

Designing Distributed Diagnosers for Complex Physical Systems

Indranil Roychoudhury, Gautam Biswas, Xenofon Koutsoukos and Sherif Abdelwahed

Department of EECS & ISIS

Vanderbilt University, Nashville, TN 37235. USA.

Email: gautam.biswas@vanderbilt.edu

Abstract

Online diagnosis methods require large computationally expensive diagnosis tasks to be decomposed into sets of smaller tasks so that time and space complexity constraints are not violated. This paper defines the distributed diagnosis problem in the *Transcend* qualitative diagnosis framework, and then develops heuristic algorithms for generating a set of local diagnosers that solve the global diagnosis problem without a coordinator. Two versions of the algorithm are discussed. The time complexity and optimality of these algorithms are compared and validated through experimental results.

1 Introduction

Online fault diagnosis is a key component for developing health monitoring and adaptive control techniques to maintain the safety and operational performance of mission-critical systems under adverse conditions. In model based diagnosis, the system model provides the basis for generating and analyzing fault hypotheses that explain the observed discrepancies in system behavior. For large systems analyzing faulty behaviors is computationally expensive, especially if the algorithm requires a single, composite model of the entire system. This motivates the need to develop methodologies for decomposing the diagnosis task into sets of smaller subtasks that reduce the overall computational and space complexity of online fault isolation. Each subtask can be built into a local diagnosis module, and results produced by the individual modules can be composed to derive a correct global diagnosis.

Previous work on distributed diagnosis has focused on systems with discrete behaviors, such as a set of interconnected processors [1], [2]. An approach for decentralized diagnosis has been presented in [3] where the local diagnosers communicate with a coordination process that assembles a global diagnosis result. The coordination process often require significant communication with the local diagnosers, which brings into question the reliability and scalability of these approaches. In one of the algorithms mentioned in [3] (protocol 3), the coordinator is redundant and local diagnosers generate globally valid diagnoses. However, the problem of *how* to construct such diagnosers is not addressed. A distributed diagnosis approach, where each local diagnoser communicates directly with other diagnosers has been presented in [4]. Initially, each diagnoser finds a set of local di-

agnoses and then communicates with its neighbors to further reduce the number of consistent diagnoses. The design problem is formulated as a distributed constraint satisfaction problem. The graph that represents the constraints between the fault hypotheses and the observations is partitioned to minimize the communication between local diagnosers. A similar approach that supports reconfigurable systems has been presented in [5]. System components are represented by local input/output automata and the partitioning is based on the physical connections of the system and not the communication requirements. A distributed diagnosis method that does not require coordination between local diagnosers has been proposed in [6], but the structure of the local models makes the local diagnosis and communication extremely complicated.

Our approach extends the *Transcend* framework [7], and develops a method for designing distributed diagnosers for continuous systems. Given a set of faults and measurements and a global system model, we construct independent diagnosers that together make the system completely diagnosable. Two diagnosers are independent if they do not have to share information in establishing unique global diagnosis results. Complete diagnosability is the ability to uniquely isolate every fault candidate in the system given a set of measurements. We propose two algorithms. The first ensures that the local diagnosers do not share measurements, i.e., there is no communication between the diagnosers. This decomposition may not always be possible, and the second algorithm relaxes the assumption to allow overlapping measurements, but still ensures that the local diagnosis results are globally valid.

The general problem of finding independent fault sets subsumes the set covering problem [8], which is known to be NP-Complete. To overcome this complexity, our algorithm starts with each fault as an independent set, and systematically merges faults in a way that each fault within a set is uniquely globally distinguishable for a set of non-overlapping measurements. Heuristics that favor balanced (i.e., equal-sized) sets help cut down on the exponential merge process. Details of the algorithms, their computational complexity, and experiments conducted on multi-tank systems and a reverse-osmosis system are presented in this paper.

