Failure Detection Framework for Stochastic Discrete Event Systems With Guaranteed Error Bounds

Jun Chen, *Member, IEEE*, and Ratnesh Kumar, *Fellow, IEEE*

***Abstract*—This paper studies the online fault detection for stochastic discrete-event systems (DESs) under partial observability of events. Prior works have only studied the verification of the stochastic diagnosability (S-Diagnosability) property. To the best of our knowledge, this is a first paper that investigates the online detection schemes and also introduces the notions of their missed detections (MDs) and false alarms (FAs). Due to the probabilistic nature of the problem, MDs and FAs are possible even for S-Diagnosable systems, and we establish that S-Diagnosability is a necessary and sufficient condition for achieving any desired levels of MD and FA rates. We also provide a detection scheme, that can achieve the specified MD and FA rates, based on comparing a suitable detection statistic, that we define, with a suitable detection threshold, that we algorithmically compute. We also algorithmically compute the corresponding detection delay bound. The detection scheme also works for non-S-Diagnosable systems, with the difference that in this case there exists a lower bound for achievable MD rate, that increases as the FA rate requirement is made more stringent by decreasing it.**

***Index Terms*—Discrete-event systems (DESs), failure diagnosis, online fault detection.**

I. Introduction

DETECTING system failures is essential prior to any fault tolerance action and is an important and challenging problem in many disciplines. In general a fault is a deviation of a system from its required or nominal behavior, such as reaching a fault state or executing a fault event, which can be classified as a permanent fault (as studied in [1]–[5]) or an intermittent fault (as studied in [6]–[8]), and needs to be detected accurately within an adequate bounded delay to ensure timely activation of any fault tolerance action. The problem of fault detection has been widely researched [9]–[15], and is recently studied in the setting of discrete-event systems (DESs) [16]–[25], distributed/decentralized systems [26]–[29], stochastic systems [30]–[32] and systems with temporal logic specification [33]–[35]. In this paper we consider stochastic DESs subjected to faults, modeled as execution of fault traces, or equivalently reachability of fault states.

The notion of stochastic diagnosability, *S-Diagnosability* [1], first proposed as AA-diagnosability in [5], requires that for any tolerable ambiguity level ρ and error bound τ, there must exist a delay bound n such that for any fault trace s, its extensions, longer than n and probability of ambiguity higher than ρ, must occur with probability smaller than τ. Reference [5] provided an observer based exponential complexity algorithm for checking a sufficient condition for S-Diagnosability; the condition was shown to be not necessary. The extension of fault diagnosis of stochastic DESs to decentralized setting has been examined in [28], [29], whereas the diagnosis problem under the added requirement that fault be detected before the violation of any safety property is studied in [3], [36]. These prior works have only studied the verification of the S-Diagnosability property; a technique for online fault detection hasn't yet been examined in literature. To the best of our knowledge, this is a first paper that investigates the online detection schemes for stochastic DESs and also introduces the notions of missed detections (MDs) and false alarms (FAs), or equivalently, false negatives and false positives, for the schemes. Due to the probabilistic nature of the detection problem, MDs and FAs are possible even for S-Diagnosable systems, and we establish that S-Diagnosability is a necessary and sufficient condition for achieving any desired levels of MD and FA rates.

Next we present a detection scheme, that can achieve the specified MD and FA rates, based on comparing a suitable detection statistic with a suitable detection threshold. The approach is that given any observation (of partially observed events), the detector recursively computes the conditional probability of the nonoccurrence of a fault and issues a "fault" decision if the probability of the nonoccurrence of a fault falls below an appropriately chosen threshold, and issues "no-decision" otherwise. For systems that possesses S-Diagnosability property, there always exists a detection threshold and a delay bound so that this detector is able to achieve any desired level of MD and FA rates. Conversely, the existence of a detector for any desired performance requirement implies that the system possesses the S-Diagnosability property. Algorithms for determining the detection scheme parameters of detection threshold and detection delay bound for the specified MD and FA rates requirement are also presented, based on the construction of an extended observer. The extended observer computes, for each observation sequence, the set of states reached in the system model, along with their probabilities and the number of post-fault transitions executed. The algorithms are guaranteed to terminate and the required number of iterations, prior to termination, are reported as part of the correctness proof of the algorithms. Our detection strategy works for S-Diagnosable system as well as non-S-Diagnosable systems in the same manner. For S-Diagnosable systems it is possible to achieve arbitrary performance requirement for FA and MD rates, while for a non-S-Diagnosable system an arbitrary performance requirement is achievable only for the FA rate, whereas a lower bound exists for the achievable MD rate that is a function of

Manuscript received December 1, 2013; revised June 22, 2014; accepted November 28, 2014. Date of publication December 18, 2014; date of current version May 21, 2015. This work was supported in part by the National Science Foundation under the grants, NSF-ECCS-0801763, NSF-ECCS-0926029, and NSF-CCF-1331390. Recommended by Associate Editor D. Hristu-Varsakelis.

The authors are with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA (e-mail: junchen@iastate.edu; rkumar@iastate.edu).

Digital Object Identifier 10.1109/TAC.2014.2382991

the FA rate, and increases as FA rate is decreased. A variant of the above mentioned algorithm is also presented to compute an upper bound for the minimum achievable MD rate for a non-S-Diagnosable system.

The rest of this paper is organized as follows. The notations and some preliminaries are presented in Section II, followed by the proposed online fault detector and its existence condition in Section III. The offline algorithms for computing the detector parameters (detection threshold and detection delay) are proposed in Section IV and implemented in Section V using an illustrative example, where the simulation results are also provided. Section VI concludes the paper. The Appendix contains some review results and the proofs of the theorems.

II. Notations and Preliminaries

A. Stochastic DESs

For an event set Σ, let $\overline{\Sigma} := \Sigma \cup \{\epsilon\}$ denote the set of events plus ϵ, the "no-event". Let Σ^* and Σ^+ denote the set of all finite length event sequences over Σ, including and excluding ϵ respectively, i.e., $\Sigma^+ = \Sigma^* - \{\epsilon\}$. A member of Σ^* is called a *trace*. Denote as $s \leq t$ if $s \in \Sigma^*$ is a prefix of $t \in \Sigma^*$, and use $|s|$ to denote the number of events in s or the length of s. A subset of Σ^* is called *language*. For $L \subseteq \Sigma^*$, its prefix-closure, denoted as $pr(L)$, is defined as $pr(L) := \{s \in \Sigma^* | \exists t \in L : s \leq t\}$. L is said to be prefix-closed (or simply closed) if $pr(L) = L$, i.e., whenever L contains a trace, it also contains all the prefixes of that trace. For $s \in \Sigma^*$ and $L \subseteq \Sigma^*$, $L \setminus s$ denotes the set of traces in L after s and is defined as $L \setminus s := \{t \in \Sigma^* | st \in L\}$.

A stochastic DES can be modeled as a *stochastic automaton* G which is denoted by $G = (X, \Sigma, \alpha, x_0)$, where X is the set of states, Σ is the finite set of events, $x_0 \in X$ is the initial state, and $\alpha : X \times \Sigma \times X \to [0, 1]$ is the (total) transition probability function [37], satisfying: $\forall x \in X, \sum_{\sigma \in \Sigma} \sum_{x' \in X} \alpha(x, \sigma, x') = 1$. G is said to be non-stochastic if $\alpha : X \times \Sigma \times X \to \{0, 1\}$, and a non-stochastic DES is said to be deterministic if $\forall x \in X, \sigma \in \Sigma, \sum_{x' \in X} \alpha(x, \sigma, x') \leq 1$. The transition probability function α can be extended from domain $X \times \Sigma \times X$ to $X \times \Sigma^* \times X$ recursively as follows: $\forall x_i, x_j \in X, s \in \Sigma^*, \sigma \in \Sigma, \alpha(x_i, s\sigma, x_j) = \sum_{x_k \in X} \alpha(x_i, s, x_k)\alpha(x_k, \sigma, x_j)$, and $\alpha(x_i, \epsilon, x_j) = 1$ if $x_i = x_j$ and 0 otherwise. Define a *transition* in G as a triple $(x_i, \sigma, x_j) \in X \times \Sigma \times X$ and define the language generated by G as $L(G) := \{s \in \Sigma^* | \exists x \in X, \alpha(x_0, s, x) > 0\}$. Note that the transition function of G is a total function, and in a graphical representation of G, a transition $(x, \sigma, x') \in X \times \Sigma \times X$ is omitted if and only if $\alpha(x, \sigma, x') = 0$.

The observations of events are filtered through an observation mask, $M : \overline{\Sigma} \to \overline{\Delta}$, satisfying $M(\epsilon) = \epsilon$, where Δ is the set of observed symbols. An event σ is said to be unobservable if $M(\sigma) = \epsilon$. The set of unobservable events is denoted as Σ_{uo} and the set of observable events is then denoted by $\Sigma - \Sigma_{uo}$. The observation mask can be extended from domain Σ to Σ^* inductively as following: $M(\epsilon) = \epsilon$ and $\forall s \in \Sigma^*, \sigma \in \overline{\Sigma}, M(s\sigma) = M(s)M(\sigma)$.

Example 1: Fig. 1(a) is an example of a stochastic automaton G. The set of states is $X = \{0, 1, 2, 3\}$ with initial state $x_0 = 0$, event set $\Sigma = \{a, b, c, f\}$. A state is depicted as a node, whereas a transition is depicted as an edge between its origin and termination states, with its event name and probability value labeled on the edge. The observation mask M is such that $M(f) = \epsilon$ and for any other event σ, $M(\sigma) = \sigma$. ■

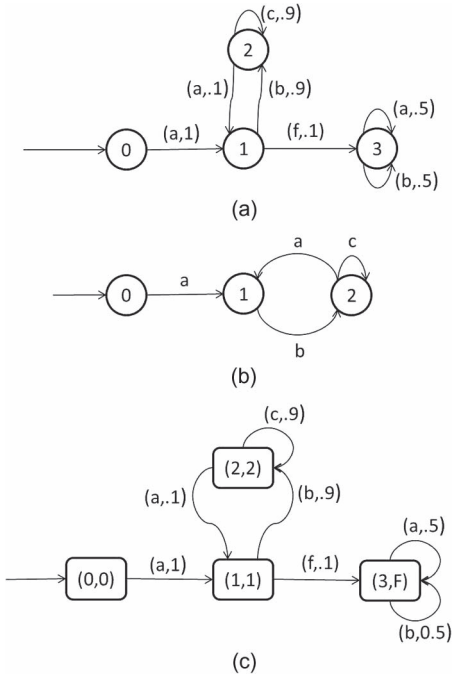

Fig. 1. (a) Stochastic automaton G; (b) deterministic nonfault specification R; (c) refined plant G^R.

