Stochastic Failure Prognosability of Discrete Event Systems

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

***Abstract*—We study the prognosis of fault, i.e., its prediction prior to its occurrence, in stochastic discrete event systems. We introduce the notion of m-steps Stochastic-Prognosability, called S_m-Prognosability, which allows the prediction of a fault at least m-steps in advance. We formalize the notion of a prognoser and also show that S_m-Prognosability is necessary and sufficient for the existence of a prognoser that can predict a fault at least m-steps prior to occurrence, while achieving any arbitrary false alarm and missed detection rates. We also provide a polynomial algorithm for the verification of S_m-Prognosability. Finally, we compare the notion of stochastic prognosability with that of stochastic diagnosability, and show that the former is a stronger notion, as can be expected.**

***Index Terms*—Discrete event systems (DESs), failure prognosis, likelihood, stochastic prognosability.**

I. Introduction

THE problem of predicting a fault prior to its occurrence is a well researched area (see for example [1]–[4]). In [2], the notion of uniformly bounded prognosability of fault was formulated for logical discrete event systems (DESs), where each fault-trace must possess a nonfault-prefix such that for all indistinguishable traces, a future fault is inevitable within a bounded delay that is uniform across all fault-traces. Such a nonfault-prefix from which a future fault is inevitable is termed an *indicator*. The notion was later extended to the decentralized setting in [3] and the requirement of the existence of a uniform bound was also removed. Reference [3] also established that the notion of prognosability is equivalent to the existence of a prognoser with no false alarm (FA) and no missed detection (MD). The issue of prognosability under a general decentralized inferencing mechanism was proposed in [5], where a prognostic decision involved inferencing among a group of local prognosers over their local decisions and their ambiguity levels, and the notion of inference-prognosability and its verification was introduced to capture the necessity and sufficiency of inferencing based decentralized prognosis. The problem of distributed prognosability under bounded-delay communications among the local prognosers was studied in [6], where the notion of joint-prognosability and its verification was proposed. Readers are referred to the above literature for more details on prognosability of logical DESs.

In order to generalize the notion of prognosability to stochastic DESs, in this paper, we introduce m-steps Stochastic-Prognosability, or simply S_m-Prognosability, which requires for any tolerance level ρ and error bound τ, there exists a reaction bound $k \geq m$, such that the set of fault-traces for which a fault cannot be predicted k steps in advance with tolerance level ρ, occurs with probability smaller than τ. We formalize the notion of a prognoser that maps observations to decisions by comparing a suitable statistic with a threshold, and show that S_m-Prognosability is a necessary and sufficient condition for the existence of a prognoser with reaction bound at least m (i.e., prediction at least m-steps prior to the occurrence of a fault) that can achieve any specified FA and MD rate requirement. In this sense, S_m-Prognosability can be viewed as a generalization of the logical prognosability, since it provides a basis for the existence and synthesis of a prognoser that can achieve a user-specified level of FA and MD. In contrast, the logical version is rather rigid, offering no further options for systems that fail to be logically prognosable, even when there may exist a prognoser that can achieve a satisfying performance as measured in terms of FA and MD rates. The introduction of false alarm and missed detection rates has provided reliability indices for prognostic systems, paving the way for their risk analysis. Also, in the logical setting, an indicator cannot visit a cycle of nonfault-states, which can be restrictive; in contrast in stochastic setting, an indicator can visit a cycle of nonfault-states as long as the cycle is not absorbing. Further, we also provide a polynomial algorithm for verifying S_m-Prognosability. We show that even the weakest form of stochastic-prognosability where the reaction bound is zero, namely, S_0-Prognosability, is stronger than stochastic-diagnosability (see [7]–[9]), meaning that whenever it is possible to predict faults (even with zero reaction bound), it is also possible to diagnose those, as can be expected.

The rest of this paper is organized as follows. The notations and some preliminaries are presented in Section II, followed by the definition of S_m-Prognosability and stochastic prognoser in Sections III and IV, respectively. Section IV also shows necessity and sufficiency of S_m-Prognosability for the existence of a m-prognoser that can fulfill any desired level of error bounds over FA and MD. Section V gives an algorithm for verifying S_m-Prognosability; Section VI includes two practical examples and the comparison with related works is provided in Section VII. The paper is concluded in Section VIII.

Manuscript received November 1, 2013; revised July 25, 2014 and November 6, 2014; accepted December 7, 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 C. Seatzu.

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.2014.2381437

0018-9286 © 2014 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

II. Notations and Preliminaries

For an event set Σ, define $\overline{\Sigma} := \Sigma \cup \{\epsilon\}$, where ϵ denotes "no-event." The set of all finite length event sequences over Σ, including ϵ, is denoted as Σ^*. A *trace* is a member of Σ^* and a *language* is a subset of Σ^*. We use $s \leq t$ to denote that $s \in \Sigma^*$ is a prefix of $t \in \Sigma^*$, $pr(s)$ to denote the set of all prefixes of s, and $|s|$ to denote the length of s or the number of events in s. For $\sim \in \{<, \leq, >, \geq, =\}$ and $n \in \mathbb{N}$, where $\mathbb{N}$ denotes the set of all nonnegative integers, define $\Sigma^{\sim n} := \{s \in \Sigma^* : |s| \sim n\}$ and denote $\Sigma^{=n}$ as Σ^n for simplicity. For $L \subseteq \Sigma^*$, its prefix-closure is defined as $pr(L) := \bigcup_{s \in L} pr(s)$, and L is said to be prefix-closed (or simply closed) if $pr(L) = L$. Given two languages L_1 and L_2, their *concatenation* is defined as $L_1 L_2 := \{st : s \in L_1, t \in L_2\}$, the set of traces in L_1 *after* L_2 is defined as $L_1 \setminus L_2 := \{t \in \Sigma^* : \exists s \in L_2, st \in L_1\}$, and the set of traces in L_1 *quotient* L_2 is defined as $L_1/L_2 := \{s \in pr(L_1) : \exists t \in L_2, st \in L_1\}$.

A stochastic DES can be modeled by a *stochastic automaton* $G = (X, \Sigma, \alpha, x_0)$, where X is the set of states, Σ is the set of events, $x_0 \in X$ is the initial state, and $\alpha : X \times \Sigma \times X \to [0, 1]$ is the transition probability function [10] satisfying $\forall x \in X$, $\sum_{\sigma \in \Sigma} \sum_{x' \in X} \alpha(x, \sigma, x') = 1$, i.e., there is no "termination" at any of the states. (Note there is no loss of generality in assuming no termination, since otherwise, one can augment the model with a newly introduced "termination-state," and transitions from each state to the termination state on a newly introduced "termination-event" that is unobservable and whose occurrence probability equals the probability of termination of the said state.) G is nonstochastic if $\alpha : X \times \Sigma \times X \to \{0, 1\}$, and a nonstochastic DES is deterministic if $\forall x \in X$, $\sigma \in \Sigma$, $\sum_{x' \in X} \alpha(x, \sigma, x') \in \{0, 1\}$, i.e., each state has at most one transition on each event. The transition probability function α can be generalized to $\alpha : X \times \Sigma^* \times X$ in a natural way by multiplying the probabilities of the individual transitions. Define the language generated by G as $L(G) := \{s \in \Sigma^* : \exists x \in X, \alpha(x_0, s, x) > 0\}$. A *component* $C = (X_C, \alpha_C)$ of G is a "subgraph" of G, i.e., $X_C \subseteq X$ and $\forall x, x' \in X_C$ and $\sigma \in \Sigma$, $\alpha_C(x, \sigma, x') := \alpha(x, \sigma, x')$, whenever the latter is defined. C is said to be a *strongly connected component* (SCC) or *irreducible* if $\forall x, x' \in X_C$, $\exists s \in \Sigma^*$ such that $\alpha_C(x, s, x') > 0$. A SCC C is said to be *closed* if for each $x \in X_C$, $\sum_{\sigma \in \Sigma} \sum_{x' \in X_C} \alpha_C(x, \sigma, x') = 1$.

To represent the limited sensing capabilities of a prognoser, we introduce an event observation mask, $M : \overline{\Sigma} \to \overline{\Delta}$, where Δ is the set of observed symbols and $M(\epsilon) = \epsilon$. An event σ is *unobservable* if $M(\sigma) = \epsilon$. The set of unobservable events is denoted as Σ_{uo}, and so the set of observable events is given by $\Sigma - \Sigma_{uo}$. The observation mask can be generalized to $M : 2^{\Sigma^*} \to 2^{\Delta^*}$ in a natural way: $\forall s \in \Sigma^*$, $\sigma \in \overline{\Sigma}$, $L \subseteq \Sigma^*$, $M(\epsilon) = \epsilon$, $M(s\sigma) = M(s)M(\sigma)$ and $M(L) = \{M(s) : s \in L\}$.

For a stochastic automaton $G = (X, \Sigma, \alpha, x_0)$ with generated language $L(G) = L$, let $K \subseteq L$ be a nonempty closed sublanguage representing a nonfault-specification for G, i.e., $L - K$ consists of behaviors that execute some fault. Then the task of prognosis is to predict the execution of any trace in $L - K$ prior to its execution, and at least m steps in advance, and with sufficient confidence. Let $K \subseteq L$ be generated by a *deterministic* automaton $R = (Q, \Sigma, \beta, q)$ such that $L(R) = K$ (from now on we interchangably use K and R to refer to the "nonfault-specification"). Then the refinement of the plant with respect to the specification, denoted as G^R, can be used to capture the fault-traces in the form of the reachability of a fault-state carrying the label F in G^R, which is given by $G^R := (X \times \overline{Q}, \Sigma, \gamma, (x_0, q_0))$, where $\overline{Q} = Q \cup \{F\}$, 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) $L(G^R) = L(G) = L$, 2) any fault-trace $s \in L - K$ transitions the refinement G^R to a fault-state (a state containing F as its second coordinate); and 3) the occurrence probability of each trace in G^R is the same as that in G, i.e., $\sum_{x \in X} \alpha(x_0, s, x) = \sum_{(x, \overline{q}) \in X \times \overline{Q}} \gamma((x_0, q_0), s, (x, \overline{q}))$.