2 Background

The *Transcend* architecture employs a model-based approach based on analysis of transients to isolate abrupt faults in process components. It combines a novel qualitative scheme with a quantitative parameter estimation scheme for fault isolation and identification in continuous systems. Dy-

dynamic system models are constructed using bond graphs [9]. Faults map to component parameters in the bond graph. *Transcend* focuses on abrupt faults that are modeled as discrete and persistent changes in component parameters. 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 a qualitative analysis of the fault transient dynamics. Specifically, the magnitude and slope of the transient residual, derived from measurements, are mapped onto $\{+, 0, -\}$ symbols (after energy-based filtering [10]) for qualitative matching against fault signatures.

2.1 Qualitative fault isolation

Fault isolation is developed on a graphical model representation, the Temporal Causal Graph (TCG), that is derived automatically from a bond graph model of the system [7].

Definition 1 A Temporal Causal Graph (TCG) is a directed graph $\langle V, L, D \rangle$. $V = E \cup F$, where V is a set of vertices, E is a set of effort variables and F is a set of flow variables in the 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. ■

The TCG captures the causality of physical effects in the system, and retains the dynamics expressed in 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.

We illustrate fault isolation in *Transcend*, as well as the new algorithm developed in this paper, using a hypothetical physical system that consists of six fluid tanks connected to each other with pipes, with a source of flow into the first tank and a pipe for draining from each of the tanks. In the bond graph model, all the pipes are modeled by resistances and tanks are modeled as capacities, making this a sixth order system. Fig. 1 shows the 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 of possible faults i.e., $F = \{C_1, \dots, C_6, R_1, \dots, R_6, R_{12}, \dots, R_{56}\}$, includes changes in all tank capacities, drain pipe resistances and connecting pipe resistances. A $+$ ($-$) superscript implies a fault that results in an increase (decrease) in parameter value. For example, C_i^- indicates a decrease in tank i capacitance. The set of measurements $M = \{e1, f2, e6, f7, e11, f12, e16, f17, e21, f22, e26, f27\}$, includes all tank pressures and flow rates through the drain pipes. Fig. 2 shows the TCG for the tank system.

Transcend 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 at and after the point of fault occurrence.

Definition 2 A fault signature of order N is an ordered 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. ■




Figure 1: Six-tank system and its bond graph model.

During hypothesis generation, the algorithm identifies the set of component parameters with a hypothesized direction of change in the parameter that explains the observed deviation. The TCG is then used to generate the fault signatures for all measurements for all of the hypothesized faults. During hypothesis refinement, as time progresses, new symbolic measurement variables that become available are matched against the fault signatures for each hypothesized candidate. If a signature becomes inconsistent with the observed behavior, the candidate is dropped. A systematic analysis of the qualitative diagnosis establishes the discriminatory power of qualitative fault signatures.

Lemma 1 In a purely qualitative framework, faults can be discriminated if the signature shows the direction of abrupt change and the direction of change immediately following the abrupt change, i.e., $(+, +)$, $(+, -)$, $(-, +)$ and $(-, -)$. However, $(+, +)$ and $(-, -)$ signatures imply an unstable system, therefore, they are unlikely.

If there is no abrupt change, then the first direction of change in the measured signal provides discriminatory information i.e., $(0..+)$ and $(0..-)$. ■

This lemma implies that a single measurement can distinguish at most four faults. This gives a measure of the maximum number of faults that can be isolated for a set of measurements. The above lemma also informs us when the qualitative fault isolation scheme can discriminate no further among fault candidates [11]. The results have been used in measurement selection algorithms [12] to find the minimum number of measurements that establish complete diagnosability given a set of faults. These ideas are exploited in a different way to establish independent fault sets in the next section.

3 Designing Distributed Diagnosis Systems

System variable dependencies in a mathematical model for continuous system behavior are expressed as continuous functions of time. Changes in any part of a continuous system propagate to all other parts, therefore, decomposition is not easily achieved by exploiting the temporal properties of event propagation. In this work, we make use of the topological properties of the bond graph model, and the TCG models derived from these bond graphs to design non-interacting diagnosers. The diagnosers satisfy the strong constraint that a local diagnosis result is globally valid. Our algorithms for partitioning fault sets are based on this property.




Figure 2: Temporal causal graph for the six-tank system.