B. Fault/Nonfault Behaviors and Refined Plant

In order to define the fault/nonfault behaviors of a stochastic automaton $G = (X, \Sigma, \alpha, x_0)$, its event set Σ is partitioned into fault events $\Sigma_f \subseteq \Sigma$ versus nonfault events $\Sigma - \Sigma_f$, where events in Σ_f are assumed to be unobservable. Then the overall behaviors of G is given by its generated language $L(G)$, whereas the set of nonfault behaviors of G is given by $K := L(G) \cap (\Sigma - \Sigma_f)^*$. The remaining behaviors $L - K$ are called the fault behaviors. Another approach to describing the fault/nonfault behaviors of a given stochastic automaton G is to specify the nonfault behaviors K in form of a *deterministic* automaton $R = (Q, \Sigma, \beta, q)$ such that $L(R) = K$ [26]. Note that β here is a binary-valued total function that is defined to be zero for fault events and so omitted in graphical representation for convenience. Then the refinement of G with respect to R, denoted as *refined plant* G^R, can be used to capture the traces violating the given specification in form of the reachability of a fault state and is given by $G^R := (Y, \Sigma, \gamma, (x_0, q_0))$, where $Y = X \times \overline{Q}$ and $\overline{Q} = Q \cup \{F\}$, and $\forall (x, \overline{q}), (x', \overline{q}') \in X \times \overline{Q}, \sigma \in \Sigma, \gamma((x, \overline{q}), \sigma, (x', \overline{q}')) = \alpha(x, \sigma, x')$ if the following holds:

$$(\overline{q}, \overline{q}' \in Q \wedge \beta(\overline{q}, \sigma, \overline{q}') > 0) \vee (\overline{q} = \overline{q}' = F) \vee \left(\overline{q}' = F \wedge \sum_{q \in Q} \beta(\overline{q}, \sigma, q) = 0 \right)$$

and otherwise $\gamma((x, \overline{q}), \sigma, (x', \overline{q}')) = 0$.

Then it can be seen that the refined plant G^R has the following properties: (1) the generated language of the refined plant G^R is the same as that generated by G, i.e., $L(G^R) = L(G)$; (2) any trace (system behavior) in $L(G)$ but not in $L(R)$ transitions the refined plant G^R to a fault state (a state containing F as its second coordinate); (3) the probability of occurrence of each trace in G^R is the same as that in G, i.e., $\sum_{x\in X}\alpha(x_0,s,x) = \sum_{y\in Y}\gamma((x_0,q_0),s,y)$.

For $y_i, y_j \in Y$ and $\delta \in \Delta$, define the set of traces originating at y_i, terminating at y_j and executing a sequence of unobservable events followed by a single observable event with observation δ as $L_{G^R}(y_i,\delta,y_j) := \{s \in \Sigma^* | s = u\sigma, M(u) = \epsilon, M(\sigma) = \delta, \gamma(y_i,s,y_j) > 0\}$. Define $\alpha(L_{G^R}(y_i,\delta,y_j)) := \sum_{s\in L_{G^R}(y_i,\delta,y_j)}\gamma(y_i,s,y_j)$ and denote it as $\mu_{i,\delta,j}$ for short, i.e., it is the probability of all traces originating at y_i, terminating at y_j and executing a sequence of unobservable events followed by a single observable event with observation δ. Also define $\lambda_{ij} = \sum_{\sigma\in\Sigma_{uo}}\gamma(y_i,\sigma,y_j)$ as the probability of transitioning from y_i to y_j while executing a single unobservable event. Then it can be seen that $\mu_{i,\delta,j} = \sum_k \lambda_{ik}\mu_{k,\delta,j} + \sum_{\sigma\in\Sigma:M(\sigma)=\delta}\gamma(y_i,\sigma,y_j)$, where the first term on the right hand side (RHS) corresponds to transitioning in at least two steps (i to intermediate k unobservably, and k to j with a single observation δ at the end), whereas the second term on RHS corresponds to transitioning in exactly one step. Thus for each $\delta \in \Delta$, given the values $\{\lambda_{ij} | i,j \in Y\}$ and $\{\sum_{\sigma\in\Sigma:M(\sigma)=\delta}\gamma(y_i,\sigma,y_j) | i,j \in Y\}$, all the probabilities $\{\mu_{i,\delta,j} | i,j \in Y\}$ can be found by solving the following matrix equation (see for example [38] for a similar matrix equation):

$$\boldsymbol{\mu}(\delta) = \boldsymbol{\lambda}\boldsymbol{\mu}(\delta) + \boldsymbol{\gamma}(\delta) \tag{1}$$

where $\boldsymbol{\mu}(\delta)$, $\boldsymbol{\lambda}$ and $\boldsymbol{\gamma}(\delta)$ are all $|Y| \times |Y|$ square matrices whose ijth elements are given by $\mu_{i,\delta,j}$, λ_{ij} and $\sum_{\sigma\in\Sigma:M(\sigma)=\delta}\gamma(y_i,\sigma,y_j)$, respectively.

Example 2: For system presented in Fig. 1(a), the deterministic nonfault specification R is given in Fig. 1(b). Then the refined plant G^R is shown in Fig. 1(c). Let the state space of G^R be $Y = \{y_1 = (0,0), y_2 = (1,1), y_3 = (2,2), y_4 = (3,F)\}$. By solving matrix (1), we get

$$\boldsymbol{\mu}(a) = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & .05 \\ 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & .5 \end{bmatrix}$$

$$\boldsymbol{\mu}(b) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & .9 & .05 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & .5 \end{bmatrix}$$

$$\boldsymbol{\mu}(c) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & .9 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

■

III. Stochastic Diagnosability and Online Detection

A. Stochastic Diagnosability of DESs

Let us recall the definition of S-Diagnosability [1] (referred as AA-diagnosability in [5]).

Definition 1: Given a stochastic DES with generated language L and a closed sublanguage K as a nonfault specification, (L,K) is said to be Stochastically Diagnosable, or simply S-Diagnosable, if

$$(\forall \tau > 0 \wedge\ \forall \rho > 0)(\exists n \in \mathbb{N})(\forall s \in L - K)$$
$$Pr(t : t \in L \setminus s, |t| \geq n, P_N(st) > \rho) < \tau$$

where $P_N : L - K \to [0,1]$ is a map that assigns to each fault trace $s \in L - K$, the probability of s being ambiguous, which is the probability of all indistinguishable nonfault traces, conditioned by the fact that ambiguity can only arise from indistinguishable system traces, and is given by

$$P_N(s) := Pr\left(u \in K | M(u) = M(s)\right) = \frac{Pr\left(u \in K : M(u) = M(s)\right)}{Pr\left(u \in L : M(u) = M(s)\right)}.$$

Note in the definition of $P_N(s)$, "|" denotes the conditioning operation. Algorithm for checking S-Diagnosability was also given in [1], which is recalled in Appendix A for reference below.

Example 3: It can be checked that system in Fig. 1(c) is S-Diagnosable. As can be seen that after the occurrence of fault if one continues to observe the system for enough number of transitions, then with high probability two consecutive a or two consecutive b will be observed, resolving the ambiguity that a fault occurred. ■

Here we present a new characterization of S-Diagnosability which states that the S-Diagnosability is lost if and only if there exists an indistinguishable pair of fault and nonfault traces such that all future observations have identical probability of being fault versus nonfault. The correctness proof is given in the Appendix.

Theorem 1: (L,K) is not S-Diagnosable if and only if

$$(\exists s \in L - K, s' \in K \text{ s.t. } M(s) = M(s'))\,(\forall o \in \Delta^*)$$
$$Pr\left(t : t \in (L-K) \setminus M^{-1}M(s), M(t) = o\right)$$
$$= Pr\left(t : t \in K \setminus M^{-1}M(s'), M(t) = o\right). \tag{2}$$

Remark 1: While the definition of S-Diagnosability applies to the set of fault traces $L - K$, Theorem 1 is symmetric with respect to fault and nonfault traces, and thus suggests that notion of diagnosability can also be defined for nonfault traces: $s \in K$ is not diagnosable if and only if there exists $s' \in (L-K) \cap M^{-1}M(s)$ such that for all future observations $o \in \Delta^*$, $Pr(M^{-1}(o) \cap K \setminus M^{-1}M(s)) = Pr(M^{-1}(o) \cap (L-K) \setminus M^{-1}M(s'))$. We denote the set of all non-diagnosable nonfault traces as $K^{nd} \subseteq K$. Clearly, for a S-Diagnosable system, $K^{nd} = \emptyset$.

B. Computation of Likelihood of No-Fault

When the system executes a trace $s \in L$, an observation $o = M(s)$ is received by a fault detector. In order to issue a fault-decision versus no-decision for the observation $o = M(s)$, we propose the detector to compute the likelihood of no-fault, and issue a fault-decision if this likelihood of no-fault is

small (i.e., below a suitable threshold), and otherwise issue no-decision. In this subsection, we present how this likelihood can be recursively computed. With a slight abuse of notation, we denote the no-fault likelihood function as, $P_N : M(L) \to [0,1]$ and define it to be the conditional probability of nonoccurrence of a fault following any observation $o \in M(L)$:

$$P_N(o) := Pr\left(u \in K | M(u) = o\right) = \frac{Pr\left(u \in K : M(u) = o\right)}{Pr\left(u \in L : M(u) = o\right)}.$$

Note that $P_N(o)$ is the probability of nonfault traces conditioned by the fact that ambiguity can only arise from the system traces that produce the observation o. In order to recursively compute P_N we proceed as follows. For a given refined plant G^R whose state space is partitioned into nonfault states versus fault states, we define a nonfault indication binary column vector $I_{nf} \in \{0,1\}^{|Y|\times 1}$, where an entry of 1 indicates a nonfault state. Also define state distribution vector $\boldsymbol{\pi} : M(L) \to [0,1]^{1\times|Y|}$, i.e., for each $o \in M(L)$, $\boldsymbol{\pi}(o)$ is the state distribution of G^R following the observation o. Then $\boldsymbol{\pi}(\cdot)$ is recursively computed as follows: $\boldsymbol{\pi}(\epsilon) = [1, 0, \ldots, 0]$, and for any $o \in M(L), \delta \in \Delta$:

$$\boldsymbol{\pi}(o\delta) = \frac{\boldsymbol{\pi}(o)\boldsymbol{\mu}(\delta)}{\|\boldsymbol{\pi}(o)\boldsymbol{\mu}(\delta)\|}$$

where $\boldsymbol{\mu}(\delta)$ is computed by solving matrix (1), and $\|\cdot\|$ is simply the sum of all vector elements. Then for an observation o, $P_N(o)$ is simply given by adding the probabilities of the nonfault states

$$P_N(o) = \boldsymbol{\pi}(o) I_{nf}$$

where note that $\boldsymbol{\pi}(o)$ and hence also $P_N(o)$ are recursively computed.

