Fault Diagnosis of Continuous Systems Using Discrete-Event Methods

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
Email: matthew.j.daigle,xenofon.koutsoukos,gautam.biswas@vanderbilt.edu

***Abstract*— Fault diagnosis is crucial for ensuring the safe operation of complex engineering systems. Although discrete-event diagnosis methods are used extensively, they do not easily apply to parametric fault isolation in systems with complex continuous dynamics. This paper presents a novel discrete-event system diagnosis approach for abrupt parametric faults in continuous systems that is based on a *qualitative* abstraction of measurement deviations from the nominal behavior. Our approach systematically generates a diagnosis model from bond graphs that is used to analyze system diagnosability and derive the discrete-event diagnoser. The proposed approach is applied to an electrical power system diagnostic testbed.**

I. Introduction

Fault diagnosis is crucial for ensuring the safe operation of complex engineering systems. Faults and degradations need to be quickly identified so corrective actions can be taken and catastrophic situations avoided. Discrete-event system (DES) methods are an important framework for event-driven diagnosis in safety-critical systems, since they comprise a well-developed theory that allows for systematic construction of computationally efficient online diagnosers.

Existing DES diagnosers [1], [2] are designed as extended observers that estimate the system state under nominal and faulty conditions. Although these methods have produced many practical diagnosis applications [1], [3]–[5], they are not suitable for systems with complex continuous dynamics. Quantizing the continuous behavior using a finite set of states and events results in large, nondeterministic models [6], [7] that degrade the performance and increase the computational requirements of the diagnosis algorithms. In the presence of faults, these models become increasingly complex, and deriving these models for different fault magnitudes is computationally inefficient.

This paper presents a novel approach to constructing DES diagnosers for isolating single, abrupt faults in continuous systems, based on a *qualitative* abstraction of the measurement deviations from the nominal behavior. The approach is derived from TRANSCEND [8], a model-based methodology for qualitative fault diagnosis in continuous systems. We derive qualitative fault signatures, which capture fault effects on measurements, and relative measurement orderings, which specify the temporal order of measurement deviations, from the system model.

Event-based diagnosis of continuous systems is also investigated in [9], [10], using temporal orders of measurement deviations to help isolate faults. These approaches, based on analytical redundancy relations, are difficult to develop for nonlinear systems and multiplicative faults. The approach does not specify how to obtain the ordering information to construct event-based diagnosers, whereas in our approach, they are derived systematically from the continuous models.

Compared to previous quantization approaches, we propose a more compact qualitative abstraction on the measurement space. Three qualitative states are defined for each measurement: above nominal, at nominal, and below nominal. Measurement deviations indicate fault occurrences and form the event set of our approach. System tracking and fault isolation are separated, so the diagnoser tracks faulty behaviors that manifest as measurement deviations. An observer using a continuous model of the system tracks nominal behavior [8]. The fault signatures and relative measurement orderings are translated into the discrete-event model of the system that represents only faulty behavior. This greatly reduces the burden of the modeling task, and provides a systematic framework for deriving the faulty behavior.

The contribution of the paper centers on: (*i*) a method for systematically constructing a labeled transition system that captures the *fault language*, which, for each fault, describes all possible sequences of measurement deviations, (*ii*) diagnosability analysis of the system and design of an event-based diagnoser, (*iii*) a spectrum of diagnoser implementations that trade off space and time efficiency, and (*iv*) a demonstration of the approach on an electrical power system diagnostic testbed.

II. Qualitative Fault Isolation

TRANSCEND [8] is designed for isolation of single, abrupt, persistent faults in continuous systems. We model the system using bond graphs [11], with faults represented as abrupt parameter value changes in the model [8]. The diagnosis model, the temporal causal graph (TCG), derived from the bond graph models of the system [8], captures effects of faults on measurements.

Abrupt faults generate transients in the dynamic system behavior. Assuming that the system output is continuous and continuously differentiable except at the points of fault occurrence, the transient response after fault occurrence can be approximated by a Taylor series expansion [12]. Measurement transients are described using the magnitude

and the derivative values of the residual signal [8]. This is the basis for establishing a signature for a fault transient. TRANSCEND abstracts these signatures using the qualitative values +, −, and 0, which imply an increase, decrease, or no change from the nominal behavior, respectively.

A fault signature is defined as the qualitative value of zeroth- through kth-order derivative changes on a measurement due to the occurrence of a fault. Only magnitude and slope of a signal can be reliably measured, so we extract two features of the observed deviations, (*i*) whether or not a discontinuity occurred, and (*ii*) the observed first-order change. Therefore, we condense higher order signatures to the magnitude change symbol and the first nonzero derivative change, e.g., 000−+−+ becomes 0−, and +−+−+−+ becomes +−. Thus, a fault signature for measurement m will be an element of the set $\Sigma_m \triangleq \{m^{+-}, m^{-+}, m^{0+}, m^{0-}, m^{+0}, m^{-0}\}$. The superscript indicates the observed deviation. The first symbol represents the immediate direction of abrupt change (a discontinuity) and the second symbol represents the slope. For +0 and −0, the 0 slope symbol implies that the fault will cause a sharp jump but no subsequent change in the slope. This will occur for sensor bias and other discrete faults. We omit signatures of ++ and −− because they represent physically unstable systems.

