Relative Measurement Orderings in Diagnosis of Distributed Physical Systems

Matthew Daigle, Xenofon Koutsoukos, and Gautam Biswas

Institute for Software Integrated Systems (ISIS)
Department of Electrical Engineering and Computer Science
Vanderbilt University
Nashville, TN 37235

Abstract

Fault diagnosis in large-scale, distributed physical systems often requires the use of a large number of measurements to achieve complete diagnosability. The computational complexity of the diagnosis algorithm increases with the number of measurements, making centralized approaches infeasible for online analysis. This paper presents an extension to the TRANSCEND framework for qualitative fault diagnosis in complex physical systems. The TRANSCEND framework is based solely on qualitative time-derivative effects. Our approach combines relative measurement orderings with the traditional fault signature approach to increase the discriminatory power of a set of measurements. The measurement orderings are based on a qualitative analysis of the temporal propagation of fault effects derived from the temporal causal graph of the system. These orderings allow for diagnosis with fewer measurements. More importantly, in large-scale systems, the orderings can be used to reduce the number of measurements used by local diagnosers, leading to more efficient algorithms. The application of the approach to large-scale, distributed systems is illustrated using a multi-tank system.

1 Introduction

Fault diagnosis in large-scale, distributed physical systems can be a challenging task because of the many ways subsystems can interact. In order to uniquely isolate all faults in such systems, a large number of measurements may be required. Moreover, in many applications, the cost of adding certain sensors may be prohibitive or just not possible physically. This motivates the need to exploit the full disciminatory power of a set of measurements so that faults can be diagnosed with the smallest number of measurements necessary.

Previous work in the area of model-based fault diagnosis produced the TRANSCEND architecture, in which faults are distinguished using qualitative time-derivative effects known as fault signatures [1]. Centralized diagnoser design is based on measurement selection algorithms, in which minimal measurement sets are determined to uniquely isolate all faults [2]. Independent local diagnosers that produce globally correct diagnoses have been designed in [3]. These algorithms distinguish faults by fault signatures alone, which do not always contain enough information to uniquely diagnose all faults.

Related work in distributed diagnosis has mostly been concerned with discrete-event systems. Failure diagnosis in discrete-event systems is based on relating faults to a set of observable events [4]. This work was extended in [5] for distributed diagnosis, where a centralized coordinator is used to generate a global diagnosis. In [6] and [7], the local diagnosers communicate directly with each other to work out a consistent global diagnosis, based on a distributed constraint satisfaction formulation of the problem.

Our work extends TRANSCEND to incorporate the notion of relative measurement orderings. Relative measurement orderings establish an ordering between measurement pairs for different faults. These orderings are based on a predicted temporal order of measurement deviation, and are systematically derived from the system model. Using this new information often increases the discriminatory power of a set of measurements, providing extra information by which faults can be distinguished, leading to more efficient diagnosers.

The paper is organized as follows. Section 2 gives the background on TRANSCEND. Section 3 formalizes the idea of relative measurement orderings. Section 4 describes the extensions to TRANSCEND, and Section 5 explains how distributed diagnosis is performed with relative measurement orderings. Section 6 gives distributed diagnoser design results for a six-tank system. Section 7 concludes the paper.

2 Background

The TRANSCEND architecture employs a qualitiative model-based approach for fault isolation in complex physical systems [1]. System models are constructed using bond graphs [8]. Faults are modeled as abrupt and persistent changes in parameter values in the bond graph model of the system. We assume the sampling rate is fast enough to capture the system dynamics. Further, we also make the single fault assumption.

Example. The application of the TRANSCEND scheme and the new algorithms presented in this work are illustrated on a multi-tank system. The system consists of a series of tanks connected by pipes. Each tank also has a drainage pipe, and the first tank has an input flow. Figure 1 shows a six-tank system and its bond graph, where tanks are respresented by capacitor components, pipes by resistor elements, common pressure points by 0-junctions, and common flow points by 1-junctions. The set of possible component faults can be mapped to the set of all parameter values in the bond graph, i.e. $F = \{C_1^-, \ldots, C_6^-, R_1^+, \ldots, R_6^+, R_{12}^+, \ldots, R_{56}^+\}$, where C_i^- is a capacitance decrease in tank i, and R_j^+ is a blockage in pipe j. The set of all possible measurements consists of all tank pressures and pipe flows, i.e. $M = \{e1, e6, e11, e16, e21, e26, f2, f7, f12, f17, f22, f27, f4, f9, f14, f19, f24\}$.

