

An Event-based Approach to Distributed Diagnosis of Continuous Systems

Matthew Daigle¹, Indranil Roychoudhury², Gautam Biswas³, and Xenofon Koutsoukos³

¹ University of California, Santa Cruz, NASA Ames Research Center, Moffett Field, CA, 94035, USA
matthew.j.daigle@nasa.gov

² SGT Inc., NASA Ames Research Center, Moffett Field, CA, 94035, USA
indranil.roychoudhury@nasa.gov

³ Institute for Software Integrated Systems, Dept. of EECS, Vanderbilt University, Nashville, TN, 37235, USA
{gautam.biswas,xenofon.koutsoukos}@vanderbilt.edu

ABSTRACT

Distributed fault diagnosis solutions are becoming necessary due to the complexity of modern engineering systems, and the advent of smart sensors and computing elements. This paper presents a novel event-based approach for distributed diagnosis of abrupt parametric faults in continuous systems, based on a qualitative abstraction of measurement deviations from the nominal behavior. We systematically derive dynamic fault signatures expressed as event-based fault models. We develop a distributed diagnoser design algorithm that uses these models for designing local event-based diagnosers based on global diagnosability analysis. The local diagnosers each generate globally correct diagnosis results locally, without a centralized coordinator, by communicating a minimal number of measurements between themselves. The proposed approach is applied to a multi-tank system, and results demonstrate a marked improvement in scalability compared to a centralized approach.

1 INTRODUCTION

The complexity of modern engineering systems warrants the adoption of fault diagnosis capabilities to ensure system safety, reliability, and availability. Faults must be quickly isolated so that mitigation or recovery actions may be taken. As systems become more complex, it is correspondingly more difficult to develop and deploy centralized diagnosis solutions. Further, such centralized schemes have single points of failure, do not scale as the size of systems increases, and have large computational and memory requirements. This, along with the increased pervasiveness of distributed, networked components, fuels the need for distributed diagnosis frameworks.

In previous work, we have developed a centralized framework for qualitative event-based diagnosis for parametric faults in continuous systems (Daigle *et al.*, 2009). Deviations of measured behavior from predicted nominal behavior, termed *fault signatures*, are

captured qualitatively using magnitude and slope symbols, forming the basis of the qualitative fault isolation scheme (Mosterman and Biswas, 1999). The orders in which these deviations manifest, termed *relative measurement orderings*, are also used for fault isolation, thus forming event-based descriptions of fault-induced behavior. This diagnostic information may be computed from the system model and used to build event-based diagnosers similar to those used for discrete-event systems (DES) (Sampath *et al.*, 1996). However, this centralized approach scales poorly, because as the number of faults and measurements increases, the possible number of event traces increases as well.

To address the problems of centralized diagnosis, we apply the distributed diagnoser design methodology presented in (Roychoudhury *et al.*, 2009) to the formal event-based framework developed in (Daigle *et al.*, 2009). The distributed diagnoser design approach of (Roychoudhury *et al.*, 2009) is based on global diagnosability analysis, where the local diagnosers are designed to provide globally correct diagnosis results, without a centralized coordinator, and by communicating a minimal number of measurements among themselves. The approach does not incorporate measurement orderings, but the addition of measurement orderings improves diagnosability, allowing the local diagnosers to be more efficient.

This paper presents, using a multi-tank system as a case study, how a global event-based diagnoser may be decomposed into several independent local event-based diagnosers, each of which leverages measurement orderings for diagnosis. We develop an algorithm for designing distributed diagnosers based on the ideas of (Roychoudhury *et al.*, 2009), but which uses measurement orderings to guide the diagnoser design process. Distributed diagnoser design results demonstrate the reduction in diagnoser size that may be obtained using this approach, resulting in, for each subsystem, a small, compact local diagnoser capable of providing globally correct diagnoses of local faults. Results demonstrate the improved scalability of the distributed approach over a centralized approach.

The paper is organized as follows. Section 2 formulates the system model. Section 3 reviews quali-

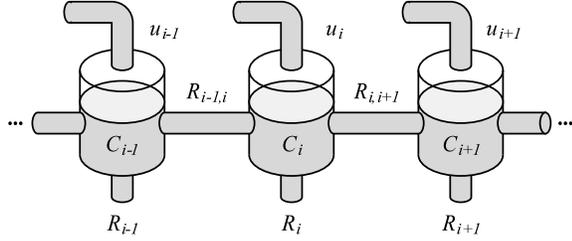


Figure 1: Tank system schematic.

tative fault isolation and event-based fault modeling, and defines diagnosability in the event-based framework. Section 4 describes the distributed diagnoser design problem. Section 5 discusses the global and local diagnoser construction, and Section 6 demonstrates the approach in simulation, and provides scalability results. Section 7 concludes the paper.

2 MODEL FORMULATION

We consider the problem of single fault diagnosis in continuous systems. We assume the system, \mathcal{S} , is described by

$$\begin{aligned}\dot{\mathbf{x}}(t) &= \mathbf{f}(\mathbf{x}(t), \boldsymbol{\theta}(t), \mathbf{u}(t)) + \mathbf{v}(t) \\ \mathbf{y}(t) &= \mathbf{h}(\mathbf{x}(t), \boldsymbol{\theta}(t), \mathbf{u}(t)) + \mathbf{n}(t),\end{aligned}$$

where $\mathbf{x}(t) \in \mathbb{R}^{n_x}$ is the state vector, $\boldsymbol{\theta}(t) \in \mathbb{R}^{n_\theta}$ is the parameter vector, $\mathbf{u}(t) \in \mathbb{R}^{n_u}$ is the input vector, $\mathbf{v}(t) \in \mathbb{R}^{n_v}$ is the process noise vector, assumed to be zero-mean Gaussian, \mathbf{f} is the state equation, $\mathbf{y}(t) \in \mathbb{R}^{n_y}$ is the output vector, $\mathbf{n}(t) \in \mathbb{R}^{n_n}$ is the measurement noise vector, assumed to be zero-mean Gaussian, and \mathbf{h} is the output equation. The dimension of a vector \mathbf{a} is denoted by n_a .

