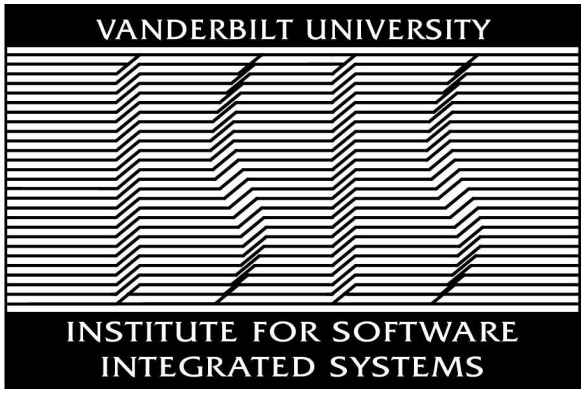


Institute for Software Integrated Systems
Vanderbilt University
Nashville, Tennessee, 37235

# A Comprehensive Diagnosis Methodology for Complex Hybrid Systems: A Case Study on Spacecraft Power Distribution Systems

Matthew Daigle, Indranil Roychoudhury, Gautam Biswas, Xenofon Koutsoukos,
Ann Patterson-Hine, and Scott Poll

**TECHNICAL REPORT**

ISIS-08-908

# A Comprehensive Diagnosis Methodology for Complex Hybrid Systems: A Case Study on Spacecraft Power Distribution Systems

Matthew Daigle†, Indranil Roychoudhury*, Gautam Biswas*, Xenofon Koutsoukos*, Ann Patterson-Hine†, and Scott Poll†

*Institute for Software Integrated Systems, Department of EECS, Vanderbilt University, Nashville, TN, USA

†NASA Ames Research Center, Moffett Field, CA, USA

matthew.j.daigle@nasa.gov, indranil.roychoudhury@vanderbilt.edu

***Abstract*—The application of model-based diagnosis schemes to real systems introduces many significant challenges, such as building accurate system models for heterogeneous systems with complex behaviors, dealing with noisy measurements and disturbances during system operation, and producing valuable results in a timely manner with limited information and computational resources. The Advanced Diagnostics and Prognostics Testbed (ADAPT), deployed at NASA Ames Research Center, is a representative spacecraft electrical power distribution system that embodies a number of these challenges for developing realistic diagnosis and prognosis algorithms. ADAPT contains a large number of interconnected components, along with a number of circuit breakers and relays that enable a number of different power distribution configurations. The system includes electrical dc and ac loads, mechanical subsystems, such as motors, and fluid systems, such as pumps. The system components are susceptible to different types of faults that include unexpected changes in parameter values, discrete faults in switching elements, and sensor faults. This paper presents Hybrid TRANSCEND, a comprehensive model-based diagnosis scheme to address these challenges. The scheme uses the hybrid bond graph modeling language to systematically develop computational models and algorithms for hybrid state estimation, robust fault detection, and efficient fault isolation. The computational methods are implemented as a suite of software tools that enables analysis and testing through simulation, diagnosability studies, and deployment on the experimental testbed. Simulation and experimental results demonstrate the effectiveness of this methodology in efficient diagnosis of heterogeneous components for an embedded system.**

***Index Terms*—Model-based diagnosis, electrical power distribution systems, distributed diagnosis.**

## I. INTRODUCTION

The increasing complexity of modern engineering systems has made the deployment of online health monitoring and diagnosis schemes a necessary and important challenge to ensure safe, reliable, and efficient operation of these systems. Model-based diagnosis schemes are the preferred approach because they allow for more general and robust diagnosis solutions [1]–[5]. However, deployment of these schemes on real systems presents significant challenges that include building accurate and reliable models, designing robust observers and fault detectors, and developing fault isolation schemes that produce valuable information in a timely manner with limited information and computational resources.

The Advanced Diagnostics and Prognostics Testbed (ADAPT), developed at NASA Ames Research Center [6], emulates spacecraft power storage and distribution systems, and is designed to provide an environment where researchers and practitioners can deal with a number of these challenges in developing and testing their diagnosis and prognosis algorithms. In developing our diagnosis schemes for ADAPT, we have to address a number of new and significant challenges in the context of a real-world application. Given that we are dealing with real system components, we have limited information and data to estimate and validate the parameters of our models. In addition, our system consists of heterogeneous subsystems that exhibit different behavior characteristics with different time constants. For example, the ADAPT system includes both dc and ac subsystems with corresponding loads. Some of the loads are purely electrical, some are electromechanical, e.g., motors and fans, and some, such as pumps, extend to the fluid domain. Another challenge is designing robust detectors that are sensitive to small fault magnitudes but maintain low false alarm rates in the presence of measurement noise and disturbances. When tracking and analyzing system behavior online, we are limited by the set of installed sensors, and the instrumentation that is in place for data collection and storage. Further, real systems exhibit a number of different kinds of faults in sensors, actuators, and the process, which may be represented as unexpected changes in model parameter values, or unexpected changes in the on/off modes of the switches that operate the system in different configurations, such as supplying power to different loads. Our model-based diagnosis approach has to account for different kinds of behaviors and different fault types in an integrated framework.

In order to address the challenges in diagnosis of real-world systems, we use the Hybrid TRANSCEND methodology, which is a comprehensive model-based approach for combined qualitative/quantitative diagnosis in hybrid systems [2]. For this work, the Hybrid TRANSCEND algorithms had to be extended to isolate discrete faults [7], and incorporate a distributed diagnosis framework [8] for developing separate diagnosers for the dc and ac subsystems. The two subsystem diagnosers operate in a coordinated fashion to produce the overall system diagnosis results in a timely manner. The overall implementation is incorporated into our Fault Adaptive

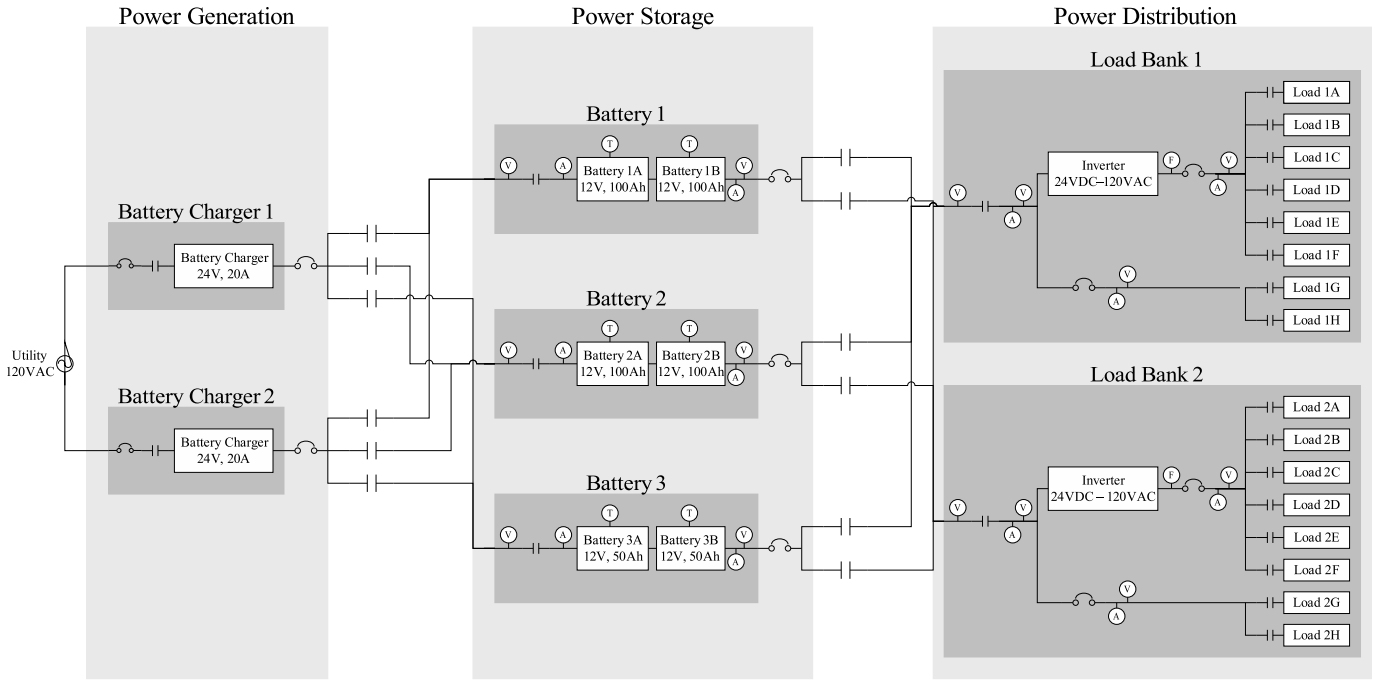


Fig. 1. Schematic diagram of ADAPT.

Control Technology (FACT) [9] paradigm, a model-integrated computing approach [10], which provides a framework for automatically synthesizing the runtime code of the deployed diagnosis system, thus making our model-based approach scalable and generalizable to a number of different applications.

