Using Factored Bond Graphs for Distributed Diagnosis of Physical Systems

Indranil Roychoudhury[1], Gautam Biswas[2] and Xenofon Koutsoukos[2]
[1] SGT Inc., NASA Ames Research Center, Moffett Field, CA 94035, USA
indranil.roychoudhury@nasa.gov
[2] ISIS, Dept. of EECS, Vanderbilt University, Nashville, TN 37235, USA
{gautam.biswas, xenofon.koutsoukos}@vanderbilt.edu

Abstract

This paper presents a distributed Bayesian fault diagnosis scheme for physical systems. Our diagnoser design is based on a procedure for factoring the global system bond graph (BG) into a set of structurally observable bond graph factors (BG-Fs). Each BG-F is systematically translated into a corresponding DBN Factor (DBN-F), which is then used in its corresponding local diagnoser for quantitative fault detection, isolation, and identification. By construction, the random variables in each DBN-F are conditionally independent of the random variables in all other DBN-Fs, given a subset of communicated measurements considered as system inputs. Each DBN-F and BG-F pair is used to derive a local diagnoser that generates globally correct diagnosis results by local analysis. Together, the local diagnosers diagnose all single faults of interest in the system. We demonstrate on an electrical system how our distributed diagnosis scheme is computationally more efficient than its centralized counterpart, but without compromising the accuracy of the diagnosis results.

1. INTRODUCTION

Online diagnosis schemes can help ensure the safe and efficient operation of real-world engineering systems. However, many model-based diagnosis schemes are centralized, and hence, expensive in memory and computational requirements. They also scale poorly to changes in system configuration, and result in single points of failure. Distributed diagnosis schemes address the drawbacks of centralized schemes [1–3]. This paper presents a distributed Bayesian scheme for diagnosing incipient (slow) and abrupt (fast) faults in system parameters of physical systems, where each local diagnoser communicates a minimal number of measurements with other diagnosers to generate globally correct diagnosis results. Together, the set of local diagnosers diagnose all single faults of interest in the system.

The basis of our diagnoser design is the procedure for *factoring* (or partitioning) the global system bond graph (BG) into multiple independent bond graph factors (BG-Fs), each of which is structurally observable [4]. These BG-Fs are designed by identifying energy storage elements in integral causality, the outgoing effort or flow of which can be computed as an algebraic function of some sensor measurements.

Once these energy storage elements are identified, they are replaced by modulated sources of flow or effort, whose modulation factors are based on the sensor readings so that their output flows or efforts equal that of the output flow or effort of the energy storage elements they replace. The integrals of the effort or flow imposed by an energy storage element in integral causality represent the state variables in the state-space model of dynamic systems. Hence, by replacing the subset of energy storage elements with modulated sources, the state variables are decoupled by considering measurements on the path (i.e., the relation) between two state variables. Given these measurements, the state variable value computation at each time step depends on the measurement values (now considered inputs) and not the other state variables. Thus each BG-F is decoupled across time from all other BG-Fs, given the subset of measurements now considered inputs.

In our work, we translate BG-Fs into Dynamic Bayesian Networks Factors (DBN-Fs), which provide a systematic method for modeling uncertainty in the behavior of dynamic systems [5]. Applying the factoring approach produces structurally observable BG-Fs, which are systematically converted to DBN-Fs using the procedure presented in [6]. Each DBN-F and BG-F pair is used to derive a local diagnoser. We implement a particle filter (PF)-based inference approach [7] on each DBN-F for fault detection, isolation and identification, and use the BG-F for qualitative fault isolation. By construction, the random variables in each DBN-F are guaranteed to be conditionally independent of the random variables in all other DBN-Fs, given a subset of communicated measurements considered as system inputs. Therefore, each local diagnoser operates independently given the measurements and generates local diagnostic hypotheses that are provably correct globally.

Our approach combines the DBN-based quantitative diagnosis scheme with a qualitative fault isolation scheme to further improve diagnosis efficiency. We demonstrate on an electrical system how our distributed diagnosis approach is computationally more efficient than its centralized counterpart, but without compromising the accuracy of diagnosis.

2. MODELING FOR DIAGNOSIS

In our work, we systematically derive the diagnosis models for fault isolation in the form of *temporal causal graphs*

(a) Schematic.

(d) Bond graph.

(e) Two-Factored bond graph with imposed derivative causality.