3.1 Complete Diagnosability and Measurements

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

Definition 3 Given a set of faults $F = \{f_1, f_2, \dots, f_l\}$ and a set of measurements $M = \{m_1, m_2, \dots, m_n\}$, a fault isolation scheme achieves complete diagnosability if it can uniquely isolate all possible single faults $f_i, i \in [1 \dots l]$ given the measurements in M . ■

In the Transcend structure this translates to $(\forall i, j \in [1, l], i \neq j)(\exists m_k \in M) \text{FS}(f_i, m_k) \neq \text{FS}(f_j, m_k)$ where $\text{FS}(f_i, m_j)$ is the fault signature for measurement m_j 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 1.

For example, consider the six-tank system in Fig. 1. Let us assume that the faults of interest are tank capacity decreases, i.e., C_1^- and C_2^- , and the outlet pipe resistance blocks, i.e., R_2^+ . 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 the three faults. This is because each fault has at least one distinguishing fault signature from the rest of the fault set. If $e6$ were the only measured variable, then the set of faults are no longer diagnosable because C_1^- and R_2^+ have the same signatures for $e6$.

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

Table 1: Fault signatures from tanks 1 and 2 for the six-tank system.

The problem of measurement selection, i.e. finding a minimal set of measurements to isolate all single faults in the system, is an instance of the set covering problem [8; 13], which is known to be NP-complete. Consequently, there is no efficient algorithm that can guarantee an optimal solution for this problem. The exponential complexity may not be an issue because measurement selection is typically executed at design time to solve the sensor placement problem.

3.2 Independence of faults for diagnosis

To derive a set of distributed noninteracting local diagnosers, we define the notion of independence of faults in a system. We start with a strong definition of independence, and develop an algorithm for partitioning the fault set F into subsets in a way that the diagnosers use no overlapping measurements. We then relax the definition of independence, and derive a new algorithm that generates local diagnosers

that provide globally correct diagnosis but allow overlapping measurements between diagnosers.

Definition 4 Two fault sets $P_1, P_2 \in F$, are said to be strongly independent for diagnosis given measurement set, M , if there exists two sets $Q_1, Q_2 \subseteq M$, such that (i) P_1 is globally diagnosable for measurement set Q_1 , (ii) P_2 is globally diagnosable for measurement set Q_2 , and (iii) $Q_1 \cap Q_2 = \emptyset$ ■

Assume $P_1 = \{C_1^-\}, P_2 = \{C_2^-\}$ and $M = \{e1, e6\}$. P_1 and P_2 are independent fault sets because $Q_1 = \{e1\}$ can uniquely isolate C_1^- and $Q_2 = \{e6\}$ can diagnose C_2^- . $P_1 = \{C_1^-\}, P_2 = \{C_2^-, R_2^+\}$ are also independently diagnosable for $M = \{e1, e6\}$ because $Q_1 = \{e1\}$ and $Q_2 = \{e6\}$ and $Q_1 \cap Q_2 = \emptyset$.

3.3 Algorithm for Partitioning the Fault Set

To develop a systematic formulation for the problem we first define a fault signature matrix. The rows in this matrix correspond to the faults $\{f_i | 1 \leq i \leq l\}$ and the columns correspond to the available measurements $\{m_j | 1 \leq j \leq n\}$. This matrix is denoted as $FSM = [\text{FS}(f_i, m_j)]_{l \times n}$, where $\text{FS}(f_i, m_j)$ is the fault signature for measurement m_j given fault f_i .