This paper discusses a number of innovative and novel features of Hybrid TRANSCEND. First, we extend our hybrid diagnosis scheme for parametric faults to a combined parametric and discrete fault scheme. Second, to deal with the limited sensors in the ac subsystem, we develop a new method for deriving parametric and discrete fault signatures for ac measurements, which exhibit fault transients that occur faster than the sampling frequency of the sensors. Given the signatures, we can analyze the diagnosability of the system. Third, we develop a comprehensive methodology for combined diagnosis of dc and ac subsystems in a hybrid systems framework, which is achieved using an extension of our previous distributed diagnosis methods [8] for continuous systems based on system diagnosability analysis. Fourth, we illustrate the effectiveness of our approach by extensive experimental studies we have conducted, some on the hardware testbed, and some in a fully-developed simulation environment called VIRTUAL ADAPT [6].

The paper is organized as follows. Section II describes the ADAPT system and the challenges that it presents for developing real world model-based diagnosis solutions. Section III presents FACT, our model-integrated diagnosis tool-suite, and Section IV describes our modeling scheme. Section V discusses our approach to tracking complex hybrid system behaviors and online fault detection. Section VI describes our integrated framework for diagnosis of the heterogeneous components of the ADAPT testbed, and Section VII discusses the details of our online fault isolation scheme. Section VIII discusses our experimental results, and Section IX provides the conclusions and our directions for future work on real-world diagnosis applications.

## II. THE ADVANCED DIAGNOSTICS AND PROGNOSTICS TESTBED

The ADAPT system schematic, shown in Fig. 1, illustrates 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.

The current testbed operational infrastructure, shown in Fig. 2, contains a `User` component, which simulates a crew member and provides commands to the testbed, an `Antagonist` component, which injects faults and spoofs sensor data sent to the `User`, and a `Test Article` component, such as a diagnoser, which receives the data and commands issued by the `User` and determines the health of the system. The `Observer` component logs all system data in order to evaluate the performance of the test articles. A common communication interface between the testbed and the various components is supported through a publish/subscribe messaging server that operates at 2 Hz.

We have identified over 170 different faults that are of interest to ADAPT. The `Antagonist` can inject discrete faults by blocking or changing user commands to the testbed, and sensor faults by spoofing sensor data. Only a subset of the faults can be injected into the system, so the remainder of the faults are synthesized using an accurate simulator that we have

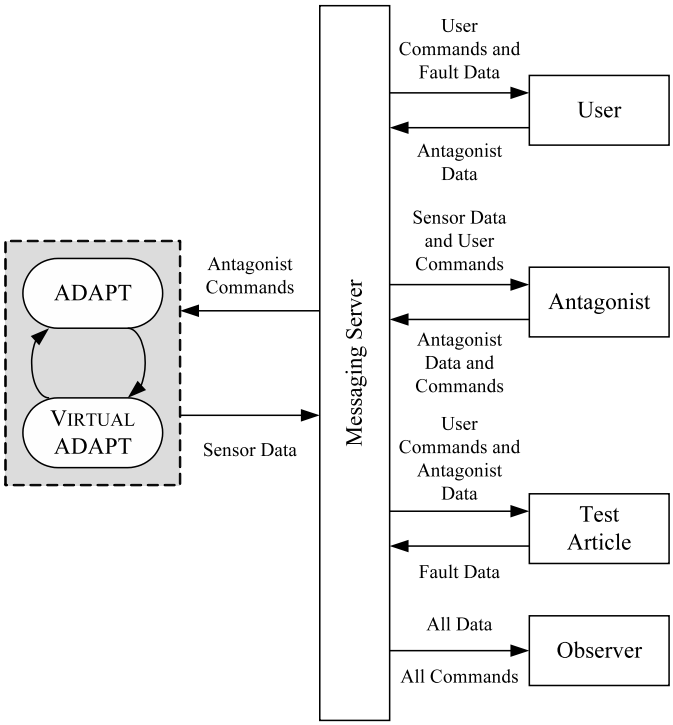


Fig. 2. Messaging architecture for ADAPT.

developed, called VIRTUAL ADAPT [6]. The `Antagonist` can use the simulator to realistically spoof sensor data based on simulated faulty scenarios. VIRTUAL ADAPT is also used as an offline version of the testbed for diagnoser design and diagnosis experiments. As indicated in Fig. 2, the simulation testbed includes the same interfaces as the actual testbed and replicates its behavior under nonfaulty and a number of faulty conditions.

### A. Challenges of Model-Based Diagnosis

Like other real-world systems, ADAPT presents a number of challenges to model-based diagnosis, including those associated with model development, system monitoring, and fault isolation.

*1) Model Development:* ADAPT can operate in a large number of configurations and contains a number of components, some of which have complex, nonlinear dynamic behaviors. The system contains over fifty relays and circuit breakers that can configure the system into different modes of operation. To avoid the computational intractability of dealing with more than $2^{50}$ modes of system operation, we require a modeling framework that specifies the system mode as a combination of the modes of a set of switching elements that operate at the component level. Other challenges to modeling include the derivation of accurate nonlinear models (e.g., the charging and discharging characteristics of the battery) from a relatively small number of available measurements, and dealing with fast-switching components (e.g., the inverter).

*2) System Monitoring:* A complete model-based approach to online diagnosis requires mechanisms for accurately tracking the dynamic behavior of the system in the presence of modeling errors, measurement noise, and disturbances in the system. This requires systematic analysis to achieve the proper tradeoff between sensitivity of detection and false alarm generation. Monitoring the behavior of ADAPT is particularly difficult, especially because the ac components operate at 60 Hz while the available measurements are only provided at a rate of 2 Hz by the data collection instrumentation.

*3) Fault Isolation:* Our fault isolation scheme has to deal with parametric process faults (e.g., unexpected changes in system parameter values, such as a resistance), additive sensor faults (e.g., bias in a sensor), and discrete faults (e.g., unexpected changes in system operating mode). ADAPT requires a combined diagnosis approach for simultaneous diagnosis of its dc and ac subsystems. The transient dynamics of faults in the dc components can be tracked at the sampling frequency used by the testbed instrumentation, i.e., 2 Hz. However, this sampling rate is insufficient to track fault transients in ac components, which occur at rates much faster than 60 Hz, the nominal operating frequency for these components.

## III. THE FAULT ADAPTIVE CONTROL TECHNOLOGY TOOL-SUITE

To address the challenges in developing useful diagnosis solutions for complex systems, we have developed the FACT tool-suite [9], which uses a model-integrated computing approach, where we can automatically synthesize simulation models, hybrid observers, and diagnoser code from component-based system models. We construct these models using graphical interfaces provided in the Generic Modeling Environment [10], which is a meta-modeling framework for specifying domain-specific modeling languages. We design model transformations for automatically synthesizing the code for the different components of the run-time application. This greatly simplifies the entire development process, from development and testing the initial prototypes to generating the diagnosers for the runtime environment.

Our approach to hybrid systems modeling is based on the hybrid bond graph (HBG) language [11]. The HBG language supports energy-based topological modeling of physical processes. System components are modeled as HBG fragments, which are connected through energy and signal ports to define the complete system behavior. Switching is defined locally at the component level, and given a particular mode, the system equations can be derived automatically from the model configuration [12]. Therefore, pre-enumeration of the complete system modes is not required during system design.

The overall FACT schematic for diagnosis applications appears in Fig. 3. System behavior is tracked dynamically using a hybrid observer, which is derived automatically from the HBG model [2]. Given the inputs and system measurements, the observer computes the estimated outputs, which provides the nominal reference for diagnosis up to the point of fault detection. The observer scheme helps to account for model uncertainty and sensor noise. Fault detection and symbol generation are defined as tests of statistical significance, implemented using the Z-test [13] and a sliding window technique [14]. Fault detectors are tuned to adjust sensitivity in order to minimize false alarms and missed detections.

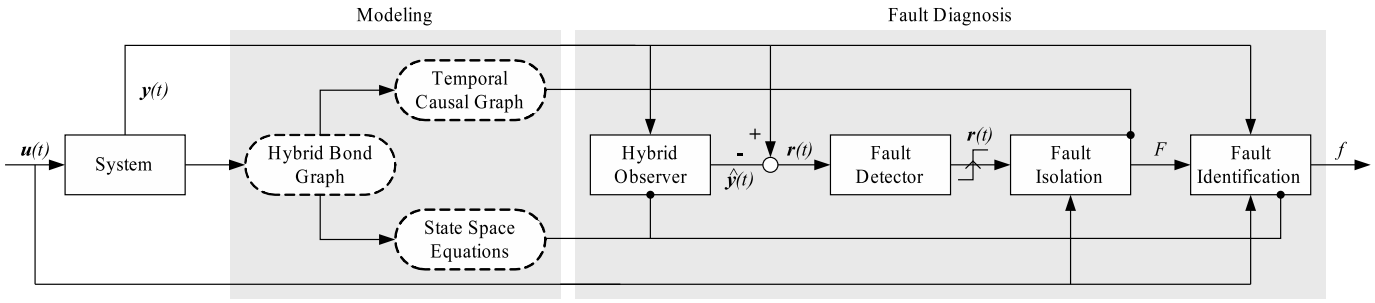


Fig. 3. Schematic diagram of FACT.