(b) Dynamic Bayesian Network.
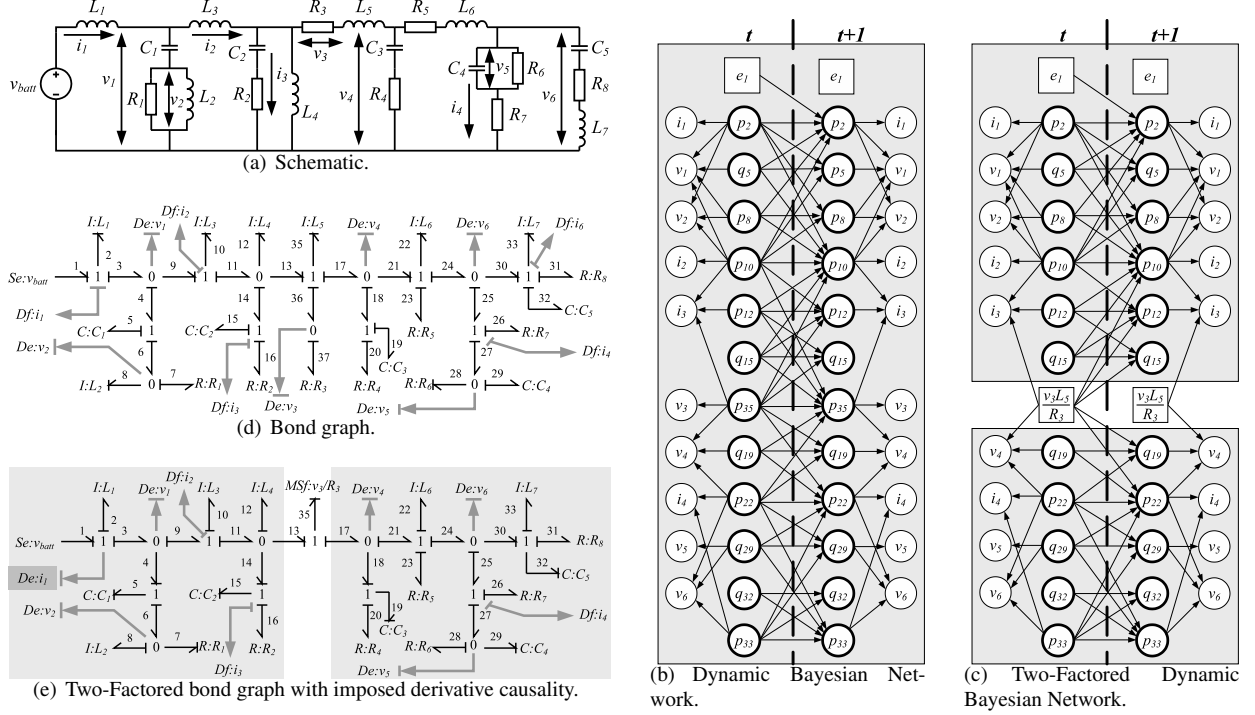
(c) Two-Factored Dynamic Bayesian Network.

Figure 1. Models of the twelfth-order electrical system.

(TCGs) [8], and DBN-Fs for fault detection, isolation, and identification [9, 10]. All of these models are derived from the system's BG model [11]. Fig. 1(d) shows the BG of a twelfth-order electrical circuit shown in Fig. 1(a). In the electrical domain, the effort variables denote voltage difference across, and flow variables denote current through, BG elements. For example, $f_2 = i_1$ denotes the current through the inductor L_1, and $e_7 = v_2$ denotes the voltage difference across resistor R_1. $e_1 = v_{batt}$ denotes the voltage imposed by the voltage supply. Effort and flows are measured using sensor elements, e.g., $De : v_2$ is a voltage sensor. For every C element in integral causality, the corresponding state variable is the *displacement* variable, q, such that $\dot{q} = f$. Similarly for every I element in integral causality, the corresponding state variable is the *momentum* variable, p, such that $\dot{p} = e$.

BGs represent a parameterized energy-based model of physical systems. Faults are defined as changes in the nominal parameter values. An *incipient fault* is a slow change in a system parameter, p (with nominal parameter value function, $p(t)$), and modeled as a linear function with a constant slope, Δ_p^i, added to the nominal component parameter value function, $p(t)$, i.e., $p^{\pm i}(t) = p(t) \pm \Delta_p^i \times (t - t_f)$, $t > t_f$, where t_f is the time of fault occurrence, and $p^{\pm i}(t)$ is the temporal profile of parameter p with an incipient fault [10]. An *abrupt fault* is modeled as an addition of a constant persistent bias term, $\Delta_p^a \cdot p(t)$, to the nominal parameter value, $p(t)$, i.e., $p^{\pm a}(t) = p(t) \pm \Delta_p^a \cdot p(t)$, $t > t_f$, where Δ_p^a is the percentage

change in the parameter expressed as a fraction, and $p^{\pm a}(t)$ is the temporal profile of parameter p with an abrupt fault [10].

A TCG, automatically derivable from a BG [8], is essentially a signal flow graph that captures the causal and temporal relations between its nodes, which represent system variables (effort and flow values), through directed *edges* and their *labels*. The direction of a TCG edge and its label are based on *causality*, which establishes the cause and effect relationships between the e_i and f_i variables of a bond i based on constraints imposed by the incident BG elements.

A DBN is a directed acyclic graph structure that represents a probabilistic discrete-time model of a dynamic system. A DBN can be defined as $D = (\mathbf{X}, \mathbf{U}, \mathbf{Y})$, where $\mathbf{X}$, $\mathbf{U}$, and $\mathbf{Y}$ are sets of stochastic random variables that denote (hidden) state variables, system input variables, and measured variables in the dynamic system, respectively [5]. Links between nodes denote causal dependencies between nodes within a time step and across time steps. Graphically, a DBN is a two-slice Bayesian network, representing a snapshot of system behavior in two consecutive time slices, t and $t + 1$. Each DBN time-slice represents the Markov process observation model, $P(\mathbf{Y}_t | \mathbf{X}_t, \mathbf{U}_t)$, while the across-time links represent the Markov state-transition model, $P(\mathbf{X}_{t+1} | \mathbf{X}_t, \mathbf{U}_t)$. DBNs exploit the conditional independence between system variables to provide a compact system representation for reasoning about dynamic systems behavior in the presence of uncertainty. Bayesian inference algorithms have been widely used
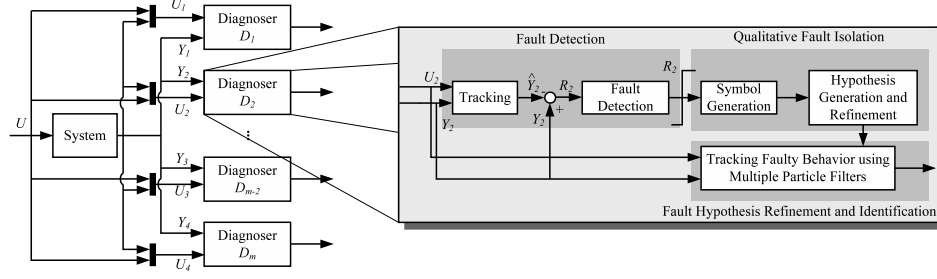
Figure 2. The distributed diagnosis architecture.