Definition 5 The distinguishing measurement set for fault $f_i \in F$ is defined by the mapping $\text{Dis}: F \rightarrow 2^M$ where $\text{Dis}(f_i) = \{M' \subseteq M | f_i \text{ is diagnosable given } M'\}$ ■

We assume $\text{Dis}(f_i) \neq \emptyset$ as all faults should be diagnosable by a global diagnoser. Our task is to find a covering of fault sets, $\{P_i | \bigcup_{1 \leq i \leq N} P_i = F\}$, such that the number of sets, N is a maximum ($N \leq l$), that satisfies

$$(\forall P_i, P_j \in F) \left[\bigcup_{f_i \in P_i} \text{Dis}(f_i) \right] \cap \left[\bigcup_{f_j \in P_j} \text{Dis}(f_j) \right] = \emptyset$$

In other words, the optimal solution to the above problem, is to find the partition of independent fault sets that covers all faults and is of maximum size. Each fault set P_i , has an associated measurement set $Q_i \subseteq M$ such that $Q_i = \bigcup_{f_j \in P_i} \text{Dis}(f_j)$, and we denote the solution to the problem as $(P_i, Q_i)_{i=1, \dots, N}$.

However, the problem of finding these different (P_i, Q_i) pairs itself is exponential as all combinations of measurements need to be considered for each fault. This makes the covering problem doubly exponential in computational complexity. To derive practical solutions to this problem, we introduce heuristic search methods.

Algorithm Description

We explain the search space for the covering problem using a search tree. The root node of this tree contains n sets (total number of measurements), where each set is a two-tuple, (P_i, Q_i) . $Q_i = m_i, 1 \leq i \leq n$, and P_i is the corresponding distinguished fault-set for m_i .¹ If a measurement cannot, by itself,

¹ A distinguished fault set is a set of faults that are diagnosable given Q_i , i.e., $Q_i = \bigcup_{f_j \in P_i} \text{Dis}(f_j)$

Fault	e1	f2	e6	f7	e11	f12	e16	f17	e21	f22	e26	f27
ζ_1	+--+--	+--+--	0+--+	0+--+	00+--	00+--	000+--	000+--	0000+--	0000+--	00000+	00000+
ζ_2	0+--+	0+--+	+--++	+--++	0+--+	0+--+	00+--	00+--	000+--	000+--	0000+--	0000+--
ζ_3	00+--	00+--	0+--+	0+--+	+--++	+--++	0+--+	0+--+	00+--	00+--	000+--	000+--
ζ_4	000+--	000+--	00+--	00+--	0+--+	0+--+	+--++	+--++	0+--+	0+--+	00+--	00+--
ζ_5	0000+--	0000+--	000+--	000+--	00+--	00+--	0+--+	0+--+	+--++	+--++	0+--+	0+--+
ζ_6	00000+	00000+	0000+--	0000+--	000+--	000+--	00+--	00+--	0+--+	0+--+	+--++	+--++
R_{12}^+	0+--+	0+--+	0+--+	0+--+	00+--	00+--	000+--	000+--	0000+--	0000+--	00000-	00000-
R_{23}^+	00+--	00+--	0+--+	0+--+	0+--+	0+--+	00+--	00+--	000+--	000+--	0000+--	0000+--
R_{34}^+	000+--	000+--	00+--	00+--	0+--+	0+--+	0+--+	0+--+	00+--	00+--	000+--	000+--
R_{45}^+	0000+--	0000+--	000+--	000+--	00+--	00+--	0+--+	0+--+	0+--+	0+--+	00+--	00+--
R_{56}^+	00000+	00000+	0000+--	0000+--	000+--	000+--	00+--	00+--	0+--+	0+--+	0+--+	0+--+

Table 2: Fault signatures for the six-tank system example.

completely diagnose any fault, its fault set is empty. A recursive procedure is then employed to generate the subsequent levels of the tree, starting with level 1, the children of the root node. Nodes at level $i+1$ in the tree are generated by creating all possible merged pairs of measurement sets from the parent nodes at level i , and computing the corresponding distinguishing fault-sets for the new measurement sets. For example, if the partition for a node is $\{(P_1, Q_1), (P_2, Q_2), (P_3, Q_3)\}$, the following nodes will be formed as children of this node: $\{(P_{12}, Q_1 \cup Q_2), (P_3, Q_3)\}$, $\{(P_{13}, Q_1 \cup Q_3), (P_2, Q_2)\}$, and $\{(P_{23}, Q_2 \cup Q_3), (P_1, Q_1)\}$. Note that P_{ij} , the distinguished fault set formed by the merger of Q_i and Q_j can include more faults than $P_i \cup P_j$ because the two sets of measurements can, in theory, uniquely diagnose more faults than the sum of the faults that each can diagnose. For example, a single measurement can at most uniquely diagnose 4 different faults, but two measurements can diagnose up to $4^2 = 16$ different faults.