FACT implements the Hybrid TRANSCEND [1], [2] methodology for combined qualitative/quantitative diagnosis in hybrid systems. Fault isolation is based on qualitative analysis of deviations from nominal behavior in the measurements caused by faults. This can be followed by a fault identification scheme to estimate the fault magnitude [2], [15]. In recent work, we have extended our qualitative diagnosis scheme to deal with both parametric and discrete faults [7], and have developed a distributed diagnosis approach for continuous systems [8]. The extended diagnosis schemes and the distributed diagnosis approach provide an innovative framework for developing a comprehensive model-based diagnosis methodology for spacecraft power distribution systems.

## IV. MODEL DEVELOPMENT

Our component-based models of hybrid physical systems are based on the HBG modeling language [11]. HBGs extend bond graphs [12], and are particularly suitable for diagnosis because they incorporate causal and temporal information required for deriving and analyzing fault transients. In bond graphs, components are represented by vertices, and bonds, drawn as half arrows, capture ideal energy connections between the components. Associated with each bond are two variables: *effort*, $e$, and *flow*, $f$, and the product $e \cdot f$ defines the rate of energy transfer through the bond. In the electrical domain, effort and flow map to voltage and current, respectively. 1-junctions represent series connections (where all $f$ values are equal and $\sum e = 0$), and 0-junctions represent parallel connections (where all $e$ values are equal and $\sum f = 0$). Component behaviors are modeled as resistances, $R$, which capture energy dissipation in the system ($e = Rf$), capacitances, $C$, ($\dot{e} = \frac{1}{C} f$) and inductances, $I$, ($\dot{f} = \frac{1}{I} e$), which capture energy storage functions, and sources of flow ($Sf$) and effort ($Se$), which model the flow of energy into and out of the system. Nonlinearities in the system are modeled as time varying parameter values, which are defined as functions of system variables using modulating elements. The constituent equations of the bond graph elements and junctions define a set of differential algebraic equations that together describe the continuous system behavior.

HBGs [11] extend bond graphs by introducing *switching junctions*, which act as ideal switches, enabling a junction to be in either the on or the off mode of operation. Off 1-junctions behave as sources of zero flow. Similarly, off 0-junctions act as sources of zero effort. When on, switching junctions behave as normal junctions. The switching behavior is defined by a *control specification* (CSPEC), modeled as a finite automaton [2], [11], where the 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.

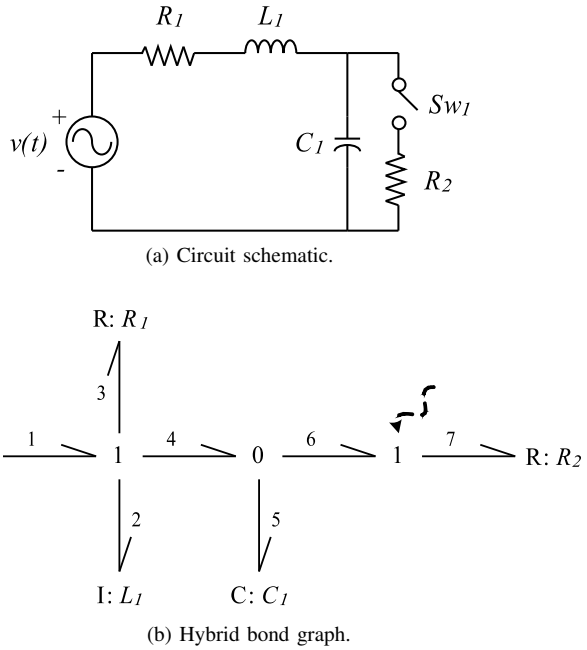


(a) Circuit schematic.

(b) Hybrid bond graph.

Fig. 4. Switched circuit example.

Consider the example electrical circuit shown in Fig. 4. The circuit consists of an ac source $Se$ with voltage, $v(t)$, resistors $R_1$ and $R_2$, inductor $L_1$, and capacitor $C_1$. The series and parallel connections in the circuit are captured using the $1-$ and $0-$junctions, respectively. The switch, $Sw_1$, is assumed to be ideal, and hence, is modeled by a switching 1-junction, representing a series connection that can be on or off. The switching junction is denoted by the dashed arrow in Fig. 4b.

In this work, we focus on the diagnosis of single, abrupt, persistent faults in hybrid systems. 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. Sensor faults are modeled as additive parametric faults. Discrete faults correspond to differences be-

tween the actual and expected state of a switching component in the HBG model, and are modeled using unobservable fault events in the CSPECs [7].

The ADAPT model includes component models for the batteries, inverters, relays, circuit breakers, dc loads that include simple circuits, and ac loads that include fans, pumps, and light bulbs. These components can be composed to form different configurations or subsystems of ADAPT. Faults can be introduced into these components by changing their nominal parameter values or by creating an unexpected change in the mode of their switching junctions. From the HBG models, we can derive a hybrid state-space formulation which forms the basis for the hybrid observer and the parameter estimation scheme, a reconfigurable block diagram which forms the basis of our simulation models, and the temporal causal graph (TCG), which forms the basis for performing qualitative fault isolation from transients.

### A. Generating Simulation Models

We use the HBG to automatically generate simulation models of the system for offline diagnosis experiments. Each mode of the HBG corresponds to a bond graph model that defines the continuous behavior within a mode. The computational model for each mode (e.g., state-space equations or signal flow graphs) can be derived systematically from this BG model using well-defined methods [12]. We have developed efficient methods for incrementally generating the computational model after a mode change occurs [16], [17], which offers significant advantages for large hybrid systems like ADAPT, because it avoids unnecessary pre-enumeration of all system modes. This scheme has been used to develop the VIRTUAL ADAPT simulation testbed [6] mentioned earlier. VIRTUAL ADAPT, implemented in Matlab Simulink [18], includes external wrappers to communicate to the messaging server of ADAPT. In more recent work in progress, this method is also being employed to efficiently regenerate the state-space equations for the observer when mode changes occur.

### B. Temporal Causal Graphs

Our model for qualitative fault diagnosis, the temporal causal graph (TCG), is derived from the bond graph model of the system. In addition to capturing the system dynamics, the model explicitly captures the propagation of both parametric and discrete fault effects on other system variables, which include the system measurements [1], [19]. 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 ($+1$) or inverse ($-1$) proportionality, integral ($dt$), and parametric dependencies (e.g., $1/R_1$). The algebraic relations imply instantaneous propagation effects, whereas the integral edges imply a delay in the propagation. Links from discrete fault events to variables may have $\pm 1$ labels and additional $N$ and $Z$ labels, if the fault causes the variable value to go from zero to nonzero or from nonzero to zero, respectively. The directionality of these edges is determined by *causality*, i.e., the preferred order for computing the effort and flow variable values. The causal directions are derived from the bond graph model [12].

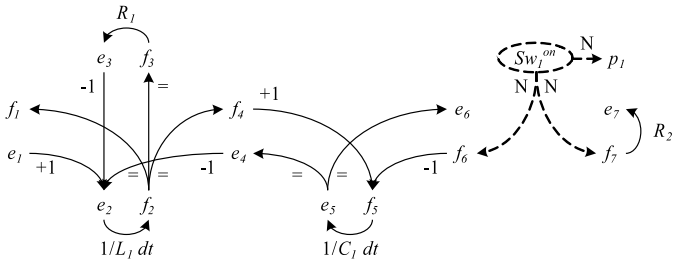


Fig. 5. Example TCG for the nominal mode where the switch is off.

The TCG for the circuit example is given in Fig. 5 for the mode where the switch is off. If the switch turns on unexpectedly (represented by fault event $Sw_1^{on}$), then the flow of current through the switch will go from zero to a nonzero value, which then affects the values of other variables in the system. A change in a parameter value caused by a fault, e.g., $R_1^+$, cannot cause such discrete changes from zero to nonzero values or vice versa.

## V. MONITORING AND FAULT DETECTION

As illustrated in the FACT diagnosis scheme in Fig. 3, the fault detector triggers the fault isolation and identification modules. The robust fault detection scheme combines a hybrid observer for tracking nominal system behavior and a statistical hypothesis testing scheme for robust fault detection. The two components take into account measurement noise and modeling errors while keeping the false alarm rate low.

### A. Hybrid Observer

The hybrid observer combines the use of an extended Kalman filter (EKF) for tracking continuous system behavior, and automata for tracking the on/off mode of every switching junction in the HBG model and making mode transitions when indicated by the CSPECs [2]. Mode changes produce a reconfiguration in the HBG model. As a result, the state-space equations are recomputed automatically, the EKF is updated, and the tracking of continuous behavior resumes. The EKF scheme assumes the modeling errors and measurement noise are uncorrelated Gaussian with zero mean, therefore, the two covariance matrices that represent the modeling error and measurement noise are assumed to be diagonal with known variance values.