Example 4: In the system of Fig. 1(c), the indication vector is given as $I_{nf} = [1,1,1,0]^T$, and the state distribution vector is initialized as $\boldsymbol{\pi}(\epsilon) = [1,0,0,0]$. If $o = aba$, then $P_N(o) = 0.783$; if $o = ababc$, then $P_N(o) = 1$; if $o = abaa$, then $P_N(o) = 0$. ■

C. Online Detection Scheme

For issuing online detection decision, we propose a detector, $D : M(L) \to \{F, \epsilon\}$ that for each observation in $M(L)$ issues either a "fault (F)" decision or "no-decision (ϵ)" by comparing the likelihood of no-fault to a suitable threshold, as follows:

$$\forall o \in M(L), [D(o) = F] \Leftrightarrow [\exists \bar{o} \leq o : P_N(\bar{o}) \leq \rho_D] \quad (3)$$

where ρ_D is the detection threshold, appropriately chosen to meet the desired FA rate requirement. Note by definition, if a detection decision is F, then it remains F for all future observations, i.e., the detector "does not change its mind," which is expected for the case of permanent faults.

Remark 2: For given detector parameters, the detection scheme (3) requires solving (1) offline for each $\delta \in \Delta$, and computing online the likelihood of no-fault upon the arrival of a new observation. The former has the complexity of $O(|\Delta| \times |X|^3 \times |\overline{Q}|^3 + |\Sigma| \times |X|^2 \times |\overline{Q}|^2) \leq O(|\Sigma| \times |X|^3 \times |\overline{Q}|^3)$, whereas the latter requires an $O(|X|^2 \times |\overline{Q}|^2)$ complexity. Since (1) can be solved offline before the initialization of the online monitoring, the online monitoring and detection has a quadratic complexity.

Note a *false alarm* occurs if the detector D issues F while the refined plant is in a nonfault state; and dually a *missed detection* occurs if the detector D fails to issue a F decision within an appropriate delay bound n_D after the occurrence of a fault. So letting P_D^{md} and P_D^{fa} denote the MD and FA rates, respectively, of a detector D, we have

$$P_D^{md} := Pr\left(st \in L - K : s \in L - K, |t| \geq n_D, P_N\left(M(st)\right) > \rho_D\right) \quad (4)$$

$$P_D^{fa} := Pr\left(s \in K : P_N\left(M(s)\right) \leq \rho_D\right). \quad (5)$$

Example 5: For the refined plant of Fig. 1(c) which is S-Diagnosable, suppose we set the threshold $\rho_D = 0.8$. Then any nonfault trace in $a(bc^+a)^*ba \subset K$ will produce false alarm $(P_N(aba) = 0.783 < \rho_D)$, and thus $P_D^{fa}|_{\rho_D=0.8} = Pr(u \in a(bc^+a)^*ba) = 47.37\%$. On the other hand if we set $\rho_D = 0.5$, then any nonfault trace in $a(bc^+a)^*baba \subset K$ will produce false alarm $(P_N(ababa) = 0.488 < \rho_D)$, and thus $P_D^{fa}|_{\rho_D=0.5} = Pr(u \in a(bc^+a)^*baba) = 4.26\%$.

Now supposing that 4.26% FA rate is acceptable, so we fix the detection threshold ρ_D to 0.5. If the detection delay bound is set to be $n_D = 3$, then any fault trace $s \in a(bc^+a)^*fbab \in L - K$ will miss detection and thus the MD rate is given by $P_D^{md}|_{\rho_D=0.5,n_D=3} = 6.58\%$. On the other hand if the detection delay bound is set to be $n_D = 4$, then any fault trace $s \in L - K$ can be detected, i.e., $P_D^{md}|_{\rho_D=0.5,n_D=4} = 0$. ■

The following theorem provides insight into the significance of the S-Diagnosability property for the purpose of online fault detection, by showing its necessity and sufficiency for the existence of an online detector that can achieve any desired levels of MD and FA rates. The correctness proof of Theorem 2 is given in Appendix.

Theorem 2: (L, K) is S-Diagnosable if and only if for any FA rate requirement $\phi > 0$ and MD rate requirement $\tau > 0$, there exist a detection threshold $\rho_D > 0$ and a delay bound n_D such that $P_D^{fa} \leq \phi$ and $P_D^{md} \leq \tau$.

The reason any FA rate can be achieved is because fewer and fewer nonfault traces produce false alarm when the detection threshold is made smaller and smaller (and so by choosing a low enough threshold, it is always possible to ensure that any given FA rate requirement can be met). The achievement of an arbitrary MD rate is ensured by the definition of S-Diagnosability, which requires that the detection statistics, namely the likelihood of no-fault, always falls below any detection threshold, no matter how small, within some bounded delay, and with arbitrarily high probability.

IV. Computation of Detection Threshold and Delay

In the this section we provide algorithms for computing the parameters ρ_D and n_D so as to achieve the desired level of MD and FA rates.

A. Algorithms for ρ_D and n_D

Here we provide a brief outline of the algorithm that computes ρ_D: In order to compute detection threshold ρ_D for a given FA rate requirement ϕ, Algorithm 1 constructs an "extended observer tree," that for each observation sequence, estimates the states (as any observer does), and organizes it in a tree form where nodes are observations tagged with the estimated states and the edges are transitions on a next new observation, with the extension that each state in the estimate is labeled by the probability of reaching it. The construction of Algorithm 1 makes the "extended observation tree" formal. These probability labels are then used to compute the probability P_N for each observation, or equivalently, each node of the extended observer tree. The tree extends to a depth so that if no detection decision are made for any of the nodes (equivalently, corresponding observations) in the tree, then the FA rate caused by the detection decisions at the future successors is upper bounded by the desired rate ϕ. The existence of such a depth is guaranteed by Theorem 3, and to ensure no detection decision for any of the nodes in T, we simply choose the detection threshold to be smaller than the minimum P_N value among all nodes of T (recall by (3) that a detection decision is only issued when the P_N value falls below the threshold).

Algorithm 1: For a given refined plant G^R and a FA rate requirement ϕ, do the following:

1) This step is just a preparatory step to identify certain classes of states before beginning to construct an extended observer tree. Identify all the states in $X \times Q$ from which a fault state in G^R is reachable, and denote this set of states as Y_1 (these are nonfault states from where fault states are reachable, and correspond to states reached by traces in K_1 defined in the proof of Theorem 2). Identify $Y_{2,3} = X \times Q - Y_1$ (these are nonfault states reached by traces in $K_2 \cup K_3$ defined in the proof of Theorem 2).
2) Iteratively construct an extended observer tree T with set of nodes, $\overline{Z} = Z \times M(L)$, where $Z = 2^{((X\times\overline{Q})\times(0,1])}$, and the depth of tree grows by 1 in each iteration until the stopping criterion is satisfied—see below. Then each node of T is of the form $\overline{z} = (z, o(\overline{z}))$, where $o(\overline{z}) \in M(L)$ is a unique observation associated with the node $\overline{z}$ and $z = \{((x_i, \overline{q}_i), p_i)\} \subseteq (X \times \overline{Q}) \times (0, 1]$ is set of state estimates, with the ith one denotes $(x_i, \overline{q}_i)$, tagged with its occurrence probability p_i. The tree T is rooted at $\overline{z}_0 = \{((0,0),1),\epsilon\}$. $\overline{z}_2 \in \overline{Z}$ is a δ-child ($\delta \in \Delta = M(\Sigma) - \{\epsilon\}$) of $\overline{z}_1 \in \overline{Z}$ if and only if $o(\overline{z}_2) = o(\overline{z}_1)\delta$ and for every $((x_2, \overline{q}_2), p_2) \in z_2$, it holds that $p_2 = \sum_{((x_1,\overline{q}_1),p_1)\in z_1} \sum_{s\in\Sigma^*: M(s)=\delta} p_1 \times \gamma((x_1, \overline{q}_1), s, (x_2, \overline{q}_2))$. *It can be seen that* $((x_2, \overline{q}_2), p_2)$ *is included in* z_2 *if and only if* $(x_2, \overline{q}_2)$ *can be reached from a state included in* z_1 *following extra observation* δ *and* p_2 *is the probability of reaching* $(x_2, \overline{q}_2)$ *from initial state following the observation* $o(\overline{z}_2)$.

   Using the probability values of states in any node $\overline{z}$ of T, we can compute the likelihood of no-fault following the observation $o(\overline{z})$, by way of adding the probabilities of the non-fault states of the node, and next normalizing over all states of the node as follows:

$$\forall \overline{z} = (z, o(\overline{z})): \quad P_N(\overline{z}) := \frac{\sum_{((x,\overline{q}),p)\in\overline{z}, \overline{q}\neq F} p}{\sum_{((x,\overline{q}),p)\in\overline{z}} p}.$$

   Then $P_N(\overline{z})$ equals $P_N(o(\overline{z}))$, and corresponds to the conditional probability of no-fault given the observation $o(\overline{z})$.

   Terminate the tree at a uniform depth so the set of leaf nodes $\overline{Z}_m \subseteq \overline{Z}$ satisfy:
   - $(\overline{z}, \overline{z}' \in \overline{Z}_m) \Rightarrow (|o(\overline{z})| = |o(\overline{z}')| =: d_1)$ (each terminal node is reached after the same number of observations, which guarantees the uniformity of the depth of T, which we denote as d_1), and
   - $\sum_{\overline{z}\in\overline{Z}_m} \sum_{((x,\overline{q}),p)\in\overline{z}:(x,\overline{q})\in Y_1} p + \sum_{\overline{z}\in\overline{Z}_m: P_N(\overline{z})\leq\rho_{\min}} \sum_{((x,\overline{q}),p)\in\overline{z}:(x,\overline{q})\in Y_{2,3}} p < \phi$, where $\rho_{\min} := \min_{\overline{z}\in\overline{Z}: P_N(\overline{z})\neq 0} P_N(\overline{z})$ (for states in Y_1 contained in terminal nodes, their added probabilities of the first term equals $Pr(K_1 \cap M^{-1}(\Delta^{>d_1}))$, which upper bounds the FA rate of their successors (see proof of Theorem 2); for the states in $Y_{2,3}$ contained in the terminal nodes having $P_N \leq \rho_{\min}$, their added probabilities of the second term equals $Pr(s \in [K_2 \cup K_3] \cap M^{-1}(\Delta^{>d_1}) : P_N(M(s)) \leq \rho_{\min})$, which upper bounds the FA rate of their successors (see proof of Theorem 2); we require the combined upper bounds to be less that ϕ, which ensures that even if all successors produce false alarm, the FA rate requirement is still met).
3) Return any $\rho_D < \rho_{\min}$. (Note that with this choice of ρ_D, all nonfault traces whose observations are included in T will have no detection decisions (and so no false alarms either), and only their extensions can have detection decisions (some of which may be false alarms). But by construction, the probability of those extensions is upper bounded by ϕ, as desired.)

