Improving Diagnosability of Hybrid Systems through Active Diagnosis ⋆

Matthew Daigle * **Xenofon Koutsoukos** ** **Gautam Biswas** **

* *University of California, Santa Cruz, NASA Ames Research Center, Moffett Field, CA 94035 (e-mail: matthew.j.daigle@nasa.gov).*
** *Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN 37235 (e-mail: {xenofon.koutsoukos,gautam.biswas}@vanderbilt.edu)*

Abstract: Fault diagnosis is key to ensuring system safety through fault-adaptive control. This task is difficult in hybrid systems with combined continuous and discrete behaviors because mode changes make diagnosability hard to achieve. Including additional sensors can improve diagnosability, but that is not always feasible. An alternative strategy is active diagnosis, where we improve the diagnosis result by executing or blocking controllable events. We present a qualitative, event-based approach to active diagnosis of hybrid systems, where we automatically synthesize event-based diagnosers for hybrid systems that can determine if the system is diagnosable through passive or active diagnosis. We apply our active diagnosis scheme to a real-world electrical power distribution system.

Keywords: active diagnosis; hybrid systems; discrete-event systems.

1. INTRODUCTION

Diagnosability relates to the ability of a diagnostic system to obtain unique diagnosis results, and, therefore, affects many aspects of the design of diagnostic systems. In hybrid systems, where the system behavior contains continuous and discrete behaviors, achieving a diagnosable system becomes more difficult. Mode changes complicate the diagnosis task, because the effects of faults may change or become masked from one mode to another. A system may be diagnosable in a single mode, but over all modes, diagnosability may be lost when mode changes occur during online fault isolation (Daigle et al., 2008). Diagnosability of hybrid systems can be improved by including additional sensors, however, this is not always feasible due to constraints in system design, cost, etc. An alternative approach to improving diagnosability is through *active diagnosis*, that is, modifying the system behavior through control actions in order to improve the diagnosis result.

Diagnosability has been well-studied in discrete-event systems, and an active diagnosis approach is developed in (Sampath et al., 1998), where controllable events that cause a loss of diagnosability are forbidden. Active diagnosis has been studied for continuous systems in (Niemann, 2005), where auxiliary inputs are used to improve fault isolation. An approach to active diagnosis of hybrid systems using analytical redundancy relations is presented in (Bayoudh et al., 2008), where active diagnosis is formulated as a conditional planning problem to find control actions that drive the system to diagnosable regions.

In previous work, we have developed an approach to qualitative fault isolation of hybrid systems that can handle both parametric (Narasimhan and Biswas, 2007) and discrete faults (Daigle, 2008). In this paper, we develop an event-based framework for active diagnosis of hybrid systems. Faults are viewed as unobservable events, and measurement deviations and controlled mode change commands constitute observable events. We design event-based diagnosers for the system, which are used to verify passive and active diagnosability properties of the system. Active diagnosis is enabled by selectively blocking or executing controllable events to avoid ambiguities in fault isolation and achieve faster, more precise diagnosis results. We apply our active diagnosis scheme to a subset of the Advanced Diagnostics and Prognostics Testbed (ADAPT) at NASA Ames Research Center, which is a complex electrical power distribution system.

2. BACKGROUND

We consider the problem of active diagnosis in hybrid systems. We represent faults as unobservable events that either (i) produce an unexpected step change in a system parameter value (parametric faults), or (ii), produce an unexpected change in system mode (discrete faults). In this paper, we assume single faults and controlled mode changes only. Autonomous modes changes and multiple faults can be incorporated in a more complex framework combining the techniques presented in (Narasimhan and Biswas, 2007; Daigle, 2008). This paper does not consider these extensions to focus on the fundamental problem of active diagnosis and diagnosability for hybrid systems.

2.1 Qualitative Fault Isolation

Our goal is quick detection and isolation of faults, therefore, we isolate faults based on the transients they produce

⋆ This work was supported in part by NSF grant CNS-0615214 and NASA NRA grant NNX07AD12A.

in the system measurements. We model the system and use a hybrid observer to estimate the state and outputs. Statistically significant differences between actual and estimated outputs signal that a fault has been detected (Biswas et al., 2003). The deviations in the measurements are then abstracted into symbolic form and matched against predicted values for fault isolation.