The observer receives updated measurements at a rate of 2 Hz, the frequency of the measurement system. However, the observer must receive measurements at kHz rates to accurately track ac system behavior, due to the controlled fast-switching behavior of the inverter. To accommodate this, we execute the model at the kHz rate, but use the observers to update state estimates at the 2 Hz data collection rate. All of the dc measurements contribute to the state update function in the EKF. For the ac subsystems, we use the instantaneous values of ac voltage and currents from the model as measured values, since the rms and phase sensor reading are based on average

computations and cannot be directly used in the EKF update functions.

### B. Fault Detection

Our fault detection scheme employs separate fault detectors for each measurement, which allows each detector to be tuned individually to achieve maximum sensitivity for a given signal. For each measurement we define the residual as $r(t) = y(t) - \hat{y}(t)$, where $y(t)$ is the measurement signal, and $\hat{y}(t)$ is the estimated output signal generated by the hybrid observer. The fault detection scheme employs the Z-test to look for nonzero residual signals [14]. The Z-test requires that the sample mean and standard deviation of a given population be known [13]. We estimate the population standard deviation and sample mean using a sliding window technique that is illustrated in Fig. 6. A small sliding window (e.g., 5 samples), $W_1$, is used to estimate the current mean $\mu_r(t)$ of a residual signal, i.e.,

$$\mu_r(t) = \frac{1}{W_1} \sum_{i=t-W_1+1}^{t} r(i).$$

The variance of the nominal residual signal is computed using a window $W_2$ preceding $W_1$, where $W_2 \gg W_1$ (e.g., 100 samples). $W_2$ is offset by $W_1$ by a buffer $W_{delay}$ (e.g., 50 samples), to ensure that $W_2$ does not contain any samples after fault occurrence:

$$\mu'_r(t) = \frac{1}{W_2} \sum_{i=t-W_1-W_{delay}-W_2+1}^{t-W_1-W_{delay}} r(i)$$

$$\sigma^2_r(t) = \frac{1}{W_2} \sum_{i=t-W_1-W_{delay}-W_2+1}^{t-W_1-W_{delay}} (r(i) - \mu'_r(t))^2.$$

Given a pre-specified confidence level, $\alpha$, (e.g., $\alpha = 95\%$) tables provide the bounds $z^-$ and $z^+$ for a two-sided Z-test. The thresholds for the fault-no fault decision, $\varepsilon^-_r(t)$ and $\varepsilon^+_r(t)$, are computed as:

$$\varepsilon^-_r(t) = z^- \frac{\sigma_r(t)}{\sqrt{W_1}} + E \quad (1)$$

$$\varepsilon^+_r(t) = z^+ \frac{\sigma_r(t)}{\sqrt{W_1}} - E, \quad (2)$$

where $E$ is a modeling error term. A computed mean value $\mu_r(t)$ that lies outside of the thresholds at time $t$ implies a fault. For practical applications, parameters $W_1$, $W_2$, and $W_{delay}$, the confidence level $\alpha$, and the modeling error term, $E$ of the fault detector have to be tuned experimentally to optimize performance (i.e., minimize false alarms while keeping detection sensitivity high) [14].

## VI. Diagnoser Design

Our approach to diagnosing faults in power distribution systems, such as ADAPT, combines schemes for diagnosis from 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 describe how we derive fault signatures for the two cases. Given the signatures, we analyze system diagnosability, and based on this analysis, develop the overall system diagnoser as two interacting diagnosers: the dc subsystem diagnoser, and the ac subsystem diagnoser, using our distributed diagnosis methodology.

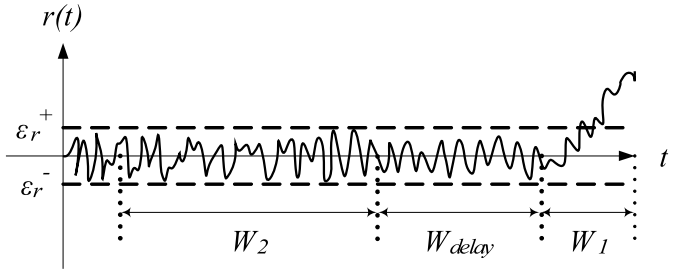


Fig. 6. Sliding windows in the fault detection scheme.

### A. Fault Signatures for DC Measurements

For the dc measurements, the fault signatures are derived from the transients that result after fault occurrence. Assuming that the system output is continuous and continuously differentiable except at the points of fault occurrence and mode changes, the transient response after a fault occurrence 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 [1], [15]. In TRANSCEND, these signatures are represented in a qualitative form: + (increase), − (decrease), and 0 (no change) in the magnitude and derivatives of the residual signal. 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 proved that the first change and subsequent slope provide all of the discriminatory evidence for qualitative fault isolation in dynamic systems [15]. Therefore, our fault signatures include two symbols: the magnitude change and slope of the residual signal. For discrete fault analyses, fault signatures have been extended to include a third symbol that indicates if a fault causes zero to nonzero or nonzero to zero value change in a measurement. Discrete faults cause mode changes at junctions, and, as a result, variable values linked to this junction may go from nonzero to zero abruptly (for a junction turning off) or go from zero to nonzero abruptly (for a junction turning on). The symbols N, Z, and X, imply zero to nonzero, nonzero to zero, or no discrete change behavior in the measurement from the estimate [7].

Fault signatures representing the transient behavior due to parametric and discrete faults are defined as follows for our three-symbol representation.

**Definition 1** (Fault Signature from Transients)**.** A *fault signature* for a fault $f$ in a system mode $q$ defines the qualitative effect in magnitude, slope, and discrete change in measurement $m$ due to the occurrence of $f$.

Fault signatures are derived for each hypothesized fault $f$ in mode $q$ by performing a forward propagation function on the TCG [1], [7]. Signatures for the circuit example are

TABLE I
FAULT SIGNATURES FROM TRANSIENTS FOR THE CIRCUIT WITH THE SWITCH ON

| Candidate | $V_{R1}$ | $I_{R_2}$ |
|---|---|---|
| $C_1^+$ | 0+,X | −+,X |
| $C_1^-$ | 0−,X | +−,X |
| $L_1^+$ | −+,X | 0−,X |
| $L_1^-$ | +−,X | 0+,X |
| $R_1^+$ | −+,X | 0−,X |
| $R_1^-$ | +−,X | 0+,X |
| $R_2^+$ | 0−,X | −+,X |
| $R_2^-$ | 0+,X | +−,X |
| $Sw_1^{off}$ | 0−,X | −⋆,Z |

given in Table I, assuming the voltage source is dc instead of ac, and variable values are nominally positive, where the measurements are the voltage across $R_1$, $V_{R_1}$, and the current through $R_2$, $I_{R_2}$. The table shows that the system is not diagnosable with the selected measurements, because faults in $L_1$ and $R_1$ cannot be distinguished in this mode.

### B. Fault Signatures for AC Measurements

As discussed, analyzing fault transients in the ac domain, where the components operate at 60Hz, would require sampling of measurements at rates greater than 120 Hz. This would make the online analysis of fault signatures computationally infeasible. Besides, as discussed earlier, the ADAPT system is equipped only with rms and phase sensors for the voltages and currents in the ac subsystem, and these sensors generate their output at 2 Hz. Therefore, we resort to a steady state analysis of the system bond graph to derive fault signatures for the ac measurements given faults in the ac components. The HBG model of the system provides the relations between the component parameter values and the ac measurements. The parameters for the $R$, $C$, and $I$ elements are replaced by their complex impedance representations in the ac domain. Given the frequency, $\omega = 2\pi f$, ($f$ is frequency in hertz) 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 $Z_L = j\omega L$. The bond graph structure then provides the series/parallel relations between the voltages across and currents through the different elements in the circuit (Kirchhoff's laws), and, like before, by combining the constitutive relations of the elements and the junction equations, we can derive the voltage and current variable relations in symbolic form. By algebraic manipulation, we get the symbolic form of the expressions for the rms and phase of these measurements. Computing the partial derivative of a measurement with respect to a particular fault variable, provides an expression for the fault signature in symbolic form. After substituting nominal values of all other parameters, if the sign of this partial derivative is always positive (negative) for the considered fault magnitudes, then the corresponding fault signature is defined to be a + (−). If the sign cannot be uniquely determined, the ambiguity is represented using a $\star$. 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 see what effects the mode changes will have.

Fault signatures for representing steady-state changes are defined as follows for our two-symbol representation.

**Definition 2** (Fault Signature by Steady-state Analysis)**.** A *fault signature* for a fault $f$ in a system mode $q$ defines the qualitative effect in magnitude and discrete change in measurement $m$ due to the occurrence of $f$.