Definition 1: A *fault signature* for a fault f and measurement m is the qualitative effect of the occurrence of f on m, and is denoted by $\sigma_{f,m} \in \Sigma_{f,m}$, where $\Sigma_{f,m} \subseteq \Sigma_m$. We denote the set of all fault signatures for fault f as Σ_f.

The fault signatures are systematically derived from the TCG using a forward propagation algorithm to predict qualitative effects of faults on measurements [8]. Ambiguities may arise in the qualitative arithmetic, resulting in a signature containing a $*$, which may manifest as either +, −, or 0. So, in general, $\sigma_{f,m}$ may not be unique.

In addition to fault signatures, the TCG captures the temporal order of measurement deviations, defined as *relative measurement orderings* [13], [14]. Relative measurement orderings refer to the intuition that fault effects will manifest in some parts of the system before others. If there are energy storage elements in the path between two sensors in the bond graph, then the path can be characterized by a strictly proper transfer function, and therefore, the energy storage element imposes a delay in the transient responses at the two sensors. If there are no energy storage elements, the relation is algebraic and no delay will be observed.

Definition 2: Consider a fault f and measurements m_i and m_j. If f manifests in m_i before 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 Ω_f.

Throughout the paper we will illustrate the diagnosis method using a circuit example, shown in Fig. 1(a). Fig. 1(b) illustrates its associated bond graph model. We assume that our input voltage, $v(t)$, is constant and positive. The derived TCG is given in Fig. 1(c). Relations between system variables are direct (+1) or inverse (-1) proportionality relations, component parameter values (e.g., R_1), or time-derivative effects (dt). For the circuit, the set of faults is assumed to be $F = \{R_1^+, R_1^-, R_2^+, R_2^-, C_1^+, C_1^-, L_1^+, L_1^-\}$, where the superscript indicates the direction of change of the parameter value. We define the measurement set as the current through L_1, the voltage across C_1, and the current through R_2, or $M = \{f_2, e_5, f_6\}$ in the bond graph model.

TABLE I
FAULT SIGNATURES AND RELATIVE MEASUREMENT ORDERINGS FOR THE CIRCUIT

Fault	f_2	e_5	f_6	Measurement Orderings
R_1^+	0−	0−	0−	$f_2 \prec e_5, f_2 \prec f_6$
R_1^-	0+	0+	0+	$f_2 \prec e_5, f_2 \prec f_6$
R_2^+	0−	0+	−+	$e_5 \prec f_2, f_6 \prec f_2, f_6 \prec e_5$
R_2^-	0+	0−	+−	$e_5 \prec f_2, f_6 \prec f_2, f_6 \prec e_5$
C_1^+	0+	−+	−+	$e_5 \prec f_2, f_6 \prec f_2$
C_1^-	0−	+−	+−	$e_5 \prec f_2, f_6 \prec f_2$
L_1^+	−+	0−	0−	$f_2 \prec e_5, f_2 \prec f_6$
L_1^-	+−	0+	0+	$f_2 \prec e_5, f_2 \prec f_6$

The fault signatures and relative measurement orderings for the circuit system are given in Table I. For example, consider R_2^+. An increase in R_2 will cause an immediate decrease in f_6. Since all subsequent paths from f_6 to any other observed variable in the system contain some edge with a dt specifier (implying an integration), then deviations in these measurements will only be detected after f_6 deviates. The measured variable e_5 will deviate next with a first-order increase. The change is opposite of f_6 because of the -1 specifier in the path, which implies an inverse relationship. The measured variable f_2 will deviate next because of the dt specifier on the path from e_5 to f_2, with a second-order decrease. This will be eventually detected as a first-order change.

III. EVENT-BASED FAULT MODELING

We combine the notion of fault signatures and relative measurement orderings into an event-based framework, where measurement deviations are symbolically abstracted to events. For a specific fault, the combination of all fault signatures and relative measurement orderings yields all the possible ways a fault can manifest. Our event set is then the set of possible measurement deviations. We denote one of these possibilities as a *fault trace*.

Definition 3: A *fault trace* for a fault f, denoted by λ_f, 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 Ω_f.

Consider C_1^+. $\lambda_{C_1^+} = e_5^{-+} f_6^{-+} f_2^{0+}$ is a valid fault trace, but $\lambda_{C_1^+} = f_2^{0+} e_5^{-+} f_6^{-+}$ is not because the measurement deviation sequence does not satisfy $\Omega_{C_1^+}$. Note also that the definition implies that fault traces are of maximal length, i.e., it includes all measurement fault signatures attributed to this fault. We group the set of all fault traces into a *fault language*, which is represented concisely by a *labeled transition system* (LTS).

Definition 4: The *fault language* of a fault f, denoted by L_f, is the set of all fault traces for f.