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 transient dynamics caused by abrupt faults. Deviations in measurement values after a fault occurrence constitute a fault signature, where predicted deviations in magnitude and higher order derivative values are mapped to symbols of the set $\{+, 0, -\}$, which correspond to deviations above normal, no deviations, and deviations below normal, respectively.

The fault isolation algorithm in TRANSCEND utilizes the Temporal Causal Graph (TCG) representation to predict fault effects. The TCG can be derived directly from the bond graph model of the system. It models the causal relations between system variables and thus causality of physical effects in the system. It specifies the signal flow graph of the system in a form where edges are labeled with a single component parameter value

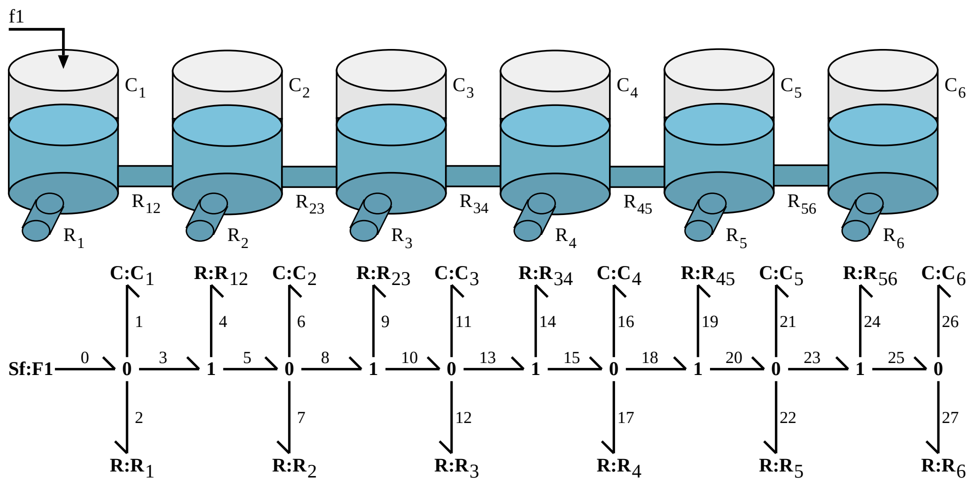

Figure 1: Six-tank system diagram and its bond graph model.

or a direct or inverse proportionality relation. Figure 2 depicts the TCG of the six-tank system.

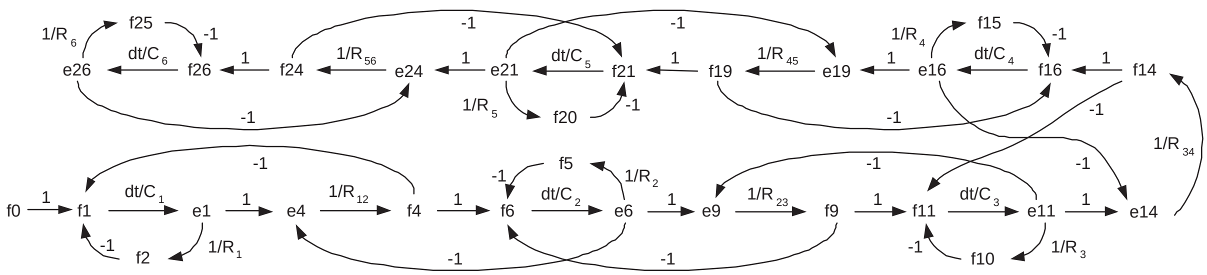

Figure 2: Six-tank system temporal causal graph.

Fault signatures are generated using a forward-propagation algorithm on the TCG. The qualitative effect of a fault, + or $-$, is propagated to all measurement vertices in the TCG. Instantaenous edges are traversed first, followed by temporal edges (specified by a *dt* label). When a temporal edge is traversed, the signature along that path goes from an i^{th} to $(i+1)^{th}$ order signature. In this way fault signatures are built up for each measurement. These give a temporal progression of the predicted qualitative changes in the signal's transient by the following lemma [9].