To illustrate the approach, we consider the circuit given in Fig. 4a. The measured signals are the voltage across $R_1$, $v_{R_1}(t)$, and the current through $R_2$, $i_{R_2}(t)$. 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, $R_1 = 1\ \Omega$, and $R_2 = 2\ \Omega$. 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, $\omega$, the impedances are $Z_{L1} = j\omega L_1$, for the inductor $L_1$, $Z_{C_1} = \frac{1}{j\omega C_1}$ for the capacitor $C_1$, $Z_{R1} = R_1$, for the resistor $R_1$, and $Z_{R_2} = R_2$, for the resistor $R_2$. Using the bond graph model as described above, we derive the symbolic expressions describing the measurements as a function of the inputs and the impedances, and compute the fault signature matrix for each mode (see Table II for fault signatures in mode $q_1$):

$$
\begin{aligned}
v_{R_1} &= \frac{vR_1}{Z_{eq}} \\
i_{R_2} &= \begin{cases} 0, & \text{for mode } q_0 \\ \dfrac{vZ_{C_1,R_2}}{Z_{eq}R_2}, & \text{for mode } q_1 \end{cases}
\end{aligned}
$$

where

$$
\begin{aligned}
Z_{C_1,R_2} &= \left(j\omega C_1 + \frac{1}{R_2}\right)^{-1} \\
Z_{eq} &= \begin{cases} j\omega L_1 + R_1 + \frac{1}{j\omega C_1}, & \text{for mode } q_0 \\ j\omega L_1 + R_1 + Z_{C_1,R_2}, & \text{for mode } q_1 \end{cases}
\end{aligned}
$$

Using these expressions, we can calculate the fault signature matrix for each mode. The signatures for mode $q_1$ are shown in Table II. In some cases, the direction of change in measurement values depends on fault magnitude. For example, $C_1^+$ will always cause a decrease in the rms value of $V_{R_1}$, but $C_1^-$ may cause either an increase or decrease in $V_{R_1}$ depending on its magnitude, as shown in Fig. 7. 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 $\star$ (see Table II). Discrete faults do not produce ambiguous signatures. For example, when the switch is on, the rms value of $I_{R_2}$ 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 $I_{R_2}$, and it will go to zero. This is represented by the fault signature −,Z.

TABLE II
FAULT SIGNATURES FOR AC MEASUREMENTS FOR THE CIRCUIT WITH THE SWITCH ON

| Fault | $V_{R_1}$ | $\phi_{V_{R_1}}$ | $I_{R_2}$ | $\phi_{I_{R_2}}$ |
|---|---|---|---|---|
| $C_1^+$ | −, X | −, X | −, X | −, X |
| $C_1^-$ | ⋆, X | +, X | +, X | +, X |
| $L_1^+$ | −, X | −, X | −, X | −, X |
| $L_1^-$ | +, X | +, X | +, X | +, X |
| $R_1^+$ | +, X | +, X | −, X | +, X |
| $R_1^-$ | −, X | −, X | +, X | −, X |
| $R_2^+$ | +, X | −, X | −, X | −, X |
| $R_2^-$ | −, X | ⋆, X | +, X | +, X |
| $Sw_1^{off}$ | −, X | +, X | +, Z | −, Z |

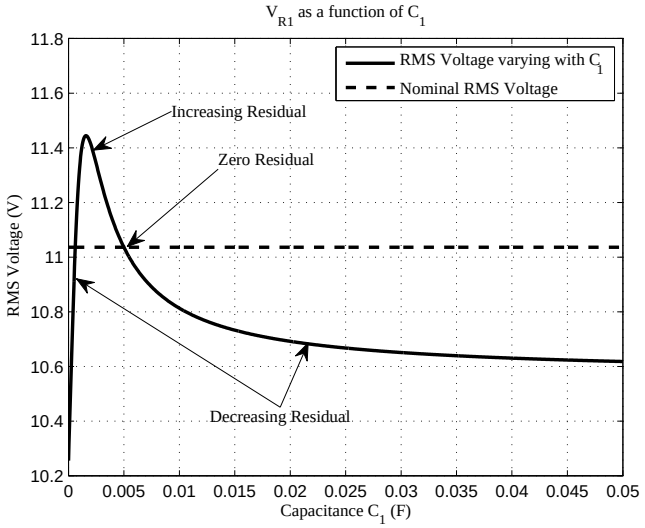


Fig. 7. $V_{R_1}$ rms value as a function of $C_1$ magnitude.

### C. Distributed Diagnoser Design

Distributed diagnosers make the overall diagnosis approach more efficient, because they partition the diagnosis task into smaller subtasks [20]–[22]. In [8], we presented an approach for designing distributed diagnosers for continuous systems whose subsystem structure is given (Algorithm 1 in [8]). In this paper, we extend this approach to hybrid systems, which allows us to decouple the diagnosers for the dc and ac subsystems of ADAPT. Our objective is to decompose the overall diagnosis task into smaller subtasks performed by local diagnosers such that the local diagnosers generate globally correct diagnosis results while minimizing the number of measurements required to be communicated amongst the local diagnosers.

To generate distributed diagnosers for hybrid systems, we require the fault signatures for each mode of the system, which are generated using the techniques previously discussed. 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 [19]. 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 $\kappa$ 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 $\widetilde{M}_i$.

**Definition 3** (Global Diagnosability)**.** A set of faults, $F_i \subseteq F$ is *globally diagnosable* by $\widetilde{M}_i \subseteq M$ if $\widetilde{M}_i$ can uniquely isolate every fault, $f \in F_i$, from all other faults in $F$ for every possible sequence of mode transitions.

Given the set of available measurements, global diagnosability is not always attainable in real-world systems, and, in fact, we will show in Section VIII that ADAPT is not globally diagnosable. We first analyze the diagnosability of the system. If the system is not globally diagnosable for a set of measurements, we define the notion of "aggregate faults". An aggregate fault 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.

Given $F_i$ and $\widetilde{M}_i$, we construct a *local diagnoser* [8], $\mathcal{D}_{F_i,\ \widetilde{M}_i,Q}$, for each subsystem. By ensuring that each $\widetilde{M}_i$ is minimal, the local diagnosers share minimal information with one another.

The procedure for designing diagnosers for a partitioned hybrid system is presented in Algorithm 1. Our design goal is to minimize the number of additional measurements, while ensuring that each subsystem is globally diagnosable. For each subsystem $S_i$, we assign to $remFaults_i$ the faults in $F_i$ that are not globally diagnosable using measurements in $M_i$. The search is simplified by defining a notion of proximity among subsystems and using this information to prioritize the selection of additional measurements for a local diagnoser. We represent the system, $S$, as a graph of connected subsystems. The proximity $d$ is defined as the minimum path length from $S_g$ to $S_h$. The search for additional measurements starts from closer subsystems. When $remFaults_i$ is not empty, we start by assigning $\widetilde{M}_i$ equal to $M_i$, and generating a working measurement set $\widetilde{M}_i^{d\leq 1}$ by pooling in measurements from all subsystems, $S_l$, at a distance $d \leq 1$ from subsystem $S_i$, i.e., $\widetilde{M}_i^{d\leq 1} = \bigcup_l M_l$. We select additional measurements from $\widetilde{M}_i^{d\leq 1} - M_i$ to reduce the number of faults in $remFaults_i$. When different measurement combinations provide the same reductions, we pick the measurement set $\widetilde{M}_i$ that adds minimal number of external measurements to $M_i$ while making the maximum number of faults in $remFaults_i$ globally diagnosable. The set $\widetilde{M}_i$ is expanded, and $remFaults_i$ is reduced to a smaller set. If $remFaults_i$ is non-empty, $d$ is incremented, and the procedure is repeated till $remFaults_i$ is empty. At this point, we have the local diagnoser $\mathcal{D}_{F_i,\ \widetilde{M}_i,Q}$. We will show the results of this algorithm on ADAPT in Section VIII, where we define as subsystems as the dc subsystem, $S_{dc}$, and the ac subsystem, $S_{ac}$.

**Algorithm 1** Designing Diagnosers for a Partitioned System

**Input:** $\kappa$ 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 $\mathcal{D}_{F,M_i,Q}$.
  $\delta = 1$; $\widehat{M}_i = M_i$
  **while** $remFaults_i \neq \varnothing$ **do**
    identify measurement set $\widehat{M}_i$ from measurements of subsystems $S_i$ at a distance $d \leq \delta$ that isolates maximal $r \in remFaults_i$, and $\widetilde{M}_i - \widehat{M}_i$ is minimal.
    $\widetilde{M}_i = \widetilde{M}_i \cup \widehat{M}_i$
    $remFaults_i = remFaults_i - r$
    **if** $remFaults_i \neq \varnothing$ **then**
      $\delta = \delta + 1$
  construct $\mathcal{D}_{F_i,\ \widetilde{M}_i,Q}$

## VII. Online Fault Isolation

The distributed diagnosers are designed offline. We use our tracking scheme for online fault detection and qualitative fault isolation, where observed measurement deviations are matched to predicted fault signatures to isolate faults. In the following, we describe our method for robust symbol generation, and the online signature matching scheme for qualitative fault isolation.

### A. Symbol Generation

For each dc measurement, we extract the magnitude and slope of the deviation, as well as the discrete change feature. For each ac measurement, we use only the first change and the discrete change behavior. The changes are abstracted symbolically to +, 0, −, N, Z, and X symbols, and the computed symbols form the observed fault signatures that are matched to predicted signatures during fault isolation.