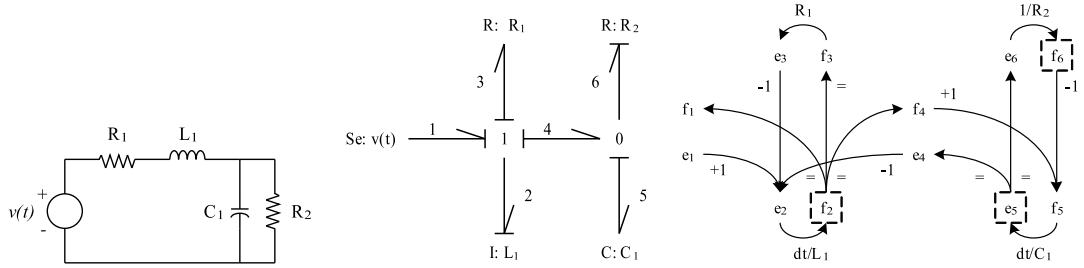

(a) Schematic. (b) Bond graph. (c) Temporal causal graph.

Fig. 1. Circuit example.

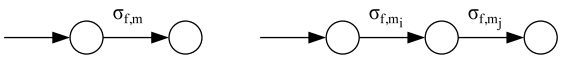

Fig. 2. Fault signature LTS representation (left) and relative measurement ordering LTS representation (right).

Definition 5: A *labeled transition system* is a tuple $\mathcal{L} = (Q, q_o, \Sigma, \delta)$ such that: Q is a set of states, $q_o \in Q$ is an initial state, Σ is a set of labels, and $\delta \subseteq Q \times \Sigma \times Q$ is a transition relation.

To systematically construct the LTS representation of a fault language, called the *fault model*, we can represent each fault signature and each relative measurement ordering as a LTS, and then compose all the information. Each fault signature $\sigma_{f,m}$ can be represented as an LTS, shown to the left of Fig. 2. It consists of only the single event corresponding to the fault signature[1]. Also, each relative measurement ordering, $m_i \prec_f m_j$, with associated signatures σ_{f,m_i} and σ_{f,m_j}, can be represented as an LTS, shown to the right of Fig. 2. It consists of the two associated signatures in the determined ordering.

Lemma 1: The LTS representation of a fault language L_f for fault f, denoted by $\mathcal{L}_f$, is the synchronous product of the individual LTS for all $\sigma_{f,m} \in \Sigma_f$ and all $m_i \prec_f m_j \in \Omega_f$.

Proof: Since the synchronous product must obey all individual ordering constraints and includes all measurement deviation events for the fault, it produces all valid measurement deviation sequences and no others. ■

Definition 6: A fault f_i is *distinguishable* from a fault f_j, denoted by $f_i \nsim f_j$, if $(\forall\ \lambda_{f_i} \in L_{f_i}, \lambda_{f_j} \in L_{f_j})\ (\neg\exists\ \lambda)\ \lambda_{f_i}\lambda = \lambda_{f_j}$.

Two faults are distinguishable if they always eventually produce different event sequences. A fault language represents all possible measurement deviation sequences for a particular fault, so if one fault can exhibit a trace that is a substring of one of some other fault's possible traces, then the faults cannot be distinguished in finite time since the fault traces are, by definition, maximal. If this is never the case, then an event must always occur which distinguishes the faults.

[1]If $\sigma_{f,m}$ is not unique, multiple edges for each possibility are needed going from the first state of the LTS to the final state. This represents the constraint that a measurement's deviation is only observed once.

We wish to obtain system *diagnosability*, based on the notion of distinguishability, with which we can *guarantee* that if measurement deviation events are generated correctly, all faults of interest can be isolated.

Definition 7: A system is *diagnosable* if $(\forall f_i, f_j \in F) f_i \neq f_j \implies f_i \nsim f_j$.

A system is diagnosable if each possible fault trace is consistent with a unique fault. If two faults are distinguishable, then they cannot manifest in the same way. Therefore, if all pairs of faults are distinguishable, then a given fault trace must match only one fault. If this holds for all faults, the system is diagnosable.

IV. Diagnoser Design

We now describe a method to systematically create such an event-based diagnoser. If the system is diagnosable, we can guarantee that the constructed diagnoser can uniquely diagnose all faults. Otherwise, ambiguities will result in the fault isolation results. First, we define formally a *diagnosis* and a *diagnoser* in our framework.

Definition 8 (Diagnosis): A *diagnosis* $d \subseteq F$ is a set of faults consistent with the observations.

Definition 9 (Diagnoser): A *diagnoser* is a tuple $\mathcal{D} = (Q, q_o, \Sigma, \delta, D, Y)$ such that: Q is a set of states, $q_o \in Q$ is an initial state, Σ is a set of labels, $\delta \subseteq Q \times \Sigma \times Q$ is a transition relation, $D \subseteq \mathcal{P}(F)$ is a set of diagnoses, and $Y : Q \to D$ is a diagnosis map.