The merge process is repeated at the next level of search till a partition $P = P_1, \dots, P_N$ is obtained such that $\bigcup P_i = F$. Note that a partition generated at level i of the tree is preferred to a partition generated at level j , where $i < j$. This is because the partition at level i will have a greater number of sets than the partition at level j . In our work, two partitions generated at the same level are given equal preference. In the future, we will introduce further criteria to differentiate among partitions generated at the same level (e.g., sets balanced by size).

A breadth-first search (BFS) solution to the fault set covering problem is exhaustive, therefore, if no solution is found in level i , the merge operator described above will be applied to all nodes at this level to generate the fault partitions at level $i+1$. If one or more goal nodes are found at a level the BFS algorithm terminates, and all solutions at this level are defined as “optimal.” Since all nodes at each level are expanded, the search space, and, the search time are doubly exponential. It may happen that the only solution found is (F, M) , and the search tree is expanded to k levels to find this solution, where $|M| = k$, the minimum number of measurements required to completely diagnose the fault set.

Our algorithm uses the BFS control structure, but applies heuristics to reduce the space time complexity of the search. If a goal node is found at a particular level, the algorithm terminates. If a goal node is not found in the current level, our algorithm chooses the best node for expansion, using a heuristic function $h = |\bigcup_{f_i \in P_i} f_i|$, the number of faults that are diagnosable in node P_i . The node with the largest h value

Algorithm 1

-
- Input:** Set of l faults $F = \{f_i | i = 1, \dots, l\}$,
Set of n measurement $M = \{M_j | j = 1, \dots, n\}$,
The Temporal Causal Graph
- Compute $\text{FSM} = [\text{FS}(f_i, m_j)]_{l \times n}$
 - Generate root
 - REPEAT:
 - Check for goal node
 - Calculate h values for each node
 - IF: Goal node is found
 - Output this node
 - Break
 - ELSE:
 - Expand node with highest h value
-

is chosen for expansion. In general, a fault may be globally diagnosable in more than one set of a partition. In this case, we assign the fault to the smaller set. This results in balanced sets within a node.

Relaxing the Definition of Independent Faults

A primary reason for not allowing overlapping measurements is to keep the communication overhead between the diagnosers to a minimum. However, if the only solution Algorithm 1 produces is (F, M) , we can relax the condition for non overlapping measurements and generate local diagnosers that generate global diagnoses.

Definition 6 *Two fault sets $P_1, P_2 \in F$, are said to be weakly independent for diagnosis given measurement set, M , if there exists two sets $Q_1, Q_2 \subseteq M$, such that (i) P_1 is globally diagnosable for measurement set Q_1 , and (ii) P_2 is globally diagnosable for measurement set Q_2 .* ■

The algorithm to derive the set of diagnosers under the weak independence condition takes on a different structure. Our goal is to allow overlapping measurements but minimize the overlap to keep the communication overhead low. We assume additional information about the system structure, i.e., the k interacting subsystems are known. Associated with each subsystem are the set of measurements M_i , and the set of faults F_i , such that $\bigcup_{1 \leq i \leq k} M_i = M$ and $\bigcup_{1 \leq i \leq k} F_i = F$. Further, Also $\forall i, j$ such that $i \neq j$, $M_i \cap M_j = \emptyset$ and $F_i \cap F_j = \emptyset$. We use this to derive k local diagnosers, one for each subsystem.

Each subsystem may have faults that cannot be diagnosed

uniquely by the measurements in its subsystem. For each such fault, the distinguishing measurement set is calculated and then the least number of additional measurements required to make this fault uniquely diagnosable is added to the measurements of the subsystem the fault belongs to. It is clear that allowing overlap of measurements results in an exponential algorithm as all possible ways in which the measurements can be grouped is by itself exponential. We can further restrict the search space if we give preference to selecting measurements from neighboring subsystems, if this information is available. In this case, the minimal overlap problem can be solved in polynomial time on the average. However, in the worst case, this algorithm will still be exponential.