We denote a measurement as m , which is a time-varying signal of $\mathbf{y}(t)$ obtained from an associated sensor. A measurement set is denoted as M .

We consider single, abrupt, parametric faults, where faults are modeled as unexpected step changes in system parameter values. We name faults by the associated parameter and the direction of change, i.e., θ^+ denotes a fault defined as an increase in the value of parameter θ , and θ^- denotes a fault defined as a decrease in the parameter value. We denote a fault as f and a set of faults as F .

Throughout the paper, we will use a multi-tank system as a running example. The tanks are connected serially as shown in Fig. 1, and we will consider a variable number of tanks. For tank i , u_i denotes the input flow, C_i denotes the capacitance, and R_i denotes the drain pipe resistance. For tanks i and j , $R_{i,j}$ denotes the resistance of the connecting pipe. For an n -tank system, the pressure of tank $i = 1$ is described by

$$\dot{p}_i = \frac{1}{C_i} \left(u_i - \frac{1}{R_i} (p_i) - \frac{1}{R_{i,i+1}} (p_i - p_{i+1}) \right),$$

of tanks $i = 2, \dots, n-1$ by

$$\begin{aligned}\dot{p}_i &= \frac{1}{C_i} \left(u_i + \frac{1}{R_{i-1,i}} (p_{i-1} - p_i) \right. \\ &\quad \left. - \frac{1}{R_i} (p_i) - \frac{1}{R_{i,i+1}} (p_i - p_{i+1}) \right),\end{aligned}$$

and of tank $i = n$ by

$$\dot{p}_i = \frac{1}{C_i} \left(u_i - \frac{1}{R_i} (p_i) - \frac{1}{R_{i-1,i}} (p_{i-1} - p_i) \right).$$

The complete fault set consists of $\{C_i^-, C_i^+, R_i^-, R_i^+ : i = 1, \dots, n\} \cup \{R_{i,i+1}^-, R_{i,i+1}^+ : i = 1, \dots, n-1\}$. The complete measurement set is defined as $\{q_i : i = 1, \dots, n\}$, where q_i describes the output flow of tank i , i.e.,

$$q_i = \frac{1}{R_i} (p_i).$$

3 QUALITATIVE EVENT-BASED DIAGNOSIS FRAMEWORK

We develop an event-based, qualitative diagnosis framework. Faults are viewed as unobservable events, manifesting as persistent abrupt changes in system parameter values. These faults cause transients in the system behavior, causing deviations in observed measurement values from nominal measurement values. In this section, we first review the theoretical framework for qualitative fault isolation, followed by a formal framework for event-based fault modeling.

3.1 Qualitative Fault Isolation

Measurement deviations from nominal values caused by faults are abstracted using qualitative $+$, $-$, and 0 values to form *fault signatures* (Mosterman and Biswas, 1999). Fault signatures represent these deviations as the immediate change in magnitude and the first nonzero derivative change.

Definition 1 (Fault Signature). A *fault signature* for a fault f and measurement m is the qualitative magnitude and slope change in m caused by the occurrence of f , and is denoted by $\sigma_{f,m} \in \Sigma_{f,m}$.

In general, ambiguities may exist in the fault signatures, so $\sigma_{f,m}$ may not be unique. A fault signature is written as $s_1 s_2$, where s_1 is the qualitative magnitude change and s_2 is the qualitative slope change, e.g., $+-$.

We also capture the temporal order of measurement deviations, termed *relative measurement orderings* (Daigle *et al.*, 2007b), based on the intuition that fault effects will manifest in some parts of the system before others. Measurement orderings are based on analysis of the transfer functions from faults to measurements (Daigle *et al.*, 2007b).

Definition 2 (Relative Measurement Ordering). If fault f manifests in measurement m_i before measurement m_j , then we define a *relative measurement ordering* between m_i and m_j for fault f , denoted by $m_i \prec_f m_j$. We denote the set of all measurement orderings for f as $\Omega_{f,M}$.

The fault signatures and measurement orderings can be computed automatically from a system model. One method is to use a temporal causal graph (TCG) representation that is derived from the system model, along with a forward propagation algorithm to predict qualitative effects of faults on measurements and their possible sequences of deviations (Daigle, 2008).

The fault signatures and measurement orderings for a three-tank system with $F = \{C_1^-, C_2^-, C_3^-, R_1^+, R_2^+$,

Table 1: Fault Signatures and Relative Measurement Orderings for the Three-tank System

Fault	q_1	q_2	q_3	Measurement Orderings
C_1^-	+−	0+	0+	$q_1 \prec q_2, q_1 \prec q_3, q_2 \prec q_3$
C_2^-	0+	+−	0+	$q_2 \prec q_1, q_2 \prec q_3$
C_3^-	0+	0+	+−	$q_2 \prec q_1, q_3 \prec q_1, q_3 \prec q_2$
R_1^+	−+	0+	0+	$q_1 \prec q_2, q_1 \prec q_3, q_2 \prec q_3$
R_2^+	0+	−+	0+	$q_2 \prec q_1, q_2 \prec q_3$
R_3^+	0+	0+	−+	$q_2 \prec q_1, q_3 \prec q_1, q_3 \prec q_2$
R_{12}^+	0+	0−	0−	$q_2 \prec q_3$
R_{23}^+	0+	0+	0−	$q_2 \prec q_1$

$R_3^+, R_{12}^+, R_{23}^+$ and $M = \{q_1, q_2, q_3\}$ are shown in Table 1. For example, a decrease in the capacitance of tank 1, denoted by C_1^- , causes a discontinuous increase in the tank 1 output flow, q_1 , followed by a smooth decrease, denoted by the signature +−. This is followed by smooth increases in q_2 and then q_3 . The tanks provide natural delays of the propagation of fault effects, which manifest in the computed orderings.

