Revised Test for Stochastic Diagnosability of Discrete-Event Systems

Jun Chen, *Member, IEEE*, Christoforos Keroglou, Christoforos N. Hadjicostis, *Senior Member, IEEE*, and Ratnesh Kumar, *Fellow, IEEE*

***Abstract*—This paper provides revisions to the algorithms presented by Chen *et al.*, 2013 for testing diagnosability of stochastic discrete-event systems. Additional new contributions include PSPACE-hardness of verifying strong stochastic diagnosability (referred as A-Diagnosability in Thorsley *et al.*, 2005) and a necessary and sufficient condition for testing stochastic diagnosability (referred as AA-Diagnosability in Thorsley *et al.*, 2005) that involves a new notion of probabilistic equivalence.**

***Note to Practitioners*—Detecting system failures is essential prior to fault mitigation. For stochastic discrete-event systems, the property of stochastic diagnosability (S-Diagnosability) allows one to detect any system failure with arbitrarily small error bound and within bounded delay. This paper contributes by revising and extending the results in the previous work by Chen *et al.*, 2013, regarding the verification of S-Diagnosability.**

***Index Terms*—Complexity, discrete-event systems (DESs), failure diagnosis, PSPACE-hardness.**

I. Introduction

TWO notions of diagnosability for stochastic discrete-event systems (DESs) were introduced and studied in [1], and later further explored in [2]. The property of stochastic diagnosability (S-Diagnosability; referred as AA-Diagnosability in [1]) requires that given any tolerable ambiguity level ρ and error bound τ, there must exist a delay bound n, such that for any faulty trace, its extensions, longer than n and with a probability of ambiguity higher than ρ, occur with probability smaller than τ. The property of strong stochastic diagnosability (SS-Diagnosability; referred as A-Diagnosability in [1]) restricts this by having tolerance bound $\rho = 0$. Methods for verifying a sufficient (but not necessary) condition for S-Diagnosability and a sufficient and necessary condition for its stronger version were presented in [1]. In [2], improvements were proposed by presenting necessary and sufficient conditions for both S- and SS-Diagnosability. However, there were subtle oversights in the tests of [2], arising when there is nondeterminism, which are revised in this paper. Some new results are also presented.

Recall that in a graph, a strongly connected component (SCC) is a subgraph with the property that any pair of nodes of the subgraph can reach each other, through paths within the subgraph. Furthermore, an SCC is closed if the outgoing probability at any node is zero (formal definitions are given in Section II). The oversight in [2] stems from the fact that when two closed SCCs synchronize, then in the presence of nondeterminism, the result may not remain closed (the synchronized system may not stay within the SCC with probability 1).

Note that the violation of SS-Diagnosability occurs when there exist two closed SCCs in the system that are persistently ambiguous (meaning one is reached upon fault and another without fault, with indistinguishable observation, and all extensions within the faulty SCC remain ambiguous). Owing to nondeterminism, the synchronization of two such closed SCCs need not manifest as a closed SCC in the testing automaton of [2] (which is an observationally synchronized product of two system copies), and may possess outgoing transitions with nonzero probability. See Fig. 1, where the persistently closed SCCs (consisting of states $1, 2$, and 3, and $4, 5$, and 6, respectively) of the system synchronize to produce a nonclosed SCC consisting of states $[(5, F), (2, 2)]$ and $[(4, F), (1, 1)]$ of testing automaton that has nonzero probability in the outgoing transition to state $[(4, F), (2, 2)]$. As a result, the condition for non-SS-Diagnosability of our previous paper [2], namely the existence of a closed SCC in the testing automaton, is only sufficient but not necessary for non-SS-Diagnosability. So the result of [2] only provides a polynomially verifiable necessary condition for SS-Diagnosability. In fact, we prove in this paper that likely no polynomially verifiable necessary and sufficient condition can be found, by establishing the problem to be PSPACE-hard.

The same problem of incorrectly handling nondeterminism due to partial observability also results in the condition for S-Diagnosability in [2] being only necessary. In this paper, we provide a correct necessary and sufficient condition for testing S-Diagnosability by weakening the sufficient condition in [1], by utilizing a new notion of p-equivalence. The necessary and sufficient condition for non-S-Diagnosability first checks for the existence of a pair of persistently ambiguous closed SCCs, and next checks whether any such pair can remain p-equivalent, in the sense that any future observation occurs with the same distribution in both SCCs.

Besides revising the tests for S- and SS-Diagnosability of [2], the new contributions of this paper are summarized as follows.

- We prove that deciding SS-Diagnosability is PSPACE-hard.
- We provide a neccessary and sufficient condition for S-Diagnosability, and a testing algorithm, which, whenever it terminates, can conclusively test S-Diagnosability.

