

Comprehensive Diagnosis of Complex Electrical Power Distribution Systems[★]

Indranil Roychoudhury^{*} Matthew Daigle^{**} Gautam Biswas^{*}
Xenofon Koutsoukos^{*} Ann Patterson-Hine^{**} Scott Poll^{**}

^{*} *Institute for Software Integrated Systems, Department of EECS,
Vanderbilt University, Nashville, TN, USA
(email:indranil.roychoudhury@vanderbilt.edu)*

^{**} *NASA Ames Research Center, Moffett Field, CA, USA*

^{***} *University of California, Santa Cruz, CA, USA*

Abstract: Electrical power distribution systems are composed of heterogeneous components, which include continuous power sources, discrete relays, passive and active elements and loads, and fast-switching power conversion subsystems. This heterogeneity introduces significant challenges for model-based diagnosis, such as building accurate models, and generating fast and accurate diagnoses while ensuring robustness to measurement noise and modeling errors. In this paper, we present a comprehensive methodology for the diagnosis of parametric and discrete faults in electrical power distribution systems that include dc and ac components. We use a hybrid bond graph modeling language to systematically develop computational models and algorithms for hybrid state estimation, robust fault detection, and efficient fault isolation. Simulation and experimental results on a real-world electrical power distribution system demonstrate the effectiveness of our methodology.

Keywords: Model-based diagnosis; electric power systems; distributed diagnosis.

1. INTRODUCTION

The increasing complexity of modern electrical power distribution systems requires online health monitoring and diagnosis to ensure safe, efficient operation. Model-based diagnosis schemes [Mosterman and Biswas, 1999, Narasimhan and Biswas, 2007] are preferred over other approaches for their more general and robust diagnosis solutions. But, their deployment on real systems presents significant challenges, including building accurate and reliable models, designing robust observers and fault detectors, and developing fault isolation schemes that quickly produce valuable information with limited information and computational resources. Moreover, most power distribution systems are composed of heterogeneous dc and ac components that operate at different time constants, thus introducing additional challenges for correct diagnosis.

One such system is the Advanced Diagnostics and Prognostics Testbed (ADAPT) [Poll *et al.*, 2007] at NASA Ames Research center. ADAPT emulates spacecraft electrical power distribution systems, and is designed to provide an environment where researchers and practitioners can develop and test their diagnosis algorithms. ADAPT poses a number of significant challenges, such as the limited availability of information and data to estimate and validate our model parameters, and the installed sensors providing data sample only at the rate of 2 Hz, which cannot sufficiently capture the dynamics of some ADAPT subsystems that operate at much higher frequencies.

In order to address the challenges of diagnosis of real-world systems like ADAPT, we extend the Hybrid TRANSCEND methodology, a comprehensive model-based approach for combined qualitative/quantitative diagnosis of hybrid systems [Narasimhan and Biswas, 2007]. First, we develop a new method for deriving parametric and discrete fault signatures from steady-state analysis of ac measurements, which exhibit fault transients that occur faster than the sensor sampling rates. Second, we develop a distributed methodology for the comprehensive diagnosis of faults in dc and ac subsystems in a hybrid systems framework, which extend our previous diagnosis approach for continuous systems [Roychoudhury *et al.*, 2008b] by analyzing the diagnosability of a system over multiple modes. Our distributed diagnosis scheme addresses the drawbacks of centralized diagnosis schemes, such as a single point of failure, large computational complexity, and poor scalability. Third, we illustrate the effectiveness of our approach by experimental studies on ADAPT and a fully-developed, high-fidelity simulation environment called VIRTUAL ADAPT [Roychoudhury *et al.*, 2008c].

2. MODELING

Our component-based hybrid system models are built using the hybrid bond graph (HBG) modeling language [Mosterman and Biswas, 1998]. HBGs are particularly suitable for diagnosis because they incorporate causal and temporal information required for deriving and analyzing fault transients, and allow automatic generation of simulation models, observers, and diagnosis models [Narasimhan and Biswas, 2007]. Fig. 1(b) shows the

[★] This work was supported in part by the National Science Foundation under Grant CNS-0615214 and NASA NRA NNX07AD12A.




Fig. 1. Switched circuit example.

HBG model of the electrical circuit shown in Fig. 1. In an HBG, components are represented by vertices, and bonds, drawn as half arrows, capture ideal energy connections. Associated with each bond, i , are two variables: effort, e_i , and flow, f_i , whose product defines the rate of energy transfer through the bond. In the electrical domain, effort and flow map to voltage and current, respectively.