Measurement deviations are abstracted using qualitative `+`, `-`, and `0` values to form *fault signatures*, which represent the immediate change in magnitude and the first nonzero derivative change (Mosterman and Biswas, 1999). They also represent what is termed discrete change behavior, which describes whether the signal went from a nonzero to a zero value (`Z`), a zero to a nonzero value (`N`), or had no zero/nonzero value changes (`X`) (Daigle, 2008). In the following, we denote the set of modes as $Q = \{q_1, q_2, \ldots, q_r\}$, the set of faults as $F = \{f_1, f_2, \ldots, f_n\}$, and the set of measurement points as $M = \{m_1, m_2, \ldots, m_p\}$.

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

In addition to fault signatures, we also capture the temporal order of measurement deviations, termed *relative measurement orderings* (Daigle et al., 2007), which refer to 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., 2007).

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

The fault signatures and measurement orderings can be automatically computed from a temporal causal graph representation derived from the system model (Mosterman and Biswas, 1999; Daigle, 2008).

The fault isolation task consists of matching observed event sequences, i.e., measurement deviations ($\sigma \in \Sigma_M$) and controlled mode changes ($\sigma \in \Sigma_Q$), to predicted sequences associated with fault candidates. We define a candidate as a hypothesized fault and a hypothesized system mode.

Definition 3. (Candidate). A *candidate* c is defined as $c = (f_i, q_i)$, where $f_i \in F$ is a hypothesized fault, and $q_i \in Q$ is a hypothesized current mode. The set of all candidates is denoted as C.

We wish to find candidates that are *consistent* with the sequence of observed events. A *diagnosis* is a collection of candidates that are consistent with the observations after the time of fault occurrence, t_f.

Definition 4. (Diagnosis). At time $t \geq t_f$, a *diagnosis* $d \subseteq C$ is a set of candidates consistent with the observations made on the system during the interval $[t_f, t]$.

Fault isolation is performed incrementally, as new events are received. At each new event, the current diagnosis is refined by eliminating candidates that are inconsistent with the new event, given the previous sequence of events.

2.2 Event-based Fault Modeling

In order to characterize diagnosability in our framework, we first need to define what it means for a candidate to be consistent with a sequence of observable events. Within a given mode, a fault will produce a number of possible sequences of measurement deviations. We define the set of these traces as a *fault language*.

Definition 5. (Fault Language). The *fault language* of a fault $f \in F$ with measurements M in mode q, denoted by $L_{f,M,q}$, is the set of all traces that includes, for every $m \in M$ that will deviate due to f in q, a fault signature $\sigma_{f,m,q}$, such that the sequence of signatures satisfies $\Omega_{f,M,q}$.

When a controlled mode change occurs, the system model is updated, and a new nominal reference for symbol generation is computed. The measurement deviations produced in the new mode must match what is predicted for the fault in that mode, ignoring already deviated measurements. This defines consistency for a fault candidate. The set of consistent candidates depends also on the initial mode and the nominal sequence of mode changes leading to the expected mode at the point of fault detection. Given this, we can now define a *candidate trace*. In the following, we denote the system mode transition function by μ.

Definition 6. (Candidate Trace). An event trace $\lambda = \sigma$ is a *candidate trace* for $c = (f_i, q_i)$ and initial mode q_0, if σ is a prefix of $\lambda' \in L_{f_i,M,q_i}$ where $q_i = \mu(f_i, q_0)$. An event trace $\lambda = \lambda_i \sigma_{i+1}$ is a *candidate trace* for $c = (f_i, q_{i+1})$ and initial mode of fault occurrence q_0, if λ_i is a candidate trace for (f_i, q_i), and if $\sigma_{i+1} \in \Sigma_Q$ then $\mu(\sigma_{i+1}, q_i) = q_{i+1}$, or if $\sigma_{i+1} \in \Sigma_M$ then $q_i = q_{i+1}$ and σ_{i+1} is a prefix of $\lambda' \in L_{f_i,M-M_i,q_{i+1}}$, where M_i is the set of deviated measurements up to the ith event. A candidate trace for c with initial mode q_0 is denoted as λ_{c,q_0}.

In other words, given a candidate trace, an extension of that trace by a measurement deviation event will also be a candidate trace for the same candidate, if the deviation is still consistent with the candidate (i.e., consistent with the fault language in the new mode). An extension of the trace by a mode change event, however, will be a candidate trace for a different candidate, namely, the one defined by changing the mode of the old candidate to the new mode.

Clearly, there may be an infinite number of candidate traces because controlled mode changes may keep occurring indefinitely. However, we are only concerned with *maximal* traces, i.e., those for which all measurements that will deviate in the current mode have deviated.