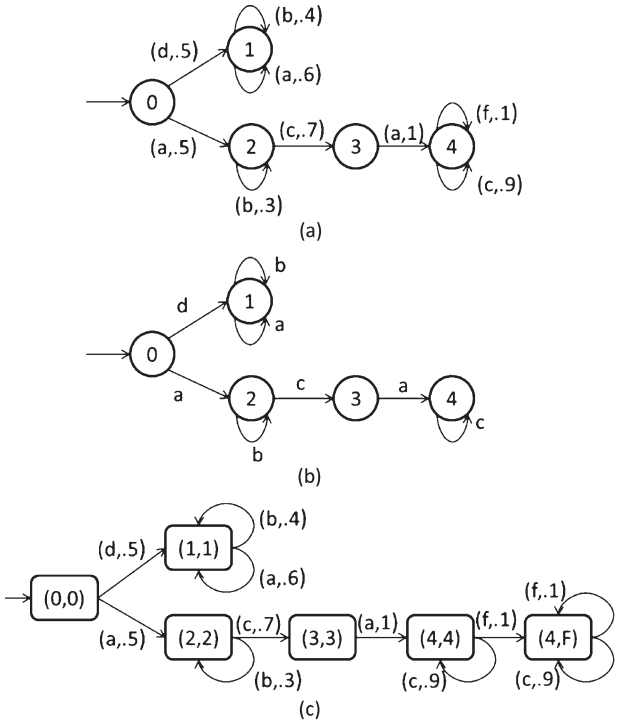

Fig. 1. (a) Stochastic automaton G. (b) Nonfault specification R. (c) Refinement G^R.

Example 1: Fig. 1(a) is an example of a stochastic automaton G. The set of states is $X = \{0, 1, 2, 3, 4\}$ with initial state $x_0 = 0$, and event set $\Sigma = \{a, b, c, d, f\}$. 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 $M(\{d, f\}) = \{\epsilon\}$ and $M(\sigma) = \sigma$ for $\sigma \in \Sigma - \{d, f\}$. The nonfault-specification is given in Fig. 1(b). Therefore, $L - K = \{ab^*cac^*f\}\Sigma^* \cap L$ and the refinement G^R is shown in

Fig. 1(c). As can be seen, all traces in $L-K$ transitions G^R to the only fault-state $(4,F)$. In G^R there are two closed SCCs, one is formed by the nonfault-state (1,1) and its selfloop transitions whereas the other is formed by the fault-state $(4,F)$ and its selfloop transitions. ■

For $x_i, x_j \in X$ and $\sigma \in \Sigma - \Sigma_{uo}$, define the set of traces originating at x_i, terminating at x_j and executing a sequence of unobservable events followed by a single observable event σ as $L_G(x_i,\sigma,x_j) := \{s \in \Sigma^* : s = u\sigma, M(u) = \epsilon, \alpha(x_i,s,x_j) > 0\}$. Define $\alpha(L_G(x_i,\sigma,x_j)) := \sum_{s\in L_G(x_i,\sigma,x_j)} \alpha(x_i,s,x_j)$ as the occurrence probability of traces in $L_G(x_i,\sigma,x_j)$ and denote it as $\mu_{i,\sigma,j}$ for short. Also define $\lambda_{ij} = \sum_{\sigma\in\Sigma_{uo}} \alpha(x_i,\sigma,x_j)$ as the probability of transitioning from x_i to x_j while executing a single unobservable event. Then it can be seen that $\mu_{i,\sigma,j} = \sum_m \lambda_{im}\mu_{m,\sigma,j} + \alpha(x_i,\sigma,x_j)$, where the first term on the right-hand side (RHS) involves transitioning in at least two steps via some intermediate state, whereas the second RHS term involves transitioning directly in exactly one step. Thus, for each $\sigma \in \Sigma - \Sigma_{uo}$, given the values $\{\lambda_{ij} | i,j \in X\}$ and $\{\alpha(x_i,\sigma,x_j) | i,j \in X\}$, all the probabilities $\{\mu_{i,\sigma,j} | i,j \in X, \sigma \in \Sigma - \Sigma_{uo}\}$ can be found by solving the following matrix equation (see for example [11] for a similar matrix equation):

$$\boldsymbol{\mu}(\sigma) = \boldsymbol{\lambda}\boldsymbol{\mu}(\sigma) + \boldsymbol{\alpha}(\sigma) \tag{1}$$

where $\boldsymbol{\mu}(\sigma)$, $\boldsymbol{\lambda}$, and $\boldsymbol{\alpha}(\sigma)$ are all $|X| \times |X|$ square matrices whose ijth elements are given by $\mu_{i,\sigma,j}$, λ_{ij}, and $\alpha(x_i,\sigma,x_j)$, respectively. In the presence of partial observability, we define $L_G(x_i, M(\sigma), x_j) := \cup_{\sigma'\in\Sigma:M(\sigma')=M(\sigma)} L_G(x_i,\sigma',x_j)$, i.e., it is the set of all traces originating at x_i, terminating at x_j and executing a sequence of unobservable events followed by a single observable event that has the same mask value $M(\sigma)$. Then their occurrence probability is given by $\alpha(L_G(x_i,M(\sigma),x_j)) := \sum_{\sigma'\in\Sigma:M(\sigma')=M(\sigma)} \mu_{i,\sigma',j}$.

III. Prognosability of Stochastic DESs

In this section, we formalize the notion of prognosability, called *m-steps Stochastic-Prognosability*, or simply *S_m-Prognosability*, for stochastic DESs, and provide necessary and sufficient conditions for the verification of S_m-Prognosability. In the next section, we show that for finite-state systems, S_m-Prognosability is necessary and sufficient for the existence of a prognoser that can predict a fault at least m-steps prior to occurrence, while achieving any arbitrary false alarm and missed detection rates.

Let L be a nonempty closed language and $K \subseteq L$ be a nonempty closed language representing a nonfault-specification. In order to be able to make a prognostic decision, we define the *n-step prognostic probability of no-fault* following an observation $o \in M(L)$ as (where N in the subscript denotes "no-fault"):

$$P_N^n(o) := \frac{Pr\left(\left\{M^{-1}(o)\right\}\Sigma^n \cap K\right)}{Pr\left(\left\{M^{-1}(o)\right\}\Sigma^n \cap L\right)} = \frac{Pr\left(\left\{M^{-1}(o)\cap K\right\}\Sigma^n \cap K\right)}{Pr\left(M^{-1}(o)\cap L\right)} \tag{2}$$

and the *least prognostic probability of no-fault* following $o \in M(L)$ as

$$P_N^*(o) := \min_{n\in\mathbb{N}} P_N^n(o) = \frac{\min_{n\in\mathbb{N}} Pr\left(\left\{M^{-1}(o)\right\}\Sigma^n \cap K\right)}{Pr\left(\left\{M^{-1}(o)\right\}\cap L\right)}. \tag{3}$$

Note $P_N^n(o)$ is the probability, following the observation o, that the system does not execute a fault in the next n steps; and $P_N^*(o)$ is the least probability, following the observation o, that the system does not execute a fault over all finite-step futures. Note in the denominator of (2), we used the fact that probability of all extensions of length n, beyond the traces in $M^{-1}(o)$, is the same as the probability of traces in $M^{-1}(o)$, for there is no termination at any of the states. As a result, the denominator is constant with respect to n, and the minimum only applies to the numerator in (3).

To help formalize the prognosability for stochastic DESs, we introduce the notions of *boundary* fault-traces whose all strict prefixes are nonfault, *m-steps interior* nonfault-traces for which a fault can occur in the next $(m+1)$th step while no fault can occur within the next m steps, *persistent* nonfault-traces whose all extensions are nonfault, *indicator* nonfault-traces for which a future fault is guaranteed with arbitrary confidence and *nonindicator* nonfault-traces that are not the indicator traces.

Definition 1: Given a pair (L,K) of closed languages with $K \subseteq L$, we define the set of

- *boundary* fault-traces as, $\partial := \{s \in L-K : pr(s) - \{s\} \subseteq K\}$;
- *m-steps interior* nonfault-traces of K with respect to L (where $m \geq 0$) as, $\partial_m^- := \{s \in K : \{s\}\Sigma^{\leq m} \cap (L-K) = \emptyset, \{s\}\Sigma^{m+1} \cap \partial \neq \emptyset\}$;
- *persistent* nonfault-traces of K with respect to L as, $\aleph := \{s \in K : \forall n \in \mathbb{N}, \{s\}\Sigma^n \cap (L-K) = \emptyset\} = \{s \in K : \forall n \in \mathbb{N}, Pr(\{s\}\Sigma^n \cap K) = Pr(s)\}$;
- *indicator* nonfault-traces of K with respect to L as, $\mathfrak{J} := \{s \in K : \forall \rho > 0, \exists n \in \mathbb{N}, Pr(\{s\}\Sigma^n \cap K) \leq \rho\}$;
- *nonindicator* nonfault-traces of K with respect to L as, $\Upsilon := K - \mathfrak{J}$.

Note that $\Upsilon = \{s \in K : \exists \rho > 0, \forall n \in \mathbb{N}, Pr(\{s\}\Sigma^n \cap K) > \rho\}$, and since $\aleph$ is obtained by replacing ρ with $Pr(s)$ in the right-hand side of this equality, it follows that $\aleph \subseteq \Upsilon$. Also note that $\aleph$ is "extension-closed" in the sense that if it possesses $s \in K$, then it also possesses all extensions $t \in L$ with $s \leq t$.

Example 2: For system in Fig. 1, $L - K = ab^*cac^*f(c+f)^*$, and the set of boundary fault-traces is $\partial = ab^*cac^*f$, and so $\partial_2^- = ab^*$, $\partial_1^- = ab^*c$. The set of persistent nonfault-traces is given by $\aleph = d(a+b)^*$ as all its extensions are nonfault. The set of indicator traces is $\mathfrak{J} = \{a\}\Sigma^* \cap K$, and the set of nonindicator traces is $\Upsilon = \{\epsilon\} \cup \{d\}\Sigma^* \cap L = \{\epsilon\} \cup d(a+b)^*$. ■

Next we introduce the definition of S_m-Prognosability which requires that, for any threshold value $\rho > 0$ and error bound $\tau > 0$, there exists a reaction bound $k \geq m$, such that the set of boundary fault-traces, that are either shorter than k in length or for which a prognostic decision cannot be made k steps

in advance with confidence level ρ, occurs with probability smaller than τ.

Definition 2: A pair (L, K) of closed languages with $K \subseteq L$ is said to be *m-steps Stochastically-Prognosable*, or simply *S_m-Prognosable*, if

$$(\forall \tau, \rho > 0)(\exists k \geq m) \\ Pr\left(s \in \partial : [|s| \leq k] \vee \left[\forall u \in s/\Sigma^{>k}, P_N^*\left(M(u)\right) > \rho\right]\right) < \tau \quad (4)$$

where P_N^* is as defined by (2) and (3).