3.2 Event-based Fault Modeling

Fault signatures combined with relative measurement orderings provide event-based information for diagnosis. For a given fault, the combination of all fault signatures and measurement orderings yields all the possible ways a fault can manifest in the measurements. We denote each of these possibilities as a *fault trace*.

Definition 3 (Fault Trace). A *fault trace* for a fault f over measurements M , denoted by $\lambda_{f,M}$, is a string of length $\leq |M|$ that includes, for every $m \in M$ that will deviate due to f , a fault signature $\sigma_{f,m}$, such that the sequence of fault signatures satisfies $\Omega_{f,M}$.

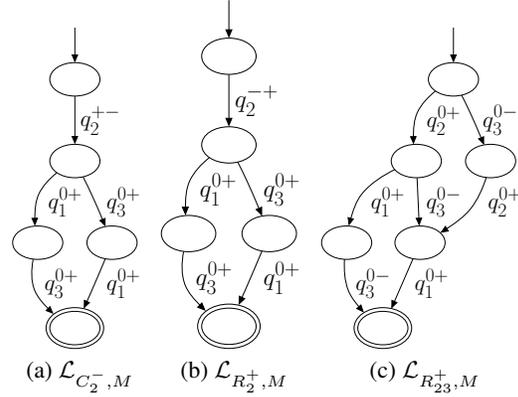
Note that the definition implies that fault traces are of maximal length, i.e., a fault trace includes deviations for all measurements affected by the fault. We group the set of all fault traces into a *fault language*. The *fault model*, defined by a finite automaton, concisely represents the fault language of a fault.

Definition 4 (Fault Language). The *fault language* of a fault $f \in F$ with measurement set M , denoted by $L_{f,M}$, is the set of all fault traces for f over measurements M .

Definition 5 (Fault Model). The *fault model* for a fault $f \in F$ with measurement set M , is the finite automaton that accepts exactly the language $L_{f,M}$, and is given by $\mathcal{L}_{f,M} = (S, s_0, \Sigma, \delta, A)$ where S is a set of states, $s_0 \in S$ is an initial state, Σ is a set of events, $\delta : S \times \Sigma \rightarrow S$ is a transition function, and $A \subseteq S$ is a set of accepting states.

The finite automata representation allows for the composition of the fault signatures and relative measurement orderings into fault models. The possible fault signatures and measurement orderings can be composed automatically to form the fault models based on the synchronization operation (Daigle *et al.*, 2009).

Selected fault models for a three-tank system are shown in Fig. 2. For example, as seen in $\mathcal{L}_{C_2^-}$, the

Figure 2: Fault models for some faults of the three-tank system, where $M = \{q_1, q_2, q_3\}$.

fault C_2^- may manifest as the fault traces $q_2^{+−} q_1^{0+} q_3^{0+}$ and $q_2^{+−} q_3^{0+} q_1^{0+}$, as implied by the fault signatures and measurement orderings.

3.3 Diagnosability

With the formal fault isolation framework defined, we may now establish the notions of *distinguishability* and *diagnosability* in this framework. Using these definitions, we can then formally define the distributed diagnoser design problem. Distinguishability between faults is characterized as follows.

Definition 6 (Distinguishability). With measurements M , a fault f_i is distinguishable from a fault f_j , denoted by $f_i \approx_M f_j$, if f_i always eventually produces effects on the measurements that f_j cannot.

Under our framework, one fault will be distinguishable from another fault if it cannot produce a fault trace that is a prefix (denoted by \sqsubseteq) of a trace that can be produced by the other fault¹. If this is not the case, then when that trace manifests, the first fault cannot be distinguished from the second.

We define a system in our framework as follows.

Definition 7 (System). A *system* \mathcal{S} is tuple $(F, M, L_{F,M})$, where $F = \{f_1, f_2, \dots, f_n\}$ is a set of faults, M is a set of measurements, and $L_{F,M} = \{L_{f_1,M}, L_{f_2,M}, \dots, L_{f_n,M}\}$ is the set of fault languages.

If a system is diagnosable, then we can make guarantees about the unique isolation of every fault in the system.

Definition 8 (Diagnosability). A system $\mathcal{S} = (F, M, L_{F,M})$ is *diagnosable* if $(\forall f_i, f_j \in F) f_i \neq f_j \implies f_i \not\approx_M f_j$.

If \mathcal{S} is diagnosable, then every pair of faults is distinguishable using the measurements in M . So, each fault trace we observe can be linked to exactly one fault, meaning that we can uniquely isolate all faults of interest. If \mathcal{S} is not diagnosable, then ambiguities

¹A fault trace λ_i is a prefix of fault trace λ_j if there is some (possibly empty) sequence of events λ_k that can extend λ_i such that $\lambda_i \lambda_k = \lambda_j$.

will remain after fault isolation, i.e., after all possible measurement deviations have been observed.

4 DISTRIBUTED DIAGNOSER DESIGN

Given a system that is *diagnosable*², our objective is to decompose the overall diagnosis task into smaller sub-tasks performed by local diagnosers with the following properties: (i) all single faults of interest in the system can be diagnosed, and (ii) the local diagnosis results are *globally* correct. These two properties eliminate the need for a centralized coordinator.

The system \mathcal{S} is split into n subsystems, where each fault is assigned to exactly one subsystem, and each subsystem gets a subset of the complete measurement set. The subsystem definitions are provided by the user as input.

Assumption 1. $\mathcal{S} = (F, M, L_{F,M})$ is split into $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n$, where $\mathcal{S}_i = (F_i, M_i, L_{F_i, M_i})$, such that (i) $F = F_1 \cup F_2 \cup \dots \cup F_n$, (ii) $\forall i \neq j \in [1, n], F_i \cap F_j = \emptyset$, and (iii) $\forall i M_i \subseteq M$.

Subsystems may be locally diagnosable. A locally diagnosable subsystem is one in which its own faults can be uniquely isolated using its own measurements.

