POLITECNICO DI TORINO
Repository ISTITUZIONALE

Fast Spread in Controlled Evolutionary Dynamics

(Article begins on next page)

24 April 2024

Fast Spread in Controlled Evolutionary Dynamics

Lorenzo Zino, *Member, IEEE,* Giacomo Como, *Member, IEEE,* and Fabio Fagnani, *Member, IEEE*

Abstract—We study a controlled evolutionary dynamics that models the spread of a novel state in a network where the exogenous control aims to quickly spread the novel state. We estimate the performance of the system by analytically establishing upper and lower bounds on the expected time needed for the novel state to replace the original one. Such bounds are expressed as functions of the control policy adopted and of the network structure, and establish fundamental limitations on the system's performance. Leveraging these results, we classify network structures depending on the possibility of achieving a fast spread of the novel state (i.e., complete replacement in a time growing logarithmically with the network size) using simple open-loop control policies. Finally, we propose a feedback control policy that using little knowledge of the network and of the system's evolution at a macroscopic level allows for a substantial speed up of the spreading process, guaranteeing fast spread on topologies where simple open-loop control policies are not sufficient. Examples and simulations corroborate our findings.

Index Terms—Network systems; Feedback control; Spreading processes; Evolutionary dynamics; Diffusion of innovation.

I. Introduction

IN the last decades, the study of spreading processes in network systems has significantly advanced. Increasingly refined models have been proposed, with applications spanning from epidemic spreading to the diffusion of innovation and the adoption of social norms. The theoretical analysis of these models has guided the development of accurate control policies. The literature on epidemics offers a paradigmatic example, whereby the understanding of how the network of interactions influences the spread of a disease [2]–[5] has paved the way for the design and study of control policies to mitigate or stop epidemic outbreaks [6]–[9]. Similar, for decision-making in social systems, the extensive analysis of the voter model [10] and of dynamical interaction models in game-theoretic frameworks [11], [12] have inspired the design of techniques to optimally place spreaders in a network [13], [14] and to control networks of imitative agents [15].

In this paper, we focus on a dynamical network system that models the spread of novel states in a population. It first appeared in the literature in [16] under the name of *evolutionary dynamics*, where it is used to model the competition between a novel species and the original one in a geographic region. In

this formulation, the network's nodes represent locations and interconnections are determined by proximity. Locations are assumed to be small, so it is reasonable to assume that each location is fully occupied by a single species, as in [16]. This is modeled by equipping each node with a binary state. The spreading process follows a probabilistic rule and takes place through pairwise interactions between adjacent nodes, which yield a competition between the species present in the two nodes to place their offspring. This model naturally accounts for possible biases that favor one species against the other, thus reflecting the presence of an *evolutionary advantage*. Besides evolutionary competition, this model may also find other valuable applications, e.g., diffusion of innovation in social systems. In this setting, the two states may represent two alternative technologies, and the nodes are users interacting and exchanging information on a social network. As a result of these pairwise interactions, one of the two individuals may get convinced to adopt the technology used by the other one, thus via an imitative/contagion mechanism. Evolutionary advantage captures an intrinsic bias, thus accounting, e.g., for a quality or cost difference between the two products.

The literature on evolutionary dynamics [16]–[20] focuses on understanding how the network structure influences the probability that the novel state diffuses in the network (termed fixation probability) and the duration of such a spreading process. However, few analytical results have been established. To the best of our knowledge, the fixation probability has been analytically computed for specific network topologies [17], [18], while most of the results are based on simulations [16], [19], [20], and no control policies have been studied. In a preliminary work [1], we proposed a new formalism for evolutionary dynamics, which presents two main novelties: i) the spreading process is modeled through a link-based (instead of a node-based as in [16]) activation mechanism, and ii) an exogenous control action is incorporated to model the controlled introduction of the novel state in some nodes. This change of perspective and the explicit introduction of a control action have enabled us to gain new analytical insight into the expected duration of the spreading process. In [1], we introduced a first feedback control policy to speed up the spreading process, but its feasibility was limited to very specific networks and it presented two drawbacks: it was sensible to small data errors and very costly from the control viewpoint. An improved policy was proposed in [21] and tested on a real-world case study via numerical simulations.

Here, we undertake a fundamental analysis of controlled evolutionary dynamics in great generality, encompassing both open-loop and feedback control policies. Control actions are evaluated on the basis of their cost —in terms of the number of nodes where the control acts and of its total energy— and the spreading time of the process. Our main contribution is

the establishment of fundamental limitations between these quantities, from which we derive an array of ready-to-use tools to estimate the expected spreading time and control cost for specific control policies and network structures. In particular, this allows for performing a full theoretical analysis of the feedback control law presented in [21].

Formally, we cast the controlled evolutionary dynamics as a Markov process [22]. Under mild assumptions on the network structure and on the form of the exogenous control, the process has a unique absorbing state coinciding with the complete spread of the novel state. Our study focuses on the estimation of the expected time needed for the Markov process to reach such an absorbing state, in the presence of an evolutionary advantage. The exogenous control makes this Markov process nonhomogeneous, thereby hindering a direct application of classical results for homogeneous Markov processes [22]–[24]. Our analysis is based on two key properties of the dynamics, which allow us to establish performance guarantees and fundamental limitations on the controlled evolutionary dynamics. First, the process presents natural monotonicity properties with respect to the initial configuration and the evolutionary advantage of the novel state. Second, the speed of the spreading process is proportional to the size of the boundary separating the set of nodes in the two different states, similar to what was observed in network epidemic models [8].