for diagnosis of dynamic systems represented as DBNs. The system DBN is constructed from its TCG in integral causality using the method given in [12]. Fig. 1(b) shows the DBN for our example circuit, where thick-lined circles denote state variables, thin-lined circles denote observed variables, and squares denote input variables.

3. DISTRIBUTED DIAGNOSIS APPROACH

Our diagnosis approach, whose architecture is shown in Fig. 2, has three primary components [10]: (*i*) fault detection, (*ii*) qualitative fault isolation (Qual-FI), and (*iii*) quantitative fault hypothesis refinement and identification (Quant-FHRI).

A fault is *detected* when the measurement residual, i.e. the difference between the observed (faulty) and estimated (nominal) values of any measurement is determined to be statistically significant using a Z-test [13]. A PF-based observer implemented on the "nominal" DBN-F for each local diagnoser estimates the nominal system behavior. Nominal DBN-Fs are generated from the nominal BG-Fs that are obtained from the nominal global BG, and include only the state and measurement variables as random variables. The system parameters are deterministic. A PF is a sequential Monte Carlo sampling method for Bayesian filtering that approximates the belief state of a system using a weighted set of samples, or particles [7]. Each sample, or particle, consists of a value for each state variable, and describes a possible system state. As more observations are obtained, each particle is moved stochastically to a new state, and the weight of each particle is readjusted to reflect the likelihood of that observation given the particle's new state.

The Qual-FI module is activated once a fault is detected, and the *symbol generator* starts encoding if the magnitude and slope of each measurement has increased, decreased, or remained unchanged from nominal, using qualitative '+', '−', or '0' symbols [13], respectively. In parallel, all fault hypotheses, i.e., the abrupt and incipient parameter changes that explain the observed deviations, are identified by propagating the first observed measurement-deviation backwards along the TCG. Then, for each fault hypothesis, the fault signature, i.e., its expected effect on magnitude and slope of the mea-

surement residuals (encoded as qualitative '+', '−', and '0' symbols), is determined, by propagating the parameter deviation forward along the TCG. For example, the fault signature of a fault, say p^{+a}, for a measurement m_1 can be $(+-)$, denoting a discontinuous increase followed by a gradual decrease in m_1 if fault p^{+a} occurs. The true fault can be isolated by removing fault hypotheses whose signatures are inconsistent with the observed measurement symbols. The Quant-FHRI scheme is invoked when the fault hypothesis set is refined to a pre-defined size, k, or a pre-specified s simulation timesteps have elapsed.

The Quant-FHRI performs both fault hypothesis refinement and identification if multiple fault hypotheses remain when Quant-FHRI is initiated. If however, the Qual-FI has refined the set of hypotheses to a singleton, Quant-FHRI performs the task of fault identification only. For each fault hypothesis that remains at the time Quant-FHRI is initiated, a faulty DBN-F is generated by extending the nominal DBN-F to include the fault parameter as a stochastic variable in the DBN-F, as explained in [10]. A PF is then implemented on each DBN-F fault model to track the faulty observed behavior, taking as input the measurements from time $t_d - \Delta_t^{max}$, where $\Delta_t^{max} \geq t_d - t_f$ is the maximum delay possible between the time of fault occurrence, t_f, and the time of fault detection, t_d. We start tracking the faulty DBN-Fs from time $t_d - \Delta_t^{max}$ because once the fault has occurred, and till the magnitude of the fault is correctly identified, the system model is unknown to us. If we started tracking faulty measurements from time t_d onwards, our tracking will be off, since we would possibly start at wrong initial state, especially for the unknown fault variable. The idea is that if Δ_t^{max} is set to be at least as large as $t_d - t_f$, then the time of actual fault occurrence will be between $t_d - \Delta_t^{max}$ and t_d, and our initial state setting will be closer to the correct value, since it is based on state estimates from a known system model.

For each PF, a Z-test is used to determine if the deviation of a measurement estimated by the PF from the corresponding actual observation is statistically significant. As more observations are obtained, ideally the PF using the correct fault model will eventually converge to the observed measurements, while the observations estimated using the in-

correct fault models will gradually deviate from the observed measurements. Since even the correct fault model will need some time before the particles start converging to the observed measurement values, we need to delay the invocation of the Z-tests for s_d time steps, as otherwise, they may indicate a deviation from observed measurements at the very onset for all fault models. A fault hypothesis is removed from consideration if: (*i*) the Qual-FI drops that fault candidate, or (*ii*) the measurements estimated by that fault model significantly deviates from the observed faulty measurements. We typically assume that the particles for the true fault model will converge to the observed measurements within s_d time steps of its invocation. Since the fault magnitude is included as a stochastic variable in every fault model, the magnitude of the true fault (i.e., the % bias, Δ_p^a, or, the slope, Δ_p^i) is considered to be that estimated by the PF for the true fault model.