A robust method based on the Z-test is used for computing the symbolic features of the residual signal. If the measurement residual, $r(t) = y(t) - \hat{y}(t)$, is greater than $\varepsilon_r^+(t)$ (less than $\varepsilon_r^-(t)$), we assign a + (−) to the magnitude value for the residual.

The calculation of the slope of a measurement deviation starts with the estimation of the initial residual value, $\mu_{r_0}(t_d)$, after fault detection by computing the average of the residual samples over a small window ,$W_3$, i.e.

$$\mu_{r_0}(t_d) = \frac{1}{W_3} \sum_{i=t_d}^{t_d+W_3-1} r(t_d + i).$$

Again using the Z-test, the slope of the residual (i.e., the measurement) is determined over another small, but larger window (e.g., 15 samples) after the end of the smaller window [14]. The mean value of the measured and expected values of the signal after fault detection are given by:

$$\mu_{r_d}(t_d{+}t) = \begin{cases} \dfrac{\left(\sum_{i=t_d}^{t_d+W_n-1} r(t_d+i)\right)}{W_n} - \mu_{r_0}, & W_n > W_3 \\ 0, & W_n \leq W_3 \end{cases}.$$

It is assumed that the variance of the residual does not change due to the occurrence of the fault, i.e., $\sigma_r^2(t) = \sigma_r^2(t_d)$ for all $t \geq t_d$. The variance of $\mu_{r_d}$ is $\sigma_{r_d}^2(t_d{+}t) \approx \sigma_r^2/W_n$, while the variance of $\mu_{r_0}$ is $\sigma_{r_0}^2 \approx \sigma_r^2/W_3$. That is, the uncertainty of the initial residual value depends on the noise and $W_3$, while the uncertainty of the mean estimate depends on the noise and the number of samples used in the calculations. Using a confidence value $\alpha$ and the corresponding $z^+$ and $z^-$ values, the + slope symbol is generated when:

$$\mu_{r_d} > z^+\sigma_r \left(\frac{1}{\sqrt{W_3}} + \frac{1}{\sqrt{W_n}}\right).$$

Similarly, the − slope symbol is generated when

$$\mu_{r_d} < -z^+\sigma_r \left(\frac{1}{\sqrt{W_3}} + \frac{1}{\sqrt{W_n}}\right).$$

The size of the window used to calculate the mean, $W_n$, is increased until the symbol is successfully generated, or $W_n$ becomes larger than a pre-specified limit, at which the slope is reported as 0, implying that the true slope is either zero or unknown but very small.

To compute the discrete change feature, we do not use the residual, but use the observed and estimated values of the signal, by again computing the mean of the measured signal, $y(t)$, and the mean of the estimate, $\hat{y}(t)$, over a small window, $W_{dc}$:

$$\begin{aligned} \mu_y(t_d) &= \frac{1}{W_{dc}} \sum_{i=t_d}^{t_d+W_{dc}-1} y(i) \\ \mu_{\hat{y}}(t_d) &= \frac{1}{W_{dc}} \sum_{i=t_d}^{t_d+W_{dc}-1} \hat{y}(i), \end{aligned}$$

where $t_d$ is the time of fault detection. We wish to determine whether each signal belongs to a population with zero mean, and choose the variance of the population to be the variance of the residual $r(t)$, $\sigma_r^2(t)$, as a good approximation of the true variance of the zero-mean distribution. Here the thresholds are computed as:

$$\begin{aligned} \varepsilon_{y_{dc}}^+ = \varepsilon_{\hat{y}_{dc}}^+ &= z^+\frac{\sigma_r(t_d)}{\sqrt{W_{dc}}} + E_{dc} \\ \varepsilon_{y_{dc}}^- = \varepsilon_{\hat{y}_{dc}}^- &= z^-\frac{\sigma_r(t_d)}{\sqrt{W_{dc}}} + E_{dc}, \end{aligned}$$

where $E_{dc}$ is a modeling error term. These thresholds are the same for fault detection (Equations 1 and 2), only they are computed for $y(t)$ and $\hat{y}(t)$ rather than $r(t)$. If $\mu_y(t_d)$ is outside its bounds, we say it is nonzero, otherwise we say it is zero. Similarly, if $\mu_{\hat{y}}(t_d)$ is outside its bounds, we say it is nonzero, otherwise we say it is zero. If the estimate is nonzero and the measurement is zero, we report Z, and if the estimate is zero and the measurement is nonzero, we report N, otherwise, we report X.

### B. Distributed Fault Isolation

Observed fault signatures computed using symbol generation are matched to predicted fault signatures to isolate faults. Each local diagnoser, e.g., the dc and ac diagnosers, obtains the symbols for its own measurements. Inconsistent faults are eliminated, and consistent faults retained in the set of

hypothesized faults. A globally correct diagnosis result is reached when: (*i*) *all* measurements for a local diagnoser have deviated and the fault hypothesis set is reduced to a singleton fault set, or, (*ii*) a local diagnoser's hypothesis set is reduced to a singleton but all of its measurements have not deviated, *and* all other diagnosers produce a *null hypothesis*, i.e., their candidate sets are empty [8].

Mode changes are handled using the techniques presented in [2]. If a controlled mode change occurs, such as a relay turning on or off, the faults signatures for the new mode are used, and consistent faults must match future measurement deviations for the current mode. If an inconsistency is obtained, autonomous mode changes are hypothesized, such as circuit breakers tripping, then faults in the hypothesized modes should be consistent with the predictions in the hypothesized modes.

## VIII. Experimental Results

We choose a subset of components in the ADAPT system to demonstrate our approach. This subset includes one of the lead-acid batteries, two DC loads, an inverter, and two ac loads. The models of the dc components can be found in [19], and the models of the ac components can be found in [23]. A schematic of the subsystem is given in Fig. 8. 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 two dc loads connected to the battery are purely electrical, while the ac loads include a light bulb and a large fan. In addition, we also consider four relays, two of which connect the dc loads to the battery, while the remaining two connect the ac loads to the inverter. 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, $\phi$, the temperature of the light bulb, $T_{bulb}$, the rotational speed of the fan, $\omega_{fan}$, the currents through the two dc loads, $I_{L1}$ and $I_{L2}$, and the battery voltage and current, $V_B$ and $I_B$.

The fault signatures for the mode where all loads are active is given in Table III. From the table, we can see that the system is not globally diagnosable, because $Sw_3^{off}$ and $R_{bulb}^{+}$ cannot be distinguished. We form an aggregate fault from these two individual faults to apply the diagnoser design algorithm described in Section VI. To do this, we consider two subsystems (see Fig. 8), (*i*) the dc subsystem, which contains the battery, the two dc loads and the two relays, $Sw_1$ and $Sw_2$, and (*ii*) the ac subsystem, which contains the inverter, the ac loads, and the relays connecting these loads to the inverter, $Sw_3$ and $Sw_4$. The dc subsystem fault list, $F_{dc}$, includes resistances of the dc loads, $R_{L1}$ and $R_{L2A}$, the capacitance and resistance of the battery, $C_0$ and $R_1$, and faults in the switches, $Sw_1$ and $Sw_2$. The dc measurements, $M_{dc}$, include $I_{L1}$, $I_{L2}$, $V_B$, 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 in the switches, $Sw_3$, and $Sw_4$. The ac measurements, $M_{ac}$, include $V_{rms}$, $I_{rms}$, $\phi$, $T_{bulb}$, and $\omega_{fan}$.

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 III. If a dc fault occurs, no deviations will be observed on any of the ac measurements, therefore, the ac diagnoser will not isolate any ac 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 $I_{rms}$, (this is 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 automonmous 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., $\widetilde{M}_{dc} = \{V_B, I_B, I_{L1}, I_{L2}, V_{rms}, I_{rms}\}$. 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.

### A. Simulation Results

We first present diagnosis results obtained on the simulation testbed Virtual ADAPT. For this study, we did not have access to the phase sensor readings, therefore, we investigate the effectiveness of our extensions to ac diagnosis in simulation. We used the simulation model to provide the nominal reference for fault detection and symbol generation. For this set of experiments, we inject faults into the configuration where both ac loads and the first dc load are all online. For the fault detectors, we selected $W_1 = 5$, $W_2 = 100$, $W_{delay} = 50$, $W_3 = 3$, $W_n = 20$, and $\alpha = 99.97\%$. We chose $E = 0$ for all measurements except $I_B$, where $E = 0.2$, and $\phi$, where $E = 0.0001$.

The results are summarized in Table IV. In the table, $t_f$ represents the time of fault occurrence, $t_d - t_f$ is the delay in fault detection, and $t_i - t_f$ is the time to isolate the fault, which is given as the point at which a diagnoser last reduces its fault set. 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. Note that the fault $R_{bulb}^{+}$, and increase in the bulb resistance, and $Sw_3^{off}$, a fault where $Sw_3$ is stuck off, could not be distinguished, which was predicted using diagnosability analysis. In the following, we step through the reasoning of the distributed diagnosers for two interesting scenarios.