To illustrate the implications of our findings, we analyze network families with different characteristics. For highly connected networks, we prove that the evolutionary advantage and the monotonicity properties of the dynamics generate positive feedback, yielding fast spread, even in the absence of any exogenous control (except for the initial seeding): the novel state reaches complete spread in a time that grows logarithmically in the network size. On the other hand, we show that the poorly connected topologies hinder such positive feedback, as the system may enter and remain stuck in bottlenecks where the spreading process is significantly slowed down. It is in these cases that an exogenous control is the most valuable in reinforcing the spreading process.

To this aim, we design and study a feedback control law that compensates for such slowdowns by concentrating the control efforts when the system enters configurations whose boundary is too small to guarantee a fast spread. The proposed feedback policy needs limited a-priori knowledge on the topology of the network for its implementation and is driven by two macroscopic one-dimensional observables: the number of nodes in the novel state and the size of the boundary separating the set of nodes in the two different states. Despite such a low complexity structure, our control policy exhibits good performance that is guaranteed by analytical results. We apply the proposed policy to stochastic block models (SBMs), which are representative of many real-world networks [25]. For them, while simple open-loop control policies fail in yielding fast spread, our feedback control law succeeds. We believe that the capability of incorporating and studying control architectures in spreading processes and, more generally, in network dynamics is a crucial step from the application viewpoint.

The rest of the paper is organized as follows. Section II introduces the model. Section III presents the main results.

Section IV analyzes open-loop control policies, in particular constant policies, for three fundamental network examples. Section V, we propose and study a feedback control policy. Section VI concludes the paper.

Notation. $\mathbb{R}_+$ stands for the set of nonnegative reals. The n-dimensional all-0 and all-1 vectors are denoted by $\mathbf{0}$ and $\mathbf{1}$, respectively; $\delta^{(i)}$ is a vector with a 1 in its ith entry and 0 elsewhere; for $\mathcal{S} \subseteq \{1, \ldots, n\}$, $\delta^{(\mathcal{S})} = \sum_{i \in \mathcal{S}} \delta^{(i)}$. The transpose of a matrix M is denoted as $M^\top$. An inequality $\boldsymbol{x} \geq \boldsymbol{y}$ between two vectors is meant to hold true entrywise, i.e., $x_i \geq y_i$, for all i. Finally, we use the notation $\mathbb{1}_{\mathcal{A}}$ for the indicator function of a set $\mathcal{A}$ and $f(t_0^-) := \lim_{t \nearrow t_0} f(t)$ and $f(t_0^+) := \lim_{t \searrow t_0} f(t)$ for the left and right limits.

II. Model

We describe the network as a finite weighted undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, whereby $\mathcal{V} = \{1, \ldots, n\}$ is the node set, $\mathcal{E}$ is the set of undirected links, and $W = W^\top$ in $\mathbb{R}_+^{n \times n}$ is the (symmetric) weight matrix, with $W_{ij} = W_{ji} > 0$ if and only if $\{i, j\} \in \mathcal{E}$. Throughout, we shall assume that $\mathcal{G}$ is connected, i.e., W is irreducible. Each node i in $\mathcal{V}$ is characterized by a binary state

$$X_i(t) = \left\{ \begin{array}{ll} 0 & \text{if at time } t \text{ node } i \text{ is in the original state,} \\ 1 & \text{if at time } t \text{ node } i \text{ is in the novel state,} \end{array} \right.$$

evolving in continuous time according to the following two mechanisms.

Spreading mechanism. Each undirected link $\{i, j\}$ in $\mathcal{E}$ is equipped with an independent rate-W_{ij} Poisson clock modeling the interactions between nodes i and j. If the clock ticks at time t and both nodes have the same state, nothing happens. Otherwise, a conflict takes place and the winning state occupies both nodes. Conflicts are solved in a probabilistic way: the novel state wins a conflict (independently of the others) with fixed probability $\beta \in [0, 1]$, yielding $X_i(t^+) = X_j(t^+) = 1$, otherwise the original state wins and $X_i(t^+) = X_j(t^+) = 0$. The parameter β captures the *evolutionary advantage* of the original state ($\beta < 1/2$) or of the novel state ($\beta > 1/2$). For simplicity, we assume that β is uniform over edges and that every conflict ends with a winner. These assumptions may be partially relaxed, as we shall discuss in Remarks 3 and 7.

Exogenous control. We fix a locally integrable function $U : \mathbb{R}_+ \to \mathbb{R}_+^n$ whose ith entry $U_i(t)$ represents the control rate at which the novel state is enforced at node i at time t. The node will be occupied by the novel state (i.e., $x_i(t^+) = 1$) according to a rate-$U_i(t)$ Poisson clock, independent of others.