4 Performance Evaluation

We first analyze the heuristic algorithm under strong independence and then present experimental results of this algorithm on two multi-tank systems and Reverse Osmosis System. This is followed by the result obtained from the six-tank system under the weak independence condition.

To analyze the time complexity of the heuristic algorithm under strong independence, assume $|F| = l$ and $|M| = n$. The root node starts with n sets. For each measurement set Q_i , we identify the set of faults P_i diagnosable by the measurements in Q_i . The faults in P_i have unique fault signatures for the measurements in Q_i and they are computed by traversing the columns of the fault signature matrix, FS , that correspond to the measurements in Q_i . This operation can be computed in $O(l^2n)$ time. After combining the sets P_i , we merge all pairs of Q_i 's to obtain the measurement sets of the children nodes. Therefore we have $\binom{n}{2}$ nodes in the next level and each node will have $(n-1)$ measurement sets (Q_i 's). Computing the P_i 's for each node at this level is also $O(l^2n)$. Since we are expanding only one node, we will have only $\binom{n-1}{2}$ children. The number of nodes generated is $\binom{n}{2} + \binom{n-1}{2} + \binom{n-3}{2} + \dots + \binom{2}{2} = O(n^3)$ as there are at most n levels. Hence the total complexity is $O(l^2n^4)$, which polynomial in the number of faults and the number of measurements.

We apply the heuristic partitioning algorithm under strong independence to the six-tank system example of Fig. 1, and compare the results obtained with the breadth-first search scheme results. For this example, we chose the fault set $F = \{C_1^-, \dots, C_6^-, R_{12}^+, \dots, R_{56}^+\}$, which included all the tank capacitances and the connecting pipe resistances. The set of measurements considered was $M = \{e1, f2, e6, f7, e11, f12, e16, f17, e21, f22, e26, f27\}$, which included all tank pressures and flow rate through the connecting pipes. The fault signature table for the fault and measurement set are shown in Table 2.

The breadth first search came up with several possible "optimal" covers, which included $(\{C_1\}, \{e1\}), (\{R_{23}, R_{34}, R_{45}, R_{56}\}, \{e6, e11, e16, e21, e26\}), (\{C_2, R_{12}\}, \{f7\}), (\{C_3\}, \{f12\}), (\{C_4\}, \{f17\}), (\{C_5\}, \{f22\}), (\{C_6\}, \{f27\}); (\{C_1\}, \{e1\}), (\{C_2, R_{12}\}, \{e6\}), (\{C_3, R_{23}\}, \{e12, f7\}), (\{C_4, R_{34}\}, \{e16, f12\}), (\{C_5, R_{45}\}, \{e21, f16\}), (\{C_6\}, \{f27\}), (\{R_{56}\}, \{e26, f22\})$ and many more which we have not listed here. The superscripts have been dropped for improving readability.

The output of our heuristic algorithm for the same example

was $(\{C_1\}, \{e1\}), (\{C_2, R_{12}\}, \{e6\}), (\{C_3, R_{23}\}, \{e12, f7\}), (\{C_4, R_{34}\}, \{e16, f12\}), (\{C_5, R_{45}\}, \{e21, f16\}), (\{C_6\}, \{f27\}), (\{R_{56}\}, \{e26, f22\})$. The solution generated by the heuristic algorithm has seven sets, therefore, it is one of the optimal solutions generated by the BFS algorithm. When one compares the number of node expansions required to generate the solutions, the BFS search searched involved 183074 node expansions, and our algorithm derived its solution with 203 node expansions. We have run a number of other experiments with the six tank system, and in almost all cases, the heuristic covering algorithm found an optimal solution expanding 1% of the nodes that were generated by the BFS algorithm. This demonstrated that the heuristic algorithm is efficient and generates acceptable solutions. This algorithm is scales up well for large systems such as a ten-tank system as well. We ran additional experiments on 10 and 20 tank systems to demonstrate how well the partitioning algorithm scales up. The results were similar to the six tank system and not reported here.