Definition 9 (Local Diagnosability). A subsystem $\mathcal{S}_i = (F_i, M_i, L_{F_i, M_i})$ is *locally diagnosable* if $(\forall f_i \in F_i, f_j \in F_i) f_i \neq f_j \implies f_i \approx_{M_i} f_j$. We say two faults $f_i \in F_i$ and $f_j \in F_i$ are *locally distinguishable* if $f_i \not\approx_{M_i} f_j$.

Local diagnosability is not sufficient for local diagnosers to achieve globally correct diagnoses. The problem is that for \mathcal{S}_i , there may be some $f_i \in F_i$ and for \mathcal{S}_j , some $f_j \in F_j$, such that f_j produces the same effects on M_i as f_i does. The result is that, if f_j occurs local diagnoser i will say that f_i has occurred. In general, we may have faults in a subsystem that are distinguishable from faults local to the subsystem, but which may not be distinguishable from faults *outside* the subsystem. For the local diagnosers to achieve globally correct local diagnoses, the subsystems must satisfy a notion of *global diagnosability*.

Definition 10 (Global Diagnosability). A subsystem $\mathcal{S}_i = (F_i, M_i, L_{F_i, M_i})$ belonging to system $\mathcal{S} = (F, M, L_{F,M})$ is *globally diagnosable* if $(\forall f_i \in F_i, f_j \in F) f_i \neq f_j \implies f_i \approx_{M_i} f_j$. We say two faults $f_i \in F_i$ and $f_j \in F$ are *globally distinguishable* if $f_i \not\approx_{M_i} f_j$.

That is, a subsystem \mathcal{S}_i is *globally diagnosable* if all the faults F_i are distinguishable from every other fault $f \in F$ using only the measurements in M_i . If the subsystems can be structured such that each subsystem \mathcal{S}_i is globally diagnosable, then each local diagnoser can independently generate local diagnoses which are globally correct.

For example, consider the three-tank system defined earlier, with $F = \{C_1^-, C_2^-, C_3^-, R_1^+, R_2^+, R_3^+, R_{12}^+,$

$R_{23}^+\}$ and $M = \{q_1, q_2, q_3\}$. Let us define a subsystem for each tank, where for $i = 1, \dots, n-1$, \mathcal{S}_i is defined by $F_i = \{C_i^-, R_i^+, R_{i,i+1}^+\}$ and $M_i = \{q_i\}$, and for $i = n$, \mathcal{S}_i is defined by $F_i = \{C_i^-, R_i^+\}$ and $M_i = \{q_i\}$. Consider tank 1. If $0+$ is observed for q_1 , then that may be the result of local fault R_{12}^+ or any of the nonlocal faults (see Table 1). Clearly, \mathcal{S}_1 is not globally diagnosable. Note that it is locally diagnosable, as the three local faults each produce a different effect on the sole measurement of the subsystem, q_1 .

Different design problems may be defined which determine partitions of the fault set F and/or the assignment of measurements to subsystems (Roychoudhury *et al.*, 2009). In each case, the end result must be a set of globally diagnosable subsystems. In this paper, we focus on the problem where \mathcal{S} is already partitioned into subsystems, but each \mathcal{S}_i may not be globally diagnosable. We define the distributed diagnoser design problem as determining, for each \mathcal{S}_i , the minimal number of measurements to pull in from other subsystems to achieve global diagnosability. Formally, the problem can be defined as follows.

Problem (Partitioned System Diagnoser Design). Given n subsystems, where $\mathcal{S}_i = (F_i, M_i, L_{F_i, M_i})$, construct, for each subsystem, a measurement set $M_i^+ \subseteq M$ such that (i) $M_i^+ - M_i$ is minimal, and (ii) $\mathcal{S}_i' = (F_i, M_i^+, L_{F_i, M_i^+})$ is globally diagnosable.

This problem is a variation of the *measurement selection* problem, which is an instance of the set covering problem, known to be NP-complete (Narasimhan *et al.*, 1998). Our goal, while designing the local diagnosers, is to minimize the sharing of measurements across subsystems in order to limit the size of the local diagnosers and their communication requirements. We simplify the measurement search using measurement orderings as a guide, based on the intuition that measurements that deviate before others are more helpful. Further, these measurements provide the fastest diagnosis. To do this, for each fault that is not globally distinguishable, we determine the measurements that deviate first by looking at the measurement orderings, and this set of measurements over all the globally indistinguishable faults forms the current working measurement set, i.e., measurements with which we try to resolve global diagnosability. This heuristic simplifies the search process, but the algorithm is still exponential in the general case, where $O(2^{|M|})$ measurement sets must be considered for a single subsystem. The heuristic reduces the number of measurements to consider at each iteration, so only $O(2^{|M_i^+|})$ combinations end up being considered, where typically, $|M_i^+| \ll |M|$. The introduction of the heuristic trades off optimality of the diagnoser design for algorithm efficiency. Additional heuristics may also be used, e.g., the subsystem distance heuristic presented in (Roychoudhury *et al.*, 2009).

The distributed diagnoser design procedure is given as Algorithm 1. For a diagnosable system \mathcal{S} , for each \mathcal{S}_i , we first determine, using diagnosability analysis, the set of faults $F_i^* \subseteq F_i$ which are not globally distinguishable using M_i . At each iteration, for each fault that is not globally distinguishable using the current

²If the system \mathcal{S} is not diagnosable, we can define aggregate faults, where an aggregate fault is a set of faults that are indistinguishable from each other. The diagnosis methodology can be applied to the modified fault set that includes the aggregate faults (Roychoudhury *et al.*, 2009).