HBG components include resistances, R (e.g., R_1), which capture energy dissipation, i.e., $e_i = Rf_i$, capacitances, C (e.g., C_1), and inductances, I (e.g., L_1), which capture energy storage, i.e., $e_i = \frac{1}{C}f_i$ and $f_i = \frac{1}{I}e_i$, respectively, and sources of flow (Sf) and effort (Se) (e.g., $v(t)$), which capture the energy flow into and out of the system. 1-junctions represent series connections (points of common flow), and 0-junctions represent parallel connections (points of common effort). Hybrid behaviors are captured using *switching junctions* (denoted by the dashed arrow in Fig. 1(b)), which act as ideal switches, enabling a junction to be either on or off. Off 1-junctions and 0-junctions behave as sources of zero flow and effort, respectively. The switching behavior is defined by a finite automaton, the *control specification* (CSPEC), whose state determines whether the junction is on or off. The overall system mode is defined implicitly by the individual states of all the CSPECs, and this provides a concise representation of the hybrid system model.

In this work, we focus on the diagnosis of single, abrupt, persistent faults in hybrid systems. Extensions of our diagnosis approach for diagnosing incipient faults have been presented in [Roychoudhury et al., 2008a], but are beyond the scope of this paper. We classify faults into two categories: (i) *parametric faults*, and (ii) *discrete faults*. Parametric faults, which represent partial failures or degradations in system components, manifest as abrupt changes in the HBG model parameter values. A parametric fault in parameter, P , is denoted by P^+ or P^- , representing an increase or decrease in the parameter value, respectively. Discrete faults correspond to differences between the actual and expected state of a switching component in the HBG model, and are modeled using unobservable fault events in the CSPECs [Daigle, 2008].

Our model for qualitative fault diagnosis, the temporal causal graph (TCG), is derived from the current mode of the HBG system model, and explicitly captures the propagation of both parametric and discrete fault effects on system variables and outputs [Mosterman and Biswas, 1999, Daigle, 2008]. The TCG is essentially a signal flow graph whose nodes are system variables or discrete fault events. The labeled edges represent the qualitative relationships between the variables, i.e., equality, direct or inverse proportionality, integral, and parametric dependencies. Links from discrete fault events also include labels

that describe if the fault causes the variable value to go from zero to nonzero or vice versa, respectively.

3. DIAGNOSER DESIGN

Our approach to diagnosing faults in power distribution systems, such as ADAPT, combines schemes for diagnosis using transients in the dc subsystems and changes in steady-state values for ac measurements, such as rms and phase values of voltages and currents. We derive *fault signatures* for the both dc and ac measurements, which capture the effect of a fault on a measurement at its time of occurrence. Fault isolation is performed by comparing predicted fault signatures to the observed deviations from nominal behavior of the system. Given the signatures, we analyze system diagnosability, which is then used to develop the overall system diagnoser as two interacting diagnosers: the dc subsystem diagnoser, and the ac subsystem diagnoser.

3.1 Fault Signatures

In general, fault signatures are linked to the transients that occur after fault occurrence. Assuming continuous and continuously differentiable output except at the points of fault occurrence and mode changes, the transient response after a fault can be approximated by a Taylor series expansion, which is defined by the changes in magnitude and higher order derivatives in the signal at the point of fault occurrence [Mosterman and Biswas, 1999]. These qualitative signatures describe increase (+), decrease (-), and no change (0) in the magnitude and derivatives of the residual signal. A * denotes an ambiguity is determination of the sign. If a fault produces an immediate change in the residual, i.e., a discontinuity at the point of fault occurrence, then the magnitude symbol will be + or -, otherwise it will be 0. In previous work, we have proven that for parametric faults, the first change and subsequent slope provide all of the discriminatory evidence for qualitative fault isolation in dynamic systems [Manders et al., 2000]. Therefore, our fault signatures include two symbols: the magnitude change and slope of the residual signal. For discrete fault diagnosis, fault signatures have been extended to include a third symbol that indicates if a fault causes zero to nonzero (N), nonzero to zero (Z), or no zero/nonzero value changes (X) in a measurement. Note that the fault signatures depend on the direction of change in the parameter value, and not on the actual fault magnitude.