II. Notation and Background

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 transition probability function [3]. G is said to be nonstochastic if $\alpha : X \times \Sigma \times X \to \{0, 1\}$, and a nonstochastic DES is said to be

Manuscript received January 22, 2016; revised March 18, 2016; accepted March 28, 2016. Date of publication May 2, 2016; date of current version January 4, 2018. This paper was recommended for publication by Associate Editor C. Seatzu and Editor S. Reveliotis upon evaluation of the reviewers' comments. This work was supported by the National Science Foundation under Grant NSF-CCF-1331390, Grant NSF-ECCS-0801763, Grant NSF-ECCS-0926029, and Grant NSF-ECCS-1509420.

J. Chen was with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA. Current address: Idaho National Laboratory, Idaho Falls, ID 83415 USA (e-mail: jchenee2015@gmail.com).

C. Keroglou and C. N. Hadjicostis are with the Department of Electrical and Computer Engineering, University of Cyprus, Nicosia 1678, Cyprus (e-mail: keroglou.christoforos@ucy.ac.cy; chadjic@ucy.ac.cy).

R. Kumar is with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA (e-mail: rkumar@iastate.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TASE.2016.2551746

deterministic if $\forall x \in X, \sigma \in \Sigma, \sum_{x' \in X} \alpha(x, \sigma, x') \in \{0, 1\}$. The transition probability function α can be recursively extended from domain $X \times \Sigma \times X$ to $X \times \Sigma^* \times X$ 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 the language generated by G as $L(G) := \{s \in \Sigma^* : \exists x \in X, \alpha(x_0, s, x) > 0\}$. 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 observable symbols. An event σ is said to be unobservable if $M(\sigma) = \epsilon$; the set of unobservable events is denoted by 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^* in a natural way.

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 an SCC or irreducible, if $\forall x, x' \in X_C, \exists s \in \Sigma^*$, such that $\alpha_C(x, s, x') > 0$. An SCC C is said to be closed if for each $x \in X_C$, $\sum_{\sigma \in \Sigma} \sum_{x' \in X_C} \alpha_C(x, \sigma, x') = 1$. Note that for SCC that is not closed, α_C does not necessarily denote probability, as in this case, the outgoing probability from some states in that SCC may not equal 1. The states that belong to a closed SCC are recurrent states, and the remaining states (that do not belong to any closed SCC) are transient states.

For a stochastic DES $G = (X, \Sigma, \alpha, x_0)$, its nonfaulty behaviors are specified in the form of a deterministic automaton $R = (Q, \Sigma, \beta, q_0)$, such that $L(R) = K$ is the set of nonfaulty traces. Then, the remaining traces $L - K$ are called the faulty behaviors. The refinement of G with respect to R, denoted by G^R, can be used to capture the traces, violating the given specification in the form of the reachability of a faulty state and 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: $(\overline{q}, \overline{q}' \in Q \wedge \beta(\overline{q}, \sigma, \overline{q}') > 0) \vee (\overline{q} = \overline{q}' = F) \vee (\overline{q}' = F \wedge \sum_{q \in Q} \beta(\overline{q}, \sigma, q) = 0)$, and otherwise $\gamma((x, \overline{q}), \sigma, (x', \overline{q}')) = 0$.

Definition 1: Given a stochastic DES G, deterministic nonfault specification R with generated languages $L = L(G)$ and $K = L(R)$, (G, R), is said to be S-Diagnosable, if

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

where $Pr_{\text{amb}} : L - K \to [0, 1]$ is given by

$$Pr_{\text{amb}}(s) := \frac{Pr(u \in K : M(u) = M(s))}{Pr(u \in L : M(u) = M(s))}.$$

The definition of SS-Diagnosability is obtained by setting $\rho = 0$ in Definition 1.

Example 1 illustrates the omission present in the testing algorithms of [2]. As noted above, this is caused by incorrect handling of nondeterminism, under which the synchronization need not preserve the closedness of SCCs. Note the example is intentionally kept simple so the omission of [2] can be witnessed by inspection.

Example 1: Consider the stochastic plant model G and deterministic nonfault specification generator R shown in Fig. 1, where f is a fault event and unobservable. The mask function is defined as: $M(f) = M(e) = \epsilon$, $M(a_1) = M(a_2) = a$, and $M(b) = b$. The behaviors after the occurrence of f as well as e are identical under the observation mask M, and so clearly, the system is not S-Diagnosable, and hence, not SS-Diagnosable (see Definition 1). However, in the testing automaton T obtained using [2, Algorithm 1], as shown in Fig. 1(d), there is no closed ambiguous SCC or biclosed ambiguous SCC (whose existence is required by the algorithms of [2] for the violation of SS- and S-Diagnosability). ■