Algorithm 1 Distributed Diagnoser Design

Input: $\mathbb{S} = \{S_i = (F_i, M_i, L_{F_i, M_i}) : i = 1, \dots, n\}$
for all $S_i \in \mathbb{S}$ **do**
 $F_i^* \leftarrow \{f_i^* : f_i \sim_{M_i} f_j \text{ for } f_i^* \in F_i, f_j \in F, \text{ and } f_i^* \neq f_j\}$
 $M_i^+ \leftarrow M_i$
while $F_i^* \neq \emptyset$ **do**
for all $f_i^* \in F_i^*$ **do**
 $M_{f_i^*} \leftarrow \{m : \exists m', (m' \prec m) \notin \Omega_{f_i^*, M - M_i^+}\}$
end for
identify minimal $M_i^* \subseteq \bigcup_{f_i^* \in F_i^*} M_{f_i^*}$ such that
 $M_i^+ \cup M_i^*$ isolates maximal $F_i' \subseteq F_i^*$
 $M_i^+ \leftarrow M_i^+ \cup M_i^*$
 $F_i^* \leftarrow F_i^* - F_i'$
end while
construct \mathcal{D}_{F_i, M_i^+}
end for

measurement set, M_i^+ , we compute the set of measurements out of $M - M_i^+$ that may deviate first for the fault, as $M_{f_i^*}$. We then find the minimal set of measurements to add to M_i^+ from the set of measurements found in this way over all f_i^* that resolves the most globally indistinguishable faults, and add these to M_i^+ . The process repeats until S_i is globally diagnosable, resulting in the local diagnoser \mathcal{D}_{F_i, M_i^+} , whose construction is described in the next section.

It is easy to see that Algorithm 1 always succeeds in making S_i globally diagnosable, because (i) S is diagnosable, so global diagnosability for S_i can be achieved (at worst by setting $M_i^+ = M$), and (ii) the algorithm continually adds measurements to M_i^+ until S_i is globally diagnosable (and in the worst case all measurements are considered).

We apply this algorithm to the n -tank system, where for $i = 1, \dots, n-1$, S_i is defined by $F_i = \{C_i^-, R_i^+, R_{i,i+1}^+\}$ and $M_i = \{q_i\}$, and for $i = n$, S_i is defined by $F_i = \{C_i^-, R_i^+\}$ and $M_i = \{q_i\}$. For tank 1, R_{12}^+ is not globally distinguishable. From the measurement orderings, q_2 will deviate before q_3 , so $M_1^+ = \{q_2\}$. This measurement alone is sufficient to add to M_1^+ to obtain global diagnosability, so no further iteration is necessary. For tank 2, R_{23}^+ is not globally distinguishable, and both q_1 or q_3 belong to M_2^* . Measurement q_3 alone is sufficient to achieve global diagnosability. For tank 3, the subsystem is already globally diagnosable. The new measurement sets are therefore $M_1^+ = \{q_1, q_2\}$, $M_2^+ = \{q_2, q_3\}$, and $M_3^+ = \{q_3\}$.

5 DIAGNOSER IMPLEMENTATION

In this section we describe the construction of the event-based diagnosers. The goal of the event-based diagnoser is, given a sequence of measurement deviation events, to determine which faults are consistent with the observed sequence. We define formally a *diagnosis* and a *diagnoser* in our framework (Daigle *et al.*, 2009).

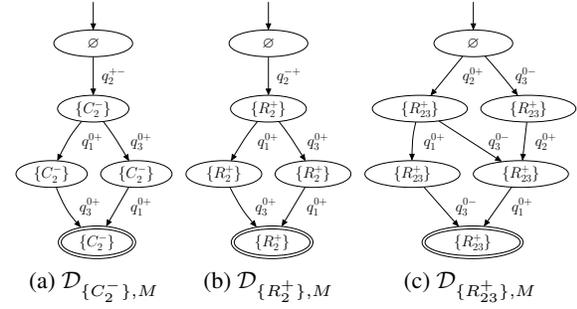


Figure 3: Diagnosers for some individual faults of the three-tank system, where $M = \{q_1, q_2, q_3\}$.

Definition 11 (Diagnosis). A *diagnosis* $d \subseteq F$ is a set of faults, each of which is consistent with the observations.

Definition 12 (Diagnoser). A *diagnoser* for a fault set F and measurement set M is a tuple $\mathcal{D}_{F, M} = (S, s_0, \Sigma, \delta, A, D, Y)$ where S is a set of states, $s_0 \in S$ is an initial state, Σ is a set of events, $\delta : S \times \Sigma \rightarrow S$ is a transition function, $A \subseteq S$ is a set of accepting states, $D \subseteq 2^F$ is a set of diagnoses, and $Y : S \rightarrow D$ is a diagnosis map.

A diagnoser is a finite automaton extended by a set of diagnoses and a diagnosis map. It takes events as inputs, which, as with fault models, correspond to measurement deviations. From the current state, a measurement deviation event causes a transition to a new state. The diagnosis for that new state represents the set of faults that are consistent with the sequence of events seen up to the current point in time.

Accepting states correspond to a fault isolation result. We say that a diagnoser *isolates* a fault if it accepts all possible valid traces for the fault and the accepting states map to diagnoses containing the fault.

Definition 13 (Isolation). A diagnoser $\mathcal{D}_{F, M}$ *isolates* fault $f \in F$ if $\mathcal{D}_{F, M}$ accepts all $\lambda_{f, M} \in L_{f, M}$ and for each $s \in A$ that accepts some $\lambda_{f, M}$, $f \in Y(s)$.

Unique isolation corresponds to system diagnosability. We say that a diagnoser *uniquely isolates* a fault if each accepting state maps to the single fault.

Definition 14 (Unique Isolation). A diagnoser $\mathcal{D}_{F, M}$ *uniquely isolates* fault $f \in F$ if $\mathcal{D}_{F, M}$ accepts all $\lambda_{f, M} \in L_{f, M}$ and for each $s \in A$ that accepts some $\lambda_{f, M}$, $\{f\} = Y(s)$.

We would like to systematically construct a diagnoser for a system S that isolates all $f \in F$, and show that if S is diagnosable, then this diagnoser uniquely isolates all $f \in F$. This procedure has been developed in previous work (Daigle *et al.*, 2009). Here, we briefly review the main points.