A diagnoser is an LTS extended by a set of diagnoses and a diagnosis map. Similar to the LTS of a fault, the labels correspond to measurement deviations. Associated with the states are diagnoses, i.e., the set of possible faults for the measurement deviations seen thus far. Like traditional DES diagnosers, diagnoser states provide estimates of the system condition. Assuming the nominal behavior can be accurately tracked with a continuous observer, our diagnoser states provide only the possible sets of faults consistent with the observed sequence of measurement deviations. That is, our diagnoser captures only the faulty system behavior.

A. Diagnosis Algorithm

The diagnoser construction procedure is shown as Algorithm 1. It is described as combining two diagnosers, but can be easily be modified to combine k diagnosers simultaneously. Diagnosers are constructed by incrementally

Algorithm 1 $\mathcal{D} \leftarrow \texttt{CreateDiagnoser}(\mathcal{D}_1, \mathcal{D}_2)$

$Q \leftarrow \varnothing$, $\delta \leftarrow \varnothing$, $D \leftarrow \varnothing$, $\Sigma \leftarrow \Sigma_1 \cup \Sigma_2$
$q_o \leftarrow (q_{o1}, q_{o2})$, $Y(q_o) \leftarrow \varnothing$, $Q_{pend} \leftarrow \{q_o\}$
while $Q_{pend} \neq \varnothing$ **do**
  $(q_1, q_2) \leftarrow \texttt{pop}(Q_{pend})$
  for all $\sigma_m \in \Sigma$ **do**
    if $m \notin H((q_1, q_2))$ **then**
      if $\delta_1(q_1, \sigma_m)$ **and** $\delta_2(q_2, \sigma_m)$ **then**
        $q' \leftarrow (\delta_1(q_1, \sigma_m), \delta_2(q_2, \sigma_m))$
        $h \leftarrow Y(\delta_1(q_1, \sigma_m)) \cup Y(\delta_2(q_2, \sigma_m))$
      else if $\delta_1(q_1, \sigma_m)$ **then**
        $q' \leftarrow (\delta_1(q_1, \sigma_m), q_2)$
        $h \leftarrow Y(\delta_1(q_1, \sigma_m))$
      else if $\delta_2(q_2, \sigma_m)$ **then**
        $q' \leftarrow (q_1, \delta_2(q_2, \sigma_m))$
        $h \leftarrow Y(\delta_2(q_2, \sigma_m))$
      else
        $q' \leftarrow \varnothing$
        $h \leftarrow \varnothing$
      if $q' \neq \varnothing$ **then**
        if $Y((q_1, q_2)) = \varnothing$ **then**
          $d \leftarrow h$
        else
          $d \leftarrow Y((q_1, q_2)) \cap h$
        if $d \neq \varnothing$ **then**
          $Q \leftarrow Q \cup \{q'\}$
          $H(q') \leftarrow H((q_1, q_2)) \cup \{m\}$
          $\delta((q_1, q_2), \sigma_m) \leftarrow q'$
          $D \leftarrow D \cup \{d\}$
          $Y(q') \leftarrow d$
          if $q' \notin Q_{pend}$ **then**
            $\texttt{push}(Q_{pend}, q')$

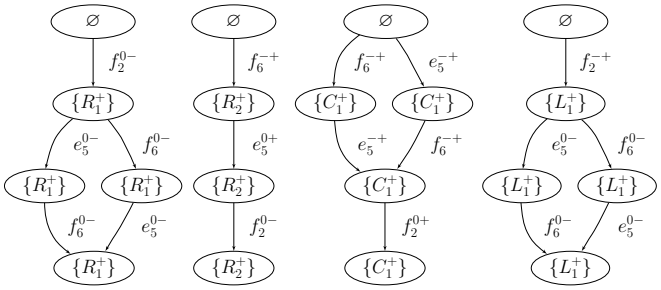

Fig. 3. Diagnosers for the individual faults of the circuit. The diagnosers for decreases in the parameter values are the same except for a reversal in the signs.

composing subdiagnosers, i.e., a diagnoser for a set of faults F_i is composed with a diagnoser for a set of faults F_j to create a new diagnoser for $F_i \cup F_j$. Initially, we begin with diagnosers for singleton fault sets. These are constructed using the individual fault models. For a single fault f, we augment $\mathcal{L}_f$ to form $\mathcal{D}_f$ by constructing the diagnosis map as mapping every state except the initial state to $\{f\}$. The initial state is mapped to the empty diagnosis $\varnothing$, because until a measurement deviation is observed, we assume the system is operating nominally. The diagnosers corresponding to the individual faults of the circuit are shown in Fig. 3.

The construction algorithm operates by tracing paths in the two given diagnosers. If the same event label is available in both current states, then we advance in both LTS, i.e., $(q_1, q_2) \xrightarrow{\sigma} (\delta(q_1, \sigma), \delta(q_2, \sigma))$. Otherwise, we advance in only one, e.g., if σ can only be taken from q_1, then $(q_1, q_2) \xrightarrow{\sigma} (\delta(q_1, \sigma), q_2)$. However, if the measurement associated with σ has already deviated along the current path (tracked using H), $\delta((q_1, q_2), \sigma)$ is set to $\varnothing$, because measurement deviations are only detected once per measurement. This also occurs if the computed diagnosis for the new state, d, is empty, because this means the current sequence of measurement deviations is inconsistent with the single fault assumption.