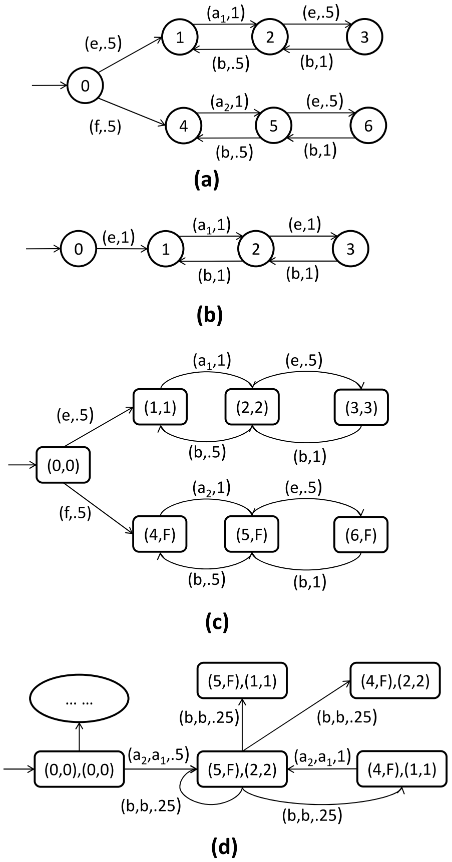

Fig. 1. Counter example, where $M(f) = M(e) = \epsilon$, $M(a_1) = M(a_2) = a$, and $M(b) = b$. (a) System G. (b) Specification R. (c) Refined system G^R. (d) Testing automaton T.

III. Corrected and New Results for SS-Diagnosability

Since closedness of SCCs is not preserved under composition for nondeterministic systems [2, eq. (2)] is incorrect, and only the necessity part of the proof to [2, Th. 1] holds. Thus, [2, Th. 1] only provides a necessary condition for SS-Diagnosability, as stated below.

Theorem 1 (Correction to [2, Th. 1]): (G, R) is SS-Diagnosable only if every closed SCC in T is unambiguous.

The corrected theorem then only provides a polynomially verifiable necessary condition for SS-Diagnosability, whereas a necessary and sufficient test for SS-Diagnosability with exponential complexity is given in [1]. We next show that SS-Diagnosability is unlikely to have a polynomial complexity algorithm by establishing its PSPACE-hardness, via a polynomial-time reduction of the Universality problem to an instance of the SS-Diagnosability problem. Since the former is PSPACE-hard, this proves that the SS-Diagnosability problem is also PSPACE-hard.

Given a nondeterministic finite automaton G_N over the alphabet Σ, the Universality problem asks if the language $L(G_N)$ contains all finite words over Σ, i.e., if $L(G_N) = \Sigma^*$ [4]. A formal definition follows.

Definition 2 [4]: Given a nondeterministic finite automaton $G_N = (X_N, \Sigma_o, \delta_N, X_N^0)$, such that the set of initial states $X_N^0 = X_N$, do we have $L(G_N) = \Sigma_o^*$?

When $|\Sigma| \geq 2$, the Universality problem with all states as initial states is known to be PSPACE-hard [4]. We now establish the aforementioned reduction. For any $G_N = (X_N, \Sigma_o, \delta_N, X_N)$,

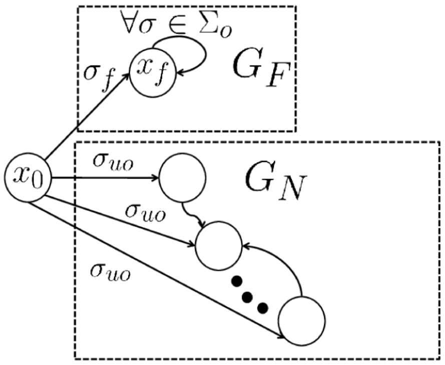

Fig. 2. Instance of SS-Diagnosability given G_N.

let $G = (X, \Sigma, \alpha, x_0)$ be such that $X = \{x_0\} \cup X_N \cup \{x_f\}$ where x_0 and x_f are new states (not in X_N) and $\Sigma = \Sigma_o \cup \{\sigma_{uo}, \sigma_f\}$ where σ_f and σ_{uo} are new events (not in Σ_o) that are unobservable, and Σ_o (events of G_N) is the set of observable events with $|\Sigma_o| \geq 2$ and $M(\sigma) = \sigma$, $\forall\sigma \in \Sigma_o$. We assign probabilities as follows.

1) $\alpha(x_0, \sigma_f, x_f) = \frac{1}{|X_N|+1}$, and $\forall\sigma \in \Sigma_o$, $\alpha(x_f, \sigma, x_f) = \frac{1}{|\Sigma_o|}$.
2) $\forall x \in X_N$, $\alpha(x_0, \sigma_{uo}, x) = \frac{1}{|X_N|+1}$.
3) $\forall x, x' \in X_N$ and $\forall\sigma \in \Sigma_o$, if $x' \in \delta_N(x, \sigma)$, then $\alpha(x, \sigma, x') = \frac{1}{\sum_{\sigma\in\Sigma_o} |\delta_N(x,\sigma)|}$.
4) $\alpha(x, \sigma, x') = 0$, for all other cases.