4. DESIGNING THE LOCAL DIAGNOSERS

The objective of our distributed diagnosis scheme is to generate globally correct diagnosis results without a centralized coordinator, and by communicating a minimal number of measurements between diagnosers. We achieve this objective by first *factoring* the system BG into *maximal* number of structurally observable [4] BG-Fs; generating a DBN-F for each BG-F; and incorporating each BG-F and DBN-F pair in its corresponding local diagnoser. The detailed procedure for generating the local diagnosers is given in [6].

These BG-Fs are designed by replacing energy-storage *I* and *C* elements by modulated sources of flow or effort, respectively. The modulated sources of effort and flow are computed as algebraic functions of observed measurements. Fig. 1(d) shows complete BG of the electrical circuit. The available measurements in this circuit are current values, $i_1, i_2, \ldots, i_4$ and voltages $v_1, v_2, \ldots, v_6$. Fig. 1(e) shows the two BG-Fs that the global BG is factored into. It is evident from Fig. 1(a) that the current through the inductor L_5 is equal to v_3/R_3. Hence, we can replace f_{35} in Fig. 1(d) with modulated $MS_f = v_3/R_3$ and create the two BG-Fs shown in Fig. 1(e). If any BG-F is structurally unobservable, it is *merged* into another BG-F, and this process is repeated till all of the BG-Fs are structurally observable. The replacement of a subset of energy storage elements with modulated source elements, effectively replaces some state variables by algebraic functions of some measured variables now considered as inputs. As the values of these state variables can now be computed at every time step, the dependence of these state variables on the values of other state variables at earlier time instances is removed, and thus, the state variables in each BG-F is decoupled across time from all other BG-Fs, given the sensor measurements.

Structural observability of a BG-F can be determined using the procedure described in [4]. A system is structurally observable if in its BG, (*i*) there exists at least one causal path for each *I* and *C* element in the preferred integral causality to a sensor element *De* or *Df*, indicating that all state-variables eventually influence a sensor, and (*ii*) inverting the causality of *every* I and C element initially in integral (preferred) causality still produces a valid causal assignment for the entire BG, indicating that a system of equations can be derived from the BG that causally relates the measurements to the state-variables, and which can be solved to determine the value of the state-variables based on the measurements. In some situations, obtaining a valid causal assignment for the BG in derivative causality may require changing a *De* or *Df* element into their dual form. Both the BG-Fs shown in 1(e) are structurally observable as they fulfill both the conditions necessary for structural observability mentioned in [6]. Note that the current sensor i_1 had to be dualized to assign derivative causality to the BG-F on the left in Fig. 1(e). Since the two BG-Fs shown in Fig. 1(e) are structurally observable, we do not require any further merging in our particular example.

On obtaining the maximal number of structurally observable BG-Fs, each BG-F is converted into a DBN-F. By construction, the random variables in each DBN-F are conditionally independent from those in other DBN-Fs given the subset of measurements now considered system inputs. Because our DBN-Fs are generated from structurally observable BG-Fs, our factored inference scheme generates accurate inference results. A DBN-F, $D_j = (\mathbf{X}_j, \mathbf{U}_j, \mathbf{Y}_j)$, is termed *conditionally independent* of other DBN-Fs, $D_k(\mathbf{X}_k, \mathbf{U}_k, \mathbf{Y}_k)$ $(k \neq j)$, given its inputs, $\mathbf{U}_j$, if every random variable in D_j is conditionally independent of all other variables in D_k given $\mathbf{U}_j$. Fig. 1(c) shows the DBN-Fs corresponding to the two BG-Fs in Fig. 1(e). Note that in the DBN-Fs, the state variable p_{35} is replaced by the input v_3L_5/R_3. Since, v_3/R_3 can be *measured* at every time step, all causal links directed into p_{35} are removed. As a result, given v_3L_5/R_3, every variable in one DBN-F is conditionally independent of the variables in the other DBN-F. Thus, the two generated DBN-Fs are conditionally independent. In the factored DBN, we do not replace state variables, such as, p_2 with i_1L_1, since this replacement does not yield any additional factors in Fig. 1(c). Moreover, we do not replace state variables p_{10} and q_{29} with i_2L_3 and v_5C_4, respectively, since we assume that inductor L_3, and capacitor C_4 can become faulty. We can see that the DBN-Fs shown in Fig. 1(c) map to the BG-Fs shown in Fig. 1(e).

Once we generate m DBN-Fs, $D_1, D_2, \ldots D_m$, from m structurally-observable BG-Fs, $B_1, B_2, \ldots, B_m$, a local diagnoser, D_i, is constructed based off the corresponding DBN-F, D_i, and BG-F, B_i. Recall, as shown in Fig. 2, our distributed diagnosis approach is implemented by each D_i, independent of the other diagnosers, as presented in the previous section: a PF applied to the nominal DBN-F, D_i, is used as the nominal observer for fault detection; a TCG generated from the cor-

responding BG-F, B_i, is used for Qual-FI; and PFs applied to faulty DBN-Fs generated by extending D_i with faulty parameters as additional state variables, are used for Quant-FHRI.

