Ulman. *Data structures and algorithms*. Addisson-Wesley, Reading, MA, 1983.

[2] Pietro Baroni, Gianfranco Lamperti, Paolo Pogliano, and Marina Zanella. Diagnosis of large active systems. *Artificial Intelligence*, 110(1):135–183, 1999.

[3] Ronald P. Bianchini and Richard W. Buskens. Implementation of on-line distributed system-level diagnosis theory. *IEEE Trans. on Computers*, 41(5):616–626, 1992.

[4] R. Debouk, S. Lafortune, and D. Teneketzis. Coordinated decentralized protocols for failure diagnosis of discrete event systems. *Discrete Event Dynamic System: Theory and Applications*, 10(1/2):33–86, January 2000.

[5] J. J. Gertler. *Fault Detection and Diagnosis in Engineering Systems*. Marcel Dekker, Inc., New York, NY, 1998.

[6] D. C. Karnopp, D. L. Margolis, and R. C. Rosenberg. *Systems Dynamics: Modeling and Simulation of Mechatronic Systems*. John Wiley & Sons, Inc., New York, third edition, 2000.

[7] J.G. Kuhl and S.M. Reddy. Fault diagnosis in fully distributed systems. In *Proceedings of the 11th IEEE International Conference on fault-Tolerant Computing)*, pages 100–105, June 1981.

[8] J. Kurien, X. Koutsoukos, and F. Zhao. Distributed diagnosis of networked embedded systems. In *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX-2002)*, pages 179–188, Semmering, Austria, May 2002.

[9] E.-J. Manders and G. Biswas. FDI of abrupt faults with combined statistical detection and estimation and qualitative fault isolation. Washington, DC, June 2003.

[10] E.-J. Manders, S. Narasimhan, G. Biswas, and P. J. Mosterman. A combined qualitative/quantitative approach for fault isolation in continuous dynamic systems. pages 1074–1079, Budapest, Hungary, June 2000.

[11] P. J. Mosterman and G. Biswas. Diagnosis of continuous valued systems in transient operating regions. 29(6):554–565, 1999.

[12] S. Narasimhan, P. J. Mosterman, and G. Biswas. A systematic analysis of measurement selection algorithms for fault isolation in dynamic systems. pages 94–101, Cape Cod, MA USA, May 1998.

[13] V. T. Paschos. A survey of approximately optimal solutions to some covering and packing problems. *ACM Computing Surveys*, 29(2):171–209, 1997.

[14] R. Su, W.M. Wohnam, J. Kurien, and X. Koutsoukos. Distributed diagnosis of qualitative systems. In *6th International Workshop on Discrete Event Systems, Zaragoza (WODES-2002)*, pages 169–174, Zaragosa, Spain, October 2002.