The first scenario represents a fault which is a $5\%$ decrease in the bulb resistance, $R_{bulb}^{-}$, at 100 s. The relevant measurement plots corresponding to this scenario are shown in Fig. 9. This change results in an increase in the phase difference, which is detected by the ac diagnoser at 100.5 s, and it

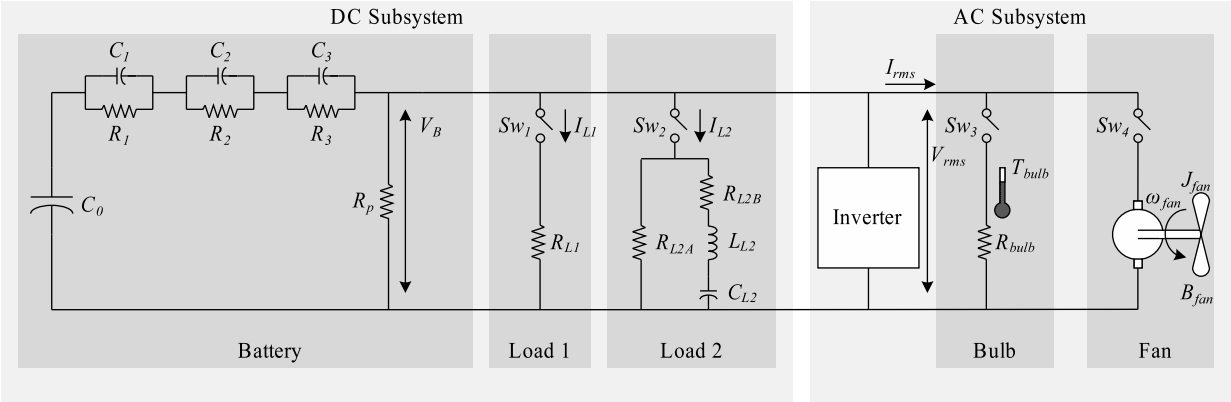


Fig. 8. Selected subset of ADAPT.

TABLE III
FAULT SIGNATURES FOR THE MODE WITH ALL LOADS ON

| Fault | DC Measurements | | | | AC Measurements | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | $V_B$ | $I_B$ | $I_{L1}$ | $I_{L2}$ | $V_{rms}$ | $I_{rms}$ | $\phi$ | $T_{bulb}$ | $\omega_{fan}$ |
| $C_0^-$ | +⋆,X | +⋆,X | +⋆,X | +⋆,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $R_1^+$ | 0−,X | 0−,X | 0−,X | 0−,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $R_{L1}^+$ | 0⋆,X | −⋆,X | −⋆,X | 0⋆,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $R_{L1}^-$ | 0⋆,X | +⋆,X | +⋆,X | 0⋆,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $R_{L2A}^+$ | 0⋆,X | −⋆,X | 0⋆,X | −⋆,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $R_{L2A}^-$ | 0⋆,X | +⋆,X | 0⋆,X | +⋆,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $Sw_1^{off}$ | 0⋆,X | −⋆,X | −⋆,Z | 0⋆,X | 00,X | 00,X | 00,X | 00,X | 00,X |
| $Sw_2^{off}$ | 0⋆,X | −⋆,X | 0⋆,X | −⋆,Z | 00,X | 00,X | 00,X | 00,X | 00,X |
| $R_{bulb}^+$ | 0⋆,X | −⋆,X | 0⋆,X | 0⋆,X | 0,X | −,X | −,X | −,X | 00,X |
| $R_{bulb}^-$ | 0⋆,X | +⋆,X | 0⋆,X | 0⋆,X | 0,X | +,X | +,X | +,X | 00,X |
| $J_{fan}^-$ | 0⋆,X | +⋆,X | 0⋆,X | 0⋆,X | 0,X | 0,X | −,X | 0,X | −+,X |
| $B_{fan}^+$ | 0⋆,X | +⋆,X | 0⋆,X | 0⋆,X | 0,X | 0,X | +,X | 0,X | 0−,X |
| $Sw_3^{off}$ | 0⋆,X | −⋆,X | 0⋆,X | 0⋆,X | 0,X | −,X | −,X | −,X | 00,X |
| $Sw_4^{off}$ | 0⋆,X | +⋆,X | 0⋆,X | 0⋆,X | 0,X | −,X | +,Z | 0,X | 0−,X |

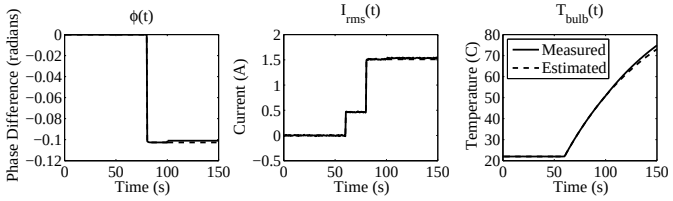


Fig. 9. $R_{bulb}^-$ fault, where $R_{bulb}$ decreases by 5%.

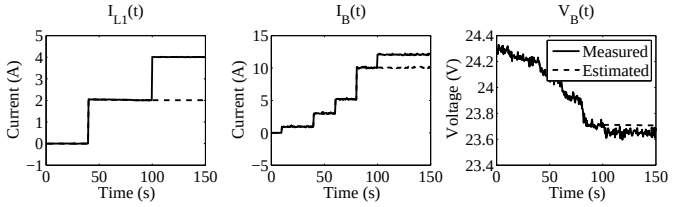


Fig. 10. $R_{L1}^-$ fault, where $R_{L1}$ decreases by 50%.

generates $\{R_{bulb}^-, B_{fan}^+, Sw_3^{off}\}$ as the initial fault candidate set. At 101.5 s, the ac diagnoser detects an increase in the $T_{bulb}$. Since only $R_{bulb}^-$ is consistent with the observed increase in $T_{bulb}$, all other candidates are dropped by the ac diagnoser, and a unique candidate is obtained. The dc diagnoser later observes the increase in $I_{rms}$, and since no faults in the dc subsystem can cause an increase in the rms inverter current, it eliminates all faults.

For a second scenario, a $50\%$ decrease in the Load 1 resistance, $R_{L1}^-$, is injected at 100 s. As shown in Fig. 10, this fault causes the Load 1 current and the battery current to increase discontinuously. Both these changes are detected at 100.0 s, and results in the dc diagnoser generating $\{C_0^-, R_{L1}^-\}$ as the fault candidates. At 103.0 s, it is determined that neither measurement exhibited any discrete change behavior, which does not affect the current candidate list. At 104.0 s, it is determined that the change in $I_B$ is a discontinuity, and that $V_B$ decraesed. The fault $C_0^-$ is dropped since it would cause the battery voltage to increaseinstead, and $R_{L1}^-$ is isolated as the true fault.

We have also studied in simulation the effect of fault magnitude and sensor noise on fault detection times and the fault isolation results. With $R_{L1}^-$, for example, the fault was detected in less than $0.5$ s, on average, for magnitudes of at least $5\%$ with the different levels of noise. Full details for the different faults in the dc subsystem can be found in [19].

TABLE IV
DIAGNOSIS RESULTS FROM SIMULATION EXPERIMENTS

| Fault | DC Diagnoser | | | AC Diagnoser | | |
|---|---|---|---|---|---|---|
| | $t_d - t_f$ | $t_i - t_f$ | Result | $t_d - t_f$ | $t_i - t_f$ | Result |
| $C_0^-$ at $-1\%$ | 0.5 | 12.0 | $\{C_0^-\}$ | N/A | N/A | $\varnothing$ |
| $R_1^+$ at $+200\%$ | 1.5 | 10.5 | $\{R_1^+\}$ | N/A | N/A | $\varnothing$ |
| $R_{L1}^+$ at $+50\%$ | 0.0 | 4.5 | $\{R_{L1}^+\}$ | N/A | N/A | $\varnothing$ |
| $R_{L1}^-$ at $-50\%$ | 0.0 | 4.0 | $\{R_{L1}^-\}$ | N/A | N/A | $\varnothing$ |
| $Sw_1^{off}$ | 0.0 | 3.0 | $\{Sw_1^{off}\}$ | N/A | N/A | $\varnothing$ |
| $R_{bulb}^+$ at $+50\%$ | 0.5 | 0.5 | $\varnothing$ | 0.5 | 0.5 | $\{R_{bulb}^+, Sw_3^{off}\}$ |
| $R_{bulb}^-$ at $-5\%$ | N/A | N/A | $\varnothing$ | 0.5 | 1.5 | $\{R_{bulb}^-\}$ |
| $J_{fan}^-$ at $-50\%$ | N/A | N/A | $\varnothing$ | 0.0 | 0.0 | $\{J_{fan}^-\}$ |
| $B_{fan}^+$ at $+50\%$ | N/A | N/A | $\varnothing$ | 0.5 | 4.0 | $\{B_{fan}^+\}$ |
| $Sw_3^{off}$ | 0.5 | 0.5 | $\varnothing$ | 0.5 | 0.5 | $\{R_{bulb}^+, Sw_3^{off}\}$ |
| $Sw_4^{off}$ | 0.5 | 0.5 | $\varnothing$ | 0.5 | 0.5 | $\{Sw_4^{off}\}$ |

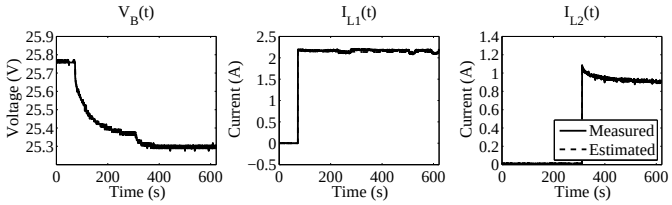


Fig. 11. Nominal system operation.