The diagnosis for the new state is formed by composing the current diagnosis with the hypothesis set. The hypothesis set, h, is the set of faults consistent with the current event. It is formed as the union of the diagnoses of the diagnoser states advanced to via the event σ. The new diagnosis for the composed diagnoser state is constructed as the intersection of the current diagnosis and the hypothesis set. For example, if $\{f_i, f_j\}$ is the current diagnosis and the hypothesis set is $\{f_i\}$ then the new diagnosis is $\{f_i\}$, which means that only f_i is consistent with the current event sequence.

The final composed diagnoser for the circuit is illustrated in Fig. 4. For example, consider the fault trace $f_6^{-+} e_5^{0+} f_2^{0-}$. For f_6^{-+} occurring as the first measurement deviation, only C_1^+ or R_2^+ could have occurred, given the known fault signatures and relative measurement orderings. Therefore, the new diagnosis is $\{C_1^+, R_2^+\}$. For e_5^{0+} occurring next, of our current faults, only R_2^+ is consistent, therefore our new diagnosis is the intersection of $\{C_1^+, R_2^+\}$ and $\{R_2^+\}$, which is $\{R_2^+\}$. At this point we obtain a unique fault. The only possible measurement deviation from here is f_2^{0-} which must still be consistent with $\{R_2^+\}$.

Theorem 1: The diagnoser constructed by Algorithm 1 for fault sets F_1 and F_2 represents all valid single fault traces for the faults in F_1 and F_2 and associates correct diagnoses with the states.

Proof: By definition, the diagnoser for a single fault f is correct because it represents L_f, so represents all possible fault traces of f, and every state (except the initial state) of $\mathcal{L}_f$ is consistent with f occurring. Assume that diagnosers $\mathcal{D}_1$ and $\mathcal{D}_2$ are correct. Then they represent all possible fault traces for fault sets F_1 and F_2, respectively. At the initial state, if an event σ happens which can only happen in one of the diagnosers, $\mathcal{D}_i$, then the diagnosis is $Y_i(\delta(q_{oi}, \sigma))$, because it must be consistent with faults in F_i that are consistent with σ. If σ can occur in both diagnosers, then the diagnosis is $Y_1(\delta_1(q_{o1}, \sigma)) \cup Y_2(\delta_2(q_{o2}, \sigma))$ because either a fault in F_1 occurred or a fault in F_2 occurred, and the diagnosis must be consistent with any fault in $F_1 \cup F_2$ consistent with σ. Assume that for a given $(q_1, q_2) \neq q_o$ the diagnoses are correct for the event sequences leading up to (q_1, q_2). Then if an event σ happens which can only happen in one of the diagnosers, $\mathcal{D}_i$, then the diagnosis is $Y((q_1, q_2)) \cap Y_i(\delta_i(q_i, \sigma))$, because it must be consistent with the previous diagnosis faults in F_i consistent with σ. If σ can occur in both diagnosers, then the diagnosis is $Y((q_1, q_2)) \cap (Y_1(\delta_1(q_1, \sigma)) \cup Y_2(\delta_2(q_2, \sigma)))$ because it must be consistent with the previous diagnosis and faults in $F_1 \cup F_2$

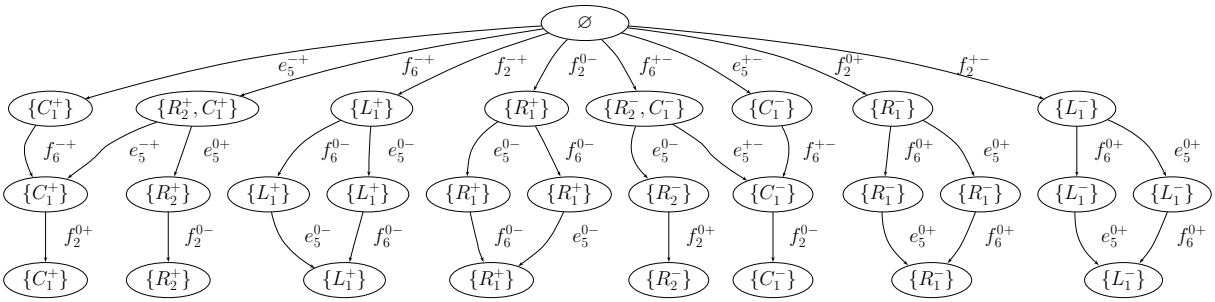

Fig. 4. Single fault diagnoser for the circuit.

consistent with σ. Therefore, for any state q, $\delta(q, \sigma)$ has a correct diagnosis. So, for any two diagnosers, the resulting diagnoser is correct. ■

The diagnoser for the circuit example, shown in Fig. 4, illustrates certain properties of our approach. Since all the leaves have diagnoses with a unique single fault, then the system is diagnosable. All possible sequences of measurement deviations corresponding to single faults occurrences are captured in the diagnoser, and lead to unique diagnoses, therefore, the system is diagnosable. We can also see that a unique diagnosis is obtained after only two of the three measurements deviate, so one measurement is redundant for single fault diagnosis of the selected faults.