Definition 7. (Maximal Candidate Trace). A candidate trace λ_{c,q_0} for $c = (f_i, q_i)$ is *maximal* if $L_{f,M-M_i,q_i} = \varnothing$, where M_i is the set of deviated measurements for λ_{c,q_0}.

Now, we can define the language of a candidate c with respect to an initial mode of fault occurrence q_0 as the set of maximal candidate traces for c starting in q_0.

Definition 8. (Candidate Language). The *candidate language* for candidate c, measurements M, and initial mode of fault occurrence q_0, denoted as L_{c,M,q_0}, is the set of all maximal candidate traces λ_{c,q_0}.

The candidate language consists of all consistent maximal traces for the candidate. A maximal trace is consistent

with a candidate if the mode of the candidate can be reached via the sequence of controlled mode changes in the trace, and the measurement deviations within the trace match the fault in the intermediate modes. In this framework diagnosability reduces to ensuring that for any two candidates, it is not possible that a maximal candidate trace for one candidate is a prefix of a maximal trace for the other candidate (Daigle, 2008; Daigle et al., 2008).

3. DIAGNOSERS

We construct from our fault models an event-based diagnoser, which maps observed sequences to consistent fault hypotheses. It is constructed as an extended finite automaton, and can be used for offline diagnosability analysis.

We wish to construct a diagnoser for a given system. In our framework, a *system* can be defined as follows.

Definition 9. (System). A *system* $\mathcal{S}$ is defined as (F, M, Q, $L_{F,M,Q}$), where $F = \{f_1, f_2, \ldots, f_n\}$ is a set of faults, $M = \{m_1, m_2, \ldots, m_p\}$ is a set of measurements, $Q = \{q_1, q_2, \ldots, q_r\}$ is a set of modes, and $L_{F,M,Q}$ is the set of fault languages for each fault in each mode, i.e., $L_{F,M,Q} = \{L_{f,M,q} : f \in F, q \in Q\}$.

We formally define a *diagnoser* as follows.

Definition 10. (Diagnoser). A *diagnoser* for a fault set F, measurements M, and modes Q, is defined as $\mathcal{D}_{F,M,Q} = (S, s_0, \Sigma, \delta, A, D, Y)$, where S is a set of states, $s_0 \in S$ is the 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, $D \subseteq 2^C$ 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. The initial state corresponds to the fault-free initial mode of the system. A diagnoser takes events as inputs, which correspond to measurement deviations $\sigma \in \Sigma_M$ and controlled mode changes $\sigma \in \Sigma_Q$. From the current state, an event causes a transition to a new state. The diagnosis for that new state represents the set of candidates that are consistent with the sequence of events seen up to the current point in time.

A diagnoser for a system $\mathcal{S}$ that captures all valid candidate traces for $\mathcal{S}$, denoted as $\mathcal{D}_{\mathcal{S}}$, can be constructed automatically from the fault models, as described in (Daigle, 2008; Daigle et al., 2008). This procedure is not described here, but intuitively, it is constructed as a finite automaton that captures all the candidate languages. A state that corresponds to a trace λ includes in its diagnosis all candidates for which λ is a candidate trace. Maximal candidate traces correspond to accepting states in the automaton.

We would like to achieve unique isolation for all candidates. To achieve this, we require that the diagnoser accepts all possible valid traces for the candidate and that the corresponding accepting states uniquely determine c.

Definition 11. (Unique Isolation). A diagnoser $\mathcal{D}_{F,M,Q}$ *uniquely isolates* a candidate c if it accepts all $\lambda \in L_{c,M,q_0}$, and for each $s \in A$ that accepts some $\lambda \in L_{c,M,q_0}$, $\{c\} = Y(s)$.

We can relate unique isolation to *passive diagnosability*.

Definition 12. (Passive Diagnosability). A system $\mathcal{S}$ is passively diagnosable if and only if, after a fault is detected, a singleton diagnosis can always be obtained in finite time.

We have shown previously that the system $\mathcal{S}$ is diagnosable, i.e., no candidate can produce a maximal trace that is a prefix of a candidate trace for some other candidate, if and only if the diagnoser $\mathcal{D}_{\mathcal{S}}$ uniquely isolates all candidates (Daigle et al., 2008). This can be determined from the following result, restated from (Daigle et al., 2008).