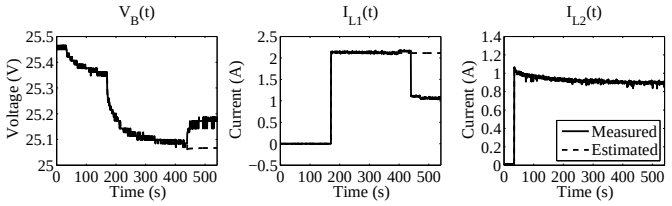


Fig. 12. $R_{L1}^+$ fault, where $R_{L1}$ increases by $100\%$.

### B. Testbed Results

We have also performed online experiments on the ADAPT testbed. In online experiments, 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 load fault and a switch fault, and refer the reader to [19] for additional experiments. In these experiments, we consider a configuration that includes the battery discharging to the two dc loads only. The nominal behavior of the system in shown in Fig. 11.

For a first scenario, a $100\%$ increase in the Load 1 resistance, $R_{L1}^+$, is manually injected at 439.5 s in mode $q_{11}$. The measured and estimated outputs are shown in Fig. 12. The increase in resistance causes a discontinuous drop in the current, detected at 440.0 s. Since the slope has not yet been computed, the possible fault candidates are $\{R_1^+, R_{L1}^+, Sw_1^{off}\}$. At 441.0 s, an increase is detected in $V_B(t)$. Since $I_{L1}^-$ cannot affect $V_B(t)$, it is dropped. $R_1^+$ is also dropped because it would have decreased, and not increased, the battery voltage. At 442.5 s, it is determined that no discrete change in $I_{L1}(t)$ occurred, so $R_{L1}^+$ is isolated as the true fault.

We now investigate an unexpected switch fault. At 375.5 s,

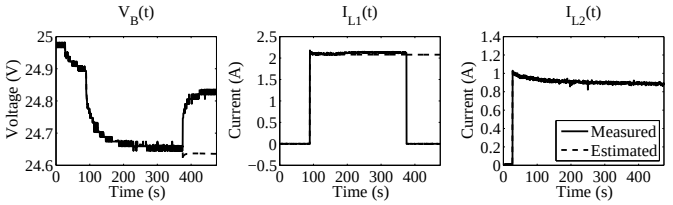


Fig. 13. $Sw_1$ turns off.

$Sw_1$ turns off without a command. The measured and estimated outputs are shown in Fig. 13. As a result of the fault, $I_{L1}(t)$ goes immediately to zero, and $V_B(t)$ increases as a result of less current being drawn. The fault is detected at 376.0 s, and the symbol generator reports a decrease in $I_{L1}(t)$. The initial fault hypotheses are then $\{R_1^+, R_{L1}^+, Sw_1^{off}\}$. At 376.5 s, the increase in $V_B(t)$ is detected, so the diagnosis reduces to $\{R_{L1}^+, Sw_1^{off}\}$. At 378.5 s, the symbol generator determines that $I_{L1}^-$ went to zero, and therefore $Sw_1^{off}$ is isolated as the true fault. Except for the discrete change behavior of $I_{L1}(t)$, the switch fault produces the same qualitative signatures as the load resistance fault. Therefore, it is clear that the additional symbol is necessary to discriminate between the faults.

## IX. CONCLUSIONS

Applying model-based diagnosis techniques to real-world systems provides many challenges, including the building of accurate system models. The modeling task is complicated because details of component models are often unavailable, interactions between components are not fully documented, and sufficient data may not be available to estimate the parameters of the model. We faced these issues when modeling a number of components of the ADAPT system – the battery, the inverter, the fan, and the pump.

Our FACT tools greatly facilitate synthesizing the different modules of the diagnosis system and the VIRTUAL ADAPT testbed. However, setting the parameters of the observer and fault detectors are also critical tasks for accurate system monitoring, avoiding false alarms, and correct symbol generation. Coming up with the right parameter values involves running a number of systematic experiments. In some cases it is

hard to guarantee false alarm rates because the nature of the modeling errors and measurement noise may be unknown. In our work, assuming Gaussian distributions, and estimating the measurement noise variance online worked well.

We extended our traditional hybrid diagnosis approach to include steady-state analysis for ac systems, which provided us with fault signatures for ac and dc measurements. Based on the signatures, we performed diagnosability analysis of the system and designed distributed diagnosers for the heterogeneous 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


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

[2] S. Narasimhan and G. Biswas, "Model-based diagnosis of hybrid systems," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 37, no. 3, pp. 348–361, May 2007.

[3] J. de Kleer and B. C. Williams, "Diagnosing multiple faults," *Artificial Intelligence*, vol. 32, pp. 97–130, 1987.

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

[5] F. Zhao, X. Koutsoukos, H. Haussecker, J. Reich, and P. Cheung, "Monitoring and fault diagnosis of hybrid systems," *IEEE Trans. on Systems, Man, and Cybernetics, Part B*, vol. 35, no. 6, pp. 1225–1240, 2005.

[6] S. Poll, A. Patterson-Hine, J. Camisa, D. Nishikawa, L. Spirkovska, D. Garcia, D. Hall, C. Neukom, A. Sweet, S. Yentus, C. Lee, J. Ossenfort, I. Roychoudhury, M. Daigle, G. Biswas, X. Koutsoukos, and R. Lutz, "Evaluation, selection, and application of model-based diagnosis tools and approaches," in *AIAA Infotech@Aerospace 2007 Conference and Exhibit*, May 2007.

[7] M. Daigle, X. Koutsoukos, and G. Biswas, "An integrated approach to parametric and discrete fault diagnosis in hybrid systems," in *HSCC 2008*, ser. LNCS. Springer-Verlag, 2008, vol. 4981, pp. 614–617.

[8] I. Roychoudhury, G. Biswas, and X. Koutsoukos, "Designing distributed diagnosers for complex continuous systems," *IEEE Transactions on Automation Science and Engineering*, 2008.

[9] E.-J. Manders, G. Biswas, J. Ramirez, N. Mahadevan, J. Wu, and S. Abdelwahed, "A model-integrated computing tool-suite for fault adaptive control," in *Proceedings of the Fifteenth International Workshop on Principles of Diagnosis*, June 2004.

[10] G. Karsai, J. Sztipanovits, A. Ledeczi, and T. Bapty, "Model-integrated development of embedded software," in *Proceeding of the IEEE*, vol. 91, no. 1, Jan. 2003, pp. 145–164.

[11] P. J. Mosterman and G. Biswas, "A theory of discontinuities in physical system models," *Journal of the Franklin Institute*, vol. 335B, no. 3, pp. 401–439, Jan. 1998.

[12] 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.

[13] R. E. Kirk, *Statistics: An Introduction*. Fort Worth: Harcourt Brace, 1999.

[14] 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.

[15] 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.

[16] M. Daigle, I. Roychoudhury, G. Biswas, and X. Koutsoukos, "Efficient simulation of component-based hybrid models represented as hybrid bond graphs," in *HSCC 2007*, ser. LNCS, A. Bemporad, A. Bicchi, and G. Butazzo, Eds. Springer-Verlag, 2007, vol. 4416, pp. 680–683.

[17] I. Roychoudhury, M. Daigle, G. Biswas, X. Koutsoukos, , and P. J. Mosterman, "A method for efficient simulation of hybrid bond graphs," in *International Conference on Bond Graph Modeling and Simulation (ICBGM 2007)*, Jan. 2007, pp. 177–184.

[18] MATLAB/Simulink. [Online]. Available: http://www.mathworks.com/

[19] M. Daigle, "A qualitative event-based approach to fault diagnosis of hybrid systems," Ph.D. dissertation, Vanderbilt University, 2008.

[20] R. Debouk, S. Lafortune, and D. Teneketzis, "Coordinated decentralized protocols for failure diagnosis of discrete event systems," *Discrete Event Dynamic System: Theory and Applications*, vol. 10, no. 1/2, pp. 33–86, January 2000.

[21] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanella, "Diagnosis of large active systems," *Artificial Intelligence*, vol. 110, no. 1, pp. 135–183, 1999.

[22] Y. Pencole and M.-O. Cordier, "A formal framework for the decentralised diagnosis of large scale discrete event systems and its application to telecommunication networks," *Artificial Intelligence*, vol. 164, no. 1-2, pp. 121 – 170, 2005.

[23] 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*, to appear.