Note that (if we ignore probabilities): 1) G can be seen as the union of $G_N = (X_N, \Sigma_o, \delta_N, X_N)$ and the singleton state automaton $G_F = (\{x_f\}, \Sigma_o, \delta_f, x_f)$, whose language $L(G_F) = \Sigma_o^*$ and 2) there exists a transition (with unobservable event) from x_0 to each state in X_N and a transition (also with unobservable event) to x_f.

To construct a diagnosability problem, we let R be such that $K = L(R) = \sigma_{uo}L(G_N) \subseteq \sigma_{uo}\Sigma_o^*$. Then, $L - K = \sigma_f\Sigma_o^*$. Moreover, X_N can be seen as the set of states of G that is consistent with the nonfaulty behavior of G, and x_f can be seen as the singleton state that is consistent with the faulty behavior of G, i.e., all nonfaulty traces in K will transition G to states in G_N while all faulty traces in $L - K$ will transition G to the singleton state x_f.

Theorem 2 shows that every instance of the language universality problem of $G_N = (X_N, \Sigma_o, \delta_N, X_N)$ with all initial states can be reduced to an instance of SS-Diagnosability problem (as it was described in the previous paragraph). Thus, the SS-Diagnosability problem is PSPACE-hard.

Theorem 2: Verification of SS-Diagnosability is PSPACE-hard.

Proof: We argue that $L(G_N) = \Sigma_o^*$ if and only if G (as in Fig. 2 and as described earlier) is not SS-Diagnosable.

($\rightarrow$) If $L(G_N) = \Sigma_o^*$, then for all $s \in L - K = \sigma_f\Sigma_o^*$, there exists $s' \in \sigma_{uo}\Sigma_o^* = K$, such that $M(s) = M(s')$. Therefore, for all extensions t of the faulty trace σ_f, $Pr_{\text{amb}}(s_f t) > 0$. Let $\tau < 1$, then for all n, $Pr(t : t \in L\backslash s, |t| \geq n, Pr_{\text{amb}}(s_f t) > 0) = 1 > \tau$. Hence G is not SS-Diagnosable.

($\leftarrow$) If $L(G_N) \neq \Sigma_o^*$, i.e., $L(G_N) \subset \Sigma_o^*$, then there exists $s \in \Sigma_o^*$, such that $s \notin L(G_N)$. Notice that $\forall u \in \Sigma_o^*$, $us \notin L(G_N)$, since $\delta_N(X_N, us) = \delta_N(\delta_N(X_N, u), s) \subseteq \delta_N(X_N, s) = \emptyset$. For any faulty trace $s_f \in L - K = \sigma_f\Sigma_o^*$, there exists an extension $u_1 s$, such that $u_1 s \in L\backslash s_f$ and $Pr_{\text{amb}}(s_f u_1 s) = 0$. Let $n_1 = |u_1 s|$ and $p_1 = 1 - Pr(u_1 s) < 1$. For other extensions $t \in L\backslash s_f \cap \Sigma^{n_1}$, such that $Pr_{\text{amb}}(s_f t) > 0$, $s_f t$ has at least one extension $u_2 s$, such that $Pr_{\text{amb}}(s_f t u_2 s) = 0$. Let $n_2 := |t u_2 s|$ and $p_2 = 1 - Pr(u_2 s) < 1$. In general, any ambiguous extensions of s_f would have at least one unambiguous extension. Let n_k be the length of the kth shortest unambiguous extension of s_f, and so

$$\begin{aligned} &Pr(t : t \in L\backslash s_f, |t| \geq n_k, Pr_{\text{amb}}(s_f t) > 0) \\ &\quad \leq \prod_{i=1}^{k} Pr(t_i : t_i \in L\backslash s_f t_1 \ldots t_{i-1}, |t_i| = n_i - n_{i-1}, \\ &\qquad\qquad Pr_{\text{amb}}(s_f t_1 \ldots t_i) > 0) \\ &\qquad \times Pr(t \in L\backslash s_f t_1 \ldots t_k, Pr_{\text{amb}}(s_f t_1 \ldots t_k t) > 0) \\ &\quad \leq \prod_{i=1}^{k}(1 - p_i). \end{aligned}$$

Since $\forall i$, $1 - p_i < 1$, the above quantity approaches (if not equals) 0 as k increases, equivalently as n_k increases. Therefore, for any $\tau > 0$, there should exist $n_k > 0$, such that

$$Pr(t : t \in L\backslash s_f, |t| \geq n_k, Pr_{\text{amb}}(s_f t) > 0) < \tau.$$

Hence, G is SS-Diagnosable.