Theorem 13. A system $\mathcal{S}$ with diagnoser $\mathcal{D}_{\mathcal{S}} = (S, I, \Sigma, \delta, A, D, Y)$ is diagnosable if and only if for all $s \in A$, $Y(s)$ is a singleton.

Therefore, we can inspect $\mathcal{D}_{\mathcal{S}}$ to determine if the system is diagnosable, simply by examining each accepting state, and ensuring that only a single candidate is included in its diagnosis. Under this notion of diagnosability, a unique diagnosis result will always be obtained *independent of* which mode changes occur during fault isolation. However, such a condition may not be achievable in a hybrid system. The effects of faults may change from mode to mode. In a new mode, it is possible that no more measurement deviations will occur to resolve any ambiguity, or the measurement deviations that will occur are not enough to distinguish the faults. If the diagnoser has no control over which controlled mode change events are issued, we cannot, in general, make any restrictions about when a mode change event will be issued. Thus, diagnosability in this sense is conservative and corresponds to passive diagnosis. It may be possible, however, to avoid ambiguous diagnosis results if certain mode changes are blocked or executed during fault isolation. Therefore, diagnosability can be improved by adopting an active diagnosis methodology.

4. ACTIVE DIAGNOSIS

In general, active diagnosis requires both prevention and execution of actions. In (Sampath et al., 1998), actions which may drive the system into undiagnosable regions are blocked, so that eventually, the actions that are executed will lead the system to a state where the fault is isolated. The system will only be truly diagnosable if it can be guaranteed that such actions will eventually occur. In (Bayoudh et al., 2008), actions are executed to drive the system toward diagnosable regions. However, external actions, e.g., those from a system operator, may inadvertently drive the system into an undiagnosable region. Therefore, in our approach to active diagnosis, we adopt a combination of the above perspectives, where the diagnoser is used to determine which controllable events (in our case, controlled mode changes) should be blocked, and which should be executed, in order to guarantee unique fault isolation results.

4.1 Active Diagnosability

First, we would like to determine if, through active diagnosis, a unique fault isolation result can always be achieved. If so, we say the system is *actively diagnosable*.

Definition 14. (Active Diagnosability). A system $\mathcal{S}$ is actively diagnosable if and only if, after a fault is detected, a singleton diagnosis can always be obtained by blocking and/or executing controllable events during fault isolation.

If unique results are ensured without modifying control actions, then the system is also passively diagnosable. Passive diagnosability is, therefore, a sufficient condition for active diagnosability. As discussed, event-based diagnosers can be used to determine if the system is passively diagnosable. However, active diagnosis, and, hence, active diagnosability, involves active control of the system. The diagnoser captures all physically possible controlled mode changes, but these are often constrained by the controller. We assume the possible controllable trajectories can be represented as a function of the current diagnosis using a finite automaton termed the *mode controller*.

Definition 15. (Mode Controller). The *mode controller* is defined as $\mathcal{C} = (S, s_0, \Sigma, D, \delta)$, where S is a set of states, $s_0 \in S$ is the initial state, Σ is a set of events, $D \subseteq 2^C$ is the set of diagnoses, and $\delta : S \times \Sigma \times D \rightarrow S$ is a transition function.

The mode controller receives events representing controlled mode changes and, conditioned on the current diagnosis, moves to a new state, which determines which future controlled mode changes are allowed. It captures the possible mode sequences that are safe to execute given the previous sequence of mode changes and the current diagnosis. Our goal is to construct a *constrained diagnoser* based on the passive diagnoser and the mode controller which we can use to verify active diagnosability.

The constrained diagnoser, denoted as $\mathcal{D}_{\mathcal{S}}^{\mathcal{C}}$, can be constructed by the following composition.

Definition 16. ($\mathcal{D}_{\mathcal{S}}^{\mathcal{C}}$). Given $\mathcal{D}_{\mathcal{S}} = (S_1, s_{01}, \Sigma_1, \delta_1, A_1, D_1, Y_1)$ and $\mathcal{C} = (S_2, s_{02}, \Sigma_2, D_2, \delta_2, Y_2)$, $\mathcal{D}_{\mathcal{S}}^{\mathcal{C}}$ is defined as $(S, s_0, \Sigma, \delta, A, D, Y)$, where

- $S = S_1 \times S_2$,
- $s_0 = (s_{01}, s_{02})$,
- $\Sigma = \Sigma_1 \cup \Sigma_2$,
- $\delta((s_1, s_2), \sigma)$ is $(\delta_1(s_1, \sigma), \delta_2(s_2, \sigma, Y_1(s_1)))$ if both transition states are not $\varnothing$,
- A is all $(s_1, s_2) \in S$ where $s_1 \in A_1$,
- $D = D_1 \cup D_2$, and
- $Y((s_1, s_2)) = Y_1(s_1)$.