The following theorem guarantees the correctness of Algorithm 1. Correctness proof is given in the Appendix.

Theorem 3: There exists $d_1 \in \mathbb{N}$ such that Algorithm 1 terminates with tree depth d_1 and returns a threshold ρ_D under which the overall FA rate is upper bounded by ϕ.

Note as the tree depth is increased, the set of traces contained in the tree, and hence their probability, also grows. Since no detection decision is issued for traces in the tree, they don't produce any false alarms, and hence the false alarm rate is upper bounded by the probability of traces not included in the tree. By increasing the tree depth, we can essentially guarantee that this upper bound is as small as desired.

Example 6: For the system G^R shown in Fig. 1(c), $Y_1 = \{(0,0),(1,1),(2,2)\}$ and $Y_{2,3} = \emptyset$. We construct the extended observer tree for the computation of detection threshold; the first 4 steps of which are as shown in Fig. 2, where $P_N(\overline{z}_0) = P_N(\overline{z}_1) = 1$, $P_N(\overline{z}_2) = 0.9474$, $P_N(\overline{z}_3) = 0$, $P_N(\overline{z}_4) = 0.7826$, $P_N(\overline{z}_5) = 1$, $P_N(\overline{z}_6) = P_N(\overline{z}_7) = P_N(\overline{z}_8) = 0$. Selecting any $\rho_D < \min_{\overline{z}\in\overline{Z}: P_N(\overline{z})\neq 0} P_N(\overline{z}) = 0.7826$, the FA rate is upper bounded by $\sum_{\overline{z}\in\overline{Z}_m} \sum_{((x,\overline{q}),p)\in\overline{z}:(x,\overline{q})\in Y_1} p = 0.09 + 0.81 = 0.9$. Algorithm 1 would proceed to a next step unless this FA rate is found to be acceptable. ■

Having provided an algorithm to compute the detection threshold ρ_D that meets the FA rate requirement ϕ, we next present an algorithm to compute the delay bound n_D to satisfy the given MD rate requirement τ. Here we provide a brief outline of the algorithm: In order to compute delay bound n_D, Algorithm 2 constructs a refined version of the extended observer tree that for each observation sequence estimates the

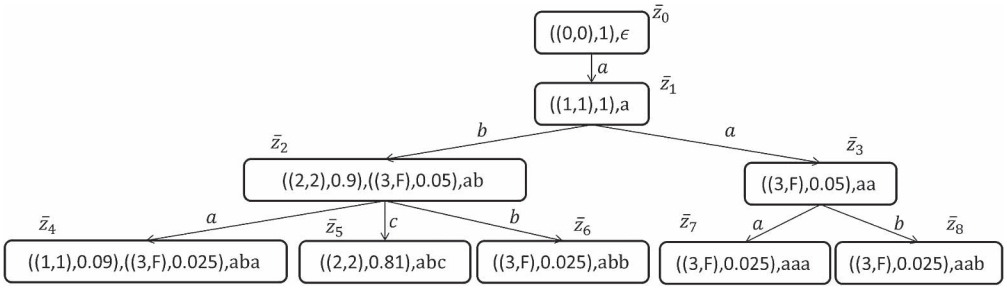

Fig. 2. Part of an extended observer tree for Example 6.

states and their probabilities, *with the refinement that keeps track of the number of post fault transitions executed for each state in the estimated state set*. The tree extends to a depth so that if no missed detections occur for any of the nodes in the tree, then the MD rate caused by the future successors is upper bounded by the desired rate τ. For S-Diagnosable systems, the existence of such a depth is guaranteed by Theorem 4, and to ensure no missed detection for any of the nodes in T, we simply choose n_D to be greater than the maximum number of post fault transitions among all nodes of T.

Algorithm 2: For a given refined plant G^R, a detection threshold ρ_D and a MD rate requirement τ, do the following:

1) Iteratively construct a refined extended observer tree T with set of nodes, $\overline{Z} = Z \times M(L)$, where $Z = 2^{((X\times\overline{Q})\times(0,1])}$ ($\mathbb{N} = \{0, 1, 2, \ldots\}$), and the depth of T grows by 1 in each iteration until the stopping criterion is satisfied—see below. Similar to Algorithm 1, each node of T is of the form $\overline{z} = (z, o(\overline{z}))$, where $z = \{((x_i, \overline{q}_i), p_i, n_i)\} \subseteq (X \times \overline{Q}) \times (0, 1] \times \mathbb{N}$, $o(\overline{z}) \in M(L)$ and the additional term n_i counts the number of post-fault transitions in reaching $(x_i, \overline{q}_i)$. The tree T is rooted at $\overline{z}_0 = \{((0, 0), 1, 0), \epsilon\}$. $\overline{z}_2 \in \overline{Z}$ is a δ-child ($\delta \in \Delta = M(\Sigma) - \{\epsilon\}$) of $\overline{z}_1 \in \overline{Z}$ if and only if $o(\overline{z}_2) = o(\overline{z}_1)\delta$, and for every $((x_2, \overline{q}_2), p_2, n_2) \in z_2$, it holds that $p_2 = \sum_{((x_1,\overline{q}_1),p_1,n_1)\in z_1} \sum_{s\in\Sigma^*: M(s)=\delta, \#\text{post}-\text{fault}(s,(x_1,\overline{q}_1))+n_1=n_2} p_1 \times \gamma((x_1, \overline{q}_1), s, (x_2, \overline{q}_2))$. Here "#post-fault" counts the number of events in s beyond a fault as follows: if $\overline{q}_1 = F$, it returns the value $|s|$, and otherwise it returns the number of transitions executed in s after a fault state is reached. *It can be seen that $((x_2, \overline{q}_2), p_2, n_2)$ is included in z_2 if and only if $(x_2, \overline{q}_2)$ can be reached from a state included in z_1 following extra observation δ, p_2 is the probability of reaching $(x_2, \overline{q}_2)$ from initial state following observation $o(\overline{z}_2)$ and n_2 is the number the post fault transitions executed.*

   For each node $\overline{z} = (z, o(\overline{z}))$, define the likelihood of no-fault given the observation $o(\overline{z})$ as in Algorithm 1:

$$P_N(\overline{z}) := \frac{\sum_{((x,\overline{q}),p,n)\in z, \overline{q}\neq F} p}{\sum_{((x,\overline{q}),p,n)\in z} p}.$$

   Terminate a branch of the tree if a detection decision has been made (P_N value smaller than ρ_D), and terminate the rest of the tree at a uniform depth so the set of leaf nodes $\overline{Z}_m \subseteq \overline{Z}$ satisfy:
   - $P_N(\overline{z}) \leq \rho_D$ (for these nodes detection decision can be issued, implying these nodes will have no missed detections), or
   - $\sum_{\overline{z}\in\overline{Z}_m: P_N(\overline{z})>\rho_D} \sum_{((x,\overline{q}),p,n)\in\overline{z}:(x,\overline{q})\in Y_1\vee\overline{q}=F} p < \tau$ (for these nodes, no detection decision will be issued since $P_N(\overline{z}) > \rho_D$, and by the choice of n_D in step 2 below there is no missed detection yet; so their added probabilities upper bounds the MD rate of their future successors, and the stopping criterion requires this to be below the desired value τ).

2) Return any $n_D > \max_{((x,\overline{q}),p,n)\in z, \overline{z}\in\overline{Z}} n$, and let d_2 denote the depth of tree T. Note that with this choice of n_D all fault traces, whose observations are included in T, will not miss detection. So clearly that the MD rate P_D^{md} is upper bounded by $\overline{P_D^{md}}$ given by:

$$\overline{P_D^{md}} := \sum_{\overline{z}\in\overline{Z}_m: P_N(\overline{z})>\rho_D} \sum_{((x,\overline{q}),p,n)\in\overline{z}:(x,\overline{q})\in Y_1\vee\overline{q}=F} p. \quad (6)$$

The following theorem guarantees the correctness of Algorithm 2. Correctness proof is given in the Appendix.

Theorem 4: For S-Diagnosable systems, there exists $d_2 \in \mathbb{N}$ such that Algorithm 2 terminates with tree depth d_2 and returns a delay bound n_D under which the overall MD rate is upper bounded by τ.

Note as before, as the tree depth is increased, the set of traces contained in the tree, and hence their probability, also grows. For all traces included in the tree, S-Diagnosability guarantees that a correct detection decision is issued within a bounded delay bound, and so any missed detection can only occur for those traces not included in the tree. So the MD rate is upper bounded by the probability of traces not included in the tree. By increasing the tree depth, we can essentially guarantee that this upper bound is as small as desired, and then read the detection delay of the traces included in the tree for which detection decision is made (i.e., whose P_N values are smaller than the detection threshold).