The next lemma states that we can always choose the reaction bound k in Definition 2 to equal m, thereby simplifying the definition a bit.

Lemma 1: A pair (L, K) of closed languages with $K \subseteq L$ is *S_m-Prognosable* if and only if

$$(\forall \tau, \rho > 0) \\ Pr\left(s \in \partial : [|s| \leq m] \vee \left[\forall u \in s/\Sigma^{>m}, P_N^*\left(M(u)\right) > \rho\right]\right) < \tau. \quad (5)$$

Proof: The sufficiency is obvious by choosing $k = m$. Now to see the converse, assume (5) is not true, i.e., $\exists \tau > 0$, $\rho > 0$, s.t. $Pr(s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]) \geq \tau$. Since we have for all $k \geq m$, $\{s \in \partial : [\forall u \in s/\Sigma^{>k}, P_N^*(M(u)) > \rho] \vee [|s| \leq k]\} \supseteq \{s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]\}$, and hence $Pr(s \in \partial : [\forall u \in s/\Sigma^{>k}, P_N^*(M(u)) > \rho] \vee [|s| \leq k]) \geq Pr(s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]) \geq \tau$. Therefore, according to Definition 2, (L, K) is not S_m-Prognosable. Hence, the necessity also holds. ■

Denote $\ell(\partial) = \min\{|s|, s \in \partial\}$ as the length of the shortest fault-trace in $L - K$. Then the following theorem provides a necessary and sufficient condition for S_m-Prognosability requiring the reaction bound m to be smaller than the length of the shortest fault-trace, $\ell(\partial)$, and every boundary fault-trace in ∂ to possess a nonfault-prefix which is more than m-steps shorter and is unambiguously an indicator.

Theorem 1: A pair (L, K) of closed languages with $K \subseteq L$ is S_m-Prognosable if and only if $m < \ell(\partial)$ and

$$(\forall s \in \partial)\left(\exists u \in s/\Sigma^{>m}\right)\left(M^{-1}M(u) \cap K \subseteq \mathfrak{J}\right). \quad (6)$$

Proof: (Sufficiency) For any $s \in \partial$, let $u \in s/\Sigma^{>m}$ be such that $M^{-1}M(u) \cap K \subseteq \mathfrak{J}$. Then

$$P_N^n\left(M(u)\right) = \frac{Pr\left(\left\{M^{-1}M(u) \cap K\right\}\Sigma^n \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)} \\ = \frac{\sum_{u' \in M^{-1}M(u) \cap K} Pr\left(\{u'\}\Sigma^n \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)}.$$

For any $\rho > 0$, define $\rho_{u'} := \rho Pr(u') > 0$ for each $u' \in M^{-1}M(u) \cap K$. Then since $M^{-1}M(u) \cap K \subseteq \mathfrak{J}$, for each $u' \in M^{-1}M(u) \cap K$, exists $n_{u'} \in \mathbb{N}$ such that $Pr(\{u'\}\Sigma^{n_{u'}} \cap K) \leq \rho_{u'}$. Let $d := \max_{u' \in M^{-1}M(u) \cap K} n_{u'}$. Note that d here is a finite integer even if $M^{-1}M(u)$ is an infinite set (resulted by unobservable loops). To see this, let $u_1 = u_{11}u_{12}$ and $u_2 = u_{11}\sigma_1 \ldots \sigma_k u_{12}$ such that $\sigma_1 \ldots \sigma_k$ is an unobservable loop. Then we have $Pr(\{u_2\}\Sigma^{n_{u_1}} \cap K) = Pr(\sigma_1 \ldots \sigma_k)Pr(\{u_1\}\Sigma^{n_{u_1}} \cap K) < \rho Pr(\sigma_1 \ldots \sigma_k)Pr(u_1) = \rho Pr(u_2) = \rho_{u_2}$, and thus $n_{u_2} \leq n_{u_1}$. Therefore, to find d, we only need to consider $u' \in M^{-1}M(u) \cap K$ such that u' doesn't contain any unobservable loop, making d finite. Therefore,

$$P_N^d\left(M(u)\right) = \frac{\sum_{u' \in M^{-1}M(u) \cap K} Pr\left(\{u'\}\Sigma^d \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)} \\ \leq \frac{\sum_{u' \in M^{-1}M(u) \cap K} \rho_{u'}}{Pr\left(M^{-1}M(u) \cap L\right)} \\ = \frac{\sum_{u' \in M^{-1}M(u) \cap K} \rho Pr(u')}{Pr\left(M^{-1}M(u) \cap L\right)} \\ = \frac{Pr\left(M^{-1}M(u) \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)}\rho \leq \rho. \text{ Hence,}$$

$$P_N^*\left(M(u)\right) \leq P_N^d\left(M(u)\right) \leq \rho.$$

Also since $m < \ell(\partial)$ implies $\{s \in \partial : |s| \leq m\} = \emptyset$, we have for all $\rho > 0$ and $\tau > 0$, $Pr(s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]) = 0 < \tau$. According to Lemma 1, (L, K) is S_m-Prognosable.

(Necessity) When $m \geq \ell(\partial)$, let $s \in \partial$ be such that $|s| = \ell(\partial) \leq m$. Obviously for any $\tau \leq Pr(s)$, $Pr(s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]) \geq Pr(s \in \partial : |s| \leq m) \geq Pr(s) \geq \tau$ for all $\rho > 0$. Therefore, (L, K) is not S_m-Prognosable. When $m < \ell(\partial)$, but (6) is not true, let $s \in \partial$ be such that $(\forall u \in s/\Sigma^{>m})(M^{-1}M(u) \cap K \cap \Upsilon \neq \emptyset)$. Then for any $u \in s/\Sigma^{>m}$ and $u' \in M^{-1}M(u) \cap K \cap \Upsilon$

$$P_N^n\left(M(u)\right) = \frac{Pr\left(\left\{M^{-1}M(u) \cap K\right\}\Sigma^n \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)} \\ \geq \frac{Pr\left(\{u'\}\Sigma^n \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)}.$$

Since $u' \in \Upsilon$, there exists $\rho_{u'} > 0$ such that $\forall n \in \mathbb{N}$, $Pr(\{u'\}\Sigma^n \cap K) > \rho_{u'}$. Therefore, for any $n \in \mathbb{N}$

$$P_N^n\left(M(u)\right) \geq \frac{Pr\left(\{u'\}\Sigma^n \cap K\right)}{Pr\left(M^{-1}M(u) \cap L\right)} \\ > \frac{\rho_{u'}}{Pr\left(M^{-1}M(u) \cap L\right)} =: \rho_u$$

and hence

$$P_N^*\left(M(u)\right) = \min_{n \in \mathbb{N}} P_N^n\left(M(u)\right) > \rho_u.$$

Thus, for any $u \in s/\Sigma^{>m}$, there exists $\rho_u > 0$ such that $P_N^*(M(u)) > \rho_u$. Therefore, for any $0 < \rho < \min_{u \in s/\Sigma^{>m}} \rho_u$ and $0 < \tau < Pr(s)$, $Pr(s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]) \geq Pr(s) > \tau$. Hence, (L, K) is not S_m-Prognosable, according to Lemma 1. ■

Example 3: For system in Fig. 1, $\ell(\partial) = 4$, so by Theorem 1, the system cannot be S_m-Prognosable with $m \geq 4$. According to Example 2, the set of indicator traces is $\mathfrak{J} = \{a\}\Sigma^* \cap K$, and the set of nonindicator traces is $\Upsilon = \{\epsilon\} \cup d(a + b)^*$, while the set of boundary fault-traces is $\partial = ab^*cac^*f$. One can check that for any $s \in \partial$, there exists $u \in s/\Sigma^{>1} \subseteq \{ab^*c\}\Sigma^* \cap K$ such that $M^{-1}M(u) \cap K \subseteq \mathfrak{J}$. Therefore, by Theorem 1,

(L, K) is S_1-Prognosable. On the other hand, for $s = acaf \in \partial$, $u = a \in s/\Sigma^{>2}$ is such that $M^{-1}M(u) \cap K \cap \Upsilon = \{da\} \neq \emptyset$. Therefore, by Theorem 1, (L, K) is not S_2-Prognosable. ■

The following corollary is directly obtained from Theorem 1, and captures the expected property that prognosability continues to hold even with smaller reaction bound.

Corollary 1: Given a pair (L, K) of closed languages with $K \subseteq L$, if (L, K) is S_m-Prognosable, then (L, K) is $S_{m'}$-Prognosable for all nonnegative $m' \leq m$, whereas if (L, K) is not S_m-Prognosable, then (L, K) is not $S_{m'}$-Prognosable for all $m' \geq m$.

For a S_m-Prognosable system, Theorem 1 requires that each boundary fault trace possess a more than m-steps shorter prefix that is unambiguously an indicator. We can strengthen this theorem by requiring that *exactly* the $(m+1)$-shorter prefix possess the said property. This requires the result of the next lemma stating that indicators are "extension-closed" (nonfault-extensions of indicators are also indicators), while nonindicators are prefix-closed (prefixes of nonindicators are also nonindicators).

Lemma 2: For a pair (L, K) of closed languages with $K \subseteq L$, it holds that $\mathfrak{J}\Sigma^* \cap K \subseteq \mathfrak{J}$, and $pr(\Upsilon) \subseteq \Upsilon$.

Proof: Let $s \in \mathfrak{J}$ be arbitrary, i.e., $\forall \rho > 0$, $\exists n \in \mathbb{N}$ s.t. $Pr(\{s\}\Sigma^n \cap K) \leq \rho$. Since for any $t \in K \setminus s$, $Pr(\{st\}\Sigma^l \cap K) \leq Pr(\{s\}\Sigma^{l+|t|} \cap K)$, we have $\forall \rho > 0$, $\exists l = n - |t| \in \mathbb{N}$ s.t. $Pr(\{st\}\Sigma^l \cap K) \leq Pr(\{s\}\Sigma^{l+|t|} \cap K) = Pr(\{s\}\Sigma^n \cap K) \leq \rho$. According to Definition 1, $st \in \mathfrak{J}$, i.e., $\forall s \in \mathfrak{J}$, $t \in K \setminus s$, $st \in \mathfrak{J}$. Therefore, $\mathfrak{J}\Sigma^* \cap K \subseteq \mathfrak{J}$.