Lemma 1. *For a measurement, lower order effects manifest before higher order effects in response to an abrupt fault in the system.*

Fault isolation in Transcend consists of generating observed fault signatures from measurement residuals. These are compared to predicted fault signatures to discriminate between the fault hypotheses. Diagnosis in Transcend is based only on fault signatures as the discriminating information. Before diagnosers can be designed, diagnosability must be ensured. A system is diagnosable if all faults of interest can be uniquely isolated with the given measurement set using the available discriminatory information.

Example. For the six-tank system, assuming the set of available measurements to be only the tank pressures, i.e. $M = \{e1, e6, e11, e16, e21, e26\}$, and the fault set to include all possible faults, the system is not diagnosable based on the fault signatures alone. Table 1 shows some of the fault signatures for the tank system. All the drainage pipe faults $\{R_1^+, \ldots, R_6^+\}$ cannot be discriminated between since they produce the same effects on all pressure measurements (the fault signatures are the same). However, from intuition one would expect that the effects will manifest themselves at different times in different measurements, because of the system's structure. If the first tank's drainage pipe becomes blocked (R_1^+), then its pressure ($e1$) will rise before the fault propogates to the last tank, causing its pressure ($e6$) to deviate above normal. It is this notion that yields the idea of relative measurement orderings, formalized in the following section.

Fault	e1	e2	e3	e4	e5	e6
C_1^-	$+-$	$0+$	$0+$	$0+$	$0+$	$0+$
C_2^-	$0+$	$+-$	$0+$	$0+$	$0+$	$0+$
R_1^+	$0+$	$0+$	$0+$	$0+$	$0+$	$0+$
R_2^+	$0+$	$0+$	$0+$	$0+$	$0+$	$0+$
R_{12}^+	$0+$	$0-$	$0-$	$0-$	$0-$	$0-$
R_{23}^+	$0+$	$0+$	$0-$	$0-$	$0-$	$0-$

Table 1: Fault signatures for the six-tank system, with leading zeros removed.

3 Relative Measurement Orderings

Relative measurement orderings refer to the intuition that the effects of faults will manifest in some parts of the system before others.

Definition 1. *Consider a fault f and measurements m_1 and m_2; if the fault manifests in m_1 before m_2 then we can define a relative measurement ordering between m_1 and m_2 for fault f, denoted as $m_1 \prec_f m_2$.*

This information can be systematically derived from the TCG of the system, since the TCG captures both the causal ordering and the temporal effects of an abrupt parameter change. Relative measurement orderings are based on the notion of a *fault path* in a TCG.

Definition 2. *A fault path for a fault f and measurement m is a path in the TCG which begins at the fault f and ends at the measurement m.*

The set of all fault paths from f to m is denoted by $FP_{f,m}$. The order of a fault path is defined as the number of temporal edges in the path. A minimum order fault path is a path in $FP_{f,m}$ which contains the minimum number of temporal edges needed to reach m from f. More than one fault path of a specific order may exist for f and m, since there are often multiple paths from one vertex to another in the TCG.

Definition 3. *The minimum order fault path set for f and m is the set of all minimum order fault paths, and is denoted as $FP^*_{f,m}$.*

A fault path represents the temporal propagation of a fault to a specific measurement variable in the system. For a certain fault, there are multiple fault paths leading to a measurement. Since lower order effects of faults manifest themselves first (by Lemma 1), only the minimum order fault path sets are useful in determining relative measurement orderings. For this purpose, we define a method of comparing fault paths.

Definition 4. *For $p \in FP_{f,m_1}$ and $p' \in FP_{f,m_2}$, $p \sqsubset p'$ if all temporal edges in p exist in p' in the same ordering, and the order of p is less than the order of p'. If so, we say p is a temporal subpath of p', denoted as $p \sqsubset p'$*

Theorem 1. *If for every $p' \in FP^*_{f,m_2}$ there exists $p \in FP^*_{f,m_1}$ such that $p \sqsubset p'$, then we have the relative measurement ordering $m_1 \prec_f m_2$.*