5. IMPLEMENTING THE LOCAL DIAGNOSERS

Once designed, each local diagnoser, D_i, implements a PF on the nominal DBN-F, D_i, to estimate the nominal system behavior. Each of these particle filters takes as inputs, U_i, and estimates X_i based on the observed measurements, Y_i. Note that a diagnoser D_i's inputs, U_i, may include some of the inputs to the global system, i.e., $U_i \cap U \neq \varnothing$, as well as some measurements now considered inputs, i.e., $U_i \cap Y \neq \varnothing$. Two diagnosers D_j, D_k *communicate*, or share, a measurement $Y \in Y$ if $Y \in U_j \wedge Y \in U_k$, i.e., measurement Y is an input to both D_j and D_k. Since there is a one-to-one correspondence between DBN-F, D_i, and local diagnoser, D_i, only measurements $(\cup_i U_i) - U$ are shared, or communicated, between the particle filter-based observers for each D_i. Further, the particle filter for the DBN-F D_i is designed to use $a\frac{|X_i|}{|X|}$ particles, where a is a user-specified parameter. Given m DBN-Fs, we know that $\sum_i |X_i| < |X|$, where X is the total number of state states in the complete system. Once a Z-test indicates statistically significant deviation of estimate nominal behavior from the observed behavior, a fault is detected in a local diagnoser, D_i. Once a fault is detected by D_i, fault hypotheses are generated by backward propagation of observed measurement deviation along the TCG obtained from the the BG-F, B_i, corresponding to DBN-F, D_i. The same TCG is also used for generating fault signatures. Typically, once Qual-FI refines the fault hypothesis set to a smaller set, Quant-FHRI is initiated on the remaining fault hypotheses. A faulty DBN-F is generated for each remaining fault hypothesis by including the fault parameter as an additional state variable in D_i. Quant-FHRI then uses PFs implemented on the different faulty DBN-Fs to track the faulty behavior and isolate and identify the true fault.

Recall the computational architecture of our distributed Bayesian fault diagnosis approach shown in Fig. 2. Each local diagnoser, D_i, essentially implements a combined qualitative-quantitative scheme for diagnosing incipient and abrupt faultsŚince, a local diagnoser operates independently of other diagnosers, the complexity of tracking using each DBN-F is less that that of tracking using the global DBN. Also, since the inference algorithms on the different factors are executed simultaneously, the total complexity of inference reduces to the complexity of inference of the particle filter with the maximum number of particles. The reduction of complexity is based on the assumptions that the sensors associated with measurements used to modulate the source elements in BG-Fs will not be faulty, and the components whose parameters are used in the algebraic functions are assumed not to

Table 1. Fault Signatures for Local Diagnosers D_1 and D_2

Diagnoser D_1						Diagnoser D_2				
Fault	i_1	i_2	i_3	v_1	v_2	Fault	i_4	v_4	v_5	v_6
$C_2^{-a}, C_2^{-i}, R_2^{+a}, R_2^{+i}$	0−	0−	0−	0+	0+	C_3^{-a}, R_4^{+a}	0+	+−	0+	0+
L_2^{-a}	0+	0−	0−	−+	−+	C_3^{-i}, R_4^{+i}	0+	0+	0+	0+
L_2^{-i}	0+	0−	0−	0−	0−	C_4^{-a}	0−	0+	+−	+−
L_3^{-a}	0+	+−	+−	−*	−+	$C_4^{-i}, R_6^{+a}, R_6^{+i}$	0−	0+	0+	0+
L_3^{-i}	0+	0+	0+	0−	0−	L_7^{-a}	−+	0−	0−	−*
L_4^{-a}	0+	0−	+−	0−	0−	L_7^{-i}	0−	0−	0−	0−
L_4^{-i}	0+	0+	0−	0−	0−	R_7^{+a}	0−	0+	0−	+−
						R_7^{+i}	0−	0+	0−	0+

fail. Therefore, there is a trade-off for robustness to gain efficiency.

Also, since the local diagnosers incorporate conditionally independent DBN-Fs, we can argue that the local diagnosers generate globally correct diagnosis results without any global coordinator. Given that the fault DBN-Fs are constructed based off conditionally independent DBN-Fs, and local diagnosers share minimal measurements with one another, a fault, $\phi \in F_j$, is only detected by diagnoser D_j, and all other diagnosers, D_k, $k \neq j$, will not detect the fault. Hence, they are not activated, even though the effect of fault ϕ may propagate to all other factors. Basically, a diagnoser D_i is activated when it detects a fault. In general, let us assume that the observer in diagnoser D_i uses the state space equations $\hat{X}_{i_{t+1}} = G_i(X_{i_t}, U_{i_t})$, and $\hat{Y}_{i_t} = H_i(X_{i_t}, U_{i_t})$. Now, a fault in BG-F, B_k implies that functions G_k and H_k do not correctly represent the actual system any more. As a result, $\hat{Y}_k \not\approx Y_k$, and a fault is eventually detected by D_k. The effects of a fault in B_k can propagate to another BG-F B_j, $j \neq k$, through their shared inputs, $(U_j \cap U_k) - U$, iff B_k and B_j communicate at least one measurement, but, since we adopt the single-fault assumption, and since by construction, two BG-Fs can never share any parameters, the state space representations G_j and H_j of all other BG-Fs, B_j, $j \neq k$, will correctly represent the actual system dynamics of each BG-F. Hence, $\hat{Y}_j \approx Y_j$, i.e., the observers in other diagnosers will correctly track the faulty measurement, and hence no fault will be detected. Consequently, the diagnoser does not get activated unless a fault is detected.