Fault Signatures for DC Measurements For the dc measurements, fault signatures are derived for each hypothesized fault in a given mode by performing a forward propagation of fault effects throughout the TCG [Mosterman and Biswas, 1999, Daigle, 2008]. Signatures for the circuit example are given in Table 1, assuming the voltage source is dc instead of ac, and variable values are nominally positive, where the measurements are V_{R_1} , the voltage across R_1 , and I_{R_2} , the current through R_2 . The table shows that the system is not diagnosable with the selected measurements, because faults in L_1 and R_1 produce the same signatures.

Table 1. Fault signatures for the circuit with the switch On

Fault	DC Signatures		AC Signatures			
	V_{R1}	I_{R2}	V_{R1}	$\phi_{V_{R1}}$	I_{R2}	$\phi_{I_{R2}}$
C_1^+	0+,X	-+,X	-,X	-,X	-,X	-,X
C_1^-	0-,X	+-,X	*,X	+,X	+,X	+,X
L_1^+	-+,X	0-,X	-,X	-,X	-,X	-,X
L_1^-	+-,X	0+,X	+,X	+,X	+,X	+,X
R_1^+	-+,X	0-,X	+,X	+,X	-,X	+,X
R_1^-	+-,X	0+,X	-,X	-,X	+,X	-,X
R_2^+	0-,X	-+,X	+,X	-,X	-,X	-,X
R_2^-	0+,X	+-,X	-,X	*,X	+,X	+,X
Sw_1^{off}	0-,X	-,Z	-,X	+,X	+,Z	-,Z

Fault Signatures for AC Measurements Fault transients in the ac domain, where the components usually operate at 60 Hz, would require sampling at rates greater than 120 Hz, making online analysis of fault signatures computationally infeasible. Further, the ADAPT system is equipped with rms and phase sensors that operate at only 2-Hz for ac voltage and current measurements. Therefore, from practical considerations, the ac fault signatures represent steady state deviations in the measurements. These fault signatures are derived by computing the partial derivative of the steady-state expression for a measurement with respect to a given fault variable to determine the sign of the measurement change.

This analysis starts by deriving the symbolic expressions relating faults to the measurements using the HBG model of the system. The parameters for the R , C , and L elements are replaced by their complex impedance representations in the ac domain. Given the frequency, ω , in radians, the impedance of a resistance, R , is $Z_R = R$, a capacitor, C , is $Z_C = \frac{1}{j\omega C}$, and an inductor, L , is $Z_L = j\omega L$. By combining the constitutive relations of the elements and the junction equations derived from the bond graph, we can generate the voltage and current variable relations in symbolic form. By algebraic manipulation, we get the symbolic form of the expressions for the ac measurements as a function of a given fault. After substituting nominal values of all other parameters, if the sign of this partial derivative is always positive, or negative, for the considered fault magnitudes, then the corresponding fault signature is defined to be a $+$, or $-$, respectively. If the sign cannot be uniquely determined, the ambiguity is represented using $*$. Since discrete faults represent changes in system mode, we determine the signatures by simply computing the rms and phase values for the different configurations, and then comparing them to compute the effects the mode changes will have.

Example To illustrate, we consider the circuit given in Fig. 1(a). The measured signals are the voltage across R_1 , $v_{R1}(t)$, and the current through R_2 , $i_{R2}(t)$. The measurements include both rms values and phase difference relative to the source voltage for both measured signals. We assume that the source voltage $v(t)$ is 120 V rms at 60 Hz, and the parameters have nominal values of $C_1 = 0.005$ F, $L_1 = 0.03$ H, $R_1 = 1$ Ω , and $R_2 = 2$ Ω . The switch implies the system can operate in two mode configurations. We need to analyze the effects of faults in both modes: q_0 , where the switch is off, and q_1 , where the switch is on. Given the frequency, ω , the impedances are $Z_{L1} = j\omega L_1$,

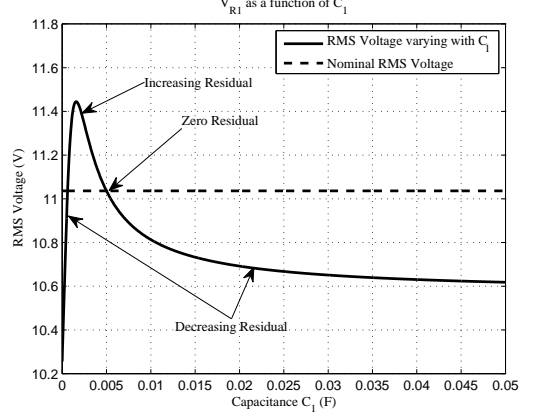


Fig. 2. V_{R1} rms value as a function of C_1 magnitude.