Proof. In the signal flow graph for the TCG, let r_1 be the measurement vertex corresponding to m_1, r_2 the vertex corresponding to m_2, and r_f correspond to the successor vertex of the edge with fault parameter f. The transfer functions from r_f to r_1, $R_1(s)$ and from r_f to r_2, $R_2(s)$, can be derived using Mason's rule. Assume for every $p' \in FP^*_{f,m_2}$ there exists $p \in FP^*_{f,m_1}$ such that $p \sqsubset p'$. Then each minimum order path from r_f to r_2 must go through r_1 or a vertex which can be expressed as $r_1 \cdot G$, where G is some constant gain. $R_2(s)$ is a sum of terms which each correspond to different forward paths from r_f to r_2. By Lemma 1, terms that correspond to forward paths of non-minimum order can be removed to produce $R_2'(s)$. Similarly, $R_1'(s)$ can be produced. Because every minimum order path from r_f to r_2 goes through a vertex $r_1 \cdot G$, $R_1'(s)$ must appear as a factor in each term of $R_2'(s)$, therefore $R_2'(s) = H(s)R_1'(s)$, where $H(s)$ is a proper transfer function. The order of m_1 is less than the order of m_2 by the definition of the $\sqsubset$ relationship, so the number of poles for $R_1'(s)$ must be less than the number of poles for $R_2'(s)$. Therefore $H(s)$ must introduce more poles than zeros to $R_2'(s)$, and, therefore, $H(s)$ is strictly proper. From $H(s)$ we can discretize using the given sampling rate of the system to get $H(z)$. Since $H(s)$ is strictly proper, $H(z)$ is, therefore $r_2'(k) = f(r_1'(k-1))$. Since $r_2'(k)$ depends only on past values of $r_1'(k)$, with appropriately selected detection thresholds[1], a deviation resulting from fault f will appear first in m_1 and then in m_2, thus $m_1 \prec_f m_2$. □

Therefore, for a given fault, we can say that it manifests in measurement m_1 before measurement m_2 if for all minimum order fault paths of m_2, there is a minimum order fault path for m_1 the fault will traverse before completely traversing the given fault path of m_2. The transient due to the fault is slower for m_2 than for m_1, therefore the fault will manifest first in m_1 and then in m_2. Using this information, we can distinguish pairs of faults.

Definition 5. *An ordering set for a fault f, R_f, is the set of all ordering relations for fault f.*

Definition 6. *A conflict between ordering sets R_{f_1} and R_{f_2} for measurement set M exists if there are two measurements $m_i, m_j \in M$ such that $\{m_i \prec_{f_1} m_j\} \in R_{f_1}$ and $\{m_j \prec_{f_2} m_i\} \in R_{f_2}$.*

For a given measurement set and for each fault, we can derive a set of fault signatures and also a set of ordering relations. Signatures alone have been used to distinguish between different faults in [1, 9]. However, the set of ordering relationships can also be used as distinguishing information for fault isolation. Therefore, the discriminatory power of a set of measurements is enhanced by using both fault signatures and relative measurement orderings. Two faults can be discriminated between if they have different fault signatures or if they have conflicts in their ordering sets. Futher, these two notions are independent

[1]This guarantees that for some time $|r_1(k)|$ will be greater than $|r_2(k)|$, after that time $|r_2(k)|$ may overtake $|r_1(k)|$ depending on the gain of $H(z)$. Therefore thresholds must be small enough such that deviations will cross them before that time.

Measurement	Order	Minimum order fault path set
$e4$	3	$\{\{dt/C_3, dt/C_2, dt/C_1\}\}$
$e7$	2	$\{\{dt/C_3, dt/C_2\}\}$
$e23$	1	$\{\{dt/C_3\}\}$
$e20$	2	$\{\{dt/C_3, dt/C_4\}\}$
$e26$	3	$\{\{dt/C_3, dt/C_4, dt/C_5\}\}$
$e12$	4	$\{\{dt/C_3, dt/C_4, dt/C_5, dt/C_6\}\}$

Table 2: Minimum order fault path sets for R_3^+, with instantaneous edges removed.

and can be combined to distinguish among fault hypotheses. The implementation can therefore check for distinguishability using the information in any order.