B. Online Diagnoser Implementation

This event-based diagnosis framework leads to three different implementations of the online diagnosis approach that trade off space and time complexity.

1) Implementation as a global LTS: Time complexity is in favor of the precomputed global diagnoser (Fig. 4). It needs only to wait for measurement deviations to occur, transition to the next state, and output the current diagnosis associated with the state. Using appropriate data structures, these operations can be achieved in constant time.

The complete diagnoser has, in the worst case, $O(|M|!)$ possible fault traces, and thus $O(|M|!)$ states. Therefore, this approach is not space-efficient. If many temporal orderings exist, then the number of possible fault traces reduces significantly, and the global diagnoser approach may be feasible.

2) Implementation as an LTS for each fault: In this approach, we only precompute the individual fault diagnosers (Fig. 3). Each fault has $O(|M|!)$ possible fault traces, but if there are many temporal orderings, this may also be reduced for many faults.

For online diagnosis, each diagnoser is traced simultaneously. When a diagnoser becomes blocked, i.e., there is no available event to to take from the current state, then it is no longer tracked, because it is no longer consistent with the observed measurement deviations. The candidate set is formed by taking the union of the faults in the current states of each active diagnoser, i.e., those faults that are still consistent with the observed measurement deviations. This operation has time complexity $O(|F|)$.

3) Implementation without explicit event fault models: If faults have few possible traces, then both of the above approaches should be both space-efficient and time-efficient. If few orderings are available, then the diagnosers approach size $O(|M|!)$, therefore, these approaches may not be feasible given the space requirements of the system. For diagnosis without using LTS-based diagnosers, we store only the fault signatures and relative measurement orderings for each fault (Table I), requiring $O(|F||M|^2)$ space.

Given a current diagnosis of d_{i-1} and an event σ_i occurring, we can check which faults are consistent with σ_i. The hypothesis set h_i consists of those faults. If $i = 1$, then the new diagnosis d_i is simply h_i. Otherwise, the new diagnosis must be consistent with d_{i-1} and with the new information, i.e., $d_i = d_{i-1} \cap h_i$. Therefore, given d_{i-1}, the new diagnosis can be computed simply as the subset of faults in d_{i-1} consistent with σ_i. This corresponds to only constructing the *path* of the global diagnoser relating to the particular fault trace we are observing.

For online diagnosis, we form the hypothesis set corresponding to the current measurement deviation by looking through the fault signatures and measurement orderings, thus taking $O(|F||M|^2)$ time. We then compute the new diagnosis, which is a function of the size of the current diagnosis and the current hypothesis set. In the worst case the hypothesis set consists of all faults, so it is $|F|$ in size. A diagnosis can be as large as $|F|$ also. The intersection of the diagnosis and hypothesis set then takes at worst $O(|F|)$ time. In practice, this time complexity is reduced because as measurements deviate, less faults are being considered.

V. Case Study

We demonstrate the proposed diagnosis framework with experiments from the Advanced Diagnostics and Prognostics Testbed (ADAPT) [15] deployed at NASA Ames. The testbed is functionally representative of a spacecraft's electrical power system, consisting of power generation, storage, and distribution subsystems. For our diagnosis experiments, we consider a subset of ADAPT that involves a battery discharging to two parallel DC loads as shown in Fig. 5.

A. Modeling Faults

The battery model describes an electric circuit equivalent. The capacitance of the battery is modeled using a large

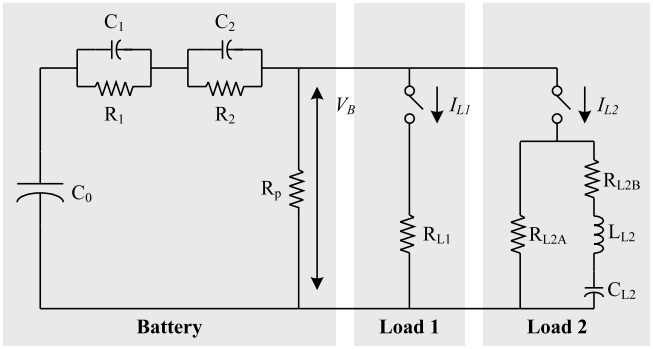

Fig. 5. Schematic diagram for the battery system.