6. EXPERIMENTAL RESULTS

This section presents experimental results of our distributed diagnosis scheme applied to the electrical system shown in Fig. 1(a). Two local diagnosers, D_1 and D_2 are designed for the two DBN-Fs shown in Fig. 1(c). These two diagnosers communicate voltage measurement v_3 between each other. Table 1 shows the possible faults that must be diagnosed by each diagnoser, and the fault signatures for the measurements available to each diagnoser.

We present an experimental run for diagnosing an abrupt fault in C_2, C_2^{-a}, with $\Delta_{C_2}^a = -0.9$, introduced at time, $t = 100$ s. As shown in Fig. 3, a negative deviation is noticed in measurement i_3 at $t = 101.4$ s. Based on this deviation, the
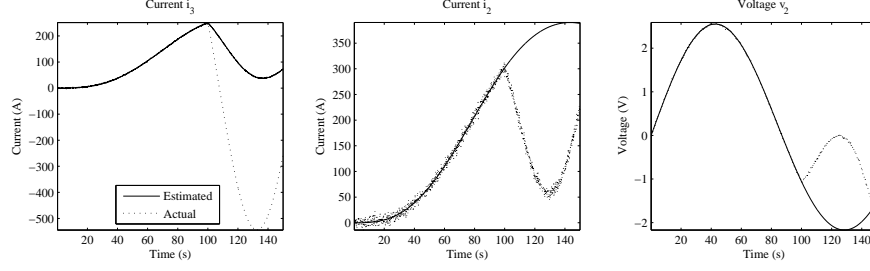
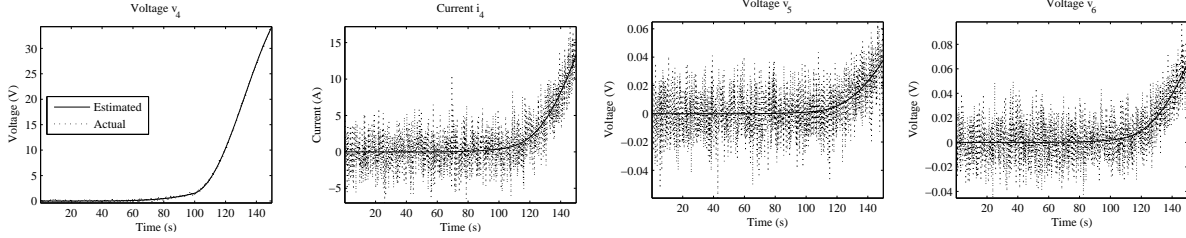**Figure 3.** Detection of C_2^{-a} fault by diagnoser D_1.

Figure 4. Tracking observations in the presence of C_2^{-a} fault by diagnoser D_2.

fault hypothesis set, $\{C_2^{-i}, C_2^{-a}, R_2^{+a}, R_2^{+i}, L_2^{-i}, L_2^{-a}, L_4^{-i}, L_4^{-a}\}$, is generated. At $t = 102.9$ s, a negative deviation is observed in measurement i_2, which is inconsistent with the fault signatures of faults L_4^{-a} and L_4^{-i} for i_2 (see Table 1). Hence, L_4^{-a} and L_4^{-i} are removed from the fault hypothesis set. Similarly, L_2^{-a} is becomes inconsistent at time $t = 102.9$ s, when the deviation in i_3 is confirmed to be a gradual decrease. At $t = 106.7$ s, the fault hypothesis set is refined to $\{C_2^{-i}, C_2^{-a}, R_2^{+a}, R_2^{+i}\}$ when v_2 shows a positive deviation. Since these fault hypotheses cannot be further refined through Qual-FI alone, Quant-FHRI is initiated. The second diagnoser does not detect any fault. We start tracking the observed measurements from time $t = 97.5$ s, and instantiate two PFs, one using a DBN-F model for fault C_2^{-i}/C_2^{-a} (see Fig. 5(a)), and the other using a DBN-F model for fault R_2^{+i}/R_2^{+a} (Figs. 5(b)), with parameters C_2 and R_2 introduced as additional state variables in the nominal DBN-Fs shown in Fig. 1(c).

Fig. 6 shows the sum of mean squared estimation errors obtained using the two different fault models. This error is identified to be statistically significant at $t = 101.3$ s and $t = 101.9$ s, for the C_2^{-a}/C_2^{-i} and R_2^{+a}/R_2^{+i} DBN-Fs, respectively. However, as noted earlier, the PF estimates made using the correct C_2 fault model eventually converges to the observed measurements at $t = 105.0$ s (which is within $s_d = 150$ s), while the estimates using the R_2 DBN fault model do not converge even after $s_d = 150$ s. Hence the true fault is isolated to be $C_2^{\pm a}/C_2^{\pm i}$ fault at $t = 251.3$ s. We run a window-based Z-test on the difference between the known nominal parameter value and the estimated state variable to determine whether the fault is an abrupt or incipient fault in C_2. At $t = 171.7$ s, the statistical test shows that the estimated parameter evolves

in a (-0) manner, implying it is an abrupt fault, and that it converges. By taking a mean of the values for 20 time steps after the abrupt fault is isolated, we obtain $\Delta_{C_2}^a = -0.897$. The estimate for the faulty parameter is shown in Fig. 6(c).