Example. In the six-tank system with the measurement set $M = \{e1, e6, e11, e16, e21, e26\}$, faults $\{R_1^+, \ldots, R_6^+\}$ cannot be distinguished using only the fault signatures because their signatures are the same for all measurements. Utilizing relative measurement orderings adds to the discriminatory power of the set M. Table 2 shows $FP^*_{R_3^+}$. It follows that the ordering set for this fault is $\{(e7 \prec e4), (e23 \prec e4), (e23 \prec e7), (e23 \prec e20), (e23 \prec e26), (e20 \prec e26), (e23 \prec e12), (e20 \prec e12), (e26 \prec e12)\}$[2]. The ordering set is computed using an algorithm described in Section 4. The effects of the fault R_3^+ (pipe blockage) will manifest in the third tank before all other tanks, in the second tank before the first, in the fourth before the fifth, and in the fifth before the sixth. The effect is similar for the other drainage pipe faults, for example R_4^+ manifests first in tank 4 and then in tank 3. We can now discriminate between drain pipe blockage faults by knowing that for each of them, their effects manifest first in the tank to which they belong. By using this new information, the system can now be shown to be completely diagnosable for the given measurement set (Section 4).

4 Distributed Diagnosis

Extending the TRANSCEND architecture to incorporate relative measurement orderings requires modifying the prediction algorithm used to derive the fault signatures, so that orderings are also determined. The prediction algorithm propagates the effects of a fault throughout the TCG, considering instantaneous edges before temporal edges, thus exploring the vertices in increasing derivative order. The algorithm is augmented to keep track of the fault paths, only keeping the minimum order fault paths for each vertex. This is a straightforward modification of the algorithm presented in [1].

After the prediction step, a second algorithm is used for each fault candidate to construct its ordering set. The algorithm does this by comparing the minimum order fault path sets between pairs of measurements. Fault paths are compared using the $\sqsubset$ operator. If for every $p' \in FP^*_{f,m_2}$ there exists $p \in FP^*_{f,m_1}$ $p \sqsubset p'$, then we have the relative measurement ordering $m_1 \prec_f m_2$, and we add it to the fault's ordering set, R_f.

We extend TRANSCEND to include relative measurement orderings as discriminatory information in the diagnosability algorithm as well. It operates by comparing fault signatures for pairs of faults. If there is no measurement for which the two faults have different signatures, the faults are indistinuishable using signatures. This procedure is

[2]The subscript R_3^+ is omitted.

Algorithm 1 Determine Diagnosability

Input: set of fault candidates F each with fault signature set S_f ordered by measurements, and ordering set R_f for measurement set M
Output: set of indistinguishable faults F^I
for all $f \in F$ **do**
  for all $f' \in F$ **do**
    if $f \neq f'$ **then**
      if $S_f = S_{f'}$ and not conflict$(R_f, R_{f'})$ **then**
        $F^I = F^I \cup \{f, f'\}$
      end if
    end if
  end for
end for
return F^I

augmented to check for conflicts in ordering sets, such that when two faults cannot be distinguished by signatures alone, their ordering sets are compared. If a conflict exists in the ordering sets, then the faults can be distinguished using this new information. The new diagnosability algorithm is shown as Algorithm 1. If a pair of faults is found which are indistinguishable using both sets of information, the faults are added to the set of indistinguishable faults.

Using Algorithm 1, we can determine if, given fault, measurement, signature, and ordering sets, all faults are distinguishable from each other. This is accomplished by comparing all pairs of faults using their signatures and ordering sets. The following propostion shows that this can be done efficiently in polynomial time.

Proposition 1. *The diagnosability algorithm runs in $O(x^2y^2)$ time, given x faults and y measurements.*

Proof. All pairs of faults are considered, and for each pair, there are $O(y)$ signatures to compare, taking $O(x^2y)$ time. Also for each fault pair, there are two ordering sets of size $O(y^2)$ to compare, and if ordering relationships are indexed by measurement, checking for conflicts can be done in time linear in the size of an ordering set, so $O(y^2)$ time per fault pair, taking $O(x^2y^2)$ operations to compare all ordering sets. Therefore this algorithm is in total $O(x^2y^2)$, so diagnosability can be determined in polynomial time. □

Relative measurement orderings change the monitoring algorithm of Transcend as well. The online monitoring algorithm starts with a set of fault candidates and their associated fault signatures after an initial deviation has been detected. It matches the candidates' predicted fault signatures to observed measurement deviations as they appear, dropping candidates whose signatures are inconsistent with observed transients. This algorithm is modified such that candidates are also dropped if there is an inconsistency between predicted measurement orderings and observed measurement orderings.