Example 7: For the system G^R in Fig. 1(c), and assuming detection threshold of $\rho_D = 0.7825$ as determined in Example 6, we construct the refined extended observer tree for the computation of delay bound; the first 5 steps of which are as shown in Fig. 3. Here $P_N(\overline{z}_0) = P_N(\overline{z}_1) = 1$, $P_N(\overline{z}_2) = 0.9474$, $P_N(\overline{z}_3) = 0$, $P_N(\overline{z}_4) = 0.7826$, $P_N(\overline{z}_5) = 0$, $P_N(\overline{z}_6) = 1$, $P_N(\overline{z}_7) = 0$, $P_N(\overline{z}_8) = 0.8265$, and $P_N(\overline{z}_9) = P_N(\overline{z}_{10}) = 1$. The branches of $\overline{z}_3$ and $\overline{z}_5$ terminate since the likelihood of no-fault is smaller than $\rho_D = 0.7826$, whereas the depth of the rest of the tree is 5. With $n_D = 1 + \max_{((x,\overline{q}),p,n)\in z, \overline{z}\in\overline{Z}} n = 4$, the MD rate is upper bounded by $\overline{P_D^{md}} = \sum_{\overline{z}\in\{\overline{z}_8,\overline{z}_9,\overline{z}_{10}\}} \sum_{((x,\overline{q}),p,n)\in\overline{z}:(x,\overline{q})\in Y_1} p = 0.081 + 0.0045 + 0.0125 + 0.081 + 0.729 = 0.908$. Algorithm 2

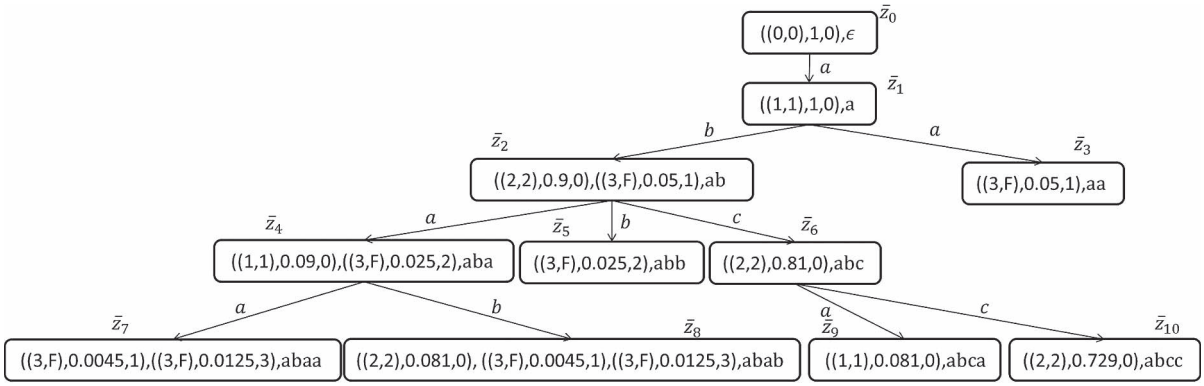

Fig. 3. Part of a refined extended observer tree for Example 7.

would proceed to a next step unless this MD rate is found to be acceptable. ■

Remark 3: Both Algorithm 1 and Algorithm 2 require the construction of an extended observer (with depths d_1 and d_2 and branching degree at most $|\Delta|$) that can have $O(|\Delta|^{d_1})$ and $O(|\Delta|^{d_2})$ nodes, respectively, and each node can have up to $|X| \times |\overline{Q}|$ elements. Therefore the complexity for *offline* computation for detection parameters ρ_D and n_D is $O(|X| \times |\overline{Q}| \times |\Delta|^d)$, where $d = \max\{d_1, d_2\}$. Note that d can depend on the system and specification models, the observation mask, and the desired bounds on MD and FA rates, and is bounded. On the other hand, as mentioned in Remark 2, the complexity of *online* monitoring is quadratic, $O(|X|^2 \times |\overline{Q}|^2)$.

B. Non-S-Diagnosable Systems

In the absence of S-Diagnosability, the termination of Algorithm 2 is not guaranteed, but a slight modification yields a terminating algorithm that finds an upper bound for the minimum achievable MD rate. In the case when the system is not S-Diagnosable, then (7) in the proof for Theorem 2 (see Appendix) may not hold for some $s \in L-K$. For given ϕ and τ, let ρ_D be chosen so that $P_D^{fa} \leq \phi$, and let $S_D^{nd} \subseteq L-K$ be the set of non-diagnosable fault traces for which there exists a MD rate $\tau' > 0$ such that the condition $Pr_D^{md}(S_D^{nd}) = Pr(st : s \in S_D^{nd}, t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) < \tau'$ is not satisfied by any $n_D \in \mathbb{N}$. Then for the traces in $(L-K)-S_D^{nd}$ there exists a detection delay bound n_D so that $\forall s \in (L-K)-S_D^{nd}$

$$Pr(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) < \tau'$$

and so the overall MD rate is upper bounded by:

$$\begin{aligned} P_D^{md} &= \sum_{s \in L-K} Pr_D^{md}(s) Pr(s) \\ &< \tau' Pr\left(L - K - S_D^{nd}\right) + Pr_D^{md}\left(S_D^{nd}\right) \\ &\leq \tau' + Pr_D^{md}\left(S_D^{nd}\right). \end{aligned}$$

Thus for non-S-Diagnosable systems, while any desired FA rate $\phi > 0$ can be always achieved by an appropriate choice of $\rho_D > 0$, a MD rate $\tau > 0$ can only be achieved if $\tau' + Pr_D^{md}(S_D^{nd}) \leq \tau$. Since n_D can be chosen to make τ' arbitrarily small, a MD rate $\tau > 0$ can be achieved if and only if $Pr_D^{md}(S_D^{nd}) < \tau$. This is captured in the following theorem, which generalizes Theorem 2 to the case of non-S-Diagnosable systems.

Theorem 5: Given a stochastic, nonfault specification-refined plant G^R with generated language L and nonfault behavior K, FA rate requirement $\phi > 0$ and MD rate requirement $\tau > 0$, there exists a detection threshold $\rho_D > 0$ such that $P_D^{fa} \leq \phi$, and for this detection threshold there exists a detection delay bound n_D such that $P_D^{md} \leq \tau$ if and only if $Pr_D^{md}(S_D^{nd}) \leq \tau$, where $S_D^{nd} \subseteq L - K$ is the set of non-diagnosable fault traces for which there exists $\tau' > 0$ such that the condition $Pr(st : s \in S_D^{nd}, t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) < \tau'$ is not satisfied by any $n_D \in \mathbb{N}$.

Remark 4: For a fixed FA rate, $Pr_D^{md}(S_D^{nd})$ is also fixed and serves as a lower bound for MD rate that the detection scheme can achieve. Note that $Pr_D^{md}(S_D^{nd})$ is a function of the FA rate requirement ϕ. When ϕ is made tighter by decreasing it, a smaller ρ_D is needed, and the resulting non-diagnosable fault traces subsume those corresponding to larger ρ_D. Therefore the minimum achievable MD rate increases as FA rate is made stringent by decreasing it.

Next we present a variant of Algorithm 2 that for a fixed threshold ρ_D computes an upper bound for $Pr_D^{md}(S_D^{nd})$. Algorithm 3 iteratively builds a refined extended observer tree T, and at each step computes an upper bound for the MD rate that either decreases or remains constant from one iteration to the next. When the latter happens, a future iteration may eventually decrease the bound, but since the optimal (least) upper bound is unknown, it is also not known how long one should continue iterating. So, to ensure termination, we adopt the heuristics of terminating the algorithm when the upper bound remains constant while n_D gets doubled.

Algorithm 3: For a given refined plant G^R and a threshold ρ_D, do the following:

1) Iteratively construct a refined extended observer tree T as in the step 1 of Algorithm 2;
2) For each depth of the tree T, set $n_D = 1 + \max_{((x,\overline{q}),p,n) \in z, \overline{z} \in \overline{Z}} n$ and compute an upper bound $\overline{P_D^{md}}$ for MD rate P_D^{md} according to (6);
3) If the upper bound $\overline{P_D^{md}}$ doesn't decrease while n_D computed in step 2 gets doubled over any two iteration steps (not necessarily consecutive), stop and return this upper bound.

V. Illustrative Example

We consider the problem of leakage detection in a two-tank system as shown in Fig. 4, which is adopted from [39]. The tanks

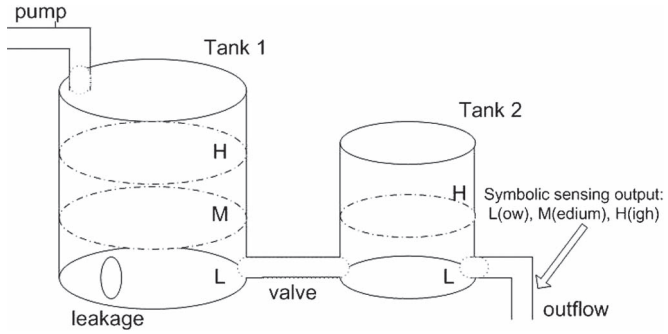

Fig. 4. Two-tank system. The fault to be detected is the leakage in the left tank with an occurrence probability of 0.05.

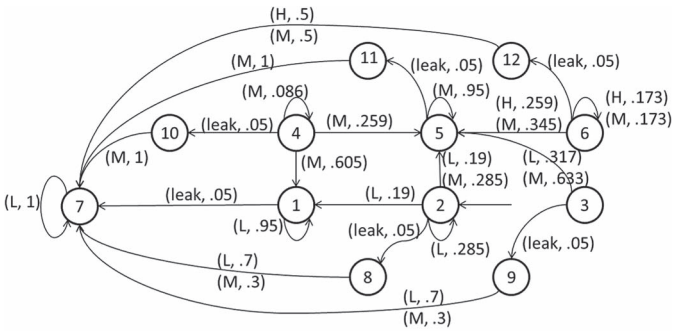

(a)

Nonfault state	Fault state	Water level in left tank	Water level in right tank
1	7	LOW	LOW
2	8	MEDIUM	LOW
3	9	HIGH	LOW
4	10	LOW	HIGH
5	11	MEDIUM	HIGH
6	12	HIGH	HIGH

(b)

Fig. 5. (a) Stochastic automaton G for the two-tank system shown in Fig. 4; (b) interpretation of states.

are connected with a valve. The water is pumped into the system in the left tank at a constant rate and outflows from the right tank. The only observation produced by this system is the symbolic sensor output (Low, Medium, High) which measures the outflow rate of the right tank at discrete times. There is a 0.05 probability that a leakage occurs in the left tank, which is to be detected. The aforementioned system is described by the stochastic automaton shown in Fig. 5(a), where the event set is $\Sigma=\{L, M, H, \text{leak}\}$, corresponding to the sensor outputs and the occurrence of leakage. All events except "leak" are fully observable, whereas "leak" is fully unobservable, i.e., $\Sigma_{uo}=\{\text{leak}\}$. The water levels in the tanks are quantized into "LOW," "MEDIUM" and "HIGH" for the left tank, and just "LOW" and "HIGH" for the right tank, and each state in the stochastic automaton denotes a combination of these water levels along with a record whether a leak occurred in past, summarized in Fig. 5(b). The system is initialized at state 2, i.e., medium level of water in the left tank and low level of water in the right tank. The states $\{1,\ldots,6\}$ are pre-fault states and states $\{i+6, i=1,\ldots,6\}$ are post-fault states, and so the nonfault specification is simply a subautomaton of the plant automaton restricted to the pre-fault states, and without the probability labels. The possibility of occurrence of leakage at each pre-fault state $i, i=1,\ldots,6$, is captured by the transition from state i to state $i+6$ labeled with the event "leak" and occurrence probability 0.05. The transitions are obtained by way of abstraction, and for further details readers are referred to [39]–[41]. It can be checked that the system is S-Diagnosable, so Theorem 2–4 apply.