This establishes the reduction with complexity of $O(|X_N|^2 \times |\Sigma_o|)$ and concludes that the verification of SS-Diagnosability is PSPACE-hard. ■

Remark 1: Indeed as reported in [5], the problems over probabilistic automata are more complex (compared with deterministic analogs), e.g., the emptiness problem for probabilistic automata is undecidable. ■

IV. Corrected and New Results for S-Diagnosability

An exponential complexity sufficient condition for S-Diagnosability was presented in [1]. In this section, we first review the sufficient condition presented in [1], followed by a correct necessary and sufficient condition for S-Diagnosability by weakening the sufficient condition in [1], by the way of utilizing the notion of p-equivalence. We start with the following notations needed for the construction of the deterministic observer that was used in [1] to check the existence of a pair of persistently ambiguous SCCs.

For $(x_i, \overline{q}_i), (x_j, \overline{q}_j) \in X \times \overline{Q}$ and $\delta \in \Delta$, define the set of traces originating at $(x_i, \overline{q}_i)$, terminating at $(x_j, \overline{q}_j)$ and executing a sequence of unobservable events, followed by a single observable event with observation δ as $L_{G^R}((x_i, \overline{q}_i), \delta, (x_j, \overline{q}_j)) := \{s \in \Sigma^* \mid s = u\sigma, M(u) = \epsilon$, and $M(\sigma) = \delta, \gamma((x_i, \overline{q}_i), s, (x_j, \overline{q}_j)) > 0\}$. Define $\alpha(L_{G^R}((x_i, \overline{q}_i), \delta, (x_j, \overline{q}_j))) := \sum_{s\in L_{G^R}((x_i,\overline{q}_i),\delta,(x_j,\overline{q}_j))} \gamma((x_i, \overline{q}_i), s, (x_j, \overline{q}_j))$ and denote it by $\mu_{i,\delta,j}$ for short, i.e., it is the probability of all traces originating at $(x_i, \overline{q}_i)$, terminating at $(x_j, \overline{q}_j)$ and executing a sequence of unobservable events, followed by a single observable event with observation δ. Also define $\pi_{ij} = \sum_{\sigma\in\Sigma_{uo}} \gamma((x_i, \overline{q}_i), \sigma, (x_j, \overline{q}_j))$ as the probability of transitioning from $(x_i, \overline{q}_i)$ to $(x_j, \overline{q}_j)$ while executing a single unobservable event. Then, it can be seen that $\mu_{i,\delta,j} = \sum_k \pi_{ik}\mu_{k,\delta,j} + \sum_{\sigma\in\Sigma:M(\sigma)=\delta} \gamma((x_i, \overline{x}_i), \sigma, (x_j, \overline{q}_j))$, where the first term on right-hand side corresponds to transitioning in at least two steps, whereas the second right-hand side term corresponds to transitioning in exactly one step. Thus, for each $\delta \in \Delta$, given the values $\{\pi_{ij} | i, j \in X \times \overline{Q}\}$ and $\{\sum_{\sigma\in\Sigma:M(\sigma)=\delta} \gamma((x_i, \overline{q}_i), \sigma, (x_j, \overline{q}_j)) | i, j \in X \times \overline{Q}\}$, all the probabilities $\{\mu_{i,\delta,j} | i, j \in X \times \overline{Q}\}$ can be found by solving the following matrix equation (see, for example, [6] for a similar matrix equation that calculates the absorbing probabilities in an absorbing Markov chain):

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

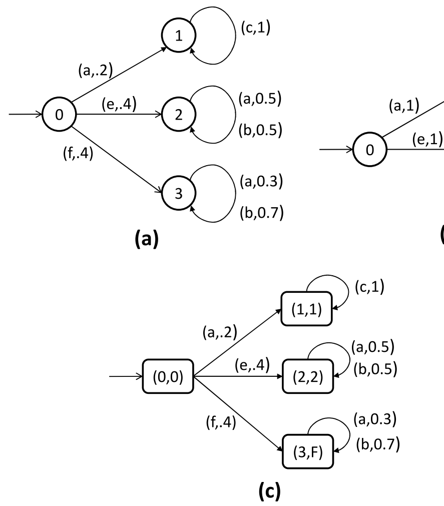

Fig. 3. (a) System G with mask function defined as: $M(f) = \epsilon$, $M(a) = a$, $M(b) = b$, and $M(c) = c$. (b) Specification R. (c) Refined system G^R.

where $\boldsymbol{\mu}(\delta)$, $\boldsymbol{\pi}$, and $\boldsymbol{\gamma}(\delta)$ are all $|X \times \overline{Q}| \times |X \times \overline{Q}|$ square matrices, whose ijth elements are given by $\mu_{i,\delta,j}$, π_{ij} and $\sum_{\sigma \in \Sigma : M(\sigma)=\delta} \gamma((x_i, \overline{q}_i), \sigma, (x_j, \overline{q}_j))$, respectively.