The constrained diagnoser is constructed by removing event sequences which are not allowed by $\mathcal{C}$ given the current diagnosis. We can use this diagnoser to determine active diagnosability of a system. Essentially, for every accepting state in the passive diagnoser where the diagnosis is not a singleton, we want to verify if (i) we can prevent entering the state by blocking some sequence of controlled mode changes (prevent the ambiguity), or (ii) we can exit that state via a sequence of controlled mode changes to a state which is not accepting (move to a state in which more measurement deviations must occur).

Theorem 17. A system $\mathcal{S}$ with constrained diagnoser $\mathcal{D}_{\mathcal{S}}^{\mathcal{C}}$ is actively diagnosable if and only if for all $s \in A$, where $d \in Y(s)$ is not a singleton, either (i) for every trace leading to s, there is some $s' \notin A$ and sequence of controlled mode change events λ_Q where $\delta(s', \lambda_Q \lambda) = s$, or (ii) there exists some sequence of controlled mode changes λ_Q where $Y(\delta(s, \lambda)) \notin A$.

Proof 1. Since the constrained diagnoser captures only controller-allowed mode change sequences, any trajectory present in the constrained diagnoser will be allowed by the controller. Since accepting states correspond to isolation results, we need to ensure that all accepting states with ambiguous diagnoses can be avoided or exited. In the first case, the trace λ_Q must be blocked, and in the second, λ_Q must be executed. If so, then the system is actively diagnosable and unique diagnosis results can be achieved.

4.2 Online Diagnosis

Offline design and diagnosability analysis of the constrained diagnoser determines whether unique diagnosis results can be obtained under active diagnosis. If the system is actively diagnosable, then there will always be an allowed sequence of controlled mode changes that can be executed, and, therefore, can be used for online diagnosis.

The overall architecture is shown in Fig. 1. Fault detection triggers symbol generation, which abstracts measurement deviations into symbolic form. Active diagnosis is split into two parts, the diagnoser, which produces diagnoses based on observed events, and the controller, which takes in high-level user commands, ε, and the current diagnoser state, and computes an appropriate trajectory of mode change commands to best satisfy the control and diagnosis objectives. The implementation of the controller is out of the scope of this paper, but different strategies can be used to best trade off information gain for diagnosis versus the control objectives, based on the system requirements. In one extreme, no diagnoser-initiated commands may be issued in order to keep control with the operator, while in the other extreme, the controller can be used to quickly reconfigure the system to the mode that provides the most information for fast diagnosis, using, e.g., conditional planning methods (Bayoudh et al., 2008) or methods based on probe or test selection (de Kleer and Williams, 1987).

5. CASE STUDY

We apply our active diagnosability framework to the Advanced Diagnostics and Prognostics Testbed (ADAPT) deployed at NASA Ames (Poll *et al.*, 2007). The testbed is functionally representative of a spacecraft's electrical power system, and consists of power generation, storage, and distribution components, including lead-acid batteries, a number of relays and circuit breakers, inverters, and various DC and AC loads. A controller (that defines $\mathcal{C}$) is implemented that restricts how the operator may configure the relays so that certain configurations, e.g., two batteries being connected in parallel, are disallowed.

We consider a subset of ADAPT to demonstrate our approach that includes a lead-acid battery, two relays, and two DC loads. The battery is modeled by an electric circuit equivalent described in (Daigle, 2008) (see Fig. 2). The battery supplies voltage to two DC loads through two relays. The selected measurements are the battery voltage, $V_B(t)$, and the currents through the relays, $I_{L1}(t)$ and $I_{L2}(t)$, i.e., $M = \{I_{L1}, I_{L2}, V_B\}$.