TABLE I
COMPUTATIONAL RESULTS OF ALGORITHM 1

FA rate ϕ	Threshold ρ_D	Tree depth d_1	# of nodes	Running time (sec.)
0.95	0.8717	2	7	0.004
0.9	0.8004	3	14	0.007
0.85	0.7473	4	25	0.016
0.8	0.7051	5	41	0.019
0.75	0.6682	6	63	0.028
0.7	0.6343	7	92	0.048
0.65	0.5722	9	175	0.074
0.6	0.5436	10	231	0.097
0.55	0.4906	12	377	0.158
0.5	0.4428	14	575	0.249
0.45	0.3996	16	833	0.383
0.4	0.3606	18	1159	0.558
0.35	0.3092	21	1793	0.984
0.3	0.2651	24	2625	1.582
0.25	0.2159	28	4089	2.908
0.2	0.1759	32	6017	5.211
0.15	0.1361	37	9177	10.64
0.1	0.0903	45	16261	32.94
0.05	0.0440	59	36050	182.3

TABLE II
COMPUTATIONAL RESULTS OF ALGORITHM 2 WITH $\rho_D = 0.044$

MD rate τ	Delay n_D	Tree depth d_2	# of nodes	Running time (sec.)
0.95	4	4	21	0.018
0.9	5	5	31	0.029
0.85	7	7	57	0.052
0.8	9	9	91	0.084
0.75	11	11	133	0.126
0.7	14	14	211	0.217
0.65	16	16	273	0.299
0.6	19	19	381	0.436
0.55	23	23	553	0.688
0.5	28	28	813	1.150
0.45	34	34	1191	1.828
0.4	43	43	1893	3.356
0.35	60	60	3661	8.068
0.3	60	60	3661	8.056
0.25	60	60	3661	8.114
0.2	60	60	3661	8.060
0.15	60	60	3661	8.038
0.1	60	60	3661	8.027
0.05	60	60	3661	8.028

We implement the proposed Algorithms 1 and 2 to compute the detection threshold ρ_D and delay bound n_D to ensure any given FA and MD rate requirements. The results are shown in Tables I and II and Fig. 6. Table I lists for various FA rates the detection threshold ρ_D returned by Algorithm 1, as well as the tree depth d_1, the number of tree nodes and the running time of the implementation of Algorithm 1 on a standard desktop PC; and the first two columns is plotted in Fig. 6(a). For example, when the FA rate is required to be under 5%, the detection threshold returned by Algorithm 1 is $\rho_D=0.044$. When we fix $\rho_D=0.044$, i.e., fix $\phi=5\%$, the delay bound n_D returned by Algorithm 2 for various MD rates is shown in Table II and Fig. 6(b); the table additionally lists for each MD rate the tree depth d_2, the number of tree nodes and the running time of the

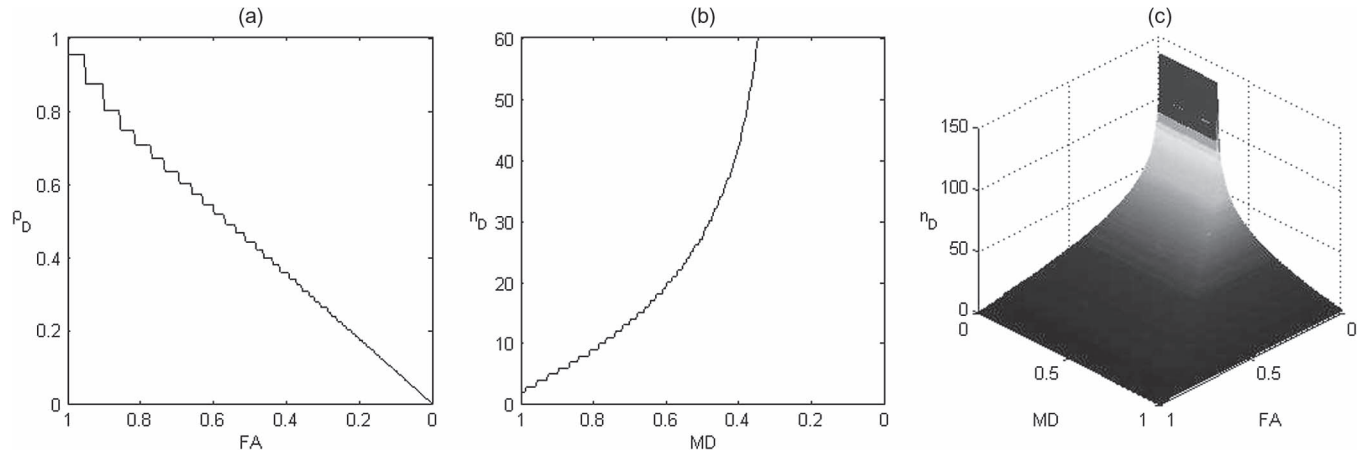

Fig. 6. Smulation results for leakage detection in two-tank system: (a) the detection threshold ρ_D as a function of ϕ; (b) the delay bound n_D as a function of τ, when $\rho_D = 0.044$ ($\phi = 5\%$); (c) n_D as a function of both ϕ and τ.

implementation of Algorithm 2 on a standard desktop PC. As can be seen, when the MD rate is required to be under 5%, the detection delay bound returned by Algorithm 2 is $n_D = 60$. If we wish to decrease the detection delay bound, then the upper bound for the MD rate will increase and possibly violate the MD rate requirement of 5%. For example if we choose $n_D = 55$, then it could only be assured that the MD rate is upper bounded by 36.24%. Recall by previous discussion, the delay bound can depend on both FA rate ϕ and MD rate τ, and this dependency is shown in Fig. 6(c). This figure along with Fig. 6(a) can be used to determine the parameters ρ_D and n_D for the specified FA and MD rates for the two-tank example. It so happens that for $n_D = 59$, the upper bound given by (6) is higher than 35%, whereas it suddenly becomes lower than 5% for $n_D = 60$. This sudden drop in upper bound explains the reason why the tree depth saturates at 60 when MD rate is decreased from 35% to 5%.

VI. Conclusion

We studied the problem of online fault detection for stochastic DESs. An online detector based on a recursive likelihood computation was proposed, whose existence for achieving any arbitrary performance requirement was shown to be equivalent to the S-Diagnosability property. Algorithms for computing the detector parameters of detection threshold and delay bound so as to achieve a given performance requirement of false alarm and missed detection rates were presented, using a proposed procedure for constructing an extended observer. The extended observer computes, for each observation sequence, the set of states reached in the system model, along with their probabilities and the number of post fault transitions executed. The algorithms are guaranteed to terminate and upper bounds on the number of iterations prior to termination were provided.

The detector has a quadratic complexity for the *online* monitoring, likelihood computation and issuing decision upon the arrival of a new observation. On the other hand, the *offline* computation of the detector parameters, namely, detection threshold and delay bound, requires constructing an extended observer whose size is exponential in the depth of the observer tree constructed, while the depth of the tree is a complex bounded function of the system and specification models, the observation mask, and the desired bounds on MD and FA rates. As can be inferred by Section V, the detector parameters of detection threshold and delay bound for various levels of MD and FA rates can be computed offline and stored in a database, and during online monitoring and detection the required set of parameters can be simply looked up each time a new level of MD and FA rates are specified.

It was also shown that our detection strategy works for S-Diagnosable as well as non-S-Diagnosable systems in the same manner. For S-Diagnosable systems it is possible to achieve arbitrary performance for FA and MD rates, while for a non-S-Diagnosable system an arbitrary performance is achievable only for the FA rate, whereas a lower bound exists for the achievable MD rate that is a function of the FA rate, and increases as FA rate is decreased. A variant of the algorithm for the S-Diagnosable case was used to compute an upper bound for the minimum achievable missed detection rate for a non-S-Diagnosable system.

Appendix

The following theorem is reproduced from [1, Theorem 3], and is used in the proofs included in this appendix. Note that s_1 (resp. s_2) denotes a trace generated in A_1 (resp. A_2).

Theorem 6 ([1]): Given two irreducible finite-state automata A_1 and A_2, where their initial state distributions are the same as their stationary state distributions, if A_1 and A_2 are not p-equivalent, then

$$(\forall \tau > 0 \wedge\ \forall \rho > 0)(\exists n \in \mathbb{N}) \\ Pr\left(s_1 : |s_1| > n, Pr\left(s_2 | M(s_1) = M(s_2)\right) > \rho\right) < \tau.$$

Next we provide complete proofs for the Theorems 1–4.

Proof of Theorem 1: (Sufficiency) If (2) is true, denote s and s' be such that (2) holds. Then for any extension t of s, it holds that, $Pr((L-K) \cap M^{-1}M(st)) = Pr(M^{-1}M(s) \cap (L-K))Pr(M^{-1}M(t) \cap (L-K) \setminus M^{-1}M(s))$ and $Pr(K \cap M^{-1}\,M\,(st))) = Pr\left(M^{-1}\,M(s) \cap K\right) Pr\left(M^{-1}M\,(t) \cap K \setminus M^{-1}\,M\,(s)\right)$. Therefore

$$\begin{aligned} P_N(st) &= \frac{Pr\left(K \cap M^{-1}M(st)\right)}{Pr\left((L-K) \cap M^{-1}M(st)\right) + Pr\left(K \cap M^{-1}M(st)\right)} \\ &= \frac{Pr\left(K \cap M^{-1}M(s)\right)}{Pr\left((L-K) \cap M^{-1}M(s)\right) + Pr\left(K \cap M^{-1}M(s)\right)} \\ &= \frac{Pr\left(K \cap M^{-1}M(s)\right)}{Pr\left(M^{-1}M(s)\right)} \end{aligned}$$

i.e., $P_N(st) = P_N(s) > 0$. Note that the above equality utilizes the fact that s and s' satisfy (2) and so $Pr(M^{-1}M(t) \cap (L - K) \setminus M^{-1}M(s)) = Pr(M^{-1}M(t) \cap K \setminus M^{-1}M(s))$. Now let $0 < \rho < P_N(s)$ and $0 < \tau < 1$. Then we have:

$$(\forall n \in \mathbb{N}) Pr(t : t \in L \setminus s, |t| \geq n, P_N(st) > \rho) = 1 > \tau.$$

It follows that the system is not S-Diagnosable.