The triple $(\mathcal{G}, \beta, U(t))$ shall be referred to as the *controlled evolutionary dynamics*, whose mechanisms induce a Markov process $X(t) = (X_1(t), \ldots, X_n(t))$ on the configuration space $\{0, 1\}^n$. The only transitions that can take place from a configuration $X(t) = \boldsymbol{x}$ are towards configurations that differ from $\boldsymbol{x}$ in a single entry. Specifically, the transition rates

$$\lambda_i^\pm(\boldsymbol{x}, t) := \lim_{h \searrow 0} \frac{1}{h} \mathbb{P}[X(t+h) = \boldsymbol{x} \pm \delta^{(i)} | X(t) = \boldsymbol{x}]$$
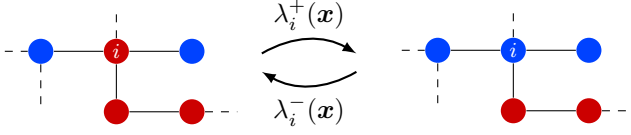
Fig. 1: Transitions of the Markov process $X(t)$. Red and blue nodes denote the original and novel state, respectively.

from $\boldsymbol{x}$ to $\boldsymbol{x} + \delta^{(i)}$ (if $x_i = 0$) and $\boldsymbol{x} - \delta^{(i)}$ (if $x_i = 1$) at time $t \geq 0$ (see Fig. 1 for an illustration) satisfy

$$
\begin{aligned}
\lambda_i^+(\boldsymbol{x}, t) &= (1 - x_i)(\beta(W\boldsymbol{x})_i + U_i(t)) \\
\lambda_i^-(\boldsymbol{x}, t) &= x_i(1 - \beta)(W(\mathbf{1} - \boldsymbol{x}))_i.
\end{aligned} \tag{1}
$$

Equation (1) states that, if node i is in state $x_i = 0$ at time t, then its transition rate to 1 is the sum of two terms: the first addend is proportional to the total weight of links towards neighbors of node i that are in state 1 at time t; the second one accounts for the control effort exerted on the node. On the other hand, if node i is in state $x_i = 1$, its transition rate to 0 is proportional to the total weight of links toward neighbors that are in state 0 at time t. In general, the process is nonhomogeneous, since the terms $U_i(t)$ can be time-varying.

The proposed framework captures key features of real-world spreading processes discussed in the introduction, such as the role of the network structure, the competition between different species or products, the possible presence of an evolutionary advantage, and the possibility to exert exogenous control on the system, and has been successfully adopted in a preliminary form in [21] to study the introduction of genetically modified mosquitoes in a geographic region to substitute species that transmit diseases. We refer to [26] for more details on the problem. Moreover, the framework has fundamental ties with existing models for network dynamics.

Remark 1: If $U(t) = \mathbf{0}$ for every $t \geq 0$, the model reduces to a voter model [10], which is biased if $\beta \neq 1/2$ [27]. A homogeneous version of the model without control and where clocks are associated with nodes instead of links has been proposed in [16]. However, the analytical results for such a model are limited to the computation of the probability that the novel state diffuses to the whole network, while more in-depth analyses are limited to specific network structures [17], [18], and the convergence time is studied only numerically.

Remark 2: An equivalent model can be obtained by replacing the exogenous control with a fictitious stubborn node s with fixed state $X_s(t) = 1$, for $t \geq 0$, and considering links with time-varying weights $W_{is} = U_i(t)/\beta$. Models with stubborn nodes have been studied in opinion dynamics [28], [29]. Differently from these works, we focus on the transient rather than the asymptotic analysis, using different tools.

Throughout the paper, we make the following assumptions.

Assumption 1: (i) $\beta > 1/2$; (ii) if $X(t) = \mathbf{0}$, then $U(t) \neq \mathbf{0}$.

Assumption 1 implies that the novel state has an evolutionary advantage on the original state and that the exogenous control is always active whenever the system is in the novel-free configuration $\boldsymbol{x} = \mathbf{0}$.

From (1), Assumption 1, and the fact that $\mathcal{G}$ is connected, it follows that $\boldsymbol{x} = \mathbf{1}$ is the unique absorbing configuration of the system and it is reachable from every other configuration. Hence, the novel state eventually spreads to the whole network almost surely in finite time [22]. From an application perspective, our interest is to shed light on the transient behavior of the system. To this aim, we introduce two performance indices: the *spreading time* and the *control cost* defined respectively as

$$
T := \inf \{t \geq 0 : X(t) = \mathbf{1}\}, \qquad J = \int_0^T \mathbf{1}^\top U(t)\mathrm{dt}.
$$

We shall then focus on the estimation of the expected values of these two indices. Since these may depend on the initial configuration distribution, we denote the conditional expected values of the spreading time and of the control cost as

$$
\mathbb{E}_{\boldsymbol{x}_0}[T] := \mathbb{E}[T \mid X(0) = \boldsymbol{x}_0], \quad \mathbb{E}_{\boldsymbol{x}_0}[J] := \mathbb{E}[J \mid X(0) = \boldsymbol{x}_0],
$$

respectively. For the purpose of this paper, we will typically be interested in the scenario with $X(0) = \mathbf{0}$, that is the novel-free configuration. Hence, most of our results will be expressed as bounds on $\mathbb{E}_{\mathbf{0}}[T]$ and $\mathbb{E}_{\mathbf{0}}[J]$.

We consider two main types of control policy:

Open-loop control policies, where the control signal $U(t)$ is predetermined. A simple example of such control policies are those with $U(t) = \boldsymbol{u}$ constant for every $t \geq 0$. We will refer to these as *constant control policies*.

Feedback control policies, where the control signal $U(t)$ is chosen as a function of the process $X(t)$ itself. Precisely, in this case we consider a function $\nu : \{0,1\}^n \to \mathbb{R}_+^n$ and we take $U(t) = \nu(X(t))$. The triple $(\mathcal{G}, \beta, \nu)$ is then called a *feedback controlled evolutionary dynamics*.