First, we construct a diagnoser, for each fault f , that isolates f , i.e., $\mathcal{D}_{\{f\}, M}$. These are shown in Fig. 3 for some of the faults of the three-tank system. They are constructed directly from the fault models $\mathcal{L}_{f, M}$, cf. Fig. 2. Because the fault model $\mathcal{L}_{f, M}$ accepts the fault language $L_{f, M}$, it is easy to show that this diagnoser isolates f .

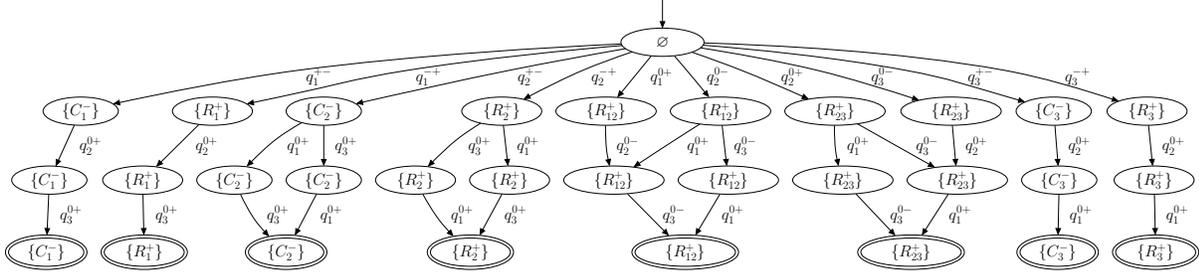


Figure 4: Three-tank system centralized diagnoser for $F = \{C_1^-, C_2^-, C_3^-, R_1^+, R_2^+, R_3^+, R_{12}^+, R_{23}^+\}$ and $M = \{q_1, q_2, q_3\}$

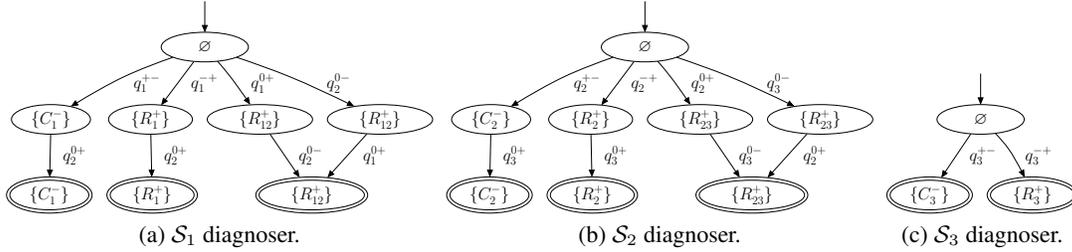


Figure 5: Local diagnosers for the three-tank system for $F_1 = \{C_1^-, R_1^+, R_{12}^+\}$, $M_1 = \{q_1, q_2\}$, $F_2 = \{C_2^-, R_2^+, R_{23}^+\}$, $M_2 = \{q_2, q_3\}$, $F_3 = \{C_3^-, R_3^+\}$ and $M_3 = \{q_3\}$.

A composition operator is then defined that composes two diagnosers, such that if each diagnoser isolates its own set of faults, the composed diagnoser will isolate the combined set of faults. We may then compose the individual diagnosers into a global diagnoser $\mathcal{D}_{F,M}$ that isolates the complete set F . We have shown that the system defined by F and M is diagnosable if and only if the diagnoser constructed in this way uniquely isolates all faults in F (Daigle *et al.*, 2009).

The resulting global diagnoser for the three-tank system described in the earlier sections is given in Fig. 4. It is clear from this figure that the system is diagnosable, as each accepting state has a unique diagnosis. In this case, a unique diagnosis is even known after only a single measurement deviation. The resulting diagnoser may be pruned to reduce diagnoser size by removing states and transitions occurring after a unique diagnosis is known (Daigle, 2008).

5.1 Local Diagnoser Implementation

The design of local diagnosers follows the same procedure as the global diagnoser, i.e., given F_i and M_i for subsystem S_i , we construct \mathcal{D}_{F_i, M_i} . The local diagnosers for the distributed diagnoser design example from the previous section are given in Fig. 5. Note that each local diagnoser except the third needs only two measurements, whereas the global diagnoser needs all three. As n increases, each local diagnoser still needs at most two measurements, whereas the global diagnoser needs all n measurements, significantly increasing its size.

In terms of scalability, the distributed diagnosis scheme clearly improves on the centralized diagnosis approach. In the worst case, the size of a diagnoser

increases factorially with the number of measurements (Daigle *et al.*, 2009). Therefore, the fewer the measurements associated with a diagnoser to achieve local and global diagnosability, the smaller a diagnoser will be. By creating local diagnosers such that each diagnoser uses only a limited number of measurements, each local diagnoser can be significantly smaller than the centralized diagnoser, and the combined size of all local diagnosers can be smaller also.

The distributed diagnosis approach works as follows. Each local diagnoser starts in its initial state. A measurement deviation event is received by all subsystems that include that measurement in their measurement set. If there is a matching event from the current state, a local diagnoser will follow that path to the next state, and remain active. If not, the local diagnoser will block, and its diagnosis result will be \emptyset . The process continues until a local diagnoser reaches an accepting state. At this point, a globally correct diagnosis is known, if each subsystem was designed to be globally diagnosable. If so, no other local diagnoser may reach an accepting state. Therefore, a globally correct diagnosis result is achieved without the use of a centralized coordinator. If the subsystems are not globally diagnosable, then two or more local diagnosers may both reach an accepting state and a coordinator is needed. We may prove this result as follows.