Example 2: For the system shown in Fig. 3, the observable event set is given by $\{a, b, c\}$. Then, $L_{G^R}((0,0), a, (1,1)) = \{a\}$ while $L_{G^R}((0,0), b, (3,F)) = \{fb\}$. Furthermore, by ordering the state space of G^R as $(0,0)$, $(1,1)$, $(2,2)$, $(3,F)$ and solving matrix equations (1), we get

$$\boldsymbol{\mu}(a) = \begin{bmatrix} 0 & 0.2 & 0.2 & 0.12 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.3 \end{bmatrix}$$

$$\boldsymbol{\mu}(b) = \begin{bmatrix} 0 & 0 & 0.2 & 0.28 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.7 \end{bmatrix}$$

$$\boldsymbol{\mu}(c) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

■

Define state distribution $\pi : M(L) \to [0,1]^{|X \times \overline{Q}|}$, which is recursively computed as: $\pi(\epsilon) = [1, 0, \ldots, 0]$, and for any $o \in M(L)$ and $\delta \in \Delta$, $\pi(o\delta) = (\pi(o)\mu(\delta)/||\pi(o)\mu(\delta)||)$ [7]. Given an SCC C of G^R, the state distribution $\pi^C(o)$ over C following an observation $o \in M(L)$ is defined as a vector with the same size as the number of states in C, whose ith element is given by $\pi_i^C(o) := (\pi_i(o)/\sum_{i \in C} \pi_i(o))$, provided the denominator is nonzero, and undefined otherwise. Given two SCCs C_1 and C_2 of G^R, a pair of distributions π^1 over C_1 and π^2 over C_2 is said to be reachable, if there exists $o \in M(L)$, such that $\pi^{C_1}(o) = \pi^1$ and $\pi^{C_2}(o) = \pi^2$.

Example 3: Supposed ab is observed for the system shown in Fig. 3. Then, it can be computed that

$$\begin{aligned} \pi(a) &= \frac{[\,1 \quad 0 \quad 0 \quad 0\,] \times \boldsymbol{\mu}(a)}{||[\,1 \quad 0 \quad 0 \quad 0\,] \times \boldsymbol{\mu}(a)||} \\ &= [\,0 \quad 0.385 \quad 0.385 \quad 0.230\,] \\ \pi(ab) &= \frac{\pi(a)\boldsymbol{\mu}(b)}{||\pi(a)\boldsymbol{\mu}(b)||} = [\,0 \quad 0 \quad 0.543 \quad 0.457\,]. \end{aligned}$$

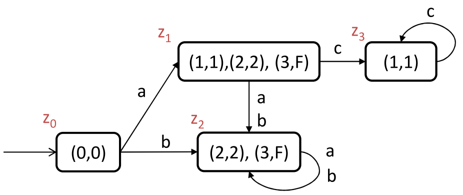

Fig. 4. Observer automaton Obs for the system in Fig. 3, whose state space is defined as $z_0 = \{(0,0)\}$, $z_1 = \{(1,1), (2,2), (3,F)\}$, $z_2 = \{(2,2), (3,F)\}$, and $z_3 = \{(1,1)\}$.

G^R, as in Fig. 3(c), has three SCCs, namely, C_1 consisting of $(2,2)$ with its self-loop transitions, C_2 consisting of $(3,F)$ with its self-loop transitions, and C_3 consisting of $(1,1)$ with its self-loop transition. Then, after observing ab, $\pi^{C_1}(ab) = [1]$ and $\pi^{C_2}(ab) = [1]$, while $\pi^{C_3}(ab)$ is undefined. ■

Given G^R, let Obs denote the deterministic observer automaton that tracks the possible system states following each observation. Obs has state space $Z \subseteq 2^{X \times \overline{Q}}$, and is initialized at node $z_0 = \{(x_0, q_0)\}$. There is a transition labeled with $\delta \in \Delta$ from node z to z' if and only if $z' = \{(x', \overline{q}') : \exists (x, \overline{q}) \in z, L_{G^R}((x, \overline{q}), \delta, (x', \overline{q}')) \neq \emptyset\}$. Given G^R and its observer automaton Obs, we construct an embedded Markov chain with state space $\{(z, (x, \overline{q})) : z \in Z, (x, \overline{q}) \in z\} \subseteq Z \times (X \times \overline{Q})$ and transition matrix Ω. Let $i = (z, (x, \overline{q}))$ and $j = (z', (x', \overline{q}'))$, then the ijth element of Ω is given by $\Omega_{ij} := \sum_{\delta \in \Delta : z \text{ transitions to } z' \text{ on } \delta} \alpha(L_{G^R}((x, \overline{q}), \delta, (x', \overline{q}')))$.