(Necessity) If (2) is not true, then for all indistinguishable pairs of fault and nonfault traces (s, s'), there exists a future observation that has different probability of being fault versus nonfault, i.e.,

$$(\forall s \in L - K, s' \in K \text{ s.t. } M(s) = M(s'))\,(\exists o \in \Delta^*)$$
$$Pr\,(t : t \in L \setminus s, M(t) = o) \neq Pr\,(t : t \in K \setminus s', M(t) = o)\,.$$

Then according to the likelihood ratio test presented in Theorem 6 (originally [1, Theorem 3]), after the occurrence of any fault trace, by comparing the number of occurrences of the minimal segment of observations that has different probability of being fault versus nonfault, the ambiguity of the occurrence of a fault decreases as the length of extension increases, i.e., there exists $n \in \mathbb{N}$ such that for all $\rho > 0$, $\tau > 0$ and $s \in L - K$, the extensions of s longer than n and having P_N larger than ρ occur with probability smaller than τ, i.e.,

$$(\forall \tau > 0 \wedge\ \forall \rho > 0)(\exists n \in \mathbb{N})(\forall s \in L - K)$$
$$Pr\,(t : t \in L \setminus s, |t| \geq n, P_N(st) > \rho) < \tau.$$

Thus we can conclude that the system is S-Diagnosable. ■

Proof for Theorem 2: (Sufficiency) For a S-Diagnosable system (L, K), we need to show the existence of ρ_D and n_D for achieving given ϕ and τ.

For finding ρ_D, first we partition the set of nonfault traces into three sub-languages, i.e., $K = K_1 \cup K_2 \cup K_3$, where K_1 is the set possessing a fault extension $(K_1 = K \cap pr(L - K))$, K_2 is the set with no fault extension and is non-diagnosable $(K_2 = K^{nd} - K_1)$, and $K_3 = K - K_1 - K_2$ is the set with no fault extension and is diagnosable. Note that if a nonfault trace has a fault extension, then it can not satisfy condition in (2) and hence is diagnosable. In other words, $K^{nd} \cap K_1 = \emptyset$. Therefore $K_2 = K^{nd} - K_1 = K^{nd}$. Also note if (L, K) is diagnosable, then $K_2 = K^{nd} = \emptyset$.

For the nonfault traces in $K_1 = K \cap pr(L - K)$ that possess a fault extension, nonfault-ness is a transient property, and so for any $\phi_1 > 0$ there exists $m_1 \in \mathbb{N}$ such that the traces in K_1 that are longer than m_1 occur with probability smaller than ϕ_1. Denote $\rho_1 = \min_{s \in K_1, |s| \leq m_1} P_N(M(s))$. Since for a nonfault trace s, $P_N(M(s)) > 0$, and since the traces of length smaller than m_1 are finite, $\rho_1 > 0$. By choosing $\rho_D < \rho_1$ we can ensure that the detector issues a decision for only the traces in K_1 that are longer than m_1. (For shorter traces, P_N value will be larger than $\rho_1 > \rho_D$, and so no decision.) Since the probability of such traces is smaller than ϕ_1, their FA rate is also smaller than ϕ_1.

For the nonfault traces in K_2 that possess no fault extensions and are non-diagnosable, there exists $m_2 \in \mathbb{N}$ such that for every trace in K_2 that is longer than m_2, further extensions will not change the P_N value (i.e., P_N will converge to a value smaller than 1; otherwise the traces would be diagnosable). Denote $\rho_2 = \min_{s \in K_2, |s| \leq m_2} P_N(M(s))$. Similar to ρ_1, we have $\rho_2 > 0$. By choosing $\rho_D < \rho_2$ we can ensure the detector issues no decision for traces in K_2 and hence no false alarm in K_2.

For the nonfault traces in K_3 that possess no fault extensions and are diagnosable, according to Theorem 6 (originally [1, Theorem 3]), for any $\phi_3 > 0$ and $\rho_3' \in (0, 1)$ there exists $m_3 \in \mathbb{N}$ such that the traces longer than m_3 and having P_N value smaller than ρ_3' occur with probability smaller than ϕ_3. Denote $\rho_3'' = \min_{s \in K_3, |s| \leq m_3} P_N(M(s))$. Similar to ρ_1 and ρ_2, we have $\rho_3'' > 0$. By choosing $\rho_D < \rho_3 = \min(\rho_3', \rho_3'')$ we can ensure that the detector issues a decision only for those traces in K_3 that are longer than m_3 and have P_N value smaller than $\rho_D < \rho_3'$. Since the probability of such traces is smaller than ϕ_3, their FA rate is smaller than ϕ_3.

Therefore for any system (*regardless* whether or not it is S-Diagnosable), if we choose ϕ_1 and ϕ_3 in such a way that $\phi_1 + \phi_3 \leq \phi$ and accordingly set $\rho_D = \min_{i=\{1,2,3\}} \rho_i$, then the overall FA rate will be given by: $P_D^{fa} \leq \phi_1 + \phi_3 \leq \phi$. Thus using our detection scheme, any FA rate can be achieved for any system (regardless of whether or not it is S-Diagnosable), while as will be seen later, this is not the case for the MD rate.

Next we need to establish the existence of n_D to meet the MD rate requirement. Since the system is S-Diagnosable, for any $\tau > 0$ and $\rho_D > 0$ that guarantees FA rate, there always exists $n_D \in \mathbb{N}$ such that $\forall s \in L - K$

$$Pr\,(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) < \tau. \tag{7}$$

With such a choice of n_D we have, $Pr_D^{md}(s) < \tau$, and so the overall MD rate is bounded by: $P_D^{md} = \sum_{s \in L-K} Pr_D^{md}(s)\, Pr(s) < \tau Pr(L - K) \leq \tau$. Thus the sufficiency of Theorem 2 holds.

(Necessity) Suppose for a system (L, K), given any $\phi > 0$ and $\tau > 0$, there exist ρ_D and n_D such that $P_D^{fa} \leq \phi$ and $P_D^{md} \leq \tau$. Letting $S_D^{md} = \{st : s \in L - K, t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D\} \subseteq L - K$ denote the set of fault traces that miss detection, according to (4), we have $Pr(S_D^{md}) = P_D^{md} \leq \tau$. Then for given $s \in S_D^{md} \subseteq L - K$, we have

$$Pr\,(st : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) \leq \tau.$$

Since the LHS is the same as $Pr(s) Pr(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D)$, for any $s \in S_D^{md}$, we have:

$$Pr\,(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) \leq \frac{\tau}{Pr(s)}.$$

Let $p = \min_{s \in S_D^{md}} Pr(s)$ and $\tau' = 2\tau/p$, then for any $s \in S_D^{md}$, we have $Pr(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) < \tau'$. Note τ can be chosen to be arbitrarily small to make τ' arbitrarily small. Furthermore for any $s \in (L - K) - S_D^{md}$, we have $Pr(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) = 0 < \tau'$. Then $\forall s \in L - K$

$$Pr\,(t : t \in L \setminus s, |t| \geq n_D, P_N(st) > \rho_D) < \tau'. \tag{8}$$

Since for any $\phi > 0$ (and hence any ρ_D) and $\tau > 0$ (and hence any $\tau' > 0$), such n_D always exists to make the above analysis true, then for any $\rho_D > 0$ and $\tau' > 0$, there exists $n_D \in \mathbb{N}$ such that (8) holds, which indicates the condition for S-Diagnosability is held. Thus the necessity of Theorem 2 holds. ■

Proof for Theorem 3: According to the proof of Theorem 2, for $i \in \{1,2,3\}$, there exists m_i such that ρ_i obtained by examining traces in K_i shorter than m_i ensures the FA rate of K_i be smaller than ϕ_i. Since $\phi_2 = 0$ (none of the traces in K_2 produce false alarm because no decision is issued for those traces), by choosing ϕ_1 and ϕ_3 such that $\phi_1 + \phi_3 \leq \phi$, the requirement of the specified FA rate is met. It follows that Algorithm 1 is guaranteed to terminate with tree depth $d_1 \leq \max_i m_i$, returning a threshold $\rho_D \leq \min_i \rho_i$ such that the overall FA rate is upper bounded by ϕ. ■

Proof for Theorem 4: In the tree of Algorithm 2, a node is deemed a leaf if the "F" decision is made upon reaching it, and otherwise the tree itself is terminated at a uniform depth so that the upper bound for the MD rate has dropped below the requirement τ. Expand (6) and we have $\overline{P_D^{md}} = \sum_{\overline{z}\in\overline{Z}_m : P_N(\overline{z})>\rho_D} \sum_{((x,\overline{q}),p,n)\in\overline{z}:(x,\overline{q})\in Y_1} p + \sum_{\overline{z}\in\overline{Z}_m : P_N(\overline{z})>\rho_D} \sum_{((x,\overline{q}),p,n)\in\overline{z}:\overline{q}=F} p$. Similar to the proof of Theorem 2, the nonfault-ness in K_1 is a transient property, and so for any $\tau_1 > 0$, there exists $m' \in \mathbb{N}$ such that $Pr(s \in K \cap pr(L-K) : |s| \geq m') < \tau_1$, and hence the first term on the RHS is less than τ_1. For S-Diagnosable systems, according to Theorem 2, for any $\tau_2 > 0$ there exists n'_D such that with this choice of delay bound, the second term on the RHS is less than τ_2. Therefore by choosing τ_1 and τ_2 such that $\tau_1 + \tau_2 \leq \tau$, Algorithm 2 is guaranteed to terminate with tree depth $d_2 \leq m' + n'_D$, returning a delay bound $n_D = 1 + \max_{((x,\overline{q}),p,n)\in z, \overline{z}\in\overline{Z}} n$ such that the overall MD rate is upper bounded by τ. ■

REFERENCES

[1] J. Chen and R. Kumar, "Polynomial test for stochastic diagnosability of discrete event systems," *IEEE Trans. Auto. Sci. Eng.*, vol. 10, no. 4, pp. 969–979, Oct. 2013.

[2] J. Chen, C. Keroglou, C. N. Hadjicostis, and R. Kumar, "Corrections to 'polynomial test for stochastic diagnosability of discrete-event systems'," *IEEE Trans. Auto. Sci. Eng.*, submitted for publication.