for the inductor L_1 , $Z_{C1} = \frac{1}{j\omega C_1}$ for the capacitor C_1 , $Z_{R1} = R_1$, for the resistor R_1 , and $Z_{R2} = R_2$, for the resistor R_2 . Using the HBG as described above, we derive the symbolic expressions describing the measurements as a function of the inputs and the impedances: $v_{R1} = \frac{v_{R1}}{Z_{eq}}$, and $i_{R2} = \{0 \text{ for mode } q_0, \frac{v_{Z_{C1,R2}}}{Z_{eq}R_2} \text{ for mode } q_1\}$, where $Z_{C1,R2} = (j\omega C_1 + \frac{1}{R_2})^{-1}$, and $Z_{eq} = \{j\omega L_1 + R_1 + \frac{1}{j\omega C_1} \text{ for mode } q_0, j\omega L_1 + R_1 + Z_{C1,R2} \text{ for mode } q_1\}$. These symbolic expressions for impedances are used to compute the fault signature matrix for each mode.

The signatures for mode q_1 are shown in Table 1. In some cases, the direction of change in measurement values depends on fault magnitude. For example, C_1^+ will always cause a decrease the rms value of V_{R1} , but C_1^- may cause either an increase or decrease in V_{R1} depending on its magnitude, as shown in Fig. 2. For its nominal value of 0.005 F, with an increase in C_1 , the measurement value always decreases, but for a decrease in magnitude, the measurement value may go above or below the nominal measurement value, so we represent the signature in this case as a $*$ (see Table 1). Discrete faults do not produce ambiguous signatures. For example, when the switch is on, the rms value of I_{R2} is 2.83 A, and when off, it is zero, therefore, when unexpectedly going from q_1 to q_0 , we will observe a decrease in I_{R2} , and it will go to zero. This is represented by the fault signature $-,Z$.

3.2 Distributed Diagnoser Design

In this paper, we extend the distributed diagnoser design approach for continuous systems (Algorithm 1 in [Roychoudhury et al., 2008b]) to hybrid systems, which allows us to decouple the diagnosers for the ADAPT ac and dc subsystems. The distributed diagnosers make the overall diagnosis approach more efficient by partitioning the diagnosis task into smaller subtasks performed by local diagnosers. The distributed diagnosers are designed to generate globally correct diagnosis results without any supervision and communicating a minimal number of measurements amongst themselves.

We start the distributed diagnoser design procedure by first generating the fault signatures for all possible faults for each mode of the system, as discussed earlier. We

assume that our fault set is complete. Since mode changes can occur during fault isolation, we also have to account for the possible interleavings of signatures for different modes. The traces formed by measurement deviations and mode change events can be represented as a finite automaton that maps states to consistent fault hypotheses [Daigle, 2008]. We denote this finite automaton as $\mathcal{D}_{F,M,Q}$, where F is the set of all possible faults, M is the set of all available measurements, and Q is the set of all system modes.

We define a subsystem $S_i = (F_i, M_i)$, where F_i is the set of faults in S_i , M_i is the set of measurements in S_i . The different F_i and M_i form partitions of the set of faults, F , and measurements, M , respectively. Given κ subsystems, $S_i = (F_i, M_i)$, $1 \leq i \leq \kappa$, and $\mathcal{D}_{F,M,Q}$, our design problem is to construct, for each subsystem, a measurement set $\widetilde{M}_i \subseteq M$ such that (i) $\widetilde{M}_i \supseteq M_i$ is minimal, and (ii) all single faults in F_i are *globally diagnosable* by measurements in \widetilde{M}_i , i.e., \widetilde{M}_i can uniquely isolate every fault, $f \in F_i$, from all other faults in F for every possible sequence of mode transitions. Given F_i and \widetilde{M}_i , we construct a *local diagnoser* [Roychoudhury et al., 2008b], $\mathcal{D}_{F_i, \widetilde{M}_i, Q}$, for each subsystem. By ensuring that each \widetilde{M}_i is minimal, the local diagnosers share minimal information with one another.