There are typically some constraints that we want to enforce on the admissible control policies. In general, the exogenous control $U(t)$ is constrained to be active only at certain specified nodes, denoted by $\mathcal{U} \subseteq \mathcal{V}$. In this case, the triple $(\mathcal{G}, \beta, U(t))$ is called an *$\mathcal{U}$-controlled evolutionary dynamics*.

We now introduce some fundamental quantities.

Definition 1: Given a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$:

(i) the *weighted boundary* [30] of a node subset $\mathcal{W} \subseteq \mathcal{V}$ is

$$
\zeta(\mathcal{W}) = \sum_{i \in \mathcal{W}} \sum_{j \notin \mathcal{W}} W_{ij};
$$

(ii) the *minimum conductance* and the *maximum expansiveness* profiles are the functions $\phi, \eta : \{0, \ldots, n\} \to \mathbb{R}_+$,

$$
\phi(a) = \min_{\mathcal{W} \subset \mathcal{V}, |\mathcal{W}| = a} \zeta(\mathcal{W}), \quad \eta(a) = \max_{\mathcal{W} \subset \mathcal{V}, |\mathcal{W}| = a} \zeta(\mathcal{W}).
$$

Since $\mathcal{G}$ is undirected, $\zeta(\mathcal{W}) = \zeta(\mathcal{V} \smallsetminus \mathcal{W})$ for all $\mathcal{W} \subseteq \mathcal{V}$. Since $\mathcal{G}$ is connected, $\phi(a) \geq \eta(a) > 0$ for $0 < a < n$.

We end this section by defining three one-dimensional stochastic processes as aggregate statistics of the process $X(t)$:

- the *number of nodes* with novel state $A(t) := \mathbf{1}^\top X(t)$;
- the size of the *boundary* of the set of nodes that are in the novel state $B(t) := X(t)^\top W(\mathbf{1} - X(t))$; and
- the *effective control rate* in nodes that are in the original state $C(t) := (\mathbf{1} - X(t))^\top U(t)$.

III. GENERAL PERFORMANCE ANALYSIS

In this section, we present a series of theoretical results on the performance of controlled evolutionary dynamics. First, we focus on deriving performance guarantees, consisting in bounds on the expected spreading time that depend on the network topology and on the form and intensity of the control. Then, we derive some fundamental limitations of the controlled evolutionary dynamics. They are expressed as control-independent lower bounds on the expected spreading time. Leveraging these general results, ready-to-use corollaries are derived for specific choices of the control policy.

A. Performance Guarantees

Here, we present a general result estimating the performance of controlled evolutionary dynamics, which will be then applied to both open-loop and feedback control policies.

The core of the proposed analysis consists in the estimation of the expected time spent by the process $X(t)$ before reaching the absorbing configuration $\mathbf{1}$. To perform such analysis, we focus on the aggregate statistics $A(t) = \mathbf{1}^\top X(t)$ counting the number of state-1 nodes in the network, and estimate the time spent by $A(t)$ before being absorbed in state n. The analysis is nontrivial because, in general, $A(t)$ is not a Markov process: each transition of $X(t)$ increases or decreases $A(t)$ by 1 unit and its rate depends on the whole vector $X(t)$ through the network weight matrix W, not just on $A(t)$ itself. Specifically, the increase and decrease rates of $A(t)$ respectively satisfy

$$\lambda_t^+ = \sum\nolimits_{i \in \mathcal{V}} \lambda_i^+(X(t), t) = \beta B(t) + C(t), \qquad (2a)$$

$$\lambda_t^- = \sum\nolimits_{i \in \mathcal{V}} \lambda_i^-(X(t), t) = (1 - \beta) B(t). \qquad (2b)$$

For $t \geq 0$, let p_t^+ and p_t^- be the conditional probabilities that, should the process $X(t)$ have a transition at time t, such transition would increase the value of $A(t)$ or decrease it, respectively, by 1 unit. Clearly, when $A(t) = 0$, we get $p_t^+ = 1$ and $p_t^- = 0$, whereas, (2) yields

$$p_t^+ = 1 - p_t^- = \frac{\lambda_t^+}{\lambda_t^+ + \lambda_t^-} = \frac{\beta B(t) + C(t)}{B(t) + C(t)} \geq \beta, \quad (3)$$

whenever $X(t) \neq \mathbf{1}$. The above is a uniform bound that allows us to estimate the conditional probability that, given that $A(t_0) = a$, the aggregate process $A(t)$ will ever go below a at some time $t > t_0$, before reaching its absorbing state $A(T) = n$. Specifically we have the following result for

$$q_a := \inf_{t_0 \geq 0} \min_{\mathbf{1}^\top \boldsymbol{x} = a} \mathbb{P}[A(t) \geq a, \forall t \geq t_0 \mid X(t_0) = \boldsymbol{x}].$$

Lemma 1: For every $1 \leq a < n$, $q_a \geq (2\beta - 1)/\beta$.
Proof: See Appendix A. ∎

Lemma 1 provides a lower bound on q_a and, ultimately, on the desired probability, depending only on the evolutionary advantage β. Using this uniform bound, we now estimate the total amount of time that the process spends in each of its nonabsorbing states, when a relation between the three aggregate statistics $A(t)$, $B(t)$, and $C(t)$ is verified. Specifically, the following is the key technical result which allows us to formulate our performance guarantees.