The algorithm that allows overlapping measurements produced the following partition: $(\{C_1, R_{12}\}, \{e1, f2, e6\}), (\{C_2, R_{23}\}, \{e6, f7, e11\}), (\{C_3, R_{34}\}, \{e11, f12, e16\}), (\{C_4, R_{45}\}, \{e16, f17, e21\}), (\{C_5, R_{56}\}, \{e21, f22, e26\}), (\{C_6\}, \{e26, f27\})$. Here we assumed each tank and the pipe connecting it to the tank to its right to belong to the same subsystem. The pressure measurement in the tank, and the flow out of each tank were the two measured variables for each subsystem. The faults included the decrease in capacitance of each tank and the increase in flow resistance in the pipe connecting each tank to its next tank. As expected, the capacitance C_i in the i^{th} tank is uniquely diagnosable by the effort variable of that tank. However, to uniquely diagnose the faults in the inter-connecting pipes, this algorithm adds the pressure variable of the adjoining tank to the measurements of subsystem i and makes all faults in F_i uniquely diagnosable.

4.1 Reverse Osmosis System

We also applied this algorithm under strong independence to design distributed diagnosers for the Reverse Osmosis (RO) system, which is a component of a Water Recovery system (WRS) constructed at the NASA Johnson Space Center as a test-bed for long duration human life support in space. Fig. 3 shows the schematic of the Reverse Osmosis system. It consists of a feed pump which brings water into the system followed a coil which acts as a reservoir for water before it is pumped in by the recirculation pump into a membrane. The membrane cleans water by a filtering action. Purified water leaves the membrane and goes to the next subsystem. Water that did not pass through the membrane goes into a recycle loop (Primary or Secondary) and is pumped into the membrane again. This system, and the Transcend FDI results are discussed in detail in [14].

In this experiment, we consider only the primary operating mode of this system. Faults of interest are the feed pump, with corresponding component parameters $\{I_{fp}, R_{fp}\}$, faults in the membrane corresponding to $\{C_{memb}, R_{memb}\}$ and faults in the recirculation pump corresponding to $\{I_{ep}, R_{ep}\}$. There are five measurements in the system; (i) the pressure at the output of recirculation pump (e37), (ii) fluid pressure at the membrane (e16) (iii) flow through the membrane (f25), (iv) pressure in the return path of the loop (e23) (v) input flow to

F-M	e37	e16	e23	f30
I_{fp}^-	0 0 - + .	0 0 + . .	0 - + . .	+ - . . .
R_{fp}^+	0 0 0 + -	0 0 0 - .	0 0 + - .	0 - + . .
C_{memb}^+	0 - + . .	- + . . .	- + . . .	0 0 + - .
R_{memb}^-	0 0 - + .	0 - + . .	0 - + . .	0 0 0 + -
I_{ep}	+ - + . .	0 + . . .	0 + . . .	0 0 + - .
R_{ep}^-	0 + - + .	0 0 + . .	0 0 + . .	0 0 0 + -

Table 3: Faults with corresponding measurement signatures for the RO system.

the RO (f30).

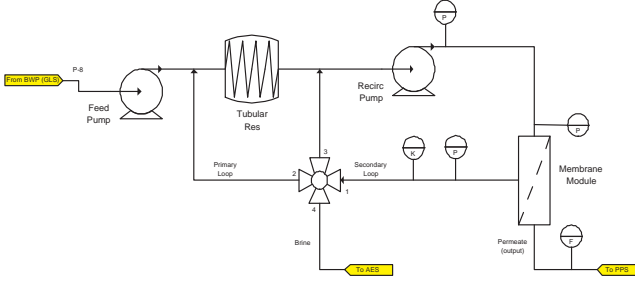


Figure 3: Reverse Osmosis system process engineering diagram.

Table 3 shows the fault signatures for the faults and measurements used in this experiment. The algorithm comes up with a scheme where three separate diagnosers can be implemented. The partition obtained by this algorithm is $(\{I_{fp}^-, R_{fp}^+\}, \{f30\}), (\{C_{memb}^+\}, \{e23\}), (\{R_{memb}^-, I_{ep}, R_{ep}^-\}, \{e37, e16\})$.