Tables 2 and 3 summarize the results of some distributed and centralized diagnosis experiments, respectively, we ran on the electrical circuit example, averaged over 5 runs. This table shows that compared to the centralized diagnosis approach, the distributed diagnosis approach results in comparable (albeit larger) parameter estimation errors. The distributed approach also takes longer to converge to correct parameter values. We attribute this difference to the proportional distri-
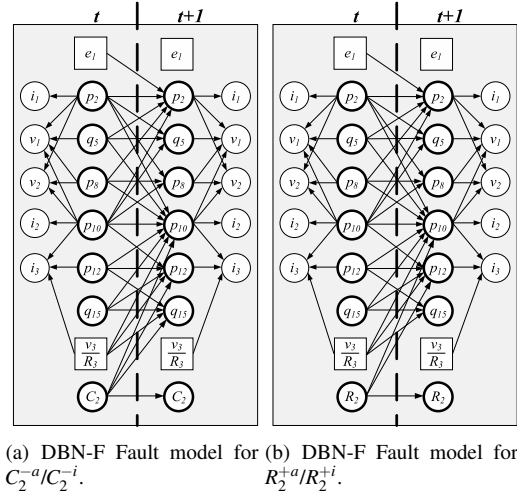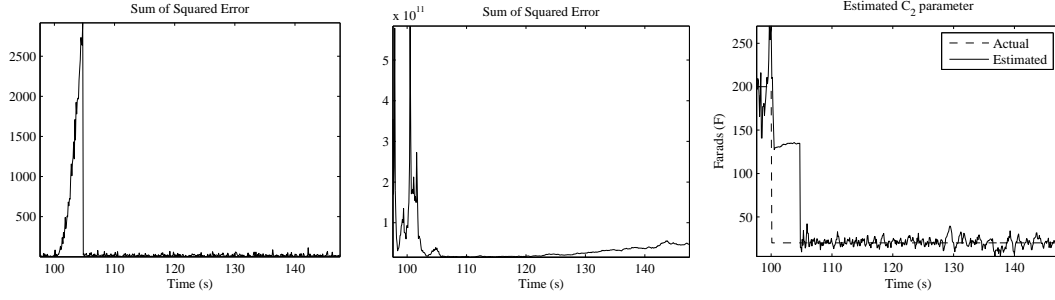

(a) DBN-F Fault model for C_2^{-a}/C_2^{-i}. (b) DBN-F Fault model for R_2^{+a}/R_2^{+i}.

Figure 5. DBN-F Fault models for distributed diagnosis experiments.

(a) Sum of mean squared estimation errors obtained by C_2 DBN fault models. (b) Sum of mean squared estimation errors obtained by R_2 DBN fault models. (c) Estimate of C_2 obtained using the C_2 DBN-F fault model.

Figure 6. Quant-FHRI using C_2 and R_2 DBN-F fault models.

Table 2. Results of Distributed Diagnosis Experiments on the Twelfth-Order Electrical Circuit with Particles Used Proportional to The Total Number of States per Factor

Fault	Magnitude	Det. Time (s)	Iso. Time (s)	Conv. Time (s)	% Mean Param. Est. Error
C_2^{-a}	−0.90	1.04	55.06	5.88	0.64
C_2^{-i}	−0.55	37.3	53.90	19.80	39.74
L_3^{-a}	−0.90	0.50	4.32	6.56	1.11
L_3^{-i}	−0.05	50.40	82.50	23.76	21.14
C_3^{-a}	−0.90	0.20	3.02	3.64	0.13
R_7^{+a}	+5.00	118.30	163.30	128.64	0.66

Table 3. Results of Centralized Diagnosis Experiments on the Twelfth-Order Electrical Circuit

Fault	Magnitude	Det. Time (s)	Iso. Time (s)	Conv. Time (s)	% Mean Param. Est. Error
C_2^{-a}	−0.90	1.26	53.80	4.76	0.27
C_2^{-i}	−0.55	32.82	53.80	11.52	6.05
L_3^{-a}	−0.90	0.50	3.98	5.08	0.49
L_3^{-i}	−0.05	49.62	82.9	12.40	5.50
C_3^{-a}	−0.90	0.2	2.8	3.26	0.12
R_7^{+a}	+5.00	196.8	377.4	115.6	0.48

bution of particles based on the size of each factor, keeping the sum total of particles used by all the PFs the same, as well as, the use of a noisy sensor to compute the value of a state variable. If the individual local diagnosers are executed on different processors, then we can increase the number of particles for each diagnoser, and our intuition is that this will improve the estimation accuracy and identification time of the local diagnosers. Also, our experimental results illustrates the accuracy versus efficiency trade-off of our distributed diagnosis scheme, because, as explained earlier, by construction, if each local diagnoser is implemented on a separate process, the worst case efficiency of our distributed diagnosis scheme is better than that of the centralized diagnosis scheme.