Example 4: For the system in Fig. 3, its deterministic observer automaton Obs is constructed, as shown in Fig. 4. Then, the embedded Markov chain has state space of size 7, ordered as follows: $[z_0, (0,0)]$, $[z_1, (1,1)]$, $[z_1, (2,2)]$, $[z_1, (3,F)]$, $[z_2, (2,2)]$, $[z_2, (3,F)]$, and $[z_3, (1,1)]$. Furthermore, the transition matrix Ω is given by

$$\Omega = \begin{bmatrix} 0 & 0.2 & 0.2 & 0.12 & 0.2 & 0.28 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Then, it can be verified that Ω has three recurrent states, namely, $[z_2, (2,2)]$, $[z_2, (3,F)]$, and $[z_3, (1,1)]$. ■

The following sufficient condition for S-Diagnosability is reproduced from [1], which requires the nonexistence of a pair of recurrent states in the same node, such that one recurrent state is faulty and the other nonfaulty.

Theorem 3 [1, Th. 4]: (G, R) is S-Diagnosable if: for every node z of Obs, if there exists $(x, \overline{q}) \in z$, such that $\overline{q} = F$ and $[z, (x, \overline{q})]$ is a recurrent state of Ω, then for all $(x', \overline{q}') \in z$, such that $[z, (x', \overline{q}')]$ is recurrent, $\overline{q}' = F$.

In the rest of this section, we extend Theorem 3 [1, Th. 4] to obtain a condition for S-Diagnosability that is both necessary and sufficient. Note when the condition of Theorem 3 is not satisfied, there exist $z \in Z$, $(x, \overline{q}) \in z$ with $\overline{q} = F$, and $(x', \overline{q}') \in z$ with $\overline{q}' \neq F$, such that both $[z, (x, \overline{q})]$ and $[z, (x', \overline{q}')]$ are recurrent in Ω. We then map the recurrent states back to G^R, and require the two SCCs of G^R that contain $(x, \overline{q})$ and $(x', \overline{q}')$, respectively, satisfy a certain property called p-equivalence, as defined as follows.

Definition 3: Two closed SCCs in G^R are said to be p-equivalent with respect to a given pair of initial distributions of the two SCCs, if each $o \in \Delta^*$ occurs with the same probability in the two SCCs (starting from their corresponding distributions).

Example 5: For the system shown in Fig. 3, it can be verified that starting from $\pi^1 = [1]$ and $\pi^2 = [1]$, C_1 and C_2 (see Example 3) have different probabilities for observation a. Thus, C_1 and C_2 are not p-equivalent with respect to $\pi^1 = [1]$ and $\pi^2 = [1]$, respectively. ■

Remark 2: As provided in [8], given two SCCs with n_1 and n_2 states and initial distributions π_1 and π_2, respectively, one can construct a basis for the $(n_1 + n_2)$-dimensional vector space that includes the $(n_1 + n_2)$-dimensional vector $[\pi_1 \; \pi_2]$. Then, the two SCCs are p-equivalent with respect to π_1 and π_2, if and only if, starting from the initial distributions π_1 and π_2, for every observation o, leading to a distribution for the two SCCs that is another basis in the $(n_1 + n_2)$-dimensional space; the generation probability of o for the two SCCs starting from π_1 to π_2, respectively, is identical. Since there exists a finite basis set, the verification of p-equivalence with respect to a given pair of initial distributions is decidable, and in fact, can be done in $O(|X|^4 \times |Q|^4 \times |\Sigma|)$. ■

Semi-Algorithm 1 provides a necessary and sufficient test for S-Diagnosability. Note that Steps 1–3 of Semi-Algorithm 1 checks the condition in Theorem 3, whereas Steps 4 and 5 checks the additional requirement of p-equivalence.

Semi-Algorithm 1: For a given pair of (G, R), do the following to check S-Diagnosability.

1) Construct the deterministic observer automaton Obs for G^R.
2) Construct the embedded Markov chain with its transition probability matrix Ω.
3) Check if there exists a node z of Obs, such that there exists $(x, \overline{q}) \in z$ with $\overline{q} = F$ and $(x', \overline{q}') \in z$ with $\overline{q}' \neq F$, and both $[z, (x, \overline{q})]$ and $[z, (x', \overline{q}')]$ are recurrent in Ω. (G, R) is S-Diagnosable, if such z does not exist. Otherwise proceed to the next step.
4) For each pair of $[z, (x, \overline{q})]$ and $[z, (x', \overline{q}')]$ that satisfy the condition in the previous step, identify the pair of SCCs C_1 and C_2 in G^R that contain $(x, \overline{q})$ and $(x', \overline{q}')$, respectively.
5) For each SCC pair C_1 and C_2 identified earlier, check if there exists a reachable distribution pair π^1 and π^2, such that C_1 and C_2 are p-equivalent with respect to π^1 and π^2, respectively. If yes, (G, R) is not S-Diagnosable; otherwise, (G, R) is S-Diagnosable.

Example 6: For the system in Fig. 3, the node z_2 of Obs has $(2, 2) \in z_2$ and $(3, F) \in z_2$, such that both $[z_2, (2, 2)]$ and $[z_2, (3, F)]$ are recurrent, while $(2, 2)$ is a nonfaulty state and $(3, F)$ is a faulty state. Hence, the condition in Theorem 3 is not satisfied, and S-Diagnosabiltiy for system in Fig. 3 cannot be determined by Theorem 3.

Consider the SCCs C_1 of G^R that consists of $(2, 2)$ with its self-loop transition and C_2 that consists of $(3, F)$ with its self-loop transition. It turns out that those two SCCs are not p-equivalent with respect to the only reachable distribution pair $\pi^1 = [1]$ and $\pi^2 = [1]$ (see Example 5). In other words, there does not exist any reachable distribution pair with respect to which C_1 and C_2 are p-equivalent. It follows from Semi-Algorithm 1 that the system is S-Diagnosable. ■

Remark 3: Steps 4 and 5 of Semi-Algorithm 1 seek the existence of a pair of distributions over a pair of closed SCCs, under whose initializations, the SCCs are p-equivalent. In general, there can be arbitrarily many reachable distributions pair, and so the decidability for checking this condition remains open at this point. However, Semi-Algorithm 1 yields a conclusive answer if it terminates (see Example 6). ■

Theorem 4 guarantees the correctness of Semi-Algorithm 1.

Theorem 4: (G, R) is S-Diagnosable if and only if: 1) condition of Theorem 3 is satisfied or 2) for every $z \in Z$, $(x, \overline{q}) \in z$ with $\overline{q} = F$, and $(x', \overline{q}') \in z$ with $\overline{q}' \neq F$, such that both $[z, (x, \overline{q})]$ and $[z, (x', \overline{q}')]$ are recurrent; the pair of SCCs in G^R that contain $(x, \overline{q})$ and $(x', \overline{q}')$, respectively, is not p-equivalent with respect to any reachable distribution pair.

Proof: Suppose condition of Theorem 3 is not satisfied. Let C_1 be the SCC of G^R containing $(x, \overline{q})$, and C_2 be the SCC of G^R containing $(x', \overline{q}')$. When C_1 and C_2 are not p-equivalent with respect to any reachable distributions pair, then for any $s_f \in L - K$ reaching C_1 and $s_n \in K$ reaching K, any $\rho, \tau > 0$, there exists $n \in \mathbb{N}$, such that $Pr(t_f \in L \backslash s_f : |t_f| \geq n, Pr_{\text{amb}}(s_f t_f) > \rho) < \tau$ [2, Th. 3]. Thus, one can conclude S-Diagnosability. On the other hand, when C_1 and C_2 are p-equivalent with respect to a reachable distribution pair π^1 and π^2, there exists $s_f \in L - K$ reaching C_1 and $s_n \in K$ reaching C_2, such that $M(s_f) = M(s_n) = o$, $\pi^1 = \pi^{C_1}(o)$, and $\pi^2 = \pi^{C_2}(o)$. Since C_1 and C_2 are p-equivalent with respect to π^1 and π^2, further extensions will not decrease the ambiguity level. Let $\rho := Pr_{\text{amb}}(s_f) > 0$. Then, for any $n \in \mathbb{N}$ and $\tau < 1$, $Pr(t_f \in L \backslash s_f : |t_f| \geq n, Pr_{\text{amb}}(s_f t_f) > \rho) = 1 > \tau$. Therefore, (G, R) is not S-Diagnosable. ■

V. Conclusion

In this paper, we corrected the oversights in [2], where in the presence of nondeterminism, when two closed SCCs synchronize, the result may not remain closed. We prove that the verification of SS-Diagnosability is PSPACE-hard, and so it is unlikely that a polynomial algorithm for verifying SS-Diagnosability exists. For testing S-Diagnosability, we provide a corrected necessary and sufficient condition that utilizes the deterministic observer to keep track of pairs of persistence ambiguous SCCs, and the notion of p-equivalence to decide S-Diagnosability. The decidability of checking this condition remains an open problem, while we provide a Semi-Algorithm, which can conclusively test S-Diagnosability, whenever it terminates.

References

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

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

[3] 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.

[4] J.-Y. Kao, N. Rampersad, and J. Shallit, "On NFAs where all states are final, initial, or both," *Theor. Comput. Sci.*, vol. 410, nos. 47–49, pp. 5010–5021, 2009.

[5] H. Gimbert and Y. Oualhadj, "Probabilistic automata on finite words: Decidable and undecidable problems," in *Proc. ICALP*, Bordeaux, France, Jul. 2010, pp. 527–538.

[6] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, Princeton, NJ, USA: Van Nostrand, 1960.

[7] J. Chen and R. Kumar, "Failure detection framework for stochastic discrete event systems with guaranteed error bounds," *IEEE Trans. Autom. Control*, vol. 60, no. 6, pp. 1542–1553, Jun. 2015.

[8] W.-G. Tzeng, "A polynomial-time algorithm for the equivalence of probabilistic automata," *SIAM J. Comput.*, vol. 21, no. 2, pp. 216–227, Apr. 1992.