Given F and M , we first analyze if every $f \in F$ is globally diagnosable by M . If the system is not globally diagnosable for a set of measurements, we define the notion of “aggregate faults” [Roychoudhury et al., 2008b], each of which includes all single faults that are not distinguishable from one other. Our diagnosis methodology treats aggregate faults as single faults, and, as a result, the reduced fault set is guaranteed to be globally diagnosable. As we will show in Section 4, ADAPT is not globally diagnosable.

The procedure for designing diagnosers for a partitioned hybrid system is presented as Algorithm 1. Our design goal is to find, for each S_i that is not globally diagnosable given M_i , the minimal set of measurements, \widetilde{M}_i required to ensure its global diagnosability. To simplify this search, we represent the system, S , as a graph of its connected subsystems, and define the notion of proximity, d , among a subsystem pair as the minimum path length from S_g to S_h . The search for additional measurements starts from closer subsystems. For each subsystem S_i , we first identify the set of faults, $remFaults_i \subseteq F_i$ that are not globally diagnosable using M_i . We identify the *minimal* set of additional measurements, \widehat{M}_i , from all subsystems at a distance $d \leq \delta$, with $\delta = 1, 2, \dots$, that isolates the *maximal* number of faults in $remFaults_i$. For every iteration of the search, we expand $\widetilde{M}_i = \widetilde{M}_i \cup \widehat{M}_i$, while reduce $remFaults_i$ to a smaller set. If $remFaults_i$ is nonempty, δ is incremented, and the procedure is repeated until it is. At this point, we have the local diagnoser $\mathcal{D}_{F_i, \widetilde{M}_i, Q}$. Since, at every iteration step, we look for measurements that diagnose some faults in $remFaults_i$ from all other faults in F , the generated diagnosers are guaranteed to generate globally correct diagnosis results without a centralized coordinator, and only communicating a minimal number of measurements, and not partial diagnosis results. We apply this algorithm to design the distributed diagnosers for ADAPT in Section 4.

Algorithm 1 Designing Diagnosers for a Partitioned System

Input: κ local subsystems, $S_i = (F_i, M_i)$, and $\mathcal{D}_{F,M,Q}$
for each S_i **do**
 identify $remFaults_i \subseteq F_i$ that are not globally diagnosable in $\mathcal{D}_{F,M_i,Q}$.
 $\delta = 1$; $\widetilde{M}_i = M_i$
 while $remFaults_i \neq \emptyset$ **do**
 identify measurement set \widehat{M}_i from measurements of subsystems S_i at a distance $d \leq \delta$ that isolates maximal $f \in remFaults_i$, and $\widetilde{M}_i - \widehat{M}_i$ is minimal.
 $\widetilde{M}_i = \widetilde{M}_i \cup \widehat{M}_i$
 $remFaults_i = remFaults_i - f$
 if $remFaults_i \neq \emptyset$ **then**
 $\delta = \delta + 1$
 construct $\mathcal{D}_{F_i, \widetilde{M}_i, Q}$

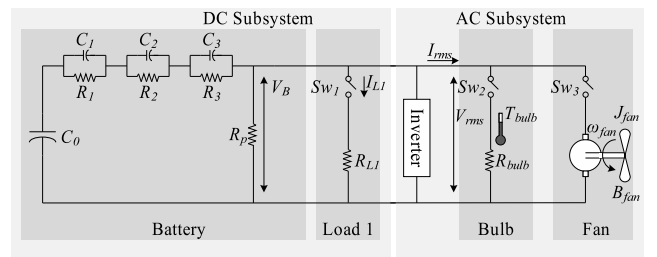


Fig. 3. Selected subset of ADAPT.

4. CASE STUDY

The ADAPT system embodies a typical functional representation of the *power generation* (two battery chargers), *power storage* (three sets of lead-acid batteries), and *power distribution* components (two inverters, a number of relays and circuit breakers, and a variety of other dc and ac loads) of a spacecraft’s electrical power system. The testbed can be commanded into different configurations, and contains sensors that measure system variables, such as voltages, currents, and temperatures.