In designing diagnosers for distributed systems, including measurement orderings can create smaller and more efficient local diagnosers. Each local diagnoser is defined by the set of faults it must be able to diagnose, and the set of measurements that are locally available to it. The goal is to determine which measurements need to be communicated from other systems in order for each local diagnoser to obtain a globally correct diagnosis. This avoids the need for a centralized coordinator. A distributed algorithm has been developed using only fault signatures to determine such a design [3]. Utilizing relative

Algorithm 2 Distributed Diagnoser Design

Input: local fault sets F_i, local measurement sets M_i, fault signatures, ordering sets, k subsystems
for subsystem $i \in 1, \ldots, k$ **do**
  identify set $remFaults_i$ such that $f \in remFaults_i$ cannot be completely distinguished using M_i (using extended diagnosability algorithm)
  for $f \in remFaults_i$ **do**
    identify minimum set of communicated measurements to globally diagnose f (using extended diagnosability algorithm)
    add this set to the local measurement set
  end for
end for

measurement orderings in many cases will allow the local diagnosers to require fewer local measurements and also require fewer measurements to be communicated.

The algorithm generates the distributed diagnoser by minimizing the number of shared measurements between subsystems. For each subsystem, if a fault is not globally diagnosable using local measurements, it searches neighboring subsystems for a minimal set of additional measurements to make the fault globally diagnosable. The goal is to achieve a unique diagnosis with minimum communication between the subsystems. This algorithm is extended to include relative measurement orderings when determining if a set of measurements can distinguish a set of faults. The modified algorithm is shown as Algorithm 2. In the worst case all combinations of measurements are considered, so the algorithm is exponential. Practically, this is done at design time so its time complexity is not of much concern.

5 Experimental Results

Evaluation of a diagnoser's design is dependent on the number of measurements it uses to diagnose a given set of faults. A diagnoser is more efficient with a smaller number of measurements because it has to make less comparisons in the online monitoring algorithm. Also, less measurements need to deviate to obtain a diagnosis, so the diagnosis will typically be achieved faster.

As previously shown, a given measurement set may not be able to completely diagnose all faults using fault signatures alone. Therefore, in order to evaluate and compare designs with and without measurement orderings, we assume the system is diagnosable using fault signatures alone, and compare the number of measurements required to obtain a global diagnosis. To meet this criteria, we assume the fault set, F, to be $\{C_1^-, \ldots, C_6^-, R_1^+, \ldots, R_6^+, R_{12}^+, \ldots, R_{56}^+\}$, and the measurement set, M, to be $\{e1, e6, e11, e16, e21, e26, f4, f9, f14, f19, f24\}$, i.e. all tank pressures and connecting pipe flows.

A centralized diagnoser design is based on finding the minimum set of measurements which provides diagnosability for the given fault and measurement sets. For the chosen F and M, without measurement orderings, all measurements are required to diagnose all faults, i.e. $M = \{e1, e6, e11, e16, e21, e26, f4, f9, f14, f19, f24\}$. With orderings, however, only the tank pressures are necessary i.e. $M = \{e1, e6, e11, e16, e21, e26\}$. In this case, using orderings reduces the size of the diagnoser.

Table 3 shows the results of a distributed diagnoser design. Communicated measurements are denoted by an asterisk. The subsystems are taken to be a tank with its respective drainage pipe and its right connecting pipe. Each subsystem is responsible for faults in its components, and has its tank pressure and connecting pipe flow rate as measurements, i.e. the local fault sets are $\{C_1^-, R_1^+, R_{12}^+\}$, $\{C_2^-, R_2^+, R_{23}^+\}$, ..., $\{C_6^-, R_6^+\}$. Without measurement orderings, the subystems each need their own tank pressures and connecting pipe flows, along with the flow rate for the previous subsystem and the tank pressure for the next subsystem. With orderings, only the tank pressures are needed, and each subsystem takes the pressure measurement from its neighboring subsystems. In this case, using orderings reduces the size of each local diagnoser.