5 Conclusions

This paper has developed a methodology for designing distributed diagnosers for complex continuous systems by partitioning the given fault set into sets of independent faults. The problem of finding both strongly as well as weakly independent sets have been addressed. The resulting local diagnosers satisfy the property that a local diagnosis result is also globally correct, therefore our approach does not need a coordinator module. By partitioning the system into strongly independent sets, we significantly reduce the computational complexity of the overall diagnosis task. As discussed earlier, the nature of continuous dynamics and the interactions between components of a physical system makes it difficult to divide the system into independent subsystems. 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.

In future work, we would like to improve the preliminary algorithm that finds sets that have the minimal overlap. This will minimize the communication requirements between diagnosers. 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-ALS program (Contract number: NCC 9-159), and the NSF EHS Program (Contact number: SGER 0208799).

The help provided by Dr. Eric Manders is gratefully acknowledged.

References

- [1] R. P. Bianchini and R. W. Buskens, "Implementation of on-line distributed system-level diagnosis theory," *IEEE Trans. on Computers*, vol. 41(5), pp. 616–626, 1992.
- [2] J. Kuhl and S. Reddy, "Fault diagnosis in fully distributed systems," in *Proceedings of the 11th IEEE International Conference on fault-Tolerant Computing*, pp. 100–105, June 1981.
- [3] R. Debouk, S. Lafortune, and D. Teneketzis, "Coordinated decentralized protocols for failure diagnosis of discrete event systems," *Discrete Event Dynamic System: Theory and Applications*, vol. 10(1/2), pp. 33–86, January 2000.
- [4] 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)*, Semmering, Austria, , pp. 179–188, May 2002.
- [5] R. Su, W. Wohnam, J. Kurien, and X. Koutsoukos, "Distributed diagnosis of qualitative systems," in *6th International Workshop on Discrete Event Systems, Zaragoza (WODES-2002)*, Zaragoza, Spain, pp. 169–174, Oct. 2002.
- [6] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanella, "Diagnosis of large active systems," *Artificial Intelligence*, vol. 110(1), pp. 135–183, 1999.
- [7] P. J. Mosterman and G. Biswas, "Diagnosis of continuous valued systems in transient operating regions," *IEEE-SMCA*, vol. 29(6), pp. 554–565, 1999.
- [8] V. T. Paschos, "A survey of approximately optimal solutions to some covering and packing problems," *ACM Computing Surveys*, vol. 29(2), pp. 171–209, 1997.
- [9] D. C. Karnopp, D. L. Margolis, and R. C. Rosenberg, *Systems Dynamics: Modeling and Simulation of Mechatronic Systems*, 3rd ed. New York: John Wiley & Sons, Inc., 2000.
- [10] E.-J. Manders and G. Biswas, "FDI of abrupt faults with combined statistical detection and estimation and qualitative fault isolation," in *SafeProcess 2003*, Washington, DC, June 2003.
- [11] E.-J. Manders, S. Narasimhan, G. Biswas, and P. J. Mosterman, "A combined qualitative/quantitative approach for fault isolation in continuous dynamic systems," in *SafeProcess 2000*, vol. 1, Budapest, Hungary, pp. 1074–1079, June 2000.
- [12] S. Narasimhan, P. J. Mosterman, and G. Biswas, "A systematic analysis of measurement selection algorithms for fault isolation in dynamic systems," in *Proc. of DX 1998*, Cape Cod, MA USA, pp. 94–101, May 1998.
- [13] A. Aho, J. Hopcroft, and J. Ulman, *Data structures and algorithms*. Reading, MA: Addison-Wesley, 1983.
- [14] G. Biswas, E.-J. Manders, J. Ramirez, N. Mahadevan, and S. Abdelwahed, "Online model-based diagnosis to support autonomous operation of an advanced life support system," *Habitation*, vol. 10(1), pp. 21–38, 2004.