We choose a subset of the ADAPT system to demonstrate our approach. This subset includes one of the lead-acid batteries, one dc load, an inverter, and two ac loads. Details of the component models can be found in [Daigle, 2008, Roychoudhury et al., 2008c]. A schematic of the subsystem is given in Fig. 3. The battery acts as a direct non-ideal voltage source for the dc loads. The inverter connected to the battery produces a constant 120 V rms, 60 Hz, sinusoidal ac output when the input voltage is in the range 21-32 V. When the voltage falls below 21 V, the inverter shuts off automatically. The dc load is a constant resistance, while the ac loads include a light bulb and a large fan. The available measurements include the rms values of inverter voltage and current, V_{rms} and I_{rms} , the phase difference between the inverter voltage and current, ϕ , the temperature of the light bulb, T_{bulb} , the rotational speed of the fan, ω_{fan} , the current through the dc load, I_{L1} , and the battery voltage and current, V_B and I_B .

Fault signatures for the mode with all loads are on are given in Table 2. We can see that the system is not globally diagnosable, because Sw_2^{off} and R_{bulb}^+ cannot be distinguished. We form an aggregate fault from these two

Table 2. Fault Signatures for the Mode with All Loads On

Fault	DC Measurements			AC Measurements				
	V_B	I_B	I_{L1}	V_{rms}	I_{rms}	ϕ	T_{bulb}	ω_{fan}
C_0^-	+-, X	+-, X	+-, X	00, X	00, X	00, X	00, X	00, X
R_1^+	0-, X	0*, X	0-, X	00, X	00, X	00, X	00, X	00, X
R_{L1}^+	0*, X	-*, X	-*, X	00, X	00, X	00, X	00, X	00, X
R_{L1}^-	0*, X	**, X	**, X	00, X	00, X	00, X	00, X	00, X
Sw_1^{off}	0*, X	-*, X	-, Z	00, X	00, X	00, X	00, X	00, X
R_{bulb}^+	0*, X	-*, X	0*, X	0, X	-, X	+, X	-, X	00, X
R_{bulb}^-	0*, X	**, X	0*, X	0, X	+, X	-, X	+, X	00, X
J_{fan}^-	0*, X	**, X	0*, X	0, X	-, X	-, X	0, X	+-, X
B_{fan}^+	0*, X	**, X	0*, X	0, X	+, X	+, X	0, X	0-, X
Sw_2^{off}	0*, X	-*, X	0*, X	0, X	-, X	+, X	-, X	00, X
Sw_3^{off}	0*, X	**, X	0*, X	0, X	-, X	-, Z	0, X	0-, X

faults to apply the diagnoser design algorithm described in Section 3. We consider two subsystems (see Fig. 3), (i) the dc subsystem, containing the battery, the two dc loads, and Sw_1 , and (ii) the ac subsystem, containing the inverter, the ac loads, and Sw_2 and Sw_3 . The dc subsystem fault list, F_{dc} , includes changes in the dc load resistance, R_{L1} , the capacitance and a parasitic resistance of the battery, C_0 and R_1 , and faults in Sw_1 . The dc measurements, M_{dc} , include I_{L1} , V_B , and I_B . The ac subsystem fault list, F_{ac} , includes faults in the inertia and resistance of the fan, J_{fan} and B_{fan} , the resistance of the light bulb, R_{bulb} , and faults Sw_2 , and Sw_3 . The ac measurements, M_{ac} , include V_{rms} , I_{rms} , ϕ , T_{bulb} , and ω_{fan} .

Using Algorithm 1, we obtain distributed diagnosers for the selected subsystems, which naturally falls out of the decoupling of the systems introduced by the inverter. The distributed diagnoser for the ac subsystem does not require any additional measurements from the dc subsystem to isolate its faults, i.e., $\widetilde{M}_{ac} = \{V_{rms}, I_{rms}, \phi, T_{bulb}, \omega_{fan}\}$. This is clear from the signatures given in Table 2. If a dc fault occurs, no deviations will be observed on any of the ac measurements, therefore, the ac diagnoser will not isolate any dc faults.

The dc subsystem, on the other hand, does require ac measurements to achieve unique isolation. Faults in the ac subsystem also cause the dc measurements to deviate. To overcome this ambiguity, the distributed diagnosis design communicates the I_{rms} measurement to the dc diagnoser. Since dc faults do not change I_{rms} , (due to the controlled behavior of the inverter) the dc diagnoser eliminates all local faults and determines the fault is in the ac subsystem when I_{rms} deviates. If it does not deviate, the dc diagnoser will isolate a dc fault and the ac diagnoser will not since it will not observe any deviations. Due to the autonomous mode change behavior of the inverter, the dc diagnoser also requires V_{rms} , because the ac measurements are affected by a dc fault, if the fault is such that it causes the inverter to shut off, i.e., $\widetilde{M}_{dc} = \{V_B, I_B, I_{L1}, V_{rms}, I_{rms}\}$. If a change occurs in V_{rms} , then a subsequent change in I_{rms} is explained by the inverter shutting off, and not an ac fault.