We consider faults in the battery, loads, relays, and sensors. Common battery faults include decrease in charge-holding capacity and resistance increases brought about by battery use and age, which manifest as a side effect of the chemical reactions. Capacitance decrease is represented as a decrease in the main capacitance parameter, C_0^-, and

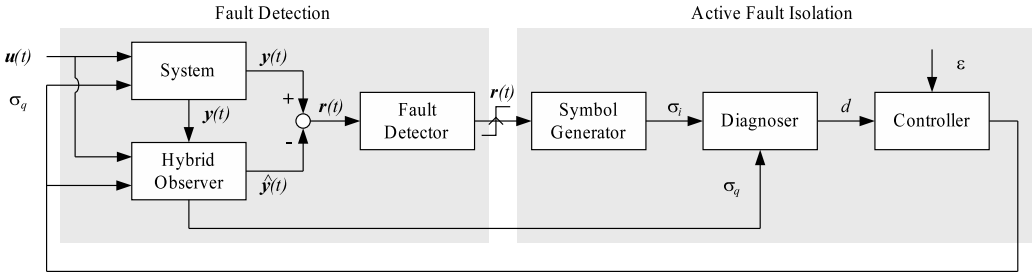

Fig. 1. Active diagnosis architecture

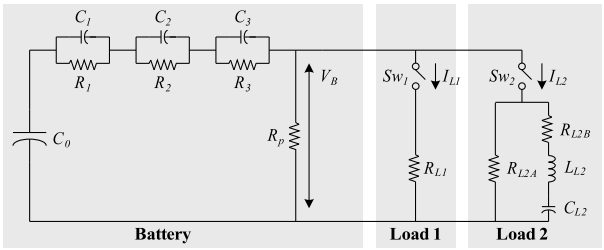

Fig. 2. Circuit equivalent for the selected subsystem.

an increase in parasitic losses by R_p^-. Faults in the system loads are represented by increases or decreases in their resistance values, R_{L1} and R_{L2A}. We also consider sensor bias faults, which produce abrupt, constant offsets 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. We represent discrete faults in Sw_1 and Sw_2 by fault events α and β, respectively, where a subscript of 0 indicates a stuck-off fault, and a subscript of 1 indicates a stuck-on fault.

5.1 Diagnosability Analysis

We denote the system mode as q_{ij} and a controlled mode change to q_{ij} as $\sigma_{q_{ij}}$, where i is the mode of Sw_1, and j is the mode of Sw_2. In this particular subsystem, $\mathcal{C}$ does not constrain the diagnoser, i.e., we allow controlled mode changes that switch the system from any one controlled mode to another. We restrict discrete faults to only occurring from expected modes where a deviation would be produced, e.g., α_1 would not produce any deviations if it occurred in a mode where Sw_1 was already on.

The fault signatures and relative measurement orderings for the faults are given in Table 1 for selected modes (q_{**} indicates the signatures and orderings are valid for any mode). The nonlinearities in the battery introduce ambiguity in the qualitative signatures, and this is denoted by the `*` symbol, e.g., a signature of `0*` may manifest as `0+` or `0-`. Since the sensors are not part of any feedback loops in the system, sensor faults affect only the measurement provided by the sensor. The other measurements are not affected, and so the corresponding fault signatures are denoted by `00`, indicating no change in the measurement from expected behavior.

Given any one mode, we find that the system is diagnosable. However, over all modes, the system is not diagnosable. Fig. 3 gives a partial diagnoser for the system that illustrates this property, with $F = \{C_0^-, R_{L1}^-\}$, $\Sigma_Q = \{\sigma_{q_{01}}, \sigma_{q_{11}}\}$, and initial mode q_{11}. If $I_{L1}^{+-,X}\sigma_{q_{01}}$

Fault	V_B	I_{L1}	I_{L2}	Measurement Orderings
(V_B^+, q_{**})	`+0,X`	`00,X`	`00,X`	$V_B \prec I_{L1}, V_B \prec I_{L2}$
(V_B^-, q_{**})	`-0,X`	`00,X`	`00,X`	$V_B \prec I_{L1}, V_B \prec I_{L2}$
(I_{L1}^+, q_{**})	`00,X`	`+0,X`	`00,X`	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
(I_{L1}^-, q_{**})	`00,X`	`-0,X`	`00,X`	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
(I_{L2}^+, q_{**})	`00,X`	`00,X`	`+0,X`	$I_{L1} \prec V_B, I_{L2} \prec I_{L1}$
(I_{L2}^-, q_{**})	`00,X`	`00,X`	`-0,X`	$I_{L1} \prec V_B, I_{L2} \prec I_{L1}$
(C_0^-, q_{11})	`+-,X`	`+-,X`	`+-,X`	$\varnothing$
(R_p^-, q_{11})	`0-,X`	`0-,X`	`0-,X`	$\varnothing$
(R_{L1}^+, q_{11})	`0*,X`	`-+,X`	`0*,X`	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
(R_{L1}^-, q_{11})	`0*,X`	`+-,X`	`0*,X`	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
(R_{L2A}^+, q_{11})	`0*,X`	`0*,X`	`-+,X`	$I_{L2} \prec V_B, I_{L2} \prec I_{L1}$
(R_{L2A}^-, q_{11})	`0*,X`	`0*,X`	`+-,X`	$I_{L2} \prec V_B, I_{L2} \prec I_{L1}$
$(\alpha_0, q_{\alpha_0 1})$	`0*,X`	`-*,Z`	`0*,X`	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
$(\alpha_1, q_{\alpha_1 1})$	`0*,X`	`+*,N`	`0*,X`	$I_{L1} \prec V_B, I_{L1} \prec I_{L2}$
$(\beta_0, q_{1\beta_0})$	`0*,X`	`0*,X`	`-*,Z`	$I_{L2} \prec V_B, I_{L2} \prec I_{L1}$
$(\beta_1, q_{1\beta_1})$	`0*,X`	`0*,X`	`+*,N`	$I_{L2} \prec V_B, I_{L2} \prec I_{L1}$