Similarly, let $s \in \Upsilon$ be arbitrary, i.e., $\exists \rho > 0$ s.t. $\forall n \in \mathbb{N}$, $Pr(\{s\}\Sigma^n \cap K) > \rho$. Then for any $u \in pr(s)$, $Pr(\{u\}\Sigma^l \cap K) \geq Pr(\{s\}\Sigma^{l-|s|+|u|} \cap K) > \rho$ for any $l - |s| + |u| \in \mathbb{N}$ and hence for any $l \in \mathbb{N}$. According to Definition 1, $u \in \Upsilon$, i.e., $\forall s \in \Upsilon$, $u \in pr(s)$, $u \in \Upsilon$. Therefore, $pr(\Upsilon) \subseteq \Upsilon$. ■

Using Lemma 2, we can strengthen Theorem 1 to obtain a new result that we employ in Section V for verifying S_m-Prognosability. The new theorem states that S_m-Prognosability holds if and only if the reaction bound $m < \ell(\partial)$, and all m-steps interior traces are distinguishable from any nonindicator trace.

Theorem 2: A pair (L, K) of closed languages with $K \subseteq L$ is S_m-Prognosable if and only if $m < \ell(\partial)$ and

$$M^{-1}M\left(\partial_m^-\right) \cap \Upsilon = \emptyset. \tag{7}$$

Proof: If $m < \ell(\partial)$ and (7) is true, then it follows from the fact that every fault-trace $s \in \partial$ possesses a nonfault-prefix $u \in \partial_m^-$ satisfying $u \in s/\Sigma^{>m}$ and Theorem 1 that (L, K) is S_m-Prognosable, and the sufficiency follows. On the other hand, if $m \geq \ell(\partial)$, then by Theorem 1, (L, K) is not S_m-Prognosable. Meanwhile if $m < \ell(\partial)$ but (7) is not true, then we can select $s \in \partial_m^-$ and $s' \in \Upsilon$ such that $M(s) = M(s')$. Then for any $u \in pr(s)$, there exists $u' \in pr(s')$ such that $M(u) = M(u')$ and $u' \in \Upsilon$ (Lemma 2), i.e., $\forall u \in pr(s)$, $M^{-1}M(u) \cap K \cap \Upsilon \neq \emptyset$. It follows from the definition of ∂_m^- that there exists $st \in \partial$ such that $st/\Sigma^{>m} = pr(s)$, and hence $\forall u \in st/\Sigma^{>m} = pr(s)$, $M^{-1}M(u) \cap K \cap \Upsilon \neq \emptyset$. According to Theorem 1, (L, K) is not S_m-Prognosable. Thus the necessity also holds. ■

Example 4: For system shown in Fig. 1, $\mathfrak{J} = \{a\}\Sigma^* \cap K$, $\Upsilon = \{\epsilon\} \cup d(a+b)^*$, $\partial_2^- = ab^*$, and $\partial_1^- = ab^*c$. Since $M^{-1}M(\partial_2^-) \cap \Upsilon = \{dab^*, ab^*\} \cap [\{\epsilon\} \cup d(a+b)^*] = dab^* \neq \emptyset$ and $M^{-1}M(\partial_1^-) = ab^*c \subseteq \mathfrak{J}$. Therefore, (L, K) is S_1-Prognosable but not S_2-Prognosable, as discussed in Example 3. ■

IV. Prognoser and Its Existence Condition

In this section, we formally define a prognoser with reaction bound at least m, called a m-prognoser, along with its FA and MD rates, and show that the notion of S_m-Prognosability introduced in the previous section acts as a necessary and sufficient condition for the existence of a m-prognoser capable of achieveing any FA and MD rates.

In order to predict a fault in advance, the prognoser computes for each $o \in M(L)$, the prognostic probability of no-fault $P_N^*(o)$ as defined by (2), (3), and compares it with an appropriately chosen threshold ρ. Whenever $P_N^*(o)$ is below this threshold, implying that there is only a small likelihood of no-fault in future, the prognoser issues a fault warning F, predicting/prognosing a future fault, and otherwise it remains silent (issues ϵ). In other words, a prognoser is formally a map, $D : M(L) \to \{F, \epsilon\}$ defined as

$$\forall o \in M(L),\ [D(o) = F] \Leftrightarrow [\exists \overline{o} \leq o : P_N^*(\overline{o}) \leq \rho] \tag{8}$$

where P_N^* is as defined by (2) and (3). Note that according to (8), once a warning is issued, it remains unchanged for the subsequent extensions.

Example 5: Consider the system G^R shown in Fig. 1. Upon receiving observation $o = abbb$, the prognoser computes $P_N^*(o) = 0.5872$ according to (2), (3), and compare it with ρ. A prognostic decision F can be issued if $\rho \leq 0.5872$. If instead the prognoser receives $o = abc$, and computes $P_N^*(o) = 0$, then for any threshold ρ, a prognostic decision F can be issued. ■

For a prognoser that aims to predict a fault at least m steps before its occurrence, a *miss detection* (MD) occurs when a fault happens while the prognoser fails to issue a warning m steps in advance, i.e., a boundary fault-trace s occurs while either $|s| \leq m$ or the prognoser is silent for its prefix s/Σ^{m+1}. On the other hand, a *false alarm* (FA) occurs when a warning is issued for a trace whose all extensions are nonfault, i.e., a trace s in $\aleph$ occurs while the prognoser issues F. Therefore, the MD rate P^{md} and the FA rate P^{fa} for a m-prognoser can be defined as

$$P^{md} = Pr(s \in \partial : [|s| \leq m] \vee \left[D\left(M\left(s/\Sigma^{m+1}\right)\right) = \epsilon\right] \tag{9}$$

$$P^{fa} = Pr\left(s \in \aleph : D\left(M(s)\right) = F\right). \tag{10}$$

Considering the fact that once the prognoser issues F, it issues F for any subsequent observations, the above equations can also be equivalently presented as

$$P^{md} = Pr\left(s \in \partial : [|s| \leq m] \vee \left[\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho\right]\right)$$

$$P^{fa} = Pr\left(s \in \aleph : \exists u \in pr(s), P_N^*\left(M(u)\right) \leq \rho\right).$$

Example 6: For the system G^R shown in Fig. 1. Suppose G^R executes $dabbb$ and produces observation $o = abbb$, then $P_N^*(o) = 0.5872$. Hence, for any m-prognoser with threshold $\rho \geq 0.5872$, traces in $\{dabbb\}\Sigma^* \cap L$ will be false alarmed. When G^R executes a trace in $ab^*cac^*f \subseteq \partial$ and produces an observation $o \in ab^*cac^*$, then $P_N^*(\overline{o})$ approaches 0 for any $\overline{o} \in ab^*c$. Therefore, for a 1-prognoser with any threshold ρ, all fault-traces can be prognosed, and hence no missed detection. However, for a 2-prognoser with $\rho = 0.3$, when G^R executes the fault-trace $abcaf$, a prognostic decision can be made only upon observing abc (since for all its prefixes, the threshold remains lower than the prognostic probability of no fault: $P_N^*(\epsilon) = 0.5,\ P_N^*(a) = 0.375,\ P_N^*(ab) = 0.444,\ P_N^*(abc) = 0$), which violates the least reaction bound $m = 2$, and hence $abcaf$ gets missed detected. ■

In order to establish a condition for the existence of a m-prognoser in terms of the property of S_m-Prognosability, we first establish the following corollary of Theorem 1 and Lemma 2.

Corollary 2: If a pair (L, K) of closed languages with $K \subseteq L$ is S_m-Prognosable, then $M^{-1}M(\Upsilon) \cap (L - K) = \emptyset$.

Proof: Suppose for contradiction that (L, K) is S_m-Prognosable and there exists $s \in \Upsilon$ such that $M^{-1}M(s) \cap (L - K) \neq \emptyset$. Let $s' \in M^{-1}M(s) \cap (L - K)$. Then for all $u' \in pr(s')$, there exists $u \in pr(s)$ such that $M(u) = M(u')$. According to Lemma 2, $u \in \Upsilon$. Therefore, $\forall u' \in pr(s') \cap K$, $M^{-1}M(u') \cap K \cap \Upsilon \neq \emptyset$. By Theorem 1, (L, K) is not S_m-Prognosable for any $m \in \mathbb{N}$, which contradicts the assumption that (L, K) is S_m-Prognosable. ■

The next lemma states that under the assumption of regularity of languages L and K, equivalently the finiteness of the state-space of G^R, no extension of an indicator can be persistently nonfault, whereas some extension of a nonindicator must be persistently nonfault. The lemma requires the finiteness of the state-space that guarantees the probability of staying in a transient state approaches 0 while the system evolves.

Lemma 3: For a pair (L, K) of closed regular languages with $K \subseteq L$, we have $\mathfrak{J}\Sigma^* \cap \aleph = \emptyset$ and $\Upsilon\Sigma^* \cap \aleph \neq \emptyset$.

Proof: Assume for contradiction that there exists $s \in \mathfrak{J}$ such that $\{s\}\Sigma^* \cap \aleph \neq \emptyset$. Let $u = \sigma_1 \ldots \sigma_n \in K \setminus s$ be such that $su \in \aleph$. Then for any $l \in \{1, \ldots, n\}$, $Pr(\{s\}\Sigma^l \cap K) \geq Pr(s\sigma_1 \ldots \sigma_l) = Pr(su)$, and for $l > n$, $Pr(\{s\}\Sigma^l \cap K) \geq Pr(\{su\}\Sigma^{l-n} \cap K) = Pr(su)$, i.e., there exists $0 < \rho < Pr(su)$ such that for any $l \in \mathbb{N}$, $Pr(\{s\}\Sigma^l \cap K) > \rho$. Therefore, $s \notin \mathfrak{J}$, a contradiction.

Similarly assume for contradiction that there exists $s \in \Upsilon$ such that $\{s\}\Sigma^* \cap \aleph = \emptyset$. Then for any $u \in L \setminus s$, it possesses a fault-extension $t \in (L - K) \setminus su$, i.e., the "nonfaulty-ness of s" is a transient property. Since the language L and K are regular and have finite state representations, for any $\rho > 0$, there exists $n \in \mathbb{N}$ such that $Pr(t \in K \setminus s, |t| \geq n) \leq \rho$, i.e., $Pr(\{s\}\Sigma^n \cap K) = Pr(s)Pr(t \in K \setminus s, |t| = n) \leq \rho Pr(s) := \rho'$ holds for any $\rho' > 0$. Hence, $s \in \mathfrak{J}$, which contradicts the assumption that $s \in \Upsilon$. ■

Remark 1: Note by Lemma 3, no extension of an indicator trace can persistently be a nonfault-trace. This requirement is weaker than the corresponding requirement for an indicator trace in the logical setting: All extensions of an indicator trace must be a fault-trace within a bounded steps [3]. A consequence of this is that, in the logical setting, an indicator trace cannot visit a cycle of nonfault-states [3], which can be restrictive. In contrast, in stochastic setting, an indicator is allowed to visit a cycle of nonfault-states as long as the cycle is non-absorbing (i.e., it has a positive exit probability, which ensures the non-persistence of remaining nonfault).