[3] W. Qiu, Q. Wen, and R. Kumar, "Decentralized diagnosis of event-driven systems for safely reacting to failures," *IEEE Trans. Auto. Sci. Eng.*, vol. 6, no. 2, pp. 362–366, Apr. 2009.

[4] J. Chen and R. Kumar, "Online failure diagnosis of stochastic discrete event systems," in *Proc. IEEE Multi-Conf. Syst. Control*, Hyderabad, India, Aug. 2013, pp. 194–199.

[5] D. Thorsley and D. Teneketzis, "Diagnosability of stochastic discrete-event systems," *IEEE Trans. Autom. Control*, vol. 50, no. 4, pp. 476–492, Apr. 2005.

[6] S. Jiang, R. Kumar, and H. E. Garcia, "Diagnosis of repeated/intermittent failures in discrete event systems," *IEEE Trans. Robot. Automat.*, vol. 19, no. 2, pp. 310–323, Apr. 2003.

[7] T. S. Yoo and H. E. Garcia, "Diagnosis of behaviors of interest in partially-observed discrete-event systems," *Syst. Control Lett.*, vol. 57, no. 12, pp. 1023–1029, 2008.

[8] T. Yoo and H. Garcia, "Stochastic event counter for discrete-event systems under unreliable observations," in *Proc. Amer. Control Conf.*, Seattle, WA, USA, Jun. 2008, pp. 1145–1152.

[9] I. Hwang, S. Kim, Y. Kim, and C. E. Seah, "A survey of fault detection, isolation, and reconfiguration methods," *IEEE Trans. Control Syst. Technol.*, vol. 18, no. 3, pp. 636–653, May 2010.

[10] H. E. Garcia and T.-S. Yoo, "Model-based detection of routing events in discrete flow networks," *Automatica*, vol. 41, no. 4, pp. 583–594, Oct. 2005.

[11] X. Zhang, M. M. Polycarpou, and T. Parisini, "A robust detection and isolation scheme for abrupt and incipient faults in nonlinear systems," *IEEE Trans. Autom. Control*, vol. 47, no. 4, pp. 576–593, Apr. 2002.

[12] C. Zhou, R. Kumar, and S. Jiang, "Keynote: Hierarchical fault detection in embedded control software," in *Proc. IEEE Int. Comp. Softw. Appl. Conf.*, Jul. 2008, pp. 816–823.

[13] R. Isermann, R. Schwarz, and S. Stolzl, "Fault-tolerant drive-by-wire systems," *IEEE Control Syst. Mag.*, vol. 27, no. 5, pp. 64–81, Oct. 2002.

[14] U. Lerner, R. Parr, D. Koller, and G. Biswas, "Bayesian fault detection and diagnosis in dynamic systems," in *Proc. Nat. Conf. Artif. Intell. (AAAI-00)*, Austin, TX, USA, Aug. 2000, pp. 531–537.

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

[16] W.-C. Lin, H. E. Garcia, and T.-S. Yoo, "A diagnoser algorithm for anomaly detection in DEDS under partial and unreliable observations: Characterization and inclusion in sensor configuration optimization," *Discrete Event Dyn. Syst.*, vol. 23, no. 1, pp. 61–91, Mar. 2013.

[17] S. H. Zad, R. Kwong, and W. Wonham, "Fault diagnosis in timed discrete-event systems," in *Proc. 38th IEEE Conf. Decision Control*, Phoenix, AZ, USA, Dec. 1999, pp. 1756–1761.

[18] G. Westerman, R. Kumar, C. Stroud, and J. Heath, "Discrete event system approach for delay fault analysis in digital circuits," in *Proc. Amer. Control Conf.*, Philadelphia, PA, USA, Jun. 1998, pp. 239–243.

[19] D. Pandalai and L. Holloway, "Template languages for fault monitoring of timed discrete-event systems," *IEEE Trans. Autom. Control*, vol. 45, no. 5, pp. 868–882, May 2000.

[20] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis, "Diagnosability of discrete-event systems," *IEEE Trans. Autom. Control*, vol. 40, no. 9, pp. 1555–1575, Sep. 1995.

[21] S. Jiang, Z. Huang, V. Chandra, and R. Kumar, "A polynomial algorithm for testing diagnosability of discrete-event systems," *IEEE Trans. Autom. Control*, vol. 46, no. 8, pp. 1318–1321, Aug. 2001.

[22] T.-S. Yoo and S. Lafortune, "Polynomial-time verification of diagnosability of partially observed discrete-event systems," *IEEE Trans. Autom. Control*, vol. 47, no. 9, pp. 1491–1495, Sep. 2002.

[23] J. Lunze, "Fault diagnosis of discretely controlled continuous systems by means of discrete-event models," *Discrete Event Dyn. Syst.*, vol. 18, no. 2, pp. 181–210, 2008.

[24] S. Bhattacharyya, Z. Huang, V. Chandra, and R. Kumar, "Discrete event systems approach to network fault management: Detection & diagnosis of faults," in *Proc. Amer. Control Conf.*, Boston, MA, USA, Jun. 30–Jul. 2, 2004, pp. 5108–5113.

[25] R. Kumar and S. Takai, "A framework for control-reconfiguration following fault-detection in discrete event systems," in *Proc. 8th IFAC Symp. Fault Detection, Supervision Safety Tech. Processes (SafeProcess)*, Mexico City, Mexico, Aug. 2012, pp. 848–853.

[26] W. Qiu and R. Kumar, "Decentralized failure diagnosis of discrete event systems," *IEEE Trans. Syst., Man, Cybern. A*, vol. 36, no. 2, pp. 384–395, Mar. 2006.

[27] R. M. G. Ferrari, T. Parisini, and M. M. Polycarpou, "Distributed fault detection and isolation of large-scale discrete-time nonlinear systems: An adaptive approximation approach," *IEEE Trans. Autom. Control*, vol. 57, no. 2, pp. 275–290, Feb. 2012.

[28] F. Liu, D. Qiu, H. Xing, and Z. Fan, "Decentralized diagnosis of stochastic discrete event systems," *IEEE Trans. Autom. Control*, vol. 53, no. 2, pp. 535–546, Mar. 2008.

[29] J. Chen and R. Kumar, "Decentralized failure diagnosis of stochastic discrete event systems," in *Proc. 9th IEEE Int. Conf. Autom. Sci. Eng.*, Madison, WI, USA, Aug. 2013, pp. 1083–1088.

[30] A. Aghasaryan, E. Fabre, A. Benveniste, R. Boubour, and C. Jard, "Fault detection and diagnosis in distributed systems: An approach by partially stochastic petri nets," *Discrete Event Dyn. Syst.*, vol. 8, no. 2, pp. 203–231, 1998.

[31] R. H. Chen, D. L. Mingori, and J. L. Speyer, "Optimal stochastic fault detection filter," *Automatica*, vol. 39, no. 3, pp. 377–390, Mar. 2003.

[32] D. Lefebvre and E. Leclercq, "Stochastic petri net identification for the fault detection and isolation of discrete event systems," *IEEE Trans. Syst., Man, Cybern. A*, vol. 41, no. 2, pp. 213–225, Mar. 2011.

[33] S. Jiang and R. Kumar, "Failure diagnosis of discrete-event systems with linear-time temporal logic specifications," *IEEE Trans. Autom. Control*, vol. 49, no. 6, pp. 934–945, Jun. 2004.

[34] R. Kumar and V. K. Garg, *Modeling and Control of Logical Discrete-Event Systems*. Boston, MA, USA: Kluwer, 1995.

[35] S. Jiang and R. Kumar, "Diagnosis of repeated failures for discrete event systems with linear-time temporal-logic specifications," *IEEE Trans. Autom. Sci. Eng.*, vol. 3, no. 1, pp. 47–59, Jan. 2006.

[36] F. Liu and D. Qiu, "Safe diagnosability of stochastic discrete event systems," *IEEE Trans. Autom. Control*, vol. 53, no. 5, pp. 1291–1296, Jun. 2008.

[37] V. K. Garg, R. Kumar, and S. I. Marcus, "A probabilistic language formalism for stochastic discrete-event systems," *IEEE Trans. Autom. Control*, vol. 44, no. 2, pp. 280–293, Feb. 1999.

[38] X. Wang and A. Ray, "A language measure for performance evaluation of discrete-event supervisory control systems," *Appl. Math. Modelling*, vol. 28, no. 9, pp. 817–833, Sep. 2004.
[39] J. Lunze and J. Schröder, "State observation and diagnosis of discrete-event systems described by stochastic automata," *Discrete Event Dyn. Syst.*, vol. 11, no. 4, pp. 319–369, 2001.
[40] J. Lunze, "Qualitative modelling of linear dynamical systems with quantised state measurements," *Automatica*, vol. 30, no. 3, pp. 417–431, 1994.
[41] J. Lunze, "On the Markov property of quantised state measuirement sequences," *Automatica*, vol. 34, no. 11, pp. 1439–1444, 1998.

Jun Chen (S'11–M'14) received the B.S. degree in electrical engineering from Zhejiang Uniersity, Hangzhou China, in 2009, and the Ph.D. degree in electrical engineering from Iowa State University, Ames, IA, USA, in 2014.

His research interests include discrete-event systems, cyber-physical systems and stochastic systems, together with their fault diagnosis and prognosis, resiliency control and information security. He is a TPC member for Chinese Control and Decision Conference since 2013.

Dr. Chen received the Research Excellence Award from Iowa State University.

Ratnesh Kumar (S'87–M'90–SM'00–F'07) received the B.Tech. degree in electrical engineering from IIT Kanpur, India, in 1987 and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Texas, Austin, TX, USA, in 1989 and 1991, respectively.

He has been a Professor of electrical and computer engineering, Iowa State University, Ames, IA, USA, since 2002. Prior to this, he held faculty position at the University of Kentucky (1991–2002) in ECE and has held visiting positions at the University of Maryland, Applied Research Laboratory (at Penn State University), NASA Ames, Idaho National Laboratory, and United Technologies Research Center. His research interests include model-based design of embedded software, web-services, networks and cyberphysical systems, sensors and their networks with application to agriculture, power systems and energy harvesting. He is or has been an Associate Editor of *ACM Transactions on Embedded Computing Systems*, the *SIAM Journal on Control and Optimization*, and the *Journal of Discrete Event Dynamical Systems*.

Dr. Kumar received the Gold Medals for Best EE undergrad and Best All Rounder at IIT Kanpur, and Best Dissertation Award at UT Austin. He was an Associate Editor of the IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION and a member of the IEEE Control Systems Society, the IEEE Robotics and Automation Systems Society, and the IEEE Workshop on Software Cybernetics.