Theorem 1. *Given a distributed diagnoser design where each subsystem S_i is globally diagnosable, then if some $f \in F$ occurs, exactly one \mathcal{D}_i will uniquely isolate it, and all remaining diagnosers will give \emptyset .*

Proof. When f occurs it will produce some trace λ , seen as $\lambda_{f, M_1}, \dots, \lambda_{f, M_n}$, to each \mathcal{D}_i . Since F is par-

tioned (Assumption 1), f belongs to exactly one F_i of D_i . Since S_i is globally diagnosable, no other $f_i \in F_i$ can produce a trace that is a prefix of λ_{f, M_i} , and since D_i is constructed correctly, it must capture λ_{f, M_i} and will uniquely isolate f . Any other D_j will observe the trace λ_{f, M_j} , and, since S_i is globally diagnosable, no fault $f_j \in F_j$ could produce a trace that is a prefix of that trace, so D_j will block, yielding \emptyset . \square

A globally correct diagnosis result may be declared earlier if a local diagnoser has not yet reached an accepting state, but has a unique diagnosis, only if all other local diagnosers have blocked. A globally correct diagnosis result may otherwise only be declared when all measurements for a subsystem have deviated (i.e., an accepting state is reached). These conditions correspond directly to those outlined in (Roychoudhury *et al.*, 2009) in the absence of the event-based framework.

6 RESULTS

As an example to demonstrate online diagnosis in this framework, consider a six-tank system, with R_{23}^+ occurring at time 10.0. The plots of q_2 and q_3 are shown in Fig. 6. At time 10.3 a 0- is detected in q_3 , using the symbol generation mechanism described in (Daigle *et al.*, 2010). Both the local diagnosers for S_2 and S_3 use this measurement and compute this symbol. Partial diagnosers (with some faults omitted) for these subsystems are shown in Fig. 7. The S_2 diagnoser moves to a state with R_{23}^+ as the sole candidate, and the S_3 diagnoser moves to a state with R_{34}^- as the sole candidate. At time 10.4, a 0+ is detected in q_2 . The S_2 diagnoser moves to an accepting state with R_{23}^+ as the sole candidate. The S_3 diagnoser does not use this measurement so takes no action. Because the S_2 diagnoser reached an accepting state, a global diagnosis has been achieved.

For the scalability analysis, we consider n -tank systems where for $i = 1, \dots, n-1$, $F_i = \{C_i^-, C_i^+, R_i^-, R_i^+, R_{i,i+1}^-, R_{i,i+1}^+\}$ and for $i = n$, $F_i = \{C_i^-, C_i^+, R_i^-, R_i^+\}$. The diagnoser design algorithm determines that for $i = 1, \dots, n-1$, $M_i^+ = \{q_i, q_{i+1}\}$, and for $i = n$, $M_i^+ = \{q_{i-1}, q_i\}$, i.e., each subsystem pulls in a measurement from an adjacent subsystem. The local diagnoser for $i = 1, \dots, n-1$ is always 13 states with 14 transitions for the non-pruned version, and 11 states and 10 transitions for the pruned version. For local diagnoser n , both the non-pruned and pruned versions have 7 states and 6 transitions.

The scalability results of the approach as compared to a centralized approach are shown in Table 2. For both non-pruned and pruned diagnosers, we report the number of states, $|S|$, and number of transitions, $|\delta|$. For the local diagnosers, we sum the number of states over each diagnoser, $\sum |S_i|$, and the number of transitions, $\sum |\delta_i|$. The sum of the local diagnoser sizes increase linearly, whereas the size of the centralized diagnoser increases exponentially, demonstrating a clear improvement in scalability. In the case of the pruned diagnosers, the centralized diagnoser size increases linearly as well, although its size is still larger than for the local diagnosers. The linear increase of the pruned

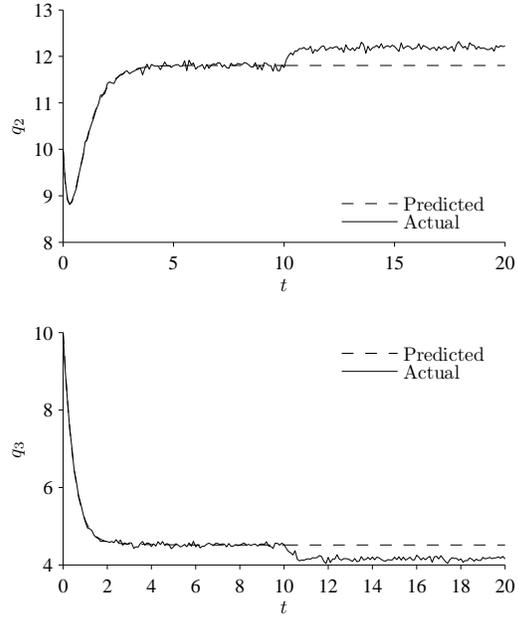


Figure 6: Six-tank predicted and observed flow outputs.

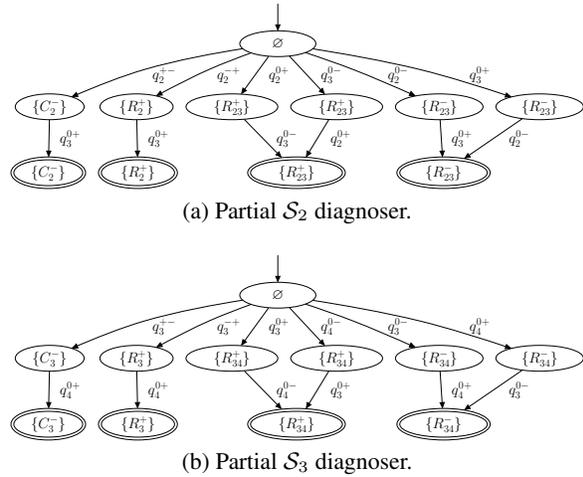


Figure 7: Some partial local diagnosers for the six-tank system

$\mathcal{D}_{F, M}$ is not a general result, but arises here because of the structure imposed by the measurement orderings.