The next lemma will be used in the proof of Theorem 3. For the sake of space, its proof is provided in the Appendix.

Lemma 4: For a pair (L, K) of S_m-Prognosable closed regular languages with $K \subseteq L$, we have

$$(\forall \rho', \phi > 0)(\exists d \in \mathbb{N})(\forall s \in \aleph) \\ Pr\left(t : t \in \aleph \setminus s, |t| \geq d, P_N^*\left(M(st)\right) < \rho'\right) < \phi \quad (11)$$

where the persistent nonfault-traces $\aleph$ is defined in Definition 1 and P_N^* is as defined by (2) and (3).

Now we are ready to present the main result of the section, which shows that for regular languages L and K, S_m-Prognosability is necessary and sufficient for the existence of a m-prognoser to satisfy any level of FA and MD rates.

Theorem 3: Consider a pair (L, K) of closed regular languages with $K \subseteq L$. Then for any FA rate $\phi > 0$ and MD rate $\tau > 0$, there exists a m-prognoser (and its associated prognostic decision threshold) defined by (8) such that the MD and FA rates defined by (9), (10) satisfy $P^{md} \leq \tau$ and $P^{fa} \leq \phi$ if and only if (L, K) is S_m-Prognosable.

Proof: (Sufficiency) Suppose (L, K) is S_m-Prognosable. Then for a nonfault-trace $s \in K - \aleph$, its extensions continuing to remain in $K - \aleph$ is a transient property. Since the language L and K are regular and have finite state representations, we have for any $\phi_1 > 0$, $\exists d_1 \in \mathbb{N}$ such that $Pr(s \in (K - \aleph) \cap \Sigma^{>d_1}) < \phi_1$. For any $s \in \aleph \cap \Sigma^{d_1}$, if we pick $\rho'_s := \min_{u \in pr(s)} P_N^*(M(u)) > 0$, we can ensure that s is not false alarmed. For any $s \in \aleph \cap \Sigma^{d_1}$, according to Lemma 4 (presented in the Appendix), for any $\phi_2 > 0$ and $\rho'_2 > 0$, there exists $d_2 \in \mathbb{N}$, such that the set of extensions of s that are longer than d_2 and have P_N^* values of their observations smaller than ρ'_2, occur with probability smaller than ϕ_2, i.e., $P^{fa}(s) < \phi_2$.

Let $d = d_1 + d_2$. If we pick $\rho' = \min_{u \in pr(s), s \in \aleph \cap \Sigma^d} P_N^*(M(u)) > 0$, $\rho < \min(\rho'_2, \rho')$ and $\phi_1 + \phi_2 < \phi$, then P^{fa} is upper bounded by

$$\begin{aligned} P^{fa} &= Pr\left(s \in \aleph : \exists u \in pr(s), P_N^*\left(M(u)\right) \leq \rho\right) \\ &= Pr\left(s \in \aleph : pr(s) \cap \Sigma^d \cap \aleph = \emptyset, \right. \\ &\qquad \left. \exists u \in pr(s), P_N^*\left(M(u)\right) \leq \rho\right) \\ &\quad + Pr\left(s \in \aleph : pr(s) \cap \Sigma^d \cap \aleph \neq \emptyset, \right. \\ &\qquad \left. \exists u \in pr(s), P_N^*\left(M(u)\right) \leq \rho\right) \\ &\leq Pr\left(s \in (K - \aleph) \cap \Sigma^{>d_1}\right) \\ &\quad + \sum_{s \in \aleph \cap \Sigma^{d_1}} Pr(s)\phi_2 < \phi_1 + \phi_2 < \phi. \end{aligned}$$

Therefore, with the above choice of ρ, an arbitrary FA rate ϕ could be achieved. Next since (L, K) is S_m-Prognosable, according to Lemma 1, with this choice of ρ, for any $\tau > 0$,

we have $P^{md} \leq Pr(s \in \partial : [\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho] \vee [|s| \leq m]) < \tau$. Therefore, the sufficiency holds.

(Necessity) To show the necessity, consider the contrapositive where (L, K) is not S_m-Prognosable. Then by Theorem 1, there are two possibilities. First, if $m \geq \ell(\partial)$, then let $s \in \partial$ be such that $|s| = \ell(\partial)$, and in which case

$$\begin{aligned} P^{md} &\geq Pr\big(s \in \partial : [|s| \leq m] \vee \big[\forall u \in s/\Sigma^{>m}, P_N^*(M(u)) > \rho\big]\big) \\ &\geq Pr\left(s \in \partial : |s| \leq m\right) \\ &\geq Pr(s). \end{aligned}$$

Therefore, a MD rate $\tau < Pr(s)$ cannot be achieved.

On the other hand, if $m < \ell(\partial)$ but (6) is not true, then exists $s \in \partial$, such that for all $u \in s/\Sigma^{>m}$, there exists $u' \in \Upsilon$ with $M(u) = M(u')$. Since $u' \in \Upsilon$, according to Lemma 3, there exists $t' \in K \setminus u'$ such that $u't' \in \aleph$. If we choose $\rho < \min_{u \in s/\Sigma^{>m}} P_N^*(u)$, then s will be missed detected, and a MD rate $\tau < Pr(s)$ cannot be achieved. On the other hand, if we choose $\rho \geq \min_{u \in s/\Sigma^{>m}} P_N^*(u)$, then $u't'$ will be false alarmed, and a FA rate $\phi < Pr(u't')$ cannot be met. Therefore, in this case, at most one of arbitrarily small FA or MD rates can be achieved, completing the contraposition argument. ■

V. Verification of S_m-Prognosability

Having established S_m-Prognosability as a central property, needed for the existence of a m-prognoser, we next provide a polynomial algorithm for the verification of S_m-Prognosability utilizing Theorem 2. We need the following definitions that identify *m-steps interior* nonfault-states from where no fault can occur within m steps but will occur at $(m+1)$th step, *indicator* nonfault-states from where a future fault is inevitable with arbitrary confidence, and *nonindicator* nonfault-states which are not indicator states.

Definition 3: Given a stochastic DES $G = (X, \Sigma, \alpha, x_0)$, deterministic nonfault-specification $R = (Q, \Sigma, \beta, q_0)$, with their refinement $G^R = (X \times \overline{Q}, \Sigma, \gamma, (x_0, q_0))$, the set of

- *m-steps interior* nonfault-states $\partial_m^-(X \times \overline{Q}) \subseteq X \times \overline{Q}$ (where $m \geq 0$) are states $(x, \overline{q})$ such that $\overline{q} \neq F$, and there exists $(x', \overline{q}')$ with $\overline{q}' = F$ and $s \in \Sigma^{m+1}$ s.t. $\gamma((x, \overline{q}), s, (x', \overline{q}')) > 0$ and for all $(x', \overline{q}')$, $s \in \Sigma^{\leq m}$, $[\gamma((x, \overline{q}), s, (x', \overline{q}')) > 0] \Rightarrow [\overline{q}' \neq F]$;
- *indicator* nonfault-states $\mathfrak{J}(X \times \overline{Q})$ are states $(x, \overline{q})$ such that $\overline{q} \neq F$ and from which the system cannot reach a closed SCC in G^R that contains a nonfault-state;
- *nonindicator* nonfault-states $\Upsilon(X \times \overline{Q})$ are states from which the system can reach a closed SCC in G^R that contains a nonfault-state.

The following lemma is immediate from Definition 1, Definition 3 and Lemma 3.

Lemma 5: Given a pair $(L = L(G), K = L(R))$ of closed regular languages with $K \subseteq L$, then for any $s \in K$,

- $[s \in \partial_m^-] \Leftrightarrow \quad [\exists (x, \overline{q}) \in \partial_m^-(X \times \overline{Q}), \gamma((x_0, q_0), s, (x, \overline{q})) > 0]$;
- $[s \in \mathfrak{J}] \Leftrightarrow [\exists (x, \overline{q}) \in \mathfrak{J}(X \times \overline{Q}), \gamma((x_0, q_0), s, (x, \overline{q})) > 0]$;
- $[s \in \Upsilon] \Leftrightarrow [\exists (x, \overline{q}) \in \Upsilon(X \times \overline{Q}), \gamma((x_0, q_0), s, (x, \overline{q})) > 0]$.

The following algorithm verifies the condition of Theorem 2.

Algorithm 1: For a given stochastic automaton $G = (X, \Sigma, \alpha, x_0)$ and a deterministic nonfault-specification $R = (Q, \Sigma, \beta, x_0)$, perform the following steps:

1) Check if the length of the shortest trace to a state $X \times \{F\}$ in G^R is smaller than m, if the answer is yes, proceed to step 2), otherwise (L, K) is not S_m-Prognosable;

2) Construct a testing automaton $T = G^R \times G^R$ such that at each step the first copy of G^R takes lead in executing transitions, whereas the second copy responds by executing an indistinguishable nonfault-trace. This automaton is denoted as $T = (Z, \Sigma \times \overline{\Sigma}, \delta, z_0)$, where

- $Z = X \times \overline{Q} \times X \times \overline{Q}$;
- $z_0 = ((x_0, q_0), (x_0, q_0))$ is the initial state;
- $\delta : Z \times \Sigma \times \overline{\Sigma} \times Z \to [0, 1]$ is defined as: $\forall ((x_1, \overline{q}_1), (x_2, \overline{q}_2)), ((x_1', \overline{q}_1'), (x_2', \overline{q}_2')) \in Z, (\sigma, \sigma') \in \Sigma \times \overline{\Sigma}$,

$$\delta\left(((x_1, \overline{q}_1), (x_2, \overline{q}_2)), (\sigma, \sigma'), ((x_1', \overline{q}_1'), (x_2', \overline{q}_2'))\right)$$
$$= \begin{cases} \gamma\left((x_1, \overline{q}_1), \sigma, (x_1', \overline{q}_1')\right), \\ \quad \text{if } (\sigma \in \Sigma_{uo}) \wedge (\sigma' = \epsilon) \\ \quad\quad \wedge ((x_2, \overline{q}_2) = (x_2', \overline{q}_2')) \wedge (\overline{q}_2' \neq F); \\ \dfrac{\gamma\left((x_1, \overline{q}_1), \sigma, (x_1', \overline{q}_1')\right) \alpha\left(L_{G^R}\left((x_2, \overline{q}_2), \sigma', (x_2', \overline{q}_2')\right)\right)}{\alpha\left(L_{G^R}((x_2, \overline{q}_2), M(\sigma))\right)}, \\ \quad \text{if } (\sigma \in \Sigma - \Sigma_{uo}) \wedge (M(\sigma) = M(\sigma')) \\ \quad\quad \wedge \left(L_{G^R}\left((x_2, \overline{q}_2), \sigma', (x_2', \overline{q}_2')\right)\right) \neq \emptyset) \\ \quad\quad \wedge (\overline{q}_2' \neq F); \\ 0 \quad \text{otherwise.} \end{cases}$$

According to the definition of δ, when the first copy of G^R executes an unobservable event, the second copy responds by ϵ (since it observes nothing); if the first copy executes an observable event σ, then the second copy responds by executing a nonfault-trace consisting of sequence of unobservable events followed by an observable event that has the same mask value as $M(\sigma)$. Note a conditioning is applied to limit the executions of the second copy to indistinguishable nonfault-traces.