Table 1. Fault signatures and relative measurement orderings for the ADAPT subsystem.

occurs, we reach an accepting state that corresponds to a diagnosis with multiple candidates. After that event, both C_0^- and R_{L1}^- are consistent. Since the state is accepting, it is possible that no new measurement deviations will occur to distinguish the faults. The resistance fault will have no visible effects on the rest of the measurements in this mode, because the source of the deviations is cut off, so we would have to wait infinitely long to verify R_{L1}^- as the true fault. We can see, however, that the system is actively diagnosable. If we prevent $\sigma_{q_{01}}$ from occurring, or change back to q_{11} if it does occur, more measurements will deviate and we can distinguish the candidate uniquely.

5.2 Experimental Results

In the following, we present an experiment to illustrate our active diagnosis methodology. We consider a fault scenario where an abrupt 33% decrease in the Load 1 resistance, R_{L1}^-, occurs. As discussed, if Sw_1 is turned off soon after the fault appears, the mode change may mask the effect of the fault on V_B, which will result in an ambiguous diagnosis. Simulated system outputs and estimated outputs from the observer are shown in Fig. 4 near the point of fault occurrence at 500.0 s in mode q_{11}. The decrease in resistance increases the current drawn by the load abruptly, and this change is detected at 500.0 s. Since the slope of the change is not yet known, the possible fault hypotheses are $\{(C_0^-, q_{11}), (R_{L1}^-, q_{11}), (I_{L1}^+, q_{11})\}$. At 501.0 s, a mode change to q_{01} is commanded. As a result, the change in $V_B(t)$ does not grow and remains hidden in the noise. Futher measurements do not deviate and the

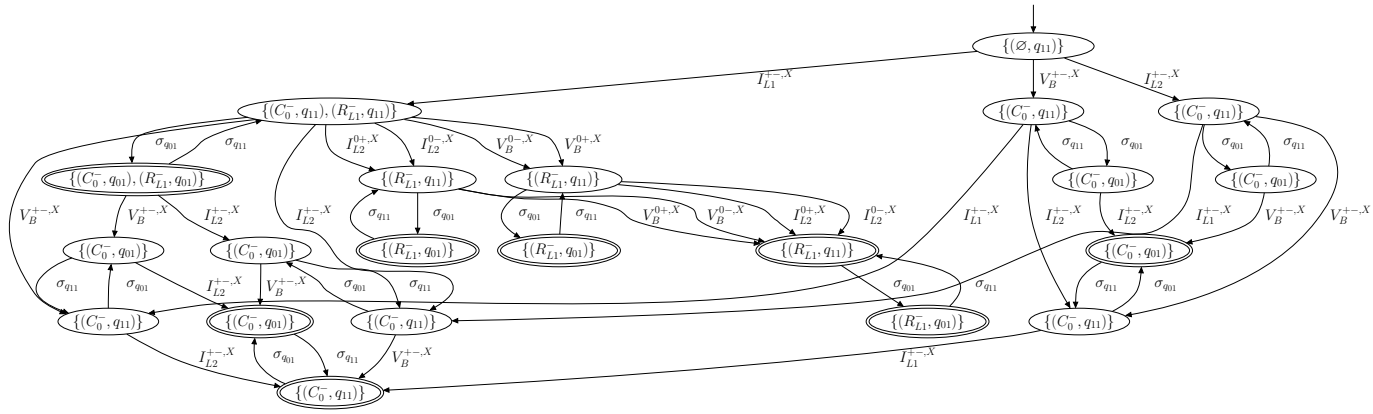

Fig. 3. Partial hybrid diagnoser for $F = \{C_0^-, R_{L1}^-\}$ and initial mode q_{11}.