TABLE II
FAULT SIGNATURES AND RELATIVE MEASUREMENT ORDERINGS FOR THE BATTERY SYSTEM

Fault	V_B	I_{L1}	I_{L2}	Measurement Orderings
C_0^-	$+\star$	$+\star$	$+\star$	$\varnothing$
R_1^+	$0-$	$0-$	$0-$	$\varnothing$
R_{L1}^+	$0\star$	$-+$	$0\star$	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
R_{L1}^-	$0\star$	$+-$	$0\star$	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
R_{L2A}^+	$0\star$	$0\star$	$+-$	$I_{L2} \prec V_B, I_{L2} \prec I_{L1}$
R_{L2A}^-	$0\star$	$0\star$	$+-$	$I_{L2} \prec V_B, I_{L2} \prec I_{L1}$
V_B^+	$+0$	00	00	$V_B \prec I_{L1}, V_B \prec I_{L2}$
V_B^-	-0	00	00	$V_B \prec I_{L1}, V_B \prec I_{L2}$
I_{L1}^+	00	$+0$	00	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
I_{L1}^-	00	-0	00	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
I_{L2}^+	00	00	$+0$	$I_{L1} \prec V_B, I_{L2} \prec I_{L1}$
I_{L2}^-	00	00	-0	$I_{L1} \prec V_B, I_{L2} \prec I_{L1}$

capacitance, C_0. Other parameters model the nonlinear behaviors and dissipative effects (see [16] for details). Battery faults include loss of charge represented by a capacitance decrease, C_0^-, and internal resistance increase, R_1^+. In the loads, faults affect the resistance values R_{L1} and R_{L2A} which can increase or decrease. In the sensors, we consider bias faults which cause abrupt changes in the measured values. Sensor faults are labeled by the measured quantity they represent, e.g. V_B^+ represents a bias fault in the battery voltage sensor.

The fault signatures and relative measurement orderings for nominal discharge operation are given in Table II. Available measurements include the battery voltage $V_B(t)$, and the load currents, $I_{L1}(t)$ and $I_{L2}(t)$. The fault models for the selected experiments are given in Fig. 6. The nonlinearities in the battery introduce ambiguity in the qualitative signatures, denoted by the $\star$ symbol. A signature of $+\star$ may manifest as $+-$ or $+0$, and a signature of $0\star$ may manifest as $0+$ or $0-$. In the fault models, all possible effects must be included. Also note that since the sensors do not feed back into the system, sensor faults affect only the measurement provided by the sensor. The other measurements are not affected, and the corresponding signature is denoted by 00, indicating no change in the measurement from expected behavior.

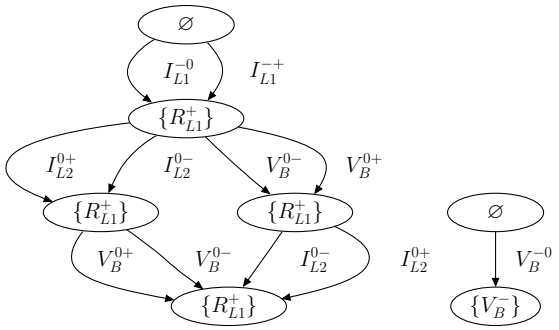

Fig. 6. Fault models for the battery system.

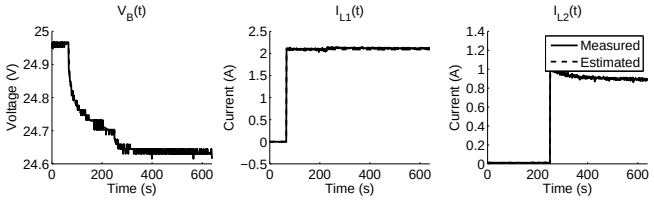

Fig. 7. Nominal system operation

B. Diagnosis Results

We investigate the diagnosis of a load fault injected in the hardware, and a sensor fault injected over the real data. A sampling rate of 2 Hz is used in all the experiments. The nominal behavior of the system in shown in Fig. 7. The system starts with both loads disconnected. First Load 1 is taken online, followed by Load 2. In the following scenarios, the fault is injected in this system configuration.

A 50% decrease in the Load 1 resistance, R_{L1}^- is injected at 417 s. The measured and estimated outputs are shown in Fig. 8. The decrease in resistance increases the current draw abruptly, and is detected at 417 s. The symbol generator reports $+0$. The first order change due to the fault is compensated for by the battery and not detected. At this point, the diagnosis is $\{R_{L1}^-, I_{L1}^+\}$. At 433 s, the fault detector observes a decrease in the battery voltage. Since I_{L1}^+ cannot produce a deviation in this measurement, it is dropped, and R_{L1}^- is isolated as the true candidate.

A positive bias of 0.5 V is injected into the voltage sensor at 500 s. The measured and estimated outputs are shown in Fig. 9. The fault detector reports an abrupt increase in battery voltage at 500 s. The symbol generator reports $+0$. The diagnosis is $\{V_B^+, C_0^-\}$. Because no further measurements deviate, C_0^- cannot be eliminated.

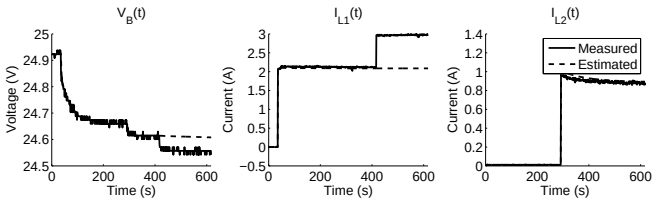

Fig. 8. R_{L1}^- fault with magnitude of 50%.