7. DISCUSSION AND CONCLUSIONS

Our approach is similar to the third protocol in [2], which generates correct results without a coordinator. But, unlike the approach presented by [1], each individual local diagnoser needs to communicate only the minimal number of measurements, and not diagnosis results, from other diagnosers to generate globally correct diagnosis results.

In [14], a combined qualitative-quantitative diagnosis approach is presented in which a look-ahead Rao-Blackwellised PF (RBPF) is first used to track system evolution, till a fault is detected, after which the consistency-based Livingstone 3 (L3) generates a set of fault candidates that are then tracked by the fault observer (another RBPF). All the fault hypotheses are included in the same model, and tracked by the fault observer. In contrast, our approach executes the qualitative and quantitative fault isolation schemes in parallel, and uses separate fault models for each fault candidate.

PFs have been used extensively for centralized system health monitoring and diagnosis applications [7, 12]. Distributed inference schemes, such as the BK algorithm [15], creates individual factors by eliminating causal links between weakly interacting subsystems. Therefore, the belief state derived from the individual factors is an approximation of the true belief state. The error in this approximation is bounded, but these bounds may not be sufficiently precise for online diagnosis, since they may result in missed alarms and less precise diagnoses in the best case scenario, and false alarms and wrong diagnosis in the worst case scenario. The Factored Particle Filtering (FPF) scheme [16] further reduces estimation errors by applying the particle filtering scheme to the BK factored inference approach. Our distributed estimation approach uses the particle filtering scheme for inference using DBNs and preserves the overall system dynamics in the factored form, and *does not* approximate the belief state. Hence, we produce accurate state estimates efficiently.

The effectiveness of our approach relies on the assumptions that the assumptions of sensors associated with measurements converted to inputs are not faulty, and components whose parameters are used in the algebraic functions do not fail. In the future, we seek to relax these assumptions. In addition, we would like to apply our diagnosis approach to a large real-world system to analyze the scalability and efficiency of our methodology. We would also like to improve the

efficiency of our diagnosis approach further by ensuring that the BG-Fs are so chosen such that minimal number of fault hypotheses remain at the end of the Qual-FI step. Finally, we would like to investigate what would be a good trade-off between the number of sensors needed for an "optimal" decomposition and the costs induced by having a large number of sensors.

REFERENCES

[1] 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.

[2] 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.

[3] I. Roychoudhury, G. Biswas, and X. Koutsoukos, "Designing distributed diagnosers for complex continuous systems," *IEEE Transactions on Automation Science and Engineering,*, vol. 6, no. 2, pp. 277–290, April 2009.

[4] C. Sueur and G. Dauphin-Tanguy, "Bond graph approach for structural analysis of MIMO linear systems," *Journal of the Franklin Institute*, vol. 328, no. 1, pp. 55–70, 1991.

[5] K. P. Murphy, "Dynamic Bayesian networks: Representation, inference, and learning," Ph.D. dissertation, University of California, Berkeley, 2002.

[6] I. Roychoudhury, G. Biswas, and X. Koutsoukos, "Factoring dynamic Bayesian networks based on structural observability," in *48th IEEE Conference on Decision and Control (CDC 2009),* to appear, 2009.

[7] D. Koller and U. Lerner, "Sampling in factored dynamic systems," in *Sequential Monte Carlo Methods in Practice*, A. Doucet, N. de Freitas, and N. Gordon, Eds. Springer-Verlag, 2001.

[8] P. J. Mosterman and G. Biswas, "Diagnosis of continuous valued systems in transient operating regions," *IEEE-SMCA*, vol. 29, no. 6, pp. 554–565, 1999.

[9] I. Roychoudhury, G. Biswas, and X. Koutsoukos, "Comprehensive diagnosis of continuous systems using dynamic bayes nets," in *Proceedings of the 19th International Workshop on Principles of Diagnosis*, 2008, pp. 151–158.

[10] I. Roychoudhury, "Distributed diagnosis of continuous systems: Global diagnosis through local analysis," Ph.D. dissertation, Vanderbilt University, Nashville, TN, 2009. [Online]. Available: http://macs.isis.vanderbilt.edu/publications

[11] D. C. Karnopp, D. L. Margolis, and R. C. Rosenberg, *Systems Dynamics: Modeling and Simulation of Mechatronic Systems*, 3rd ed. New York, NY, USA: John Wile & Sons, Inc., 2000.

[12] U. Lerner, R. Parr, D. Koller, and G. Biswas, "Bayesian fault detection and diagnosis in dynamic systems," in *Proceedings of the 17th National Conference on Artificial Intelligence*, 2000, pp. 531–537.

[13] E.-J. Manders, G. Biswas, P. J. Mosterman, L. A. Barford, and R. J. Barnett, "Signal interpretation for monitoring and diagnosis, a cooling system testbed," vol. 49, no. 3, pp. 503–509, 2000.

[14] S. Narasimhan, R. Dearden, and E. Benazera, "Combining particle filters and consistency-based approaches for monitoring and diagnosis of stochastic hybrid systems," in *Proceedings of the 15th International Workshop on Principles of Diagnosis*, 2004.

[15] X. Boyen and D. Koller, "Tractable inference for complex stochastic processes," in *Proceedings of the 14th Annual Conference on Uncertainty in Artificial Intelligence*, 1998, pp. 33–42.

[16] B. Ng and L. Peshkin, "Factored particles for scalable monitoring," in *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, 2002, pp. 370–377.