4.1 Implementation

The implemented architecture of our diagnosis scheme is shown in Fig. 4. The hybrid observer computes estimated outputs, $\hat{\mathbf{y}}(t)$, given the inputs $\mathbf{u}(t)$ and the measured out-

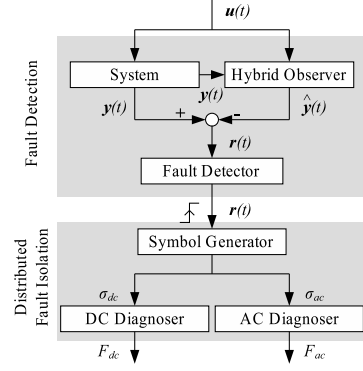


Fig. 4. Diagnosis architecture.

Table 3. Diagnosis Results from Experiments run on ADAPT

Fault	DC Diagnoser			AC Diagnoser		
	t_d (s)	t_i (s)	Result	t_d (s)	t_i (s)	Result
Experiments run in simulation						
C_0^- (-1%)	0.5	12.0	$\{C_0^-\}$	N/A	N/A	\emptyset
R_1^+ (+200%)	1.5	10.5	$\{R_1^+\}$	N/A	N/A	\emptyset
Sw_1^{off}	0.0	3.0	$\{Sw_1^{off}\}$	N/A	N/A	\emptyset
J_{fan}^- (-50%)	N/A	N/A	\emptyset	0.0	0.0	$\{J_{fan}^-\}$
B_{fan}^+ (+50%)	N/A	N/A	\emptyset	0.5	4.0	$\{B_{fan}^+\}$
Experiments run on the testbed						
R_{L1}^+ (+100%)	0.5	8.0	$\{R_{L1}^+\}$	N/A	N/A	\emptyset
R_{L1}^- (-33%)	0.5	3.5	$\{R_{L1}^-\}$	N/A	N/A	\emptyset
R_{bulb}^+ (+50%)	1.0	1.0	\emptyset	1.0	11.0	$\{R_{bulb}^+, Sw_2^{off}\}$
R_{bulb}^- (-50%)	2.5	2.5	\emptyset	2.5	2.5	$\{R_{bulb}^-, B_{fan}^+\}$
Sw_2^{off}	0.5	0.5	\emptyset	0.5	2.0	$\{R_{bulb}^+, Sw_2^{off}\}$
Sw_3^{off}	0.5	0.5	\emptyset	0.5	1.5	$\{Sw_3^{off}\}$

puts $\mathbf{y}(t)$. The difference between estimated and observed outputs forms the residual $\mathbf{r}(t)$, which, based on a statistical significance test and a sliding window technique, the fault detector uses to detect the presence of faults [Biswas et al., 2003]. When a fault is detected, the symbol generator begins to abstract the deviated measurements into fault signatures, σ . Signatures for measurements in \widetilde{M}_{dc} are sent to the dc diagnoser, and signatures for measurements in \widetilde{M}_{ac} are sent to the ac diagnoser. Each diagnoser then independently isolates faults in its own fault set.

4.2 Simulation Results

We present diagnosis results obtained on the simulation testbed Virtual ADAPT. We ran these simulation experiments to diagnose faults which were difficult or impossible to introduce in the hardware components. For this set of experiments, we inject faults into the configuration where both ac loads and the dc load are all online.

The results are summarized in Table 3. In the table, t_d is the time taken to detect a fault, and t_i is the time to isolate the fault, which is given as the point at which a diagnoser last reduces its fault set. All times in Table 3 are expressed in seconds. In all cases, the correct fault was isolated. In some cases, i.e., for C_0^- and R_1^+ , the slope had to be calculated, which took an additional amount of time.




Fig. 5. Experimental results.

4.3 Experimental Results

We have also performed experiments on ADAPT and the results are summarized in Table 3. In all cases, the diagnosis results included the true fault. Here, we have to cope with model uncertainty in addition to sensor noise, and the observer and fault detectors had to be tuned for this purpose. To demonstrate the diagnosis approach, we show the results obtained for a fault in the dc subsystem and one in the ac subsystem. Additional experimental results can be found in [Daigle, 2008].