3) Check if every state $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ with $(x_1, \overline{q}_1) \in \partial_m^-(X \times \overline{Q})$ satisfies $(x_2, \overline{q}_2) \notin \Upsilon(X \times \overline{Q})$, (L, K) is S_m-Prognosable if and only if the answer is yes.

The following theorem guarantees the correctness of Algorithm 1.

Theorem 4: A pair $(L = L(G), K = L(R))$ of closed regular languages with $K \subseteq L$ is S_m-Prognosable if and only if any fault-state can only be reached in more than m-steps in G^R and every reachable state $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ of T with $(x_1, \overline{q}_1) \in \partial_m^-(X \times \overline{Q})$ satisfies $(x_2, \overline{q}_2) \notin \Upsilon(X \times \overline{Q})$.

Proof: Obviously, we have: any fault-state can only be reached in more than m-steps if and only if $m < \ell(\partial)$. Next, by the construction of T, for any $s \in L$ and $s' \in K$, $M(s) = M(s')$ if and only if there exists $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ such that $\delta(((x_0, q_0), (x_0, q_0)), (s, s'), ((x_1, \overline{q}_1), (x_2, \overline{q}_2))) > 0$. So if every reachable state $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ with $(x_1, \overline{q}_1) \in \partial_m^-(X \times \overline{Q})$ satisfies $(x_2, \overline{q}_2) \notin \Upsilon(X \times \overline{Q})$, then by Lemma 5, every $s \in \partial_m^-$ is not ambiguous with any nonindicator trace, i.e., $M^{-1}M(\partial_m^-) \cap \Upsilon = \emptyset$. Therefore, (L, K) is S_m-Prognosable according to Theorem 2, and

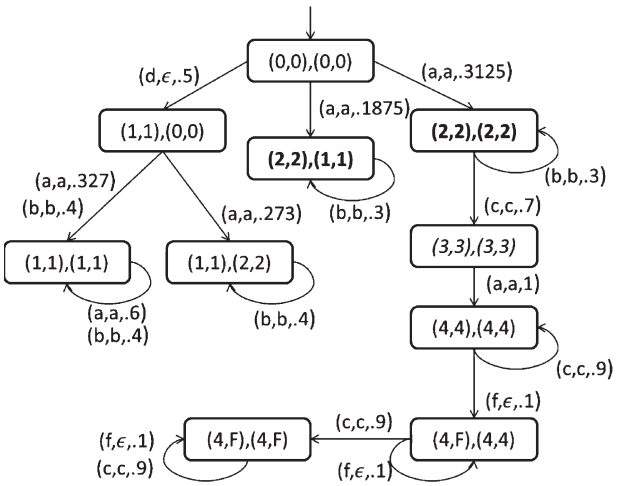

Fig. 2. Testing automaton for the system G^R shown in Fig. 1.

the sufficiency follows. On the other hand, if the theorem's condition is not satisfied, then either $m \geq \ell(\partial)$ or there exists (s, s') with $M(s) = M(s')$ and $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ such that $(x_1, \overline{q}_1) \in \partial_m^-(X \times \overline{Q})$, $(x_2, \overline{q}_2) \in \Upsilon(X \times \overline{Q})$ and $\delta(((x_0, q_0), (x_0, q_0)), (s, s'), ((x_1, \overline{q}_1), (x_2, \overline{q}_2))) > 0$. i.e., $s \in \partial_m^-$ and $s' \in \Upsilon$. Therefore, $M^{-1}M(\partial_m^-) \cap \Upsilon \neq \emptyset$. By Theorem 2, (L, K) is not S_m-Prognosable, which proves the necessity. ■

Example 7: Let us revisit the system shown in Fig. 1. According to Definition 3, $\mathfrak{J}(X \times \overline{Q}) = \{(2,2), (3,3), (4,4)\}$, $\Upsilon(X \times \overline{Q}) = \{(0,0), (1,1)\}$, $\partial_1^-(X \times \overline{Q}) = \{(3,3)\}$ and $\partial_2^-(X \times \overline{Q}) = \{(2,2)\}$. It is easy to check that $1 < 2 < \ell(\partial) = 4$. The testing automaton is shown in Fig. 2. The only state $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ such that $(x_1, \overline{q}_1) \in \partial_1^-(X \times \overline{Q})$ is labeled in *italic* and satisfies $(x_2, \overline{q}_2) \notin \Upsilon(X \times \overline{Q})$ and therefore (L, K) is S_1-Prognosable. All the states $((x_1, \overline{q}_1), (x_2, \overline{q}_2))$ such that $(x_1, \overline{q}_1) \in \partial_2^-(X \times \overline{Q})$ are labeled in bold, and there exists ((2,2),(1,1)) such that $(2,2) \in \partial_2^-(X \times \overline{Q})$ and $(1,1) \in \Upsilon(X \times \overline{Q})$. Therefore, (L, K) is not S_2-Prognosable. These are as expected from the discussion in Examples 3 and 4. ■

Remark 2: In Algorithm 1. G^R has $O(|X| \times |Q|)$ states and $O(|X|^2 \times |Q| \times |\Sigma|)$ transitions, and the testing automaton $T = G^R \times G^R$ has $O(|X|^2 \times |Q|^2)$ states and $O(|X|^4 \times |Q|^2 \times |\Sigma|^2)$ transitions. The computation of transition probabilities in T requires solving the matrix equation (1) for each $\sigma \in \Sigma - \Sigma_{uo}$ with complexity that is cubic in the number of states in G^R and linear in the number of events in G^R, namely, $O(|X|^3 \times |Q|^3 \times |\Sigma|)$. Thus, the complexity of constructing T is $O(|X|^4 \times |Q|^2 \times |\Sigma|^2 + |X|^3 \times |Q|^3 \times |\Sigma|)$. The shortest path to a fault state in G^R can be computed in $O(\sqrt{|X| \times |Q|} \times |X|^2 \times |Q| \times |\Sigma|)$ [12]. Identifying the set of m-steps interior nonfault-states in G^R can be done linearly in the size of G^R, i.e., $O(|X|^2 \times |Q| \times |\Sigma|)$, and identifying the set of indicator nonfault-states can be achieved by determining all the nonfault closed SCC in G^R using the algorithm in [13], which can be done in $O(|X|^3 \times |Q|^3)$. Therefore, the overall complexity of Algorithm 1 is $O(|X|^4 \times |Q|^2 \times |\Sigma|^2 + |X|^3 \times |Q|^3 \times |\Sigma|)$, which is polynomial in the number of states and events. Further if G is also deterministic (besides R) so that G^R has a smaller number of transitions, namely, $O(|X| \times |Q| \times |\Sigma|)$, then the verification complexity reduces to $O(|X|^2 \times |Q|^2 \times |\Sigma|^2 + |X|^3 \times |Q|^3 \times |\Sigma|)$. Furthermore, if the mask is "projection-type," the complexity further reduces due to a reduction in the number of transitions in G^R, where each state can now only have at most $|\Sigma|$ outgoing transitions, and thus the $|\Sigma|^2$ term will get replaced by $|\Sigma|$ in the complexity expression.

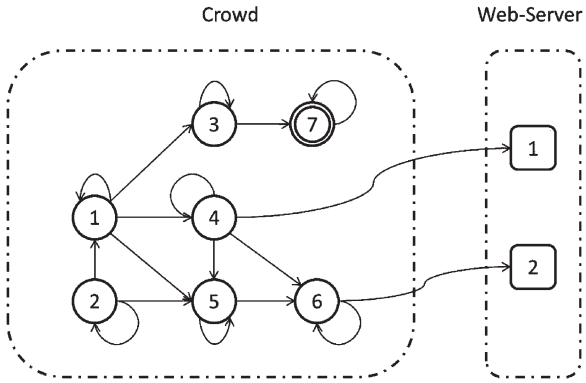

Fig. 3. Crowd with size 7 and 2 initiators.

VI. Illustrative Examples

In this section, two simple practical examples are given to illustrate our results.

Example 8: We consider the application of our results to the "Crowd" system, an anonymity protocol introduced in [14] that is used to protect the identity on the World Wide Web, which is recently studied in the stochastic DESs setting [15], [16]. When an user (called initiator) decides to send a message to a web server without revealing itself as the originator of the message, the user routes the message through a *crowd* of users (possibly itself). When a user in the crowd receives a message, it either sends the message to the web server or forwards the message to a user in the crowd (possibly itself). The above Crowd protocol is considered to be secure in hiding the identity of the originator. However, there can be a number of *corrupted* users in the crowd which can leak the information of the origin of the message, and so forwarding the message to a corrupted user is considered a fault. Further, as is customary with the analysis of Crowd ([17]), we also assume that a corrupted user does not forward a message to others. The process is depicted in Fig. 3, where the size of the crowd is taken to be 7, the possible initiators are {1,2} and the corrupted user is {7}. Now we consider the case that when a user tries to send a message to the web server and initiates a route, it also monitors the routing of that message to avoid the message being received by a corrupted user. The corresponding automaton model is given as Fig. 4, where a new initial state "0" is added from where the two initiator nodes "1" and "2" can be reached with equal probability. It is assumed that each user chooses one among its forwarding successors with a uniform probability distribution. Suppose three of the forwarding actions can be observed with the observation labels as shown, whereas the remaining forwarding actions are unobservable and so unlabeled. A fault is defined as the forwarding of a message to the