Proposition 1: Let $f : \{0, \ldots, n\} \to \mathbb{R}_+$ be such that

$$B(t) + C(t) \geq f(A(t)), \qquad \forall t \geq 0. \qquad (4)$$

For $0 \leq a < n$, let T_a be the total amount of time spent by process $A(t)$ in state a. Then,

$$\mathbb{E}[T_a \mid X(0) = \mathbf{0}] \leq \begin{cases} \dfrac{\beta}{2\beta - 1} \dfrac{1}{f(0)} & \text{if } a = 0, \\[2ex] \dfrac{1}{2\beta - 1} \dfrac{1}{f(a)} & \text{if } 0 < a < n. \end{cases}$$

$$(5)$$

Proof: For $s \geq 0$, let $T_a(s)$ denote the total time spent by $A(t)$ in a state a from time s on. Clearly, $T_a = T_a(0)$ and, for $\boldsymbol{x}_0$ in $\{0, 1\}^n$, the Markov property implies that

$$\mathbb{E}[T_a(s) \mid X(s) = \boldsymbol{x}_0, X(0) = \mathbf{0}] = \mathbb{E}[T_a(s) \mid X(s) = \boldsymbol{x}_0].$$

Let $S(t) = \inf\{h \geq 0 : X(t + h) \neq X(t)\}$ be the waiting time for the first configuration change after time t. Observe that, for $t < r \leq t + S(t)$, the total jump rate of the process $A(t)$ can be bounded, using (2) and (4), as

$$\xi(r) = \lambda_r^+ + \lambda_r^- = B(r) + C(r) \geq f(\mathbf{1}^\top \boldsymbol{x}).$$

Standard properties of Markov processes yield

$$\begin{aligned} \mathbb{E}[S(t) | X(t) = \boldsymbol{x}] &= \int_0^{+\infty} \exp\left(-\int_t^{t+s} \xi(r) \mathrm{d}r\right) \mathrm{d}s \\ &\leq \int_0^{+\infty} \exp\left(-\int_t^{t+s} f(\mathbf{1}^\top \boldsymbol{x}) \mathrm{d}r\right) \mathrm{d}s \\ &= 1/f(\mathbf{1}^\top \boldsymbol{x}). \end{aligned}$$

$$(6)$$

Now, fix $t_0 \geq 0$ and state $X(t_0) = \boldsymbol{x}_0$ such that $\mathbf{1}^\top \boldsymbol{x}_0 \leq a$. Let $t_1 \geq t_0$ be the time when $A(t)$ reaches the value a for the first time after t_0 and let $t_2 = t_1 + S(t_1)$ be the time of the first subsequent configuration change. Using the Markov property, separating the time spent in a during its first entrance from the other contributions to $T_a(t_0)$, and finally using (6) we get:

$$\begin{aligned} \mathbb{E}[&T_a(t_0) | X(t_0) = \boldsymbol{x}_0, X(t_1) = \boldsymbol{x}_1] \\ &= \mathbb{E}[T_a(t_1) | X(t_1) = \boldsymbol{x}_1] \\ &= \mathbb{E}[S(t_1) | X(t_1) = \boldsymbol{x}_1] + \mathbb{E}[T_a(t_2) | X(t_1) = \boldsymbol{x}_1)] \\ &\leq 1/f(a) + \mathbb{E}[T_a(t_2) | X(t_1) = \boldsymbol{x}_1]. \end{aligned}$$

$$(7)$$

Now, we can estimate the term $\mathbb{E}[T_a(t_2) | X(t_1) = \boldsymbol{x}_1]$ by conditioning on the direction of the jump at time t_2, getting

$$\mathbb{E}[T_a(t_2) | X(t_1) = \boldsymbol{x}_1] \leq p_{t_1}^- \theta_a + p_{t_1}^+ (1 - q_{a+1}) \theta_a, \quad (8)$$

where $\theta_a := \sup_{s \geq 0} \max_{\mathbf{1}^\top \boldsymbol{x}_0 \leq a} \mathbb{E}[T_a(s) \mid X(s) = \boldsymbol{x}_0]$. Inserting (8) into (7), we get

$$\mathbb{E}[T_a(t_0) | X(t_0) = \boldsymbol{x}_0] \leq \frac{1}{f(a)} + p_{t_1}^- \theta_a + p_{t_1}^+ (1 - q_{a+1}) \theta_a. \quad (9)$$

For $a = 0$, we have $0 = p_{t_1}^- = 1 - p_{t_1}^+$, so that Lemma 1 gives

$$\mathbb{E}[T_0(t_0) | X(t_0) = \mathbf{0}] \leq \frac{1}{f(0)} + (1 - f_1)\theta_0 \leq \frac{1}{f(0)} + \frac{1 - \beta}{\beta} \theta_0.$$

Taking the supremum over $t_0 \geq 0$ in the above, we get that $\theta_0 \leq 1/f(0) + \theta_0(1 - \beta)/\beta$, which yields the bound

$$\theta_0 \leq \frac{\beta}{2\beta - 1} \frac{1}{f(0)}. \qquad (10)$$

For $1 \le a < n$, applying (3) to the right-hand side of (9) gives

$$\mathbb{E}[T_a(t_0) \mid X(t_0) = \boldsymbol{x}_0] \le \frac{1}{f(a)} + (1 - \beta)\theta_a + \beta\frac{1 - \beta}{\beta}\theta_a$$
$$= \frac{1}{f(a)} + 2(1 - \beta)\theta_a.$$

Since this holds true for every $t_0 \ge 0$ and $\boldsymbol{x}_0 \in \{0, 1\}^n$ such that $\mathbf{1}^\top \boldsymbol{x}_0 \le a$, we get $\theta_a \le 1/f(a) + 2(1 - \beta)\theta_a$, yielding

$$\theta_a \le \frac{1}{2\beta - 1}\frac{1}{f(a)}. \tag{11}$$

Since $\mathbb{E}[T_a \mid X(0) = \mathbf{0}] \le \theta_a$, (10)–(11) yield the claim. ∎

We can now prove the fundamental result of this section.