7 CONCLUSIONS

We developed a formal framework for event-based qualitative diagnosis of continuous systems. Global and local diagnosers are automatically derived from fault signatures and relative measurement orderings, which, in turn, may be derived automatically from a system model. This results in a distributed diagnosis framework that eliminates the single point of failure associated with centralized diagnosis frameworks or distributed frameworks that require the use of a centralized coordinator, while the local diagnosers still

Table 2: Scalability Results for the Multi-tank System

Tanks	Not Pruned				Pruned			
	$ S $	$ \delta $	$\Sigma S_i $	$\Sigma \delta_i $	$ S $	$ \delta $	$\Sigma S_i $	$\Sigma \delta_i $
2	19	20	20	20	17	16	18	16
3	69	96	33	34	37	28	29	26
4	113	148	46	48	55	42	40	36
5	205	284	59	62	73	76	51	46
6	335	484	72	76	91	96	62	56
7	579	840	85	90	109	116	73	66
8	845	1264	98	104	127	136	84	76
9	1181	1812	111	118	145	156	95	86
10	1595	2500	124	132	163	176	106	96

obtain globally correct diagnoses. The approach may be naturally applied to systems with clear subsystem boundaries. The distributed approach also scales well with an increase in the number of subsystems, particularly in comparison to a centralized diagnoser.

The event-based framework presented here relates to discrete-event diagnosis methods, e.g., (Sampath *et al.*, 1996; Zad *et al.*, 2003), and also distributed discrete-event diagnosis methods such as (Debouk *et al.*, 2000). Our approach may be viewed as an implementation of Protocol 3 in (Debouk *et al.*, 2000), in which we solve the design problem to achieve the conditions for a coordinator-free approach. In (Ribot *et al.*, 2008), local diagnosers are extended with communicated events and additional sensors. We assume a diagnosable system in which sensor selection has been performed initially. The use of measurement orderings is similar to (Meseguer *et al.*, 2008; Puig *et al.*, 2005), where signatures are derived from analytical redundancy relations, but do not utilize the rich symbol framework for fault signatures used here. In (Bayouhd *et al.*, 2006), a similar approach is applied to hybrid systems, where the events are defined as changes in ARR values due to mode changes.

In future work, we will be extending the approach to multiple faults based on previous work in (Daigle *et al.*, 2007a), and to hybrid systems, based on results presented in (Daigle *et al.*, 2010). We will also investigate alternative distributed design algorithms and design heuristics.

REFERENCES

(Bayouhd *et al.*, 2006) M. Bayouhd, L. Traveé-Massuyès, and X. Olive. Hybrid systems diagnosability by abstracting faulty continuous dynamics. In *Proc. of the 17th International Workshop on Principles of Diagnosis*, pages 9–15, 2006.

(Daigle *et al.*, 2007a) M. Daigle, X. Koutsoukos, and G. Biswas. A qualitative approach to multiple fault isolation in continuous systems. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, pages 293–298, July 2007.

(Daigle *et al.*, 2007b) M. J. Daigle, X. D. Koutsoukos, and G. Biswas. Distributed diagnosis in formations of mobile robots. *IEEE Trans. on Robotics*, 23(2):353–369, April 2007.

(Daigle *et al.*, 2009) M. J. Daigle, X. Koutsoukos, and G. Biswas. A qualitative event-based approach to continuous systems diagnosis. *IEEE Transactions on Control Systems Technology*, 17(4):780–793, July 2009.

(Daigle *et al.*, 2010) M. Daigle, I. Roychoudhury, G. Biswas, X. Koutsoukos, A. Patterson-Hine, , and S. Poll. A comprehensive diagnosis methodology for complex hybrid systems: A case study on spacecraft power distribution systems. *IEEE Transactions of Systems, Man, and Cybernetics, Part A*, 4(5):917–931, September 2010.

(Daigle, 2008) M. Daigle. *A Qualitative Event-based Approach to Fault Diagnosis of Hybrid Systems*. PhD thesis, Vanderbilt University, 2008.

(Debouk *et al.*, 2000) Rami Debouk, S. Lafortune, and Demosthenis Teneketzis. Coordinated decentralized protocols for failure diagnosis of discrete event systems. *Discrete Event Dynamic Systems*, 10(1-2):33–86, 2000.

(Meseguer *et al.*, 2008) J. Meseguer, V. Puig, and T. Escobet. Fault diagnosis using a timed discrete event approach based on interval observers. In *Proc. of the 17th IFAC World Congress*, 2008.

(Mosterman and Biswas, 1999) P. J. Mosterman and G. Biswas. Diagnosis of continuous valued systems in transient operating regions. *IEEE Trans. on Systems, Man and Cybernetics, Part A*, 29(6):554–565, 1999.

(Narasimhan *et al.*, 1998) 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*, pages 94–101, Cape Cod, MA USA, May 1998.

(Puig *et al.*, 2005) V. Puig, F. Schmid, J. Quevedo, and B. Pulido. A new fault diagnosis algorithm that improves the integration of fault detection and isolation. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 3809–3814, December 2005.

(Ribot *et al.*, 2008) P. Ribot, Y. Pencolé, and M. Combacau. Design requirements for the diagnosability of distributed discrete event systems. In *Proceedings of the 19th International Workshop on Principles of Diagnosis*, pages 347–354, September 2008.

(Roychoudhury *et al.*, 2009) I. Roychoudhury, G. Biswas, and X. Koutsoukos. Designing distributed diagnosers for complex continuous systems. *IEEE Transactions on Automation Science and Engineering*, 6(2):277–290, April 2009.

(Sampath *et al.*, 1996) M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D.C. Teneketzis. Failure diagnosis using discrete-event models. *IEEE Transactions on Control Systems Technology*, 4(2):105–124, Mar. 1996.

(Zad *et al.*, 2003) S. Hashtrudi Zad, R.H. Kwong, and W.M. Wonham. Fault diagnosis in discrete-event systems: framework and model reduction. *IEEE Trans. on Automatic Control*, 48(7):1199–1212, July 2003.