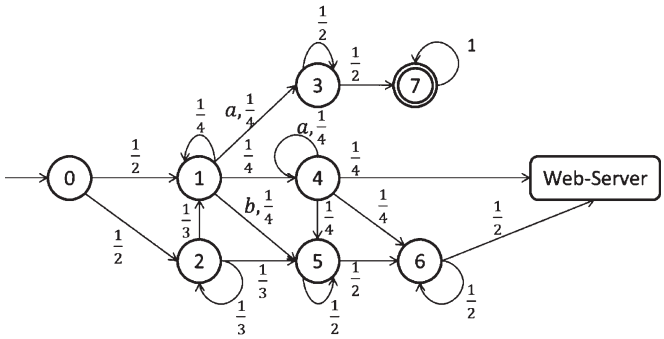

Fig. 4. Automaton for the Crowd system in Fig. 3.

corrupted user "7," i.e., the nonfault-specification can be obtained by removing the corrupted user "7" and all associate transitions. It can be checked that under this observation mask, the system is not S_m-Prognosable for any $m \geq 0$, since for any fault-trace reaching "7," all its prefixes are ambiguous with a certain nonindicator trace. To make the monitoring process meaningful, a control policy can be applied so that the self-loop of state "4" is forbidden, i.e., after receiving a message, the user "4" can only forward it to the user "5," "6" and web server. Then one can verify that the system is now S_1-Prognosable. Note in this example, neither the monitor nor the control has any affect on the corrupted user, leaving the corrupted user unaware of the existence of the monitoring or control. ■

Example 9: Consider the heating, ventilation and air conditioning (HVAC) system as examined in [2], [7], [9], which is modeled as a stochastic DES consisting of four components: a pump, a valve, a controller and a flow sensor. The model is shown in Fig. 5, which has 24 states, 11 events and 36 transitions, and is initialized at state 1. Each event in the stochastic DES has two parts, the first of which describes the motion of the controller and the second of which indicates the output of the flow sensor ("F" denotes "there is flow" and "NF" denotes "there is no flow," while no output by the flow sensor is described as ϵ, which for simplicity is omitted in Fig. 5). The unobservable events are given by $\Sigma_{uo} = \{\text{stuck\_closed}, \text{stuck\_open}\}$, which are also the fault-events Σ_f experienced by the controller; all other events are observed fully. The plant model shows the probability labels for each transitions. The deterministic nonfault-specification is obtained by excluding all the states resulted by the fault-events "stuck_closed" and "stuck_open," and is a subautomaton of the plant automaton, and without the probability labels (the definition of what constitutes a fault is independent of its occurrence probability). As can be seen, the shortest fault-trace is "stuck_closed" itself which has a length of 1. Therefore, the system cannot be S_m-Prognosable with $m \geq 1$. One can check that in this example every nonfault-trace has an extension reaching the absorbing nonfault-state "24" and hence is a nonindicator. Therefore, the system is not S_0-Prognosable. To achieve the S_0-Prognosability, one can exercise a control policy so that the system dynamics does not allow permanent idling by removing state "24" and adding a self-loop on state "10" with the same probability as transitioning to 24. Then one can verify by Algorithm 1 that the system is S_0-Prognosable. ■

VII. Comparison With Related Concepts

In this section, we will compare S_0-Prognosability with the notion of Prognosability in the logical setting [2], [3] and the notion of S-Diagnosability that are required for fault detection (as opposed to fault prediction) [7], [9], [18], [19]. To compare with the logical version of progonosability, we reproduce the definition from [3], specialized to centralized setting as follows:

Definition 4 ([3]): A pair (L, K) of closed languages with $K \subseteq L$ is said to be logically *Prognosable* if

$$(\forall s \in \partial)(\exists u \in s/\Sigma^{>0}) \left(M^{-1}M(u) \cap K \subseteq \widetilde{\mathfrak{J}} \right) \tag{12}$$

where $\widetilde{\mathfrak{J}}$ denotes the set of logical indicators and is given by $\widetilde{\mathfrak{J}} := \{s \in K : \exists n \in \mathbb{N}, L \setminus s \cap \Sigma^{\geq n} \subseteq [L - K] \setminus s\}$.

Remark 3: It is trivial to show that, for any $u \in K$, $(M^{-1}M(u) \cap K \subseteq \widetilde{\mathfrak{J}}) \Leftrightarrow (P_N^*(M(u)) = 0)$. Therefore, (12) can be equivalently written as

$$Pr\left(s \in \partial : \forall u \in s/\Sigma^{>0}, P_N^*\left(M(u)\right) > 0\right) = 0.$$

Comparing then with the definition of S_m-Prognosability under $m = 0$, so (5) can be written as

$$(\forall \tau, \rho > 0) Pr\left(s \in \partial : \forall u \in s/\Sigma^{>0}, P_N^*\left(M(u)\right) > \rho\right) < \tau.$$

It is obvious that if a system is logically Prognosable, then it is also S_0-Prognosable by definition. However, the converse is not true. For example, the system shown in Fig. 1 is S_1-Prognosable and hence is S_0-Prognosable by Corollary 1. However, it is not Prognosable since $\forall s \in \partial,\ u \in s/\Sigma^{>0},\ P_N^*(M(u)) > 0$. The stochastic version provides the flexibility of designing prognosers that can predict faults with arbitrary level of accuracy, which may be acceptable for certain applications even if 100% accuracy cannot be achieved (owing to lack of logical prognosability). Another artifact of this difference between the two notions is that, in logical setting, an indicator cannot visit a cycle of nonfault-states, which can be restrictive, but in stochastic setting, an indicator can visit a cycle of nonfault-states as long as the cycle does not form a closed SCC. In the example of Fig. 1, the prefix aca of the fault-trace $acaf$ is an indicator that ends in a non-closed cycle of nonfault-state (4,4) in G^R. While this does not violate stochastic prognosability, it ends up violating logical prognosability.

We established that stochastic-prognosability is weaker and more flexible than the logical counterpart. We next show that it is stronger than stochastic-diagnosability. The notion of S-Diagnosability, which supports fault detection after its occurrence, was introduced in [9] by the name of AA-Diagnosability and later renamed as S-Diagnosability in [7]. It requires a fault to be detected within a bounded delay of its occurrence with arbitrary level of confidence. We reproduce the definition from [7] as follows:

Definition 5 ([7]): A pair $(L = L(G), K = L(R))$ of closed regular languages with $K \subseteq L$ is said to be *Stochastically-Diagnosable*, or simply *S-Diagnosable*, if

$$(\forall \tau, \rho > 0)(\exists n \in \mathbb{N})(\forall s \in L - K)$$
$$Pr\left(t : t \in L \setminus s, |t| \geq n, Pr_{amb}(st) > \rho\right) < \tau \tag{13}$$

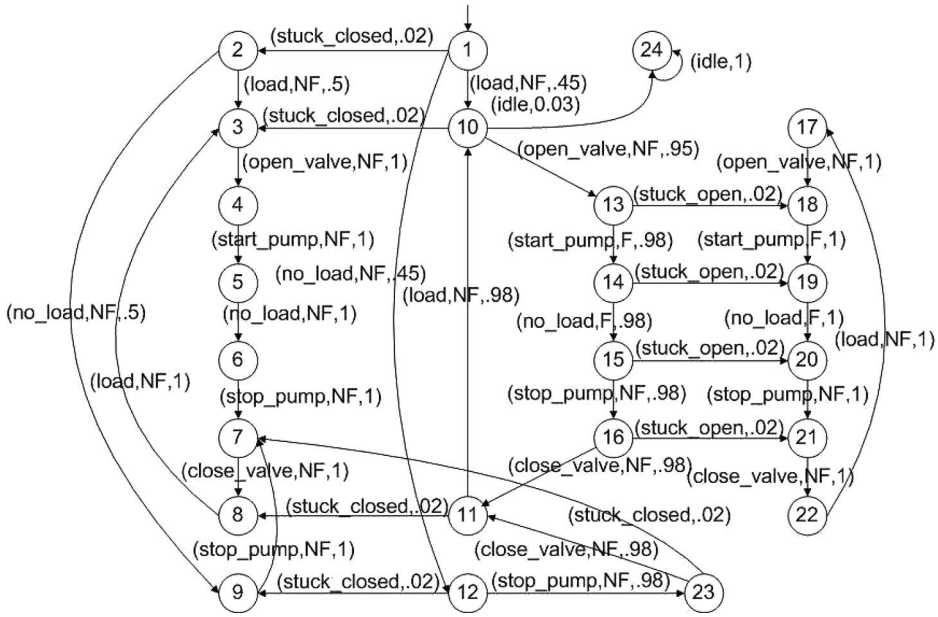

Fig. 5. Stochastic automaton G for the HVAC system under diagnosis, where each event has two parts, the first of which describes the motion of the controller and the second indicates the output of the flow sensor ("F" denotes "there is flow" and "NF" denotes "there is no flow," while no output by the flow sensor is described as ϵ which is simply omitted).

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

$$Pr_{amb}(s) := Pr\left(u \in K | M(u) = M(s)\right) \tag{14}$$

$$= \frac{Pr\left(u \in K : M(u) = M(s)\right)}{Pr\left(u \in L : M(u) = M(s)\right)}. \tag{15}$$

Note in the definition of $Pr_{amb}(s)$, "|" denotes the conditioning operation.

The next theorem establishes a necessary and sufficient condition for S-Diagnosability, which is reproduced from [8] and [18].

Theorem 5 ([8, Th. 1]): A pair $(L = L(G), K = L(R))$ of closed regular languages with $K \subseteq L$ is not S-Diagnosable if and only if

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

The next result shows that S_0-Prognosability is stronger than S-Diagnosability, meaning that whenever it is possible to predict faults, it is also possible to diagnose those, as can be expected.

Theorem 6: Given a pair $(L = L(G), K = L(R))$ of closed regular languages with $K \subseteq L$, if (L, K) is S_0-Prognosable, then it is S-Diagnosable. However, the converse need not hold.

Proof: We argue by contradiction. Assume (L, K) is S_0-Prognosable but not S-Diagnosable. Then there exist $s \in L - K$ and $s' \in K$ with $M(s) = M(s')$ satisfying the condition in Theorem 5, and all extensions of s' are nonfault-traces (otherwise the condition in Theorem 5 will not be satisfied). Then for any $n \in \mathbb{N}$, $Pr(t : t \in K \setminus s' \cap \Sigma^n) = 1$. Therefore, it follows from the definition of Υ that $s' \in \Upsilon$ (we can choose $\rho < 1$ to satisfy the definition of Υ). Considering $s \in L - K$ and $M(s) = M(s')$, we have $s \in M^{-1}M(\Upsilon) \cap L - K$, which is contradictory to Corollary 2 since (L, K) is S_0-Prognosable.