Theorem 1: Let $(\mathcal{G}, \beta, U(t))$ be a controlled evolutionary dynamics with initial configuration $X(0) = \mathbf{0}$. Then, for every $f : \{0, \ldots, n\} \to \mathbb{R}_+$ such that (4) holds true, the expected spreading time satisfies

$$\mathbb{E}_{\mathbf{0}}[T] \le \frac{\beta}{(2\beta - 1)f(0)} + \frac{1}{2\beta - 1}\sum_{a=1}^{n-1}\frac{1}{f(a)}. \tag{12}$$

Proof: It follows from Proposition 1 that

$$\mathbb{E}_{\mathbf{0}}[T] = \mathbb{E}\left[\sum_{a=0}^{n-1} T_a \mid X(0) = \mathbf{0}\right]$$
$$\le \frac{\beta}{(2\beta - 1)f(0)} + \frac{1}{2\beta - 1}\sum_{a=1}^{n-1}\frac{1}{f(a)},$$

thus proving the claim. ∎

As special case, for constant policies, we have the following.

Corollary 1: Consider a controlled evolutionary dynamics $(\mathcal{G}, \beta, \boldsymbol{u})$ under constant control policy $U(t) = \boldsymbol{u}$ with initial configuration $X(0) = \mathbf{0}$ and let ϕ be the minimum conductance profile of $\mathcal{G}$. Then, the expected spreading time satisfies

$$\mathbb{E}_{\mathbf{0}}[T] \le \frac{\beta}{(2\beta - 1)\mathbf{1}^\top \boldsymbol{u}} + \frac{1}{2\beta - 1}\sum_{a=1}^{n-1}\frac{1}{\phi(a)}. \tag{13}$$

Proof: For every $1 \le a < n$, and for all subsets $S \subset \mathcal{V}$ with $|S| = a$, we have $\phi(a) \le \zeta(S)$. Hence, if $A(t) = a$, then $B(t) + C(t) \ge \phi(a)$. If $A(t) = 0$, then $C(t) = \mathbf{1}^\top \boldsymbol{u}$. We obtain estimation (13) by applying Theorem 1 with $f(a) = \phi(a)$, for $a \ne 0$, and $f(0) = \mathbf{1}^\top \boldsymbol{u}$. ∎

Our performance guarantees can be easily extended to scenarios that include heterogeneous evolutionary advantage and conflicts with no winners, as we discuss in the following. However, to keep the presentation and the discussion of the results as simple as possible throughout the paper, we opted for keeping the simpler formulation described in Section II.

Remark 3: If the evolutionary advantage is heterogeneous across links (and/or time-varying), but always within the range $(1/2, 1]$, the result in Proposition 1 (and, consequently, in Theorem 1) can be obtained using a uniform lower bound on β, which may add a conservative multiplicative factor to the bound, but would not change its possible dependence on the network size. Similarly, if a conflict ends without any winner (i.e., with no state change) with probability $\psi \in [0, 1)$, a multiplicative constant $1/(1 - \psi)$ is added to the bounds.

B. Fundamental Performance Limitation

We present a series of results illustrating fundamental limitations of the controlled evolutionary dynamics. We start by reporting two monotonicity properties, which will be instrumental to our main results.

Lemma 2: Let $(\mathcal{G}, \beta, U(t))$ and $(\mathcal{G}, \gamma, U(t))$ be two controlled evolutionary dynamics and let $X(t)$ and $Y(t)$ be the corresponding Markov processes with $X(0) = \boldsymbol{x}_0$ and $Y(0) = \boldsymbol{y}_0$. If $\beta \le \gamma$ and $\boldsymbol{x}_0 \le \boldsymbol{y}_0$, then $\mathbb{E}_{\boldsymbol{y}_0}[T_Y] \le \mathbb{E}_{\boldsymbol{x}_0}[T_X]$ and $\mathbb{E}_{\boldsymbol{y}_0}[J_Y] \le \mathbb{E}_{\boldsymbol{x}_0}[J_X]$, where the subscripts of T and J refer to the processes they are associated with.

Proof: See Appendix B. ∎

Lemma 3: Let $(\mathcal{G}, 1, U(t))$ and $(\mathcal{G}, 1, 0)$ be two controlled evolutionary dynamics and let $X(t)$ and $Y(t)$ be the corresponding Markov processes with $X(0) = \boldsymbol{x}_0$ and $Y(0) = \boldsymbol{y}_0$. If $\boldsymbol{x}_0 \le \boldsymbol{y}_0$ and $(\mathbf{1} - \boldsymbol{y}_0)^\top U(t) = 0$, for every $t \ge 0$, then, $\mathbb{E}_{\boldsymbol{y}_0}[T_Y] \le \mathbb{E}_{\boldsymbol{x}_0}[T_X]$, where the subscripts of T and J refer to the processes they are associated with.