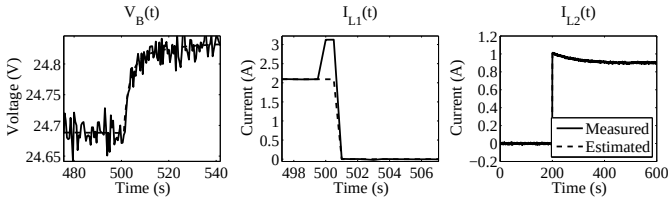

Fig. 4. R_{L1}^- fault, where R_{L1} decreases by 33% followed by $\sigma_{q_{01}}$.

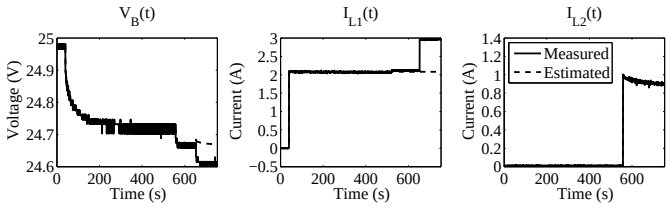

Fig. 5. R_{L1}^- fault, where R_{L1} decreases by 33%.

diagnosis will remain ambiguous unless Sw_1 is turned back on, as predicted by the diagnosability analysis.

Diagnosability analysis shows that this situation can be prevented. In the actual testbed, a 33% decrease in the Load 1 resistance, R_{L1}^-, is manually injected at 653.0 s in mode q_{11}, by abruptly changing its resistance setting. The measured and estimated outputs are shown in Fig. 5. The current increase is detected at 653.5 s, resulting in the diagnosis $\{(C_0^-, q_{11}), (R_{L1}^-, q_{11}), (I_{L1}^+, q_{11})\}$. At 654.0 s the event $\sigma_{q_{01}}$ is blocked by the active diagnoser. At 655.0 s, a decrease is detected in $V_B(t)$. Since I_{L1}^+ cannot affect $V_B(t)$, it is dropped. C_0^- is also dropped because it would have increased, and not decreased, the battery voltage. Therefore the fault is isolated as R_{L1}^-.

6. CONCLUSIONS

We presented a systematic framework for analyzing diagnosability of hybrid systems using a qualitative fault isolation approach. We create event-based diagnosers for passive and active diagnosis in hybrid systems, and show how diagnosability of the system can be determined using the diagnosers, and how using active diagnosis can improve diagnosability. If the system is actively diagnosable, unique isolation results can be achieved by blocking or executing allowable controlled mode changes during fault isolation. We demonstrated the approach on an electrical power distribution system. In future work, we will extend our approach to include autonomous mode changes, multiple faults, and more general control paradigms including fault-adaptive control.

REFERENCES

M. Bayoudh, L. Trave-Massuyes, and X. Olive. Towards active diagnosis of hybrid systems. In *Proceedings of the 19th International Workshop on Principles of Diagnosis*, pages 231–237, September 2008.

G. Biswas, G. Simon, N. Mahadevan, S. Narasimhan, J. Ramirez, and G. Karsai. A robust method for hybrid diagnosis of complex systems. In *Proc. of the 5th Symposium on Fault Detection, Supervision and Safety for Technical Processes*, pages 1125–1131, June 2003.

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

M. Daigle, X. Koutsoukos, and G. Biswas. An event-based approach to hybrid systems diagnosability. In *Proceedings of the 19th International Workshop on Principles of Diagnosis*, pages 47–54, September 2008.

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.

J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32:97–130, 1987.

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.

S. Narasimhan and G. Biswas. Model-based diagnosis of hybrid systems. *IEEE Trans. on Systems, Man and Cybernetics, Part A*, 37(3):348–361, May 2007.

H. Niemann. Fault tolerant control based on active fault diagnosis. In *Proceedings of the 2005 American Control Conference, 2005*, pages 2224–2229, 2005.

S. Poll *et al.* Evaluation, selection, and application of model-based diagnosis tools and approaches. In *AIAA Infotech@Aerospace 2007 Conf. and Exhibit*, May 2007.

M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, 43(7):908–929, 1998.