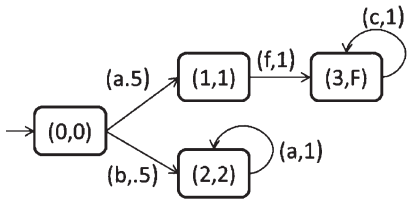

Fig. 6. Refined system G^R.

To see that the converse need not hold, we consider the system shown in Fig. 6, where $\Sigma = \{a, b, c, f\}$, $\Sigma_{uo} = \{b, f\}$ and for $\sigma \in \Sigma - \Sigma_{uo}$, $M(\sigma) = \sigma$. After the occurrence of fault-trace af, the only observations that can be produced are the traces in c^+, which are distinguishable from any nonfault-trace, and so (L, K) is S-Diagnosable. However, since $M^{-1}M(\partial_0^-) \cap \Upsilon = M^{-1}M(a) \cap \{\epsilon, ba^*\} = \{ba\} \neq \emptyset$, by Theorem 2, (L, K) is not S_0-Prognosable. ■

VIII. Conclusion

In this paper, we studied the prognosis of fault, i.e., its prediction prior to its occurrence, for stochastic discrete event systems. We formulated the notion of S_m-Prognosability for stochastic DESs, generalizing the corresponding notion from the logical setting [2], [3], and showed that it is a necessary and sufficient condition for the existence of a prognoser that can predict a fault at least m-steps prior to its occurrence, while achieving any arbitrary false alarm and missed detection

rates. (Higher accuracy of prognostic decision can be obtained by allowing shorter reaction bound.) A polynomial complexity algorithm for the verification of S_m-Prognosability was also provided, which checks on a pair of indistinguishable traces for the reachability of a pair of states, one of which is a m-steps interior nonfault-state and the other is a nonindicator state (such a pair is reachable if and only if S_m-Prognosability does not hold). The contribution of the work was further emphasized by comparing with previous related work on fault diagnosability, which was shown to be a weaker requirement than fault prognosability, as can be expected. There are several directions for future research: 1) An online recursive prognosis algorithm to compute the state distribution $\pi(o)$ resulted by an observation o so as to be able to predict a fault by checking whether $P_N^*(o) \leq \rho$, which in turn implies if $\pi(o)$ itself falls within a suitable range, and 2) algorithms for computing the decision threshold ρ and the largest possible reaction bound m for given performance requirements $\phi, \tau > 0$ for FA and MD rates. Also, an extension to the decentralized setting would be another direction for future work.

APPENDIX

Proof of Lemma 4: Since $P_N^*(M(st)) < \rho'$ if and only if $1 - P_N^*(M(st)) > 1 - \rho'$, letting $\rho := 1 - \rho'$, (11) is true if and only if

$$(\forall \rho, \phi > 0)(\exists d \in \mathbb{N})(\forall s \in \aleph) \\ Pr\left(t : t \in \aleph \setminus s, |t| \geq d, 1 - P_N^*\left(M(st)\right) > \rho\right) < \phi. \quad (16)$$

Thus, showing (11) is equivalent to showing that (16) holds. Next we show that (16) is equivalent to showing that the pair $(K, K - \aleph)$ is S-Diagnosable. First note that for any $st \in \aleph \subseteq \Upsilon$, it holds that

$$\begin{aligned} P_N^*\left(M(st)\right) &= \frac{\min_{n \in \mathbb{N}} Pr\left(\left\{M^{-1}M(st) \cap K\right\} \Sigma^n \cap K\right)}{Pr\left(M^{-1}M(st) \cap L\right)} \\ &= \frac{Pr\left(M^{-1}M(st) \cap \aleph\right)}{Pr\left(M^{-1}M(st) \cap L\right)} \\ &= \frac{Pr\left(M^{-1}M(st) \cap \aleph\right)}{Pr\left(M^{-1}M(st) \cap K\right)} \end{aligned}$$

where we have used the fact that (L, K) is S_m-Prognosable and so for $st \in \Upsilon$, $M^{-1}M(st) \cap L = M^{-1}M(st) \cap [K \cup (L - K)] = M^{-1}M(st) \cap K$ (follows from Corollary 2). Then,

$$\begin{aligned} 1 - P_N^*\left(M(st)\right) &= 1 - \frac{Pr\left(M^{-1}M(st) \cap \aleph\right)}{Pr\left(M^{-1}M(st) \cap K\right)} \\ &= \frac{Pr\left(M^{-1}M(st) \cap (K - \aleph)\right)}{Pr\left(M^{-1}M(st) \cap K\right)} \quad (17) \end{aligned}$$

which is the probability of ambiguity of st as in (15) when the pair of languages (L, K) is replaced with $(K, K - \aleph)$. Thus, we can replace $1 - P_N^*(M(st))$ in (16) with the right-hand side of (17), and in which case (16) becomes equivalent to S-Diagnosability of $(K, K - \aleph)$ as in (13).

Next we show that the pair $(K, K - \aleph)$ is indeed S-Diagnosable. Assume for contradiction that $(K, K - \aleph)$ is not S-Diagnosable. Then there exists $s \in \aleph$ and $s' \in K - \aleph$ satisfying the condition of Theorem 5. Then we have $\forall n \in \mathbb{N}$, $Pr(t : t \in [K - \aleph] \setminus s' \cap \Sigma^n) = \sum_{o \in \Delta^*} Pr(t : t \in [K - \aleph] \setminus s' \cap \Sigma^n, M(t) = o) = \sum_{o \in \Delta^*} Pr(t : t \in K \setminus s \cap \Sigma^n, M(t) = o) = Pr(t : t \in K \setminus s \cap \Sigma^n) = 1$, where the second equality follows from Theorem 5 and the last equality follows from the fact that $s \in \aleph$ (so all its extensions are in K). Thus, $\forall n \in \mathbb{N}$, $Pr(t : t \in [K - \aleph] \setminus s' \cap \Sigma^n) = 1$, i.e., $\forall n \in \mathbb{N}$, $Pr(\{s'\}\Sigma^n \cap (K - \aleph)) = 1$, implying that $\forall n \in \mathbb{N}$, $Pr(\{s'\}\Sigma^n \cap K) = 1$ (since $K - \aleph \subseteq K$), which further implies that $s' \in \aleph$. This contradicts the fact that $s' \in K - \aleph$. So the S-Diagnosability of $(K, K - \aleph)$ follows, which proves (16) and equivalently (11). ■

REFERENCES

[1] G. Vachtsevanos, F. Lewis, M. Roemer, A. Hess, and B. Wu, *Intelligent Fault Diagnosis and Prognosis for Engineering Systems*. Hoboken, NJ, USA: Wiley, 2006.

[2] S. Genc and S. Lafortune, "Predictability of event occurrences in partially-observed discrete-event systems," *Automatica*, vol. 45, no. 2, pp. 301–311, 2009.

[3] R. Kumar and S. Takai, "Decentralized prognosis of failures in discrete event systems," *IEEE Trans. Autom. Control*, vol. 55, no. 1, pp. 48–59, Jan. 2010.

[4] F. Cassez and A. Grastien, "Predictability of event occurrences in timed systems," in *Proc. 11th Int. Conf. Formal Modeling Anal. Timed Syst.*, Buenos Aires, Argentina, Aug. 2013.

[5] S. Takai and R. Kumar, "Inference-based decentralized prognosis in discrete event systems," *IEEE Trans. Autom. Control*, vol. 56, no. 1, pp. 165–171, Jan. 2011.

[6] S. Takai and R. Kumar, "Distributed failure prognosis of discrete event systems with bounded-delay communications," *IEEE Trans. Autom. Control*, vol. 57, no. 5, pp. 1259–1265, May 2012.

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

[8] J. Chen and R. Kumar, "Failure detection framework for stochastic discrete event systems with guaranteed error bounds," *IEEE Trans. Autom. Control*, 2015, to be published, DOI: 10.1109/TAC.2014.2382991.

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

[10] 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.

[11] X. Wang and A. Ray, "A language measure for performance evaluation of discrete-event supervisory control systems," *Appl. Math. Model.*, vol. 28, no. 9, pp. 817–833, Sep. 2004.

[12] A. V. Goldberg, "Scaling algorithms for the shortest paths problem," *SIAM J. Comput.*, vol. 24, no. 3, pp. 494–504, Jun. 1995.

[13] A. Xie and P. A. Beerel, "Efficient state classification of finite-state Markov chains," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 17, no. 12, pp. 1334–1339, Dec. 1998.

[14] M. K. Reiter and A. D. Rubin, "Crowds: Anonymity for web transactions," *ACM Trans. Inf. Syst. Security*, vol. 1, no. 1, pp. 66–92, Nov. 1998.

[15] A. Saboori and C. N. Hadjicostis, "Probabilistic current-state opacity is undecidable," in *Proc. 19th Int. Symp. Math. Theory Netw. Syst.*, Budapest, Hungary, Jul. 2010.

[16] B. Bérard, J. Mullins, and M. Sassolas, "Quantifying opacity," in *Proc. Int. Conf. Quantit. Eval. Syst.*, Williamsburg, VA, USA, Sep. 2010, pp. 263–272.

[17] V. Shmatikov, "Probabilistic analysis of anonymity," in *Prof. 15th IEEE Comput. Security Foundat. Workshop*, 2002, pp. 119–128.

[18] 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.

[19] 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.

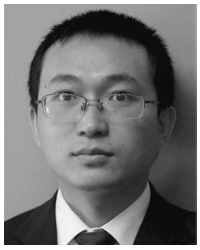

Jun Chen (S'11–M'14) received the bachelor's degree in electrical engineering from Zhejiang University, Hangzhou China, in 2009, and the Ph.D. degree in electrical engineering from Iowa State University, Ames, IA, USA, in 2014.

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, and currently is a Member of the IEEE.

Ratnesh Kumar (S'87–M'90–SM'00–F'07) received the B.Tech. degree in electrical engineering from IIT Kanpur, 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 at Iowa State University since 2002. Prior to this, he held faculty position at the University of Kentucky (1991–2002) in electrical and computer engineering 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.

Prof. Kumar is or has been an associate editor of *ACM Transactions on Embedded Computing Systems*, *SIAM Journal on Control and Optimization*, IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION, *Journal of Discrete Event Dynamical Systems*, IEEE Control Systems Society, IEEE Robotics and Automation Systems Society, and IEEE Workshop on Software Cybernetics. He received Gold Medals for Best EE undergrad and Best All Rounder at IIT Kanpur, Best Dissertation Award at UT Austin, and is a Fellow of the IEEE.