Proof: See Appendix C. ∎

Remark 4: These results have straightforward consequences. First, Lemma 2 implies that any lower bound on the expected spreading time obtained for the case when $\beta = 1$ will automatically yield a lower bound for every value of β. Second, Lemma 3 implies that, if $\beta = 1$, then it is always possible to establish a lower bound on the expected spreading time by considering an uncontrolled process with initial condition equal to 1 in all the nodes in which the exogenous control is exerted. We wish to emphasize that controlled evolutionary dynamics with $\beta < 1$ do not, in general, enjoy this property.

The previous results motivate a deeper analysis of the controlled evolutionary dynamics with $\beta = 1$. In this case, if the initial configuration is such that $\mathbf{1}^\top X(0) = k$, the Markov process $X(t)$, will always undergo exactly $n - k$ transitions before being absorbed in the all-1 configuration. Indeed, for $\beta = 1$, (1) shows that $\lambda_i^-(\boldsymbol{x}, t) = 0$ for every $\boldsymbol{x}$ in $\{0, 1\}^n$, i in $\mathcal{V}$, and $t \ge 0$. Hence, the process can only undergo transitions where the number of 1's is increasing, either driven by the spreading mechanism or by the exogenous control.

Let $0 = T_k < T_{k+1} < \cdots < T_n = T$ be the random times at which these $n - k$ jumps occur and denote by $X_h = X(T_h^+)$ the configurations of the process after the corresponding jumps, so that $X_k = X(0)$ and $X_n = \mathbf{1}$. Hence, T_h is the time of the jump from X_{h-1} to X_h. Finally, let B_h be the corresponding values for the boundary of the process when in configuration X_h, that is $B_h = \zeta(X_h)$. This time process can be described recursively as follows. We first let σ_h to be the σ-algebra generated by T_h and by the process $X(t)$ for $t \le T_h$. Given σ_{h-1} we consider two independent r.v.s t_h^{s} and t_h^{c} whose distribution functions are given, respectively, by

$$\mathbb{P}[t_h^{\mathrm{s}} \ge t \mid \sigma_{h-1}] = \exp(-B_{h-1}t), \tag{14a}$$

$$\mathbb{P}[t_h^{\mathrm{c}} \ge t \mid \sigma_{h-1}] = \exp\left(-\int_{T_{h-1}}^{T_{h-1}+t} C(s)\mathrm{d}s\right). \tag{14b}$$

We then put $t_h = \min\{t_h^{\mathrm{s}}, t_h^{\mathrm{c}}\}$. The interpretation of t_h is the following: from X_{h-1} the system can evolve either through the spreading mechanisms or through the exogenous

control exerted in a node currently in state 0; t_h^{s} and t_h^{c} are exponentially distributed r.v.s that model the random times at which the two phenomena would independently take place, according to their definition. According to the properties of Markov processes [22], the minimum of these two r.v.s models the time to wait for the next jump, i.e., $T_h = T_{h-1} + t_h$. Define

$$c_h := \int_{T_{h-1}}^{T_h} C(s)\mathrm{d}s \qquad (15)$$

to be the effective control rate exerted during the interval $[T_{h-1}, T_h]$. The following result shows that there is an exact algebraic relation that ties the average length of the interval $[T_{h-1}, T_h]$, the average effective control rate in this interval c_h, and the active boundary after the $(h-1)$-th jump B_{h-1}.

Proposition 2: For every $k \le h \le n$, it holds

$$\mathbb{E}[T_h - T_{h-1} \,|\, \sigma_{h-1}] = \frac{1 - \mathbb{E}[c_h \,|\, \sigma_{h-1}]}{B_{h-1}}, \qquad (16a)$$

$$\mathbb{E}[c_h \,|\, \sigma_{h-1}] = \mathbb{P}[t_h^{\mathrm{c}} < t_h^{\mathrm{s}} \,|\, \sigma_{h-1}]. \qquad (16b)$$

Proof: First, observe that $T_h - T_{h-1} = t_h = \min\{t_h^{\mathrm{s}}, t_h^{\mathrm{c}}\}$. Using the property of the minimum of two independent exponentially distributed r.v.s, we note that

$$\mathbb{P}[t_h \ge t \,|\, \sigma_{h-1}] = \exp\left(-\int_{T_{h-1}}^{T_{h-1}+t} \big(B_{h-1} + C(s)\big)\mathrm{d}s\right), \qquad (17)$$

which yields

$$\frac{d}{dt}\mathbb{P}[t_h \ge t \,|\, \sigma_{h-1}] = \big(B_{h-1} + C(T_{h-1}+t)\big)\mathbb{P}[t_h \ge t \,|\, \sigma_{h-1}].$$

From this and the fact that T_{h-1} and $\{C(s) \,|\, s \in [T_{h-1}, T_h]\}$ are both σ_{h-1}-measurable, we have

$$
\begin{aligned}
\mathbb{E}[c_h | \sigma_{h-1}] &= \mathbb{E}\Big[\int_{T_{h-1}}^{+\infty} C(s)\mathbb{1}_{[T_{h-1},T_h]}(s)\mathrm{d}s \,|\, \sigma_{h-1}\Big] \\
&= \int_{T_{h-1}}^{+\infty} C(s)\mathbb{E}\big[\mathbb{1}_{[T_{h-1},T_h]}(s) \,|\, \sigma_{h-1}\big]\,\mathrm{d}s \\
&= \int_{T_{h-1}}^{+\infty} C(s)\mathbb{P}[t_h \ge s - T_{h-1} \,|\, \sigma_{h-1}]\mathrm{d}s \\
&= \int_{T_{h-1}}^{+\infty} \frac{d}{ds}\mathbb{P}[t_h \ge s - T_{h-1} \,|\, \sigma_{h-1}]\mathrm{d}s \\
&\quad - B_{h-1}\int_{T_{h-1}}^{+\infty} \mathbb{P}[t_h \ge s - T_{h-1} \,|\, \sigma_{h-1}]\mathrm{d}s \\
&= 1 - B_{h-1}\mathbb{E}[t_h \,|\, \sigma_{h-1}],
\end{aligned}
\qquad (18)
$$

which yields (16a). On the other hand, it follows from (18), (17), and the fact that t_h^{c} and t_h^{s} are conditionally independent given σ_{h-1} with conditional marginal distributions as in (14),

that

$$
\begin{aligned}
\mathbb{E}[c_h \,|\, \sigma_{h-1}] &= 1 - B_{h-1}\mathbb{E}[t_h \,|\, \sigma_{h-1}] \\
&= 1 - B_{h-1}\int_0^{+\infty} \mathbb{P}[t_h \ge t \,|\, \sigma_{h-1}]\mathrm{d}t \\
&= 1 - B_{h-1}\int_0^{+\infty} e^{-\int_{T_{h-1}}^{T_{h-1}+t}(B_{h-1}+C(s))\mathrm{d}s}\mathrm{d}t \\
&= \int_0^{+\infty} B_{h-1}e^{-B_{h-1}t}\Big(1 - e^{-\int_{T_{h-1}}^{T_{h-1}+t}C(s)\mathrm{d}s}\Big)\mathrm{d}t \\
&= \mathbb{P}\big(t_h^{\mathrm{s}} - t_h^{\mathrm{c}} > 0 \,|\, \sigma_{h-1}\big) \\
&= \mathbb{P}\big(t_h^{\mathrm{c}} < t_h^{\mathrm{s}} \,|\, \sigma_{h-1}\big),
\end{aligned}
$$

thus proving (16b). ∎

If we define the control rate exerted during the interval $[T_{h-1}, T_h]$ as

$$J_h := \int_{T_{h-1}}^{T_h} \mathbf{1}^\top U(s)\mathrm{d}s, \qquad (19)$$

we can derive the following estimation.

Corollary 2: Let $(\mathcal{G}, \beta, U(t))$ be a controlled evolutionary dynamics with initial condition $X(0) = \boldsymbol{x}_0$ such that $\mathbf{1}^\top \boldsymbol{x}_0 = k > 0$. Then,

$$\mathbb{E}_{x_0}[T] \ge \sum_{h=k+1}^{n} \mathbb{E}\Big[\frac{1 - \mathbb{E}[J_h \,|\, \sigma_{h-1}]}{B_{h-1}}\Big], \qquad (20a)$$

$$\mathbb{E}_{x_0}[J] = \sum_{h=k+1}^{n} \mathbb{E}\big[\mathbb{E}[J_h \,|\, \sigma_{h-1}]\big]. \qquad (20b)$$

Proof: Summing and averaging (16a), using that $B_h \ne 0$ for $h = k, \ldots, n-1$ and that $J_h \ge c_h$ (this follows by comparing (15) and (19)), we obtain the former. The latter is obtained decomposing J into the terms in (19). ∎

Remark 5: In the case when $X(0) = \mathbf{0}$, inequality in (20a) with $k = 1$ continues to hold. Since the first jump is necessarily triggered by an exogenous activation, we can make (20a) tighter adding the term

$$\mathbb{E}[t_1] = \int_0^{+\infty} \exp\left(-\int_0^t \mathbf{1}^\top U(s)\mathrm{d}s\right)\mathrm{d}t. \qquad (21)$$

Regarding (20b), we notice that (18) yields $\mathbb{E}[c_1] = \mathbb{E}[J_1] = 1$. Relation (20b) remains an equality if we simply add 1 to the formula with $k = 1$.

The direct applicability of (20) is limited by the fact that the control efforts and the boundary evolutions can not be uncoupled when the averaging operation is taken. For $\mathcal{U}$-controlled evolutionary dynamics, Corollary 2 can be relaxed, yielding more explicit—but conservative—bounds that turn to be useful in case when $\mathcal{U}$ is sufficiently small.

Corollary 3: Let $(\mathcal{G}, \beta, U(t))$ be a $\mathcal{U}$-controlled evolutionary dynamics with initial condition $X(0) = \boldsymbol{x}_0 \le \delta^{(\mathcal{U})}$. Then,

$$\mathbb{E}_{x_0}[T] \ge \sum_{h=|\mathcal{U}|}^{n-1} \frac{1}{\eta(h)}, \qquad (22)$$

where η is defined as in Definition 1. Moreover, if the control is constant equal to $\boldsymbol{u}$ and $\boldsymbol{x}_0 = \boldsymbol{0}$,

$$\mathbb{E}_{\boldsymbol{0}}[T] \geq \frac{1}{\boldsymbol{1}^\top \boldsymbol{u}} + \sum_{h=|\mathcal{U}|}^{n-1} \frac{1}{\eta(h)}. \tag{23}$$

Proof: It follows from Lemmas 2 and 3 that it is sufficient to prove the bound for the controlled evolutionary dynamics $(\mathcal{G}, 1, 0)$ with $Y(0) = \delta^{(\mathcal{U})}$. In this case, (20a) reduces to

$$\mathbb{E}_{\boldsymbol