Subsystem	Design Without Orderings	Design With Orderings
1	$M_1 = \{e1, f4, e6^*\}$	$M_1 = \{e1, e6^*\}$
2	$M_2 = \{e6, f9, f4^*, e11^*\}$	$M_2 = \{e6, e1^*, e11^*\}$
3	$M_3 = \{e11, f14, f9^*, e16^*\}$	$M_3 = \{e11, e6^*, e16^*\}$
4	$M_4 = \{e16, f19, f14^*, e21^*\}$	$M_4 = \{e16, e11^*, e21^*\}$
5	$M_5 = \{e21, f24, f19^*, e26^*\}$	$M_5 = \{e21, e16^*, e26^*\}$
6	$M_6 = \{e26, f24^*\}$	$M_6 = \{e26, e21^*\}$

Table 3: Distributed Diagnoser Designs

Table 4 shows the results of a second distributed diagnoser design. The subsystems are taken to be all tanks with their drainage pipes as one subsystem, and each connecting pipe in its own subsystem. We consider only faults in the pipes. The subsystems are $\{\{R_1^+,\ldots,R_6^+\},\{e1,\ e6,\ e11,\ e16,\ e21,\ e26\}\}$, $\{\{R_{12}^+\},\{f4\}\}$, $\{\{R_{23}^+\},\{f9\}\}$, ..., $\{\{R_{56}^+\},\{f24\}\}$. Without measurement orderings, the subystem consisting of the tanks and drainage pipes needs the flow rates of the connecting pipes to be communicated. With measurement orderings, this communication is not required. In this case, using orderings reduces the communication requirements of the local diagnosers.

Subsystem	Design Without Orderings	Design With Orderings
1	$M_1 = \{e1, e6, e11, e16, e21, e26,$ $f4^*, f9^*, f14^*, f24^*\}$	$M_1 = \{e1, e6, e11, e16, e21, e26\}$
2	$M_2 = \{f4\}$	$M_2 = \{f4\}$
3	$M_3 = \{f9\}$	$M_3 = \{f9\}$
4	$M_4 = \{f14\}$	$M_4 = \{f14\}$
5	$M_5 = \{f19\}$	$M_5 = \{f19\}$
6	$M_6 = \{f24\}$	$M_6 = \{f24\}$

Table 4: Distributed Diagnoser Designs for a Different Subsystem Partitioning

6 Conclusions

In this paper, we have presented and analyzed the concept of relative measurement orderings for fault diagnosis. Algorithms to systematically derive the information and use it to diagnose faults have been constructed. Including this new information has been shown to increase the discriminatory power of a given set of measurements, achieving diagnosability with fewer measurements. Therefore some large systems which were not previously diagnosable now are by including measurement orderings. Also, some previously diagnosable systems are now diagnosable with fewer measurements, leadings to smaller, faster diagnosers. The benefit of using measurement orderings in diagnoser design has also been shown.

Future work will address the benefit of relative measurement orderings in the partioning method of distributed diagnoser design, where optimal partitions of the fault sets are measurement sets are determined such that no communication of measurements is needed between the distributed diagnosers. Experimental application to large-scale systems, such as a tank system and multi-robot teams, will also be explored.

References

[1] P. Mosterman and G. Biswas, "Diagnosis of continuous valued systems in transient operating regions," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 29, no. 6, pp. 554–565, 1999.

[2] 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, May 1998, pp. 94–101.

[3] I. Roychoudhury, G. Biswas, X. Koutsoukos, and S. Abdelwahed, "Designing distributed diagnosers for complex physical systems," in *Proceedings of the 16th International Workshop on Principles of Diagnosis (DX 05)*, Monterey, California, June 2005, pp. 31–36.

[4] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis, "Failure diagnosis using discrete-event models," *IEEE Transactions on Control Systems Technology*, vol. 4, no. 2, pp. 105–124, 1996.

[5] R. Debouk, S. Lafortune, and D. Teneketzis, "Coordinated decentralized protocols for failure diagnosis of discrete event systems," *Discrete Event Dynamic Systems*, vol. 10, no. 1–2, pp. 33–86, Jan 2000.

[6] J. Kurien, X. Koutsoukos, and F. Zhao, "Distributed diagnosis of networked embedded systems," in *Proceedings of the 13th InternationalWorkshop on Principles of Diagnosis (DX-2002)*, Semmering, Austria, May 2002, pp. 179–188.

[7] 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)*, Zaragosa, Spain, Oct 2002, pp. 169–174.

[8] 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.

[9] E.-J. Manders, S. Narasimhan, G. Biswas, and P. Mosterman, "A combined qualitative/quantitative approach for fault isolation in continuous dynamic systems," in *SafeProcess 2000*, vol. 1, Budapest, Hungary, June 2000, pp. 1074–1079.