First, we consider a 100% increase in the Load 1 resistance, R_{L1}^+ , injected manually at 100.0 s in the mode with all loads on. The measured and estimated outputs are shown in Fig. 5(a). The increase in resistance causes a discontinuous drop in the current, detected at 100.5 s. Since the slope has not yet been computed, the possible fault candidates are $\{R_1^+, R_{L1}^+, Sw_1^{off}\}$. At 102.5 s, the increase in V_B is detected, thus eliminating R_1^+ . At 103.5 s, it is determined that I_{L1} did not go to zero, thus eliminating Sw_1^{off} , and isolating R_{L1}^+ as the true fault. None of the measurements in the ac subsystem deviate, so the ac diagnoser does not generate any candidates.

We next consider a discrete fault where Sw_2 turns off at 100.0 s. The relevant measured and estimated outputs are shown in Fig. 5(b). At 100.5 s, an increase in V_B is detected, so the dc diagnoser generates its initial candidates as $\{C_0^-, R_{L1}^+, R_{L1}^-, Sw_1^{off}\}$. Also at 100.5 s, a decrease in I_{rms} is detected, so the initial candidates of the ac diagnoser are $\{R_{bulb}^+, J_{fan}^-, Sw_3^{off}, Sw_2^{off}\}$. Because this measurement is known to the dc diagnoser, it can eliminate all of its faults and conclude that the fault must be in the ac subsystem. At 101.0 s it is determined that the change in V_B was not a discontinuity, but the ac diagnosis remains unchanged. At 102.0 s, an increase in ϕ is detected, which reduces the fault set to $\{Sw_2^{off}, R_{bulb}^+\}$, which cannot be distinguished further, as explained earlier.

5. CONCLUSIONS

In this paper, we described how the HBG modeling paradigm provided a common framework for the generation of fault signatures for dc and ac measurements through transient and steady-state analyses, respectively. Based on the signatures generated, we performed diagnosability analysis of the system and designed distributed diagnosers for the dc and ac subsystems. In future work, we will perform additional online experiments to test our

fault detection and symbol generation strategy for a sensitivity to a variety of fault magnitudes under different sensor noise profiles. We are also improving our parameter estimation scheme for use on ADAPT, and would like to provide confidence estimates when multiple candidates are retained after fault isolation. Ongoing work is also further extending our techniques to deal with incipient faults.

REFERENCES

- 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*, pages 1125–1131, June 2003.
- M. Daigle. *A Qualitative Event-based Approach to Fault Diagnosis of Hybrid Systems*. PhD thesis, Vanderbilt University, 2008.
- E.-J. Manders, S. Narasimhan, G. Biswas, and P.J. Mosterman. A combined qualitative/quantitative approach for fault isolation in continuous dynamic systems. In *SafeProcess 2000*, volume 1, pages 1074–1079, Budapest, Hungary, June 2000.
- P. J. Mosterman and G. Biswas. A theory of discontinuities in physical system models. *Journal of the Franklin Institute*, 335B(3):401–439, January 1998.
- P.J. Mosterman and G. Biswas. Diagnosis of continuous valued systems in transient operating regions. *IEEE Trans. SMC, Part A*, 29(6):554–565, 1999.
- S. Narasimhan and G. Biswas. Model-based diagnosis of hybrid systems. *IEEE Transactions on Systems, Man and Cybernetics, Part A*, 37(3):348–361, May 2007.
- S. Poll *et al.* Evaluation, selection, and application of model-based diagnosis tools and approaches. In *AIAA Infotech@Aerospace 2007 Conference and Exhibit*, May 2007.
- I. Roychoudhury, G. Biswas, and X. Koutsoukos. Comprehensive diagnosis of continuous systems using dynamic bayes nets. In *Proc. of the 19th International Workshop on Principles of Diagnosis*, pages 151–158, 2008a.
- I. Roychoudhury, G. Biswas, and X. Koutsoukos. Designing distributed diagnosers for complex continuous systems. *IEEE Transactions on Automation Science and Engineering*, 2008b.
- I. Roychoudhury, M. Daigle, G. Biswas, and X. Koutsoukos. Efficient simulation of hybrid systems: An application to electrical power distribution systems. In *22nd European Conference on Modeling and Simulation*, pages 471–477, June 2008c.