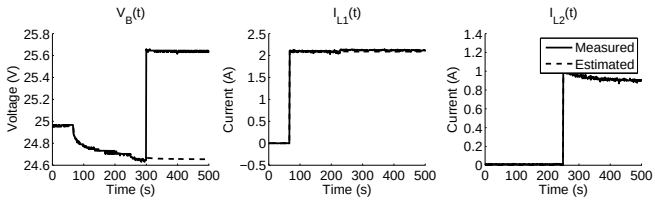

Fig. 9. V_B^+ fault with bias of 0.5.

Our diagnosis approach depends critically on the existence of a robust fault detection scheme. Currently, fault detection is achieved using an extended Kalman filter and a statistical test for evaluating the residuals [17]. The threshold parameters for the residual evaluation are manually tuned to detect substantial abrupt faults while lowering the sensitivity. Although the current results are promising, evaluating the performance of the diagnosis approach quantified by the false alarm rate is a significant issue and we plan to tackle this problem by performing additional diagnosis experiments in the experimental testbed.

VI. Conclusions

We have presented a DES modeling and diagnosis methodology applied to parametric faults in continuous-time systems. The main issue in applying DES approaches is creating a system model that captures all relevant system behavior. Quantization-based abstractions create large, nondeterministic models. On the other hand, our qualitative abstraction approach systematically creates concise discrete-event models of faulty system behavior given a continuous model of the system. In traditional DES methods, such models are often created by hand. The qualitative abstraction enables a diagnosis approach that is easily applicable to continuous systems. Diagnosis results from the electrical power systems domain demonstrated the promise of the approach.

Acknowledgments

This work was supported in part by NSF grant CNS-0615214, NASA USRA grant 08020-013, and NASA NRA grant NNX07AD12A.

References

[1] 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, Mar. 1996.

[2] S. H. Zad, R. Kwong, and W. Wonham, "Fault diagnosis in discrete-event systems: framework and model reduction," *IEEE Transactions on Automatic Control*, vol. 48, no. 7, pp. 1199–1212, July 2003.

[3] 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-02)*, Semmering, Austria, May 2002, pp. 179–188.

[4] A. Benveniste, E. Fabre, S. Haar, and C. Jard, "Diagnosis of asynchronous discrete-event systems: a net unfolding approach," *IEEE Transactions on Automatic Control*, vol. 48, no. 5, pp. 714–727, May 2003.

[5] V. Chandra, Z. Huang, and R. Kumar, "Automated control synthesis for an assembly line using discrete event system control theory," *IEEE Trans. on Systems, Man and Cybernetics, Part C*, vol. 33, no. 2, pp. 284–289, May 2003.

[6] J. Lunze, "Diagnosis of quantized systems based on a timed discrete-event model," *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, vol. 30, no. 3, pp. 322–335, 2000.

[7] X. Koutsoukos, P. Antsaklis, J. Stiver, and M. Lemmon, "Supervisory control of hybrid systems," *Proceedings of IEEE*, vol. 88, no. 7, pp. 1026–1049, July 2000.

[8] 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.

[9] V. Puig, J. Quevedo, T. Escobet, and B. Pulido, "On the integration of fault detection and isolation in model-based fault diagnosis," in *Proceedings of the 16th International Workshop on Principles of Diagnosis (DX-05)*, 2005, pp. 227–232.

[10] 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*, Dec. 2005, pp. 3809–3814.

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

[12] 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.

[13] M. Daigle, X. Koutsoukos, and G. Biswas, "Relative measurement orderings in diagnosis of distributed physical systems," in *43rd Annual Allerton Conference on Communication, Control, and Computing*, Sept. 2005, pp. 1707–1716.

[14] M. J. Daigle, X. D. Koutsoukos, and G. Biswas, "Distributed diagnosis in formations of mobile robots," *IEEE Transactions on Robotics*, vol. 23, no. 2, pp. 353–369, Apr. 2007.

[15] S. Poll, A. Patterson-Hine, J. Camisa, D. Garcia, D. Hall, C. Lee, O. Mengshoel, C. Neukom, D. Nishikawa, J. Ossenfort, A. Sweet, S. Yentus, I. Roychoudhury, M. Daigle, G. Biswas, and X. Koutsoukos, "Advanced diagnostics and prognostics testbed," in *Proceedings of the 18th International Workshop on Principles of Diagnosis*, May 2007, pp. 178–185.

[16] M. Daigle, I. Roychoudhury, G. Biswas, and X. Koutsoukos, "Efficient simulation of component-based hybrid models represented as hybrid bond graphs," Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, USA, Tech. Rep. ISIS-06-712, 2006.

[17] G. Biswas, G. Simon, N. Mahadevan, S. Narasimhan, J. Ramirez, and G. Karsai, "A robust method for hybrid diagnosis of complex systems," in *Proceedings of the 5th Symposium on Fault Detection, Supervision and Safety for Technical Processes*, June 2003, pp. 1125–1131.