

Reachability analysis in stochastic directed graphs by reinforcement learning

Original

Reachability analysis in stochastic directed graphs by reinforcement learning / Possieri, C.; Frasca, M.; Rizzo, A.. - In: IEEE TRANSACTIONS ON AUTOMATIC CONTROL. - ISSN 1558-2523. - ELETTRONICO. - 68:1(2023), pp. 462-469. [10.1109/TAC.2022.3143080]

Availability:

This version is available at: 11583/2957678 since: 2023-01-08T18:50:20Z

Publisher:

IEEE

Published

DOI:10.1109/TAC.2022.3143080

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Reachability analysis in stochastic directed graphs by reinforcement learning

Corrado Possieri, Mattia Frasca, and Alessandro Rizzo

Abstract—We characterize the reachability probabilities in stochastic directed graphs by means of reinforcement learning methods. In particular, we show that the dynamics of the transition probabilities in a stochastic digraph can be modeled via a difference inclusion, which, in turn, can be interpreted as a Markov decision process. Using the latter framework, we offer a methodology to design reward functions to provide upper and lower bounds on the reachability probabilities of a set of nodes for stochastic digraphs. The effectiveness of the proposed technique is demonstrated by application to the diffusion of epidemic diseases over time-varying contact networks generated by the proximity patterns of mobile agents.

Index Terms—Stochastic digraphs, reinforcement learning, reachability analysis.

I. INTRODUCTION

Earlier studies on networks widely relied on static graphs as a main analysis tool, that is, graphs in which the edge set does not vary in time. This modeling paradigm successfully contributed to the analysis of networked systems, even of large dimension, such as social [1] and biological [2] networks, social psychology systems [3], electrical circuits, chemical reactions [4], and industrial supply chains [5]. However, static graphs fail to capture the inherent time-varying and non-deterministic nature of the relationship between networked dynamical units. Stochastic directed graphs (digraphs) have been introduced in [6] to account for time-varying networked processes in which, at each time instant, the graph topology is constructed by selecting an edge set from a finite collection according to a given probability distribution. Typical systems that have been successfully modeled through this paradigm are chemical reactions [4], sensor networks [6], and biological systems having non-unique behaviors [7].

For such a class of dynamical processes, it is often of interest to quantify the reachability probabilities of a given set,

C. Possieri is with Istituto di Analisi dei Sistemi ed Informatica “A. Ruberti”, Consiglio Nazionale delle Ricerche (IASI-CNR), 00185 Roma, Italy (e-mail: corrado.possieri@iasi.cnr.it). M. Frasca is with Dipartimento di Ingegneria Elettrica Elettronica e Informatica, University of Catania, 95131 Catania, Italy, and also with Istituto di Analisi dei Sistemi ed Informatica “A. Ruberti”, Consiglio Nazionale delle Ricerche (IASI-CNR), 00185 Roma, Italy (e-mail: mattia.frasca@dieei.unict.it). A. Rizzo is with Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, 10129 Torino, Italy (e-mail: alessandro.rizzo@polito.it).

A.R. acknowledges financial support by Compagnia di San Paolo. M.F. acknowledges partial financial support by the Italian Ministry for Research and Education through Research Program PRIN 2017 under Grant 2017CWMP93, project VECTORS. C.P. acknowledges partial financial support by Regione Lazio through Research Program POR FESR LAZIO 2014-2020 - GRUPPI DI RICERCA 2020 under Grant GeCoWEB A0375-2020-36616, project OPENNESS.

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

which are related to the possibility of the system to assume one of the possible configurations at a certain time. This property is important to characterize the dynamical behavior of several classes of systems, such as sensor networks and power grids, as exemplified in [6], [8].

Computational methods based on the optimization of auxiliary functions have been proposed to characterize reachability probabilities in stochastic digraphs [4], [6], [7]; however, these methods may entail a high computational burden, which can make the problem intractable when the system size increases. In this technical note, we tackle this issue by computing reachability probabilities of stochastic digraphs through an approach based on dynamic programming and reinforcement learning. Transition probabilities between configurations are modeled using difference inclusions; then, the obtained model is paralleled to a suitable Markov decision process. A reinforcement learning strategy based on reward functions constructed on the basis of the Markov decision process is then exploited to establish upper and lower bounds on the reachability properties of the stochastic digraph. Our approach is validated with an application to the diffusion of epidemic diseases on mobile populations.

II. PRELIMINARIES

Let $\mathbb{Z}_n := \{1, \dots, n\}$ and let $\wp(\mathbb{Z}_n)$ denote the *power set* of \mathbb{Z}_n , i.e., the set of all the subsets of \mathbb{Z}_n . Let \mathbf{e}_i denote the i -th row of the identity matrix. The symbol $\mathbb{1}$ denotes the vector $[1 \ \dots \ 1]^\top$. The symbols \wedge , \vee , and \neg represent the logical AND, OR, and NOT operators, respectively. The symbol $\mathbb{E}[\cdot]$ denotes the expected value and the symbol \leq denotes the entry-wise \leq operator. Given a matrix M , $[M]_{i,j}$ indicates its (i, j) th entry. Given a set $\mathcal{I} \subset \mathbb{Z}_n$, the *indicator function on \mathcal{I}* is $\mathbb{1}_{\mathcal{I}}(\zeta) = 1$ if $\zeta \in \mathcal{I}$, or $\mathbb{1}_{\mathcal{I}}(\zeta) = 0$ if $\zeta \notin \mathcal{I}$.

A directed graph (briefly, a *digraph*) is a pair $\mathcal{G} := (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathbb{Z}_n$ is the *vertex set* and $\mathcal{E} \subset \mathbb{Z}_n \times \mathbb{Z}_n$ is the *edge set*. Given $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and $x \in \mathcal{V}$, the *in-neighborhood of x* and the *out-neighborhood of x* are, respectively,

$$\begin{aligned} \mathcal{N}^i(x) &:= \{y \in \mathbb{Z}_n : (y, x) \in \mathcal{E}\}, \\ \mathcal{N}^o(x) &:= \{y \in \mathbb{Z}_n : (x, y) \in \mathcal{E}\}. \end{aligned}$$

The *out-degree (in-degree)* of x is the cardinality of $\mathcal{N}^o(x)$ ($\mathcal{N}^i(x)$). A *regular digraph* is a digraph in which all vertices have the same number of out-neighbors, i.e., all the vertices have the same out-degree. A regular digraph whose vertices have out-degree k is called a k -regular digraph. Given two digraphs $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{V}, \mathcal{E}_2)$, their *union* is $\mathcal{G}_1 \cup \mathcal{G}_2 = (\mathcal{V}, \mathcal{E}_1 \cup \mathcal{E}_2)$. A *time-varying graph* is a graph whose vertex set is constant and the edge set changes in time.

A *stochastic digraph* is a time-varying digraph obtained by randomly selecting, at each time, an element from a finite collection of digraphs. Formally, a stochastic digraph is defined by a triple $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mu(\cdot))$, where:

- $\mathcal{V} := \mathbb{Z}_n$;
- $\mathcal{E} := \{\mathcal{E}_s\}_{s=1}^h$, $h \in \mathbb{Z}_{\geq 0}$, with $h < \infty$, such that $\mathcal{G}_s := (\mathcal{V}, \mathcal{E}_s)$ is a digraph for each $s \in \{1, \dots, h\}$;
- $\mu(\cdot)$ is a distribution function derived from an infinite sequence of independent, identically distributed (i.i.d.) random variables $\mathbf{w}_k : \Omega \rightarrow \{1, \dots, h\}$, $k \in \mathbb{Z}_{\geq 0}$, defined on a given probability space $(\Omega, \mathcal{F}, \mathbb{P})$ [9].

Note that the number h of admissible topologies may be much smaller than (although possibly equal to) $2^{\binom{n}{2}}$, which is the number of digraphs possible with n vertices.

The probability $\mathbb{P}(\mathbf{w}_k \in F) := \mathbb{P}(\{\omega \in \Omega : \mathbf{w}_k(\omega) \in F\})$ is well defined and independent of $k \in \mathbb{Z}_{\geq 0}$, for each $F \in \wp(\mathbb{Z}_h)$. Hence, the distribution function $\mu : \wp(\mathbb{Z}_h) \rightarrow [0, 1]$ is defined as [9, Sec. 2.1 and Sec. 11.1]

$$\mu(F) := \mathbb{P}(\mathbf{w}_k \in F), \quad (1)$$

for each $F \in \wp(\mathbb{Z}_h)$. Note that, in this setting, there is no loss of generality in considering a finite h , since the number of edge sets \mathcal{E} such that $(\mathcal{V}, \mathcal{E})$ is a digraph is finite [7].

Let us define the set-valued map $H : \mathbb{Z}_n \times \mathbb{Z}_h \rightrightarrows \mathbb{Z}_n$ of the out-neighborhoods in the instantaneous digraphs as

$$H(x, w) := \{y \in \mathbb{Z}_n : (x, y) \in \mathcal{E}_w\}. \quad (2)$$

A sequence $(\phi, \mathbf{w}) := \{(\phi_k, \mathbf{w}_k)\}_{k=0}^K$ is a *regular directed path* of a stochastic digraph \mathcal{G} starting at x if

- $\phi_0 = x$;
- $\phi_{k+1} \in \mathcal{N}_{\mathbf{w}_k}^o(\phi_k)$, for all $k \in \{0, \dots, K-1\}$.

A map \mathbf{x} from Ω to sequences in \mathbb{Z}_n , that is $\omega \mapsto \{\mathbf{x}_k(\omega)\}_{k=0}^{K(\omega)}$, where $K : \Omega \rightarrow \mathbb{Z}_{\geq 0} \cup \{\infty\}$ is a random variable denoting the length of the path, is a *stochastic directed path* from $x \in \mathbb{Z}_n$ if the following properties hold:

Path-wise feasibility: for each $\omega \in \Omega$, the sequence $\{(\mathbf{x}_k(\omega), \mathbf{w}_k(\omega))\}_{k=0}^{K(\omega)}$ is a regular directed path from x .
Causal measurability: for each $k \in \mathbb{Z}_{\geq 0}$, the mapping $\omega \mapsto \mathbf{x}_{k+1}(\omega)$ is \mathcal{F}_k -measurable, where $(\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots)$ is the minimal filtration of \mathbf{w} [9].

According to [6], the set-valued mapping $H(x, w)$ given in (2) satisfies the standing assumption of [10]–[12], and hence stochastic directed paths exist and are well defined. A stochastic directed path is *maximal* if it cannot be extended and is *complete* if it is maximal and $K(\omega) = +\infty$.

Let $\mathcal{S}(x)$ denote the set of all the maximal stochastic directed paths from x , and, given a single stochastic directed path $\mathbf{x} \in \mathcal{S}(x)$, let $\text{gph}(\mathbf{x}) = (0, \mathbf{x}_0) \times (1, \mathbf{x}_1) \times \dots \times (K, \mathbf{x}_K)$, where $K \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ denotes the length of the sequence \mathbf{x} .

Note that, although each stochastic directed path from x is a *deterministic* function from Ω to sequences in \mathbb{Z}_n , the set $\mathcal{S}(x)$ is generically not a singleton, thus making the problem of determining reachability probabilities particularly challenging.

III. COMPUTATION OF THE TRANSITION PROBABILITIES IN STOCHASTIC DIGRAPHS

In this section, we derive a series of results that, given a stochastic digraph \mathcal{G} , ultimately allow to determine the transition probabilities among its nodes, i.e., the probability of visiting other nodes at the next step given the current node.

A. Bounds for the transitions probabilities

Let us consider a stochastic digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mu(\cdot))$ and let $H : \mathbb{Z}_n \times \mathbb{Z}_h \rightrightarrows \mathbb{Z}_n$ be the map of the out-neighborhoods in the instantaneous digraphs as defined in (2). The following standing assumption is made throughout this technical note.

Assumption 1. There does not exist $(i, w) \in \mathbb{Z}_n \times \mathbb{Z}_h$ with $\mu(w) > 0$ such that $H(i, w) = \emptyset$.

Remark 1. Assumption 1 essentially requires that no node in the stochastic digraph \mathcal{G} is a sink with positive probability. If this holds, then maximal stochastic directed paths are complete. On the other hand, if this does not hold, then it is possible to define an augmented stochastic digraph \mathcal{G}' that holds the properties of the original digraph while satisfying Assumption 1. Namely, it suffices to add a further node $n+1$ and to let $H(n+1, w) = \{n+1\}$, for all $w \in \mathbb{Z}_h$, and $H(i, w) = \{n+1\}$, for all $(i, w) \in \mathbb{Z}_n \times \mathbb{Z}_h$ such that $\mu(w) > 0$ and $H(i, w) = \emptyset$. Stochastic paths of \mathcal{G}' containing the node $n+1$ represent maximal paths that are not complete.

The next theorem proves that, if a stochastic digraph satisfies Assumption 1, then it is possible to compute upper and lower bounds on the transition probabilities among its nodes.

Theorem 1. Let \mathcal{G} be a stochastic digraph satisfying Assumption 1 and let \mathbf{x} be a stochastic directed path of \mathcal{G} , i.e., $\mathbf{x} \in \mathcal{S}(x)$ for some $x \in \mathbb{Z}_n$. For each $i, j \in \mathbb{Z}_n$, the scalars

$$\ell_{i,j} := \sum_{s \in \{w \in \mathbb{Z}_h : H(i, w) = \{j\}\}} \mu(s) \in [0, 1], \quad (3a)$$

$$m_{i,j} := \sum_{s \in \{w \in \mathbb{Z}_h : H(i, w) \cap \{j\} \neq \emptyset\}} \mu(s) \in [0, 1], \quad (3b)$$

are upper and lower bounds on the transition probability $\mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i)$ from node i to node j , respectively, i.e.,

$$\ell_{i,j} \leq \mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i) \leq m_{i,j}. \quad (4)$$

Proof. Let us start by considering the special case of a stochastic digraph \mathcal{G} for which the digraphs $\mathcal{G}_s = (\mathcal{V}, \mathcal{E}_s)$, $s = 1, \dots, h$, are 1-regular. Then, for each $x \in \mathbb{Z}_n$ and each $w \in \mathbb{Z}_h$, $H(x, w)$ is a singleton, i.e., $H(x, w) = \{y\}$ for some $y \in \mathbb{Z}_n$. It follows that the transition probabilities in \mathcal{G} can be represented by a Markov chain and

$$\mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i) = \sum_{s \in \{w \in \mathbb{Z}_h : H(i, w) = \{j\}\}} \mu(s). \quad (5)$$

Thus, the bounds in (4) hold since $H(i, w) \cap \{j\} \neq \emptyset$ if and only if $H(i, w) = \{j\}$ under Assumption 1.

Let us now move to the more general case, where $\mathcal{G}_s = (\mathcal{V}, \mathcal{E}_s)$, $s = 1, \dots, h$, may not be 1-regular. In this case, $H(x, w)$ is a set-valued map. Let $\mathcal{I}_j = \{j\}$ and consider the

functions $m_j : \mathbb{Z}_{\geq 0} \times \mathbb{Z}_n \rightarrow [0, 1]$ and $\ell_j : \mathbb{Z}_{\geq 0} \times \mathbb{Z}_n \rightarrow [0, 1]$, defined as $m_j(0, \xi) := 0$ and $\ell_j(0, \xi) := 0$, for all $\xi \in \mathbb{Z}_n$, and, for all $(k, \xi) \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_n$,

$$m_j(k+1, \xi) := \sum_{w=0}^{\bar{h}-1} \max_{g \in H(\xi, w)} \{\mathbb{1}_{\mathcal{I}_j}(g), m_j(k, g)\} \mu(w),$$

$$\ell_j(k+1, \xi) := \sum_{w=0}^{\bar{h}-1} \min_{g \in H(\xi, w)} \{\mathbb{1}_{\mathcal{I}_j}(g), \mathbb{1}_{\mathcal{I}_j^c}(g) \ell_j(k, g)\} \mu(w),$$

where $\mathcal{I}_j^c := \mathbb{Z}_n \setminus \mathcal{I}_j$. Since the set \mathcal{I}_j is both closed and open [7], by [10] one has that, for all $\kappa \in \mathbb{Z}_{\geq 0}$, $\kappa \geq 1$, $m_j(\kappa, \xi) = \sup_{\xi \in \mathcal{S}(\xi)} \mathbb{P}(\exists k \in \{1, \dots, \kappa\} : \xi_k = j)$, and there exists $\bar{\xi} \in \mathcal{S}(\xi)$ such that the sup is attained. In addition, by [13], one has that $\ell_j(\kappa, \xi) = \inf_{\xi \in \mathcal{S}(\xi)} \mathbb{P}(\exists k \in \{1, \dots, \kappa\} : \xi_k = j)$, and there exists $\underline{\xi} \in \mathcal{S}(\xi)$ such that the inf is attained. This proves that (4) holds with $m_{i,j}$ and $\ell_{i,j}$ given in (3), due to the fact that the random variables w_k are i.i.d., and that $\max_{g \in H(\xi, w)} \{\mathbb{1}_{\mathcal{I}_j}(g), 0\}$ equals 1 if $H(\xi, w) \cap \{j\} \neq \emptyset$ and 0 otherwise, whereas $\min_{g \in H(\xi, w)} \{\mathbb{1}_{\mathcal{I}_j}(g), 0\}$ equals 1 if $H(\xi, w) = \{j\}$ and 0 otherwise; see [13, Prop. 3]. \square

In the special case of a stochastic digraph \mathcal{G} with 1-regular digraphs $\mathcal{G}_s = (\mathcal{V}, \mathcal{E}_s)$, $s = 1, \dots, h$, we can introduce the matrix P with entries $p_{i,j} = \mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i)$. In all the other cases, rather than P , we consider the matrices M and L , with entries given by the terms $m_{i,j}$ and $\ell_{i,j}$ in (3). Note that the matrix P is *right stochastic* [14], i.e., its rows sum to 1, whereas the matrices M and L may not satisfy this property.

Example 1. Consider the stochastic digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mu(\cdot))$, where $\mathcal{V} := \mathbb{Z}_4$, $\mathcal{E} := \{\mathcal{E}_s\}_{s=1}^2$, $\mu(\{1\}) = 2/3$, $\mu(\{2\}) = 1/3$, with $\mathcal{G}_1 = (\mathcal{V}, \mathcal{E}_1)$ and $\mathcal{G}_2 = (\mathcal{V}, \mathcal{E}_2)$ as in Figure 1.




Fig. 1. An example of stochastic digraph with $h = 2$.

First, we calculate the set-valued maps $H(x, w)$ of the out-neighborhoods, defined as in (2), for each $x \in \mathbb{Z}_n$ and $w \in \mathbb{Z}_h$:

$$\begin{aligned} H(1, 1) &= \{2, 3\}, & H(1, 2) &= \{3, 4\}, & H(2, 1) &= \{1\}, \\ H(2, 2) &= \{4\}, & H(3, 1) &= \{1, 4\}, & H(3, 2) &= \{4\}, \\ H(4, 1) &= \{2, 3\}, & H(4, 2) &= \{3\}. \end{aligned}$$

Then, expressions (3) are used to compute matrices L and M for the stochastic digraph \mathcal{G} , obtaining

$$L = \begin{bmatrix} 0 & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{3} \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \quad M = \begin{bmatrix} 0 & \frac{2}{3} & 1 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & \frac{1}{3} \\ 0 & 0 & 0 & 1 \\ 0 & \frac{2}{3} & 1 & 0 \end{bmatrix}.$$

B. Dynamics of the probability vector

Let us now consider the probability vector

$$\rho_k = [\mathbb{P}(\mathbf{x}_k = 1) \quad \cdots \quad \mathbb{P}(\mathbf{x}_k = n)],$$

whose i -th entry represents the probability of visiting the i -th node at time k . By Theorem 1, one has $\rho_k L \leq \rho_{k+1} \leq \rho_k M$, although the entries of $\rho_k L$ and $\rho_k M$ need not sum to 1.

Theorem 2. *Let \mathcal{G} be a stochastic digraph satisfying Assumption 1. Then, there exist right stochastic matrices P_1, \dots, P_ν such that the dynamics of ρ_k are given by*

$$P_k \in \{P_1, \dots, P_\nu\} =: \mathcal{P}, \quad (6a)$$

$$\rho_{k+1} = \rho_k P_k, \quad (6b)$$

with domain $\mathcal{D} := \{(\rho, P) \in [0, 1]^{1,n} \times \mathcal{P} : \rho \mathbf{1} = 1\}$.

Proof. Consider the digraphs $\mathcal{G}_s = (\mathcal{V}, \mathcal{E}_s)$, $s = 1, \dots, h$. Each of the digraphs \mathcal{G}_s can be decomposed into the union of a finite number of 1-regular digraphs, $\mathcal{G}_s = \bigcup_{i=1}^{T_s} \bar{\mathcal{G}}_{s,i}$, $s = 1, \dots, h$. Hence, letting $\bar{\mathcal{G}}_{s,i} = (\mathcal{V}, \mathcal{E}_{s,i})$ and defining $\nu = \prod_{s=1}^h T_s$, consider all the stochastic digraphs $\mathcal{G}_a = (\mathcal{V}, \mathcal{E}_a, \mu(\cdot))$, with $a = 1, \dots, \nu$, obtained letting $\mathcal{E}_a = \{(\mathcal{E}_a)_s\}_{s=1}^h$ and $(\mathcal{E}_a)_s = \mathcal{E}_{s,i}$ for some $i \in \{1, \dots, T_s\}$, $s = 1, \dots, h$. Since each \mathcal{G}_a satisfies the hypothesis that $H_a(x, w) = \{y\}$ for some $y \in \mathbb{Z}_n$, for each $(x, w) \in \mathbb{Z}_n \times \mathbb{Z}_h$, then the transition probabilities in \mathcal{G}_a can be represented through a Markov chain, whose transition matrix P_a is given by (5). Hence, letting P_1, \dots, P_ν be these matrices, the thesis follows by the definition of stochastic directed paths. \square

From Theorem 2, we derive a procedure to determine \mathcal{P} , formalized in Algorithm 1.

Algorithm 1 Computation of the set \mathcal{P}

Input: stochastic digraph \mathcal{G}

Output: set \mathcal{P} such that Eq. (6) holds

- 1: **for** $s = 1$ **to** h **do**
 - 2: determine 1-regular digraphs $\bar{\mathcal{G}}_{s,1}, \dots, \bar{\mathcal{G}}_{s,T_s}$ such that $\mathcal{G}_s := (\mathcal{V}, \mathcal{E}_s) = \bigcup_{i=1}^{T_s} \bar{\mathcal{G}}_{s,i}$ and let $\bar{\mathcal{G}}_{s,i} = (\mathcal{V}, \mathcal{E}_{s,i})$
 - 3: compute $\nu = \prod_{s=1}^h T_s$
 - 4: generate the list of all possible tuples $\mathcal{E}_1, \dots, \mathcal{E}_\nu$, $\mathcal{E}_a = \{(\mathcal{E}_a)_s\}_{s=1}^h$, with $a = 1, \dots, \nu$, whose s -th element is from $\{\mathcal{E}_{s,1}, \dots, \mathcal{E}_{s,T_s}\}$
 - 5: **for** $a = 1$ **to** ν **do**
 - 6: define the stochastic digraph $\mathcal{G}_a := (\mathcal{V}, \mathcal{E}_a, \mu(\cdot))$
 - 7: use (5) to compute the corresponding matrix P_a
 - 8: **return** $\mathcal{P} := \{P_1, \dots, P_\nu\}$
-

Note that the selection of matrices P_1, P_2, \dots, P_ν in (6a) corresponds to the choice of a random stochastic path in $\mathcal{S}(x)$. Namely, following the construction made in the proof of Theorem 2 and the definition of stochastic directed paths given in Section II, at each time instant one has to choose a path-wise feasible rule to select the node \mathbf{x}_{k+1} given \mathbf{x}_k . The selection of the matrix P_k in (6a) corresponds exactly to the choice of such a feasible rule.

The following corollary, which provides bounds on the matrices in \mathcal{P} , is a direct consequence of Theorems 1 and 2.

Corollary 1. Suppose that the assumptions of Theorem 2 hold and consider the matrices M and L with entries as in (3). Then, letting \mathcal{P} be defined as in Theorem 2, one has

$$L \preceq P \preceq M, \quad \forall P \in \mathcal{P}.$$

Example 2. Consider the stochastic digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mu(\cdot))$ given in Example 1. Following Algorithm 1, the digraphs \mathcal{G}_1 and \mathcal{G}_2 given in Figure 1 can be decomposed into the union of eight and two 1-regular digraphs, respectively. Thus, we have $T_1 = 8$, $T_2 = 2$, and $\nu = 16$. Considering the stochastic digraphs $\mathcal{G}_a = (\mathcal{V}, \mathcal{E}_a, \mu(\cdot))$, $a = 1, \dots, 16$, obtained letting $\mathcal{E}_a = \{\mathcal{E}_{1,i}, \mathcal{E}_{2,j}\}$, $i = 1, \dots, 8$, $j = 1, 2$, using (5), one has

$$\mathcal{P} = \left\{ \begin{array}{l} \begin{bmatrix} 0 & 0 & \frac{1}{3} & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & \frac{1}{3} & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \\ \begin{bmatrix} 0 & 0 & \frac{1}{3} & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \\ \begin{bmatrix} 0 & 0 & 1 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \\ \begin{bmatrix} 0 & 0 & 1 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & \frac{1}{3} \\ \frac{2}{3} & 0 & 0 & 0 \\ \frac{2}{3} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \end{bmatrix} \end{array} \right\}.$$

Note that all the matrices in \mathcal{P} are right stochastic and that their entry-wise minimum and maximum equal the corresponding entry of L and M , respectively.

C. Stochastic digraphs as Markov decision processes

So far, we have shown that the properties of stochastic digraphs can be analyzed by studying the corresponding difference inclusion (6). Note that such an inclusion can be interpreted as a Markov decision process [15] with action space $\mathcal{A} = \{1, \dots, \nu\}$ and with

$$p(j|i, a) := \mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i, \mathbf{a}_k = a) = [P_a]_{i,j}. \quad (7)$$

In this context, the action $a \in \mathcal{A}$ represents the selection of the rule employed to select the node \mathbf{x}_{k+1} given \mathbf{x}_k . However, computing the transition matrices of a large stochastic digraph via Algorithm 1 can be computationally demanding as, at step 4 of this algorithm, one has to generate the set of all the tuples of the possible edge sets.

Remark 2. The transition probabilities $\mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i, \mathbf{a}_k = a)$ can be computed more efficiently letting the action space depend on the current value of the state. Given a node i , consider the sets $H(i, w)$ with $w = 1, \dots, s$, representing all the possible sets of out-neighbors that can be reached from i . Then, the set of all the path-wise feasible rules to update node i , labeled using the index a_i , consists of a node from $H(i, 1)$, one from $H(i, 2)$ and so on, up to $H(i, s)$. The number of path-wise feasible rules that can be used to update node i is then determined by all the possible tuples that can be generated in this way and is, therefore, given by $\nu_i = \prod_{w=1}^s |H(i, w)|$. This corresponds to the cardinality of the local action space from i , i.e., $\mathcal{A}(i) := \{1, \dots, \nu_i\}$. Now, let us indicate the elements of

$H(i, w)$ as $\{j_{w,1}, \dots, j_{w,R_w}\}$ where $R_w = |H(i, w)|$, and for each $a_i \in \mathcal{A}(i)$ consider the set $(\mathcal{J}_i)_{a_i} := \{\xi_{1,a_i}, \dots, \xi_{h,a_i}\}$, representing the a_i -th tuple associated to the path-wise feasible update rule a_i , then

$$\mathcal{J}_i := \{\{j_{1,1}, \dots, j_{h,1}\}, \dots, \{j_{1,R_1}, \dots, j_{h,R_h}\}\}$$

represents the set of all tuples whose w -th element is selected from $H(i, w)$. Thus, the transition probabilities $p(j|i, a) = \mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i, \mathbf{a}_k = a)$ for $a \in \mathcal{A}(i)$ are given by

$$p(j|i, a) = \sum_{w \in \{s \in \mathbb{Z}_h : \xi_{s,a} = j\}} \mu(w). \quad (8)$$

Note that there are two clear advantages of defining $\mathcal{A}(i)$ and $p(j|i, a)$ as in (8) rather than as in (7): firstly, the set \mathcal{J}_i can be generated by just considering the value attained by the set-valued map $H(x, w)$ for $x = i$ and $w \in \mathbb{Z}_h$ and hence the probabilities $p(j|i, a)$ can be determined online without explicitly computing the matrices P_1, \dots, P_ν . Secondly, the action space $\mathcal{A}(i)$ is smaller than that of the set \mathcal{A} . Namely, its cardinality is $\nu_i = \prod_{w=1}^s |H(i, w)|$, which is in general much smaller than ν .

The following algorithm summarizes the steps of the procedure given in Remark 2.

Algorithm 2 Computation of the function $p(j|i, a)$

Input: the map of the out-neighborhoods $H(i, w)$ and the probability distribution function $\mu(\cdot)$

Output: the transition probability $p(j|i, a) = \mathbb{P}(\mathbf{x}_{k+1} = j | \mathbf{x}_k = i, \mathbf{a}_k = a)$ and the action space $\mathcal{A}(i)$

- 1: generate the list \mathcal{J}_i of all tuples whose w -th element is from $\{j_{w,1}, \dots, j_{w,R_w}\} := H(i, w)$
 - 2: calculate ν_i as $\nu_i = \prod_{w=1}^s |H(i, w)|$
 - 3: letting $(\mathcal{J}_i)_{a_i} := (\xi_{1,a_i}, \dots, \xi_{h,a_i})$ be the a_i -th tuple in \mathcal{J}_i , compute $p(j|i, a)$ as in (8)
 - 4: **return** $p(j|i, a)$ and $\mathcal{A}(i) := \{1, \dots, \nu_i\}$
-

Example 3. Consider the stochastic digraph \mathcal{G} given in Examples 1 and 2. The set-valued map H is that reported in Example 2. Using Algorithm 2, we determine the action spaces $\mathcal{A}(1) = \{1, 2, 3, 4\}$, $\mathcal{A}(2) = \{1\}$, $\mathcal{A}(3) = \{1, 2\}$, $\mathcal{A}(4) = \{1, 2\}$, and the sets $\mathcal{J}_1, \dots, \mathcal{J}_4$, given by

$$\mathcal{J}_1 = \{\{2, 3\}, \{2, 4\}, \{3, 3\}, \{3, 4\}\}, \quad \mathcal{J}_2 = \{\{1, 4\}\}, \\ \mathcal{J}_3 = \{\{1, 4\}, \{4, 4\}\}, \quad \mathcal{J}_4 = \{\{2, 3\}, \{3, 4\}\}.$$

Note that the action spaces $\mathcal{A}(i)$, $i = 1, \dots, 4$, have maximum cardinality 4, whereas \mathcal{P} has cardinality 16. Then, by using (8), one obtains the transition probabilities given in Table I.

TABLE I
TRANSITION PROBABILITIES IN \mathcal{G} .

i	a	j	$p(j i, a)$	i	a	j	$p(j i, a)$	i	a	j	$p(j i, a)$
1	1	2	$\frac{2}{3}$	1	4	3	$\frac{2}{3}$	3	1	4	$\frac{1}{3}$
1	1	3	$\frac{2}{3}$	1	4	4	$\frac{2}{3}$	3	2	4	1
1	2	2	$\frac{2}{3}$	2	1	1	$\frac{2}{3}$	4	1	2	$\frac{2}{3}$
1	2	4	$\frac{2}{3}$	2	1	4	$\frac{2}{3}$	4	1	3	$\frac{1}{3}$
1	3	3	1	3	1	1	$\frac{2}{3}$	4	2	3	1

Remark 3. The hypothesis of causal measurability is fundamental to establish the equivalence between stochastic directed paths and Markov decision processes. In fact, if such an assumption about stochastic paths is removed, i.e., $\mathbf{x}_{k+1}(\omega)$ is allowed to depend on $\mathbf{w}_{k+1}(\omega)$, then degenerate behavior may occur as shown in [7, Ex. 5]. In particular, it can be argued that removing this requirement corresponds to a non-causal selection of the action in the Markov decision process.

IV. ANALYSIS OF THE WEAK AND STRONG REACHABILITY IN STOCHASTIC DIGRAPHS

When dealing with stochastic systems with non-unique solutions, such as stochastic digraphs, one of the main challenges is to compute upper and lower bounds on the probability that a stochastic solution visits a given set. In the literature, such a task is usually carried out by means of auxiliary Lyapunov functions whose expected value is decreasing along the solutions of the system (see, e.g., [10], [13], [16] for systems with discrete state space and [6], [7] for those with continuous state space). Leveraging the correspondence established in Section III-C, the main goal of this section is to use reinforcement learning to compute upper and lower bounds on the probability that a path in the stochastic digraph visits a given set of nodes [15]. Namely, considering the Markov decision process (8), where the state correspond to a node of the stochastic digraph and the action correspond to a path-wise feasible rule to extend stochastic paths, the main objective of this section is to define the rewards

$$r(i, a, j) := \mathbb{E}(\mathbf{r}_{k+1} | \mathbf{x}_k = i, \mathbf{x}_{k+1} = j, \mathbf{a}_k = a) \quad (9)$$

to determine path-wise feasible rules that either maximize or minimize the probability of visiting a given set of nodes. In this way, the reachability properties of the stochastic digraph \mathcal{G} may be characterized. In fact, since the action $a \in \mathcal{A}(i)$ corresponds to the selection of an update rule for the stochastic path, selecting the actions that maximize a suitably defined reward $r(i, a, j)$ yields a path in $\mathcal{S}(x)$ that maximizes a given probability function. In particular, in Section IV-A, we show how to define the function $r(i, a, j)$ to obtain the weak reachability properties of the stochastic digraph \mathcal{G} , and in Section IV-B we discuss how to define $r(i, a, j)$ to determine the strong recurrence probabilities of the stochastic digraph \mathcal{G} .

First, we need to introduce some additional concepts borrowed from the dynamic programming framework [15]. Given the Markov decision process (8) and the rewards (9), let $\pi(a|x)$ be a policy that maps each state $x \in \mathbb{Z}_n$ to the probability of taking action $a \in \mathcal{A}(x)$ when the system is in the state x . We define the *state-value function* for π

$$v_\pi(k, x) := \mathbb{E}_\pi \left[\sum_{\tau=0}^{k-1} \mathbf{r}_{\tau+1} \middle| \mathbf{x}_0 = x \right],$$

which is the expected reward gathered by using π , given that the sequence $\mathbf{x} = \{\mathbf{x}_\tau\}_{\tau=0}^k$ is initialized at $\mathbf{x}_0 = x$. Letting

$$p(j, r|i, a) = \mathbb{P}(\mathbf{x}_{k+1} = j, \mathbf{r}_{k+1} = r | \mathbf{x}_k = i, \mathbf{a}_k = a)$$

be the probability that the next state and reward are j and r , respectively, given that the current state and action are i and a , respectively, by the Bellman's principle of optimality [15] we have that the function $v_\pi(k, x)$ satisfies

$$v_\pi(k, x) = \sum_{a,j,r} \pi(a|x) p(j, r|i, a) (r + v_\pi(k-1, j)). \quad (10)$$

Hence, observing that $v_\pi(1, x) = \sum_{a,j,r} \pi(a|x) p(j, r|i, a) r$, the Bellman equation (10) can be used to compute the values attained by the state-value function for π .

Let k be fixed. A policy π is defined to be better than or equal to another policy π' if $v_\pi(k, x) \geq v_{\pi'}(k, x)$ for all $x \in \mathbb{Z}_n$. By [15], an optimal policy π_* , i.e., a policy that is better than all the others, always exists. Although the optimal policy does not need to be unique, all the optimal policies share the same state-value function, called the *optimal state-value function*, which is defined as $v_*(k, x) := \max_\pi v_\pi(k, x)$.

A. Weak reachability

Let a set $\mathcal{Q} \subset \mathbb{Z}_n$, an integer $k \in \mathbb{Z}_{\geq 0}$, and a vertex $x \in \mathbb{Z}_n$ be given. Consider a random directed path $\mathbf{x} \in \mathcal{S}(x)$. The condition $\text{gph}(\mathbf{x}) \cap (\{1, \dots, k\} \times \mathcal{Q}) \neq \emptyset$ requires that the random path \mathbf{x} reaches the set \mathcal{Q} in a number of steps less than or equal to k . Hence, given $\mathcal{Q} \subset \mathbb{Z}_n$, $k \in \mathbb{Z}_{>0}$, and $\xi \in \mathbb{Z}_n$, the *weak reachability probabilities* are defined as

$$\mathfrak{R}_{\mathcal{Q}}(k, x) := \sup_{\mathbf{x} \in \mathcal{S}(x)} \mathbb{P}(\text{gph}(\mathbf{x}) \cap (\{1, \dots, k\} \times \mathcal{Q}) \neq \emptyset).$$

The value $\mathfrak{R}_{\mathcal{Q}}(k, x)$ represents the probability that a stochastic directed path starting at x that visits the set \mathcal{Q} in k or less steps exists. Given the stochastic digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mu(\cdot))$ and a set $\mathcal{Q} \subset \mathbb{Z}_n$, let us now introduce an augmented stochastic digraph $\check{\mathcal{G}}$, obtained by: i) adding two additional states, $n+1$ and $n+2$; ii) substituting in \mathcal{E}_s , $s = 1, \dots, r$, the edges (i, j) such that $j \in \mathcal{Q}$ with new edges $(i, n+1)$; and iii) adding the edges $(n+1, n+2)$ and $(n+2, n+2)$ to all \mathcal{E}_s , $s = 1, \dots, r$. Then, the next proposition holds.

Proposition 1. *Consider a stochastic digraph \mathcal{G} satisfying Assumption 1 and its corresponding augmented stochastic digraph $\check{\mathcal{G}}$. Let $\check{p}(j|i, a)$ the transition probabilities for $\check{\mathcal{G}}$, calculated by (8) and let us fix the reward function $r(i, a, j) = \mathbb{1}_{\{n+1\}}(j)$ so that $\mathbf{r}_{k+1} = \mathbb{1}_{\{n+1\}}(\mathbf{x}_{k+1})$. Then, letting v_* be the optimal state-value function corresponding to such a selection of r , for all $k \in \mathbb{Z}_{\geq 0}$ and all $x \in \mathbb{Z}_n$, one has*

$$\mathfrak{R}_{\mathcal{Q}}(k, x) = v_*(k, x). \quad (11)$$

Proof. Let ζ be a random directed path of the stochastic digraph \mathcal{G} . By [16], we have that $\mathfrak{R}_{\mathcal{Q}}(k, x) = \sup_{\zeta \in \mathcal{S}(x)} \mathbb{E} [\max_{j \in \{1, \dots, k\}} \mathbb{1}_{\mathcal{Q}}(\zeta_j)]$. Since $\check{p}(n+2|n+1, a) = 1$ for all $a \in \mathcal{A}(n+1)$ and $\check{p}(n+2|n+2, a) = 1$ for all $a \in \mathcal{A}(n+2)$ by the definition of $\check{\mathcal{G}}$, and taking into account that $\mathbf{r}_{\tau+1} = \mathbb{1}_{\{n+1\}}(\mathbf{x}_{\tau+1})$, we obtain

$$\begin{aligned} \sum_{\tau=0}^{k-1} \mathbf{r}_{\tau+1} &= \sum_{\tau=0}^{k-1} \mathbb{1}_{\{n+1\}}(\mathbf{x}_{\tau+1}) \\ &= \begin{cases} 1, & \text{if } \exists \tau \in \{1, \dots, k\} \text{ s.t. } \mathbf{x}_\tau = n+1, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Therefore, since, by the definition of $\tilde{\mathcal{G}}$, one has that $\mathbf{x}_\tau = n + 1$ if and only if the corresponding $\zeta_\tau \in \mathcal{Q}$, it results that $\mathbb{E}[\sum_{\tau=0}^{k-1} \mathbf{r}_{k+1}] = \mathbb{E}[\max_{j \in \{1, \dots, k\}} \mathbb{1}_{\mathcal{Q}}(\zeta_j)]$. Hence, the statement follows by Theorem 2 and by the definition of the optimal state-value function v_* . \square

The bound given in (11) is tight. In fact, by [10, Prop. 4], for each $(k, x) \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_n$, there exists $\bar{\mathbf{x}} \in \mathcal{S}(x)$ such that $\mathfrak{R}_{\mathcal{Q}}(k, x) = \mathbb{P}(\exists \kappa \in \{1, \dots, k\} : \bar{\mathbf{x}}_\kappa \in \mathcal{Q})$, i.e., there exists a stochastic path that attains the probability given in (11).

B. Strong recurrence

Under Assumption 1, given $\mathcal{Q} \subset \mathbb{Z}_n$, $k \in \mathbb{Z}_{>0}$, and $\xi \in \mathbb{Z}_n$, the *strong recurrence probabilities* are defined as

$$\mathfrak{V}_{\mathcal{Q}}(k, x) := \inf_{\mathbf{x} \in \mathcal{S}(x)} \mathbb{P}(\text{gph}(\mathbf{x}) \cap (\{1, \dots, k\} \times \mathcal{Q}) \neq \emptyset).$$

The value $\mathfrak{V}_{\mathcal{Q}}(k, x)$ represents the probability that all the stochastic directed paths starting at x visit the set \mathcal{Q} in k or less steps. As for the weak reachability probabilities, this function can be determined by computing the optimal value function of a Markov decision process.

Proposition 2. *Consider a stochastic digraph \mathcal{G} satisfying Assumption 1 and its corresponding augmented stochastic digraph $\tilde{\mathcal{G}}$. Let $\tilde{p}(j|i, a)$ the transition probabilities for $\tilde{\mathcal{G}}$, calculated by (8) and let us fix the reward function $r(i, a, j) = -\mathbb{1}_{\{n+1\}}(j)$, so that $\mathbf{r}_{k+1} = -\mathbb{1}_{\{n+1\}}(\mathbf{x}_{k+1})$. Then, letting v_* be the optimal state-value function corresponding to such a selection of r , for all $k \in \mathbb{Z}_{\geq 0}$ and all $x \in \mathbb{Z}_n$, one has*

$$\mathfrak{V}_{\mathcal{Q}}(k, x) = -v_*(k, x). \quad (12)$$

Proof. Let ζ be a random directed path of the stochastic digraph \mathcal{G} . By [10], we have that $\mathfrak{V}_{\mathcal{Q}}(k, x) = 1 - \sup_{\zeta \in \mathcal{S}(x)} \mathbb{E} \left[\prod_{j=1}^k \mathbb{1}_{\mathbb{Z}_n \setminus \mathcal{Q}}(\zeta_j) \right]$. Hence, by considering that

$$\prod_{j=1}^k \mathbb{1}_{\mathbb{Z}_n \setminus \mathcal{Q}}(\zeta_j) = \begin{cases} 0, & \text{if } \exists \tau \in \{1, \dots, k\} \text{ s.t. } \mathbf{x}_\tau \in \mathcal{Q}, \\ 1, & \text{otherwise,} \end{cases}$$

we have that $\prod_{j=1}^k \mathbb{1}_{\mathbb{Z}_n \setminus \mathcal{Q}}(\zeta_j) = 1 - \max_{j \in \{1, \dots, k\}} \mathbb{1}_{\mathcal{Q}}(\zeta_j)$. Therefore, the statement follows by arguments similar to those used to prove Proposition 1. \square

As for (11), the bound given in (12) is tight: by [13, Prop. 2], for each $(k, x) \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_n$, there exists $\underline{\mathbf{x}} \in \mathcal{S}(x)$ such that $\mathfrak{V}_{\mathcal{Q}}(k, x) = \mathbb{P}(\exists \kappa \in \{1, \dots, k\} : \underline{\mathbf{x}}_\kappa \in \mathcal{Q})$, i.e., there exists a stochastic path that attains the probability (12).

The results given in Propositions 1 and 2 are not surprising in light of the definition of optimal value functions, observing that node $n + 2$ is terminal, and that the only nonzero reward is gathered transitioning to such a node. Furthermore, these results are strongly linked with [8, Thm. 1 and Thm. 2]. In particular, the recursions in these statements to design optimal value functions are similar to those given in (3). However, while the objective in [8] is to design the control input for a class of hybrid systems to ensure safety, here we aim at determining the maximal and minimal probabilities of visiting a given set for stochastic digraphs. The main objective of this technical note, in fact, is to show that

the reachability properties of these discrete-time stochastic systems with non-unique solutions can be characterized via dynamic programming. Further, in the case that the dimension of the problem to be solved is too large, we can resort to stochastic optimization methods, such as the *SARSA* and *Q-learning* algorithms, thanks to the analogy with Markov decision processes established in Section III.

The reachability properties of stochastic digraphs can be also characterized by using the formulas (3), as done, e.g., in [17, Prop. 1]. However, the computational complexity of these formulas is generally larger than the one using dynamic programming or stochastic optimization methods, thus further motivating the interest in the results given in this technical note.

The following remark reviews different techniques to estimate the function $v_\pi(k, x)$ given the transition probabilities $p(j, r|i, a)$ and the reward function $r(i, a, j)$.

Remark 4. In [15], different methods are proposed to determine the function $v_*(k, x)$ given the transition probabilities $p(j, r|i, a)$ and the reward function $r(i, a, j)$. By the Bellman optimality equations, the function $v_*(k, x)$ satisfies

$$v_*(k, x) = \max_a \sum_{j, r} p(j, r|i, a) (r + v_*(k - 1, j)). \quad (13)$$

The dimension of the problem may hamper the analytical solution of the set of nonlinear equations (13). Therefore, several algorithms, such as the *policy iteration* and the *value iteration*, have been proposed to determine an approximate solution to (13). These algorithms have a worst case complexity that is polynomial in the number of states and actions. Other approaches, such as the Q-learning and the SARSA algorithms, determine an approximate solution to (13) from raw experience without a model of the MDP dynamics.

Remark 5. The functions $\mathfrak{R}_{\mathcal{Q}}(k, x)$ and $\mathfrak{V}_{\mathcal{Q}}(k, x)$ can be also computed explicitly by leveraging the construction used in the proof of Theorem 1. Namely, letting $m(k, x)$ and $\ell(k, x)$ be defined as in such a proof with \mathcal{I}_j replaced by \mathcal{Q} , by [17], for all $k \in \mathbb{Z}_{\geq 0}$ and all $x \in \mathbb{Z}_n$, one has that

$$\mathfrak{R}_{\mathcal{Q}}(k, x) = m(k, x), \quad \mathfrak{V}_{\mathcal{Q}}(k, x) = \ell(k, x).$$

Thus, in principle, these bounds on the reachability probabilities can be computed using the methods given in [4], [6], [7], whose complexity is comparable to that of solving directly the dynamic programming equations (13). On the other hand, by leveraging the correspondence between optimal state-value functions and reachability probabilities outlined in Propositions 1 and 2, the other methods reviewed in Remark 4 can be used, thus leading to smaller computational complexities, especially in the case that one is interested to characterizing infinite-horizon reachability probabilities.

V. CASE STUDY: MOBILE AGENT-BASED EPIDEMIC MODELS VIA STOCHASTIC DIGRAPHS

The main aim of this section is to show how stochastic digraphs can be used to model diffusion phenomena in populations of mobile agents. With particular reference to epidemic processes, the results of Sections III and IV can

be leveraged to determine upper and lower bounds for the temporal trends of the number of infected agents. To exemplify the potentiality of our approach, we consider a susceptible-infected-removed (SIR) epidemic process [18] that evolves over a mobile population of agents. Notably, we consider a group of N agents that move over a graph $\mathcal{M} := (\mathbb{Z}_{\mathcal{X}}, \mathcal{W})$, $\mathcal{X} \in \mathbb{Z}_{\geq 0}$. The state of the i -th agent is $\varphi_i = [\sigma_i \ \varpi_i]^\top$, where σ_i is 1 for susceptible, 2 for infected, or 3 for recovered, and $\varpi_i \in \mathbb{Z}_{\mathcal{X}}$ denotes his position in the graph \mathcal{M} . Letting $\mathcal{N}^o(\varpi)$ be the map of the out-neighborhoods of $\varpi \in \mathbb{Z}_{\mathcal{X}}$ in the graph \mathcal{M} , one may therefore describe the motion as

$$\varpi_i^+ \in \mathcal{N}^o(\varpi_i), \quad i = 1, \dots, N,$$

where ϖ_i^+ denotes the position of the i -th agent in the graph \mathcal{M} at the next time step. Each susceptible agent can transition with probability $\alpha \in [0, 1]$ to the infected state if it shares the same position with an infected agent j , i.e., $\varpi_i = \varpi_j$ for some $i \neq j$. Note that each agent can be infected by any other agent sharing the same position and that these random events are considered independent of each other. Finally, each infected agent can transition to the recovered state with probability $\beta \in [0, 1]$. Thus, in such a model, the contact graph among agents is time-varying and is influenced by the movement of agents over the graph \mathcal{M} . The next theorem shows that it is possible to construct a stochastic digraph, named **SIRM**, representing exactly the behavior of the agent-based SIR model.

Theorem 3. *Let the motion digraph \mathcal{M} , the infection probability α and the recovery probability β be given. Then, there exists a stochastic digraph **SIRM** such that its stochastic directed paths are in one-to-one correspondence with the trajectories of the agent-based SIR model with mobile agents.*

Proof. Let $\sigma = [\sigma_1 \ \dots \ \sigma_N]^\top$ denote the stack vector of all node states and let $\varpi = [\varpi_1 \ \dots \ \varpi_N]^\top$ denote the stack vector of the positions of the agents. The updates of the state σ_i are modeled by the stochastic difference inclusion $\sigma_i^+ \in G_i(\sigma, \varpi, u_i, v_i)$, where

$$G_i(\sigma, \varpi, u_i, v_i) = \begin{cases} 1, & \text{if } \sigma_i = 1 \\ & \wedge \left(\bigwedge_{\substack{j=1 \\ j \neq i}}^N \neg(\varpi_i = \varpi_j \wedge \sigma_j = 2 \wedge [u_i]_j = 2) \right), \\ 2, & \text{if } \sigma_i = 1 \\ & \wedge \left(\bigvee_{\substack{j=1 \\ j \neq i}}^N (\varpi_i = \varpi_j \wedge \sigma_j = 2 \wedge [u_i]_j = 2) \right), \\ v_i + 1, & \text{if } \sigma_i = 2, \\ 3, & \text{if } \sigma_i = 3, \end{cases}$$

and $(\mathbf{u}_i)_k : \Omega \rightarrow \{1, 2\}^N$, $(\mathbf{v}_i)_k : \Omega \rightarrow \{1, 2\}$ are two sequences of i.i.d. random variables with

$$\mathbb{P}([(u_i)_k]_j = 2) = \alpha, \quad j = 1, \dots, N, \quad (14a)$$

$$\mathbb{P}((v_i)_k = 2) = \beta, \quad (14b)$$

for $i = 1, \dots, N$. Hence, define the set-valued map

$$S_i(\sigma, \varpi, u_i, v_i) := \begin{bmatrix} G_i(\sigma, \varpi, u_i, v_i) \\ \mathcal{N}^o(\varpi_i) \end{bmatrix},$$

describing the next values of the epidemic state and of the position of the i -th agent given the values attained by the random inputs u_i and v_i . Thus, let $\bar{n} = (3\mathcal{X})^N$ and let $\bar{\Pi} : (\{1, 2, 3\} \times \mathcal{X})^N \rightarrow \mathbb{Z}_{\bar{n}}$ be a bijective function mapping each vector in $(\{1, 2, 3\} \times \mathcal{X})^N$ into an integer in $\mathbb{Z}_{\bar{n}}$. Define $\omega = [u_1^\top \ v_1 \ \dots \ u_N^\top \ v_N]^\top \in \{1, 2\}^{N^2+N}$ and consider the map $\bar{Q} : \mathbb{Z}_{\bar{n}} \times \{1, 2\}^{N^2+N} \rightarrow \mathbb{Z}_{\bar{n}}$,

$$\bar{Q}(x, \omega) = \bar{\Pi} \left(\begin{bmatrix} S_1(\bar{\Pi}^{-1}(x), u_1, v_1) \\ \vdots \\ S_N(\bar{\Pi}^{-1}(x), u_N, v_N) \end{bmatrix} \right), \quad (15)$$

where $\bar{\Pi}^{-1}$ is the inverse map of $\bar{\Pi}$. Therefore, since the set $\{1, 2\}^{N^2+N}$ is diffeomorphic to $\mathbb{Z}_{\bar{h}}$, with $\bar{h} = 2^{N^2+N}$, letting $\mathcal{V} = \mathbb{Z}_{\bar{n}}$, the stochastic directed paths of the stochastic digraph **SIRM** $(\mathcal{V}, \mathcal{E}, \mu(\cdot))$ corresponding to the map $\bar{Q}(x, \omega)$ given in (15) with the distribution $\mu(\cdot)$ induced by the sequences of i.i.d. random variables $u_1, v_1, \dots, u_N, v_N$ satisfying (14) are in one-to-one correspondence with the trajectories of the agent-based SIR model with mobile agents. \square

In the following example, we show how the stochastic digraph **SIRM** together with the reinforcement learning techniques of Section IV can be used to determine upper and lower bounds on the cumulative number of infected agents.

Example 4. Consider the agent-based SIR model with $N = 3$, $\alpha = 0.7$, $\beta = 0.3$, and \mathcal{M} as in Figure 2.




Fig. 2. Graph \mathcal{M} used in Example 4.

Using the construction in Theorem 3, we determine the stochastic digraph **SIRM** representing the agent-based SIR model. Then, using the technique given in Section III-C, we compute the transition probabilities $p(j|i, a)$ between the nodes of **SIRM**. Note that, in this setting, each node $i \in \mathcal{V}$ corresponds to a different state vector $\phi = [\phi_1^\top \ \dots \ \phi_N^\top]^\top \in (\{1, 2, 3\} \times \{1, \dots, \mathcal{X}\})^N$, and an action $a \in \mathcal{A}$ represents a motion law of all the agents.

The upper and lower bounds on the cumulative number of infected agents can be determined using the techniques described in Section IV. Namely, let $\vartheta : \mathcal{V} \rightarrow \mathbb{Z}_{\geq 0}$ be a map such that $\vartheta(i)$ is the number of agents that are either infected or recovered in the state corresponding to node i . Hence, letting $\bar{v}(i, a, j) = \vartheta(j) - \vartheta(i)$ ($\underline{v}(i, a, j) = \vartheta(i) - \vartheta(j)$), the optimal state-value function $\bar{v}_*(i)$ ($\underline{v}_*(i)$) is such that $\bar{v}_*(i) + \vartheta(i)$ ($\vartheta(i) - \underline{v}_*(i)$) is an upper (lower) bound on the cumulative number of infected agents.

Furthermore, the proposed framework can be used to compute the expected cumulative number of infected agents in the case that the agents move on the graph uniformly at random, i.e., the next position of the i -th agent is selected uniformly at random in $\mathcal{N}^o(\varpi_i)$, namely, letting P be the matrix whose

(i, j) -th entry is $p_{i,j} = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} p(j|i, a)$, which is the transition matrix of the Markov chain corresponding to a uniform random walk, letting ϕ_0 be the initial state of the agent-based SIR model, and letting $\Theta = [\vartheta(1) \ \cdots \ \vartheta(\bar{n})]^\top$, the expected cumulative number of infected agents at time k is

$$\varsigma_k = \mathbb{E}_{\Pi(\phi_0)} P^k \Theta. \quad (16)$$

Letting $\phi_0 = [2 \ 1 \ 1 \ 1 \ 1 \ 2]^\top$ the SARSA, and the Q-learning algorithms [15] have been used to estimate $\underline{v}_*(\Pi(\phi_0))$ and $\bar{v}_*(\Pi(\phi_0))$. Both algorithms have been used with learning rate equal to 0.1 over 10^4 episodes and employing an ϵ -greedy policy with $\epsilon = 0.1$. The SARSA algorithm provided the estimates $\hat{v}_*^{\text{SARSA}}(\Pi(\phi_0)) = 1.4978$ and $\hat{\bar{v}}_*^{\text{SARSA}}(\Pi(\phi_0)) = -0.7011$, whereas the Q-learning algorithm provided the estimates $\hat{v}_*^{\text{Q}}(\Pi(\phi_0)) = 1.4922$ and $\hat{\bar{v}}_*^{\text{Q}}(\Pi(\phi_0)) = -0.6986$. Figure 3 illustrates the bounds obtained using these optimal state value functions, the values ς_k gathered using (16), and the results of a Monte Carlo simulation over 10^5 samples of the agent-based SIR model.




Fig. 3. Expected number of infected agents in Example 4.

As shown in Figure 3, the expected cumulative number of infected agents determined using (16) and running a Monte Carlo simulation are in agreement. Furthermore, the values $\vartheta(i) - \underline{v}_*(i)$ and $\vartheta(i) + \bar{v}_*(i)$ are asymptotic lower and upper bounds on the expected cumulative number of infected agents.

Although in this section the attention has been focused on SIR models, the proposed techniques can be easily extended to deal with other classes of epidemic models [17], [18].

VI. CONCLUSIONS

Weak and strong reachability probabilities of a set for a stochastic digraph have been studied in this technical note. The available state-of-the-art techniques are characterized by a heavy computational burden that makes them unsuitable to efficiently tackle large graphs. In this technical note, we leverage techniques borrowed from the reinforcement learning domain to alleviate such a burden. Toward this aim, we first showed that the transition probabilities in stochastic digraphs can be modeled via a difference inclusion, which can be equivalently represented as a Markov decision process. Hence, we have discussed how to select the reward function to determine upper and lower bounds for such transition probabilities in the stochastic digraph. Crucially, we have formulated the transition probabilities as a function of an action space that depends on the current state, enabling a significant reduction

of the problem size and paving the way to the definition of an efficient algorithm for their computation.

Notably, we discussed how stochastic digraphs can be used to model epidemic spreading and how our results apply to this framework. Our results are entirely general, and can be applied without any limitation to stochastic graphs of different nature, with applications in sensor networks, biological circuits, chemical reactions, and power grids.

REFERENCES

- [1] T. Wang, H. Krim, and Y. Viniotis, "A generalized Markov graph model: Application to social network analysis," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 318–332, 2013.
- [2] F. Vecchio, F. Miraglia, G. Curcio, G. Della Marca, C. Vollono, E. Mazzucchi, P. Bramanti, and P. M. Rossini, "Cortical connectivity in fronto-temporal focal epilepsy from EEG analysis: A study via graph theory," *Clinical Neurophysiology*, vol. 126, no. 6, pp. 1108–1116, 2015.
- [3] P. Frasca, C. Ravazzi, R. Tempo, and H. Ishii, "Gossips and prejudices: Ergodic randomized dynamics in social networks," *IFAC Proceedings Volumes*, vol. 46, no. 27, pp. 212–219, 2013.
- [4] C. Possieri and A. R. Teel, "Stochastic robust simulation and stability properties of chemical reaction networks," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 1, pp. 2–12, 2018.
- [5] S. M. Wagner and N. Neshat, "Assessing the vulnerability of supply chains using graph theory," *International Journal of Production Economics*, vol. 126, no. 1, pp. 121–129, 2010.
- [6] C. Possieri and A. R. Teel, "Weak reachability and strong recurrence for stochastic directed graphs in terms of auxiliary functions," in *55th Conference on Decision and Control*, (Las Vegas, NV, USA), pp. 3714–3719, IEEE, 2016.
- [7] C. Possieri and A. R. Teel, "Asymptotic stability in probability for stochastic Boolean networks," *Automatica*, vol. 83, pp. 1–9, 2017.
- [8] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.
- [9] B. E. Fristedt and L. F. Gray, *A modern approach to probability theory*. New York, NY, USA: Springer, 2013.
- [10] A. Subbaraman and A. R. Teel, "A converse Lyapunov theorem for strong global recurrence," *Automatica*, vol. 49, no. 10, pp. 2963–2974, 2013.
- [11] S. Grammatico, A. Subbaraman, and A. R. Teel, "Discrete-time stochastic control systems: examples of robustness to strictly causal perturbations," in *52nd Conference on Decision and Control*, (Firenze, Italy), pp. 6403–6408, IEEE, 2013.
- [12] A. R. Teel, J. P. Hespanha, and A. Subbaraman, "A converse Lyapunov theorem and robustness for asymptotic stability in probability," *IEEE Transactions on Automatic Control*, vol. 59, no. 9, pp. 2426–2441, 2014.
- [13] C. Possieri and A. R. Teel, "A Lyapunov theorem certifying global weak reachability for stochastic difference inclusions with random inputs," *Systems & Control Letters*, vol. 109, pp. 37–42, 2017.
- [14] P. A. Gagniu, *Markov chains: from theory to implementation and experimentation*. Hoboken, NJ, USA: John Wiley & Sons, 2017.
- [15] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. Cambridge, MA, USA: MIT press, 2018.
- [16] A. R. Teel, "A Matrosov theorem for adversarial Markov decision processes," *IEEE Transactions on Automatic Control*, vol. 58, no. 8, pp. 2142–2148, 2013.
- [17] C. Possieri and A. Rizzo, "A mathematical framework for modeling propagation of infectious diseases with mobile individuals," in *58th Conference on Decision and Control*, (Nice, France), pp. 3750–3755, IEEE, 2019.
- [18] F. Brauer, C. Castillo-Chavez, and C. Castillo-Chavez, *Mathematical models in population biology and epidemiology*. New York, NY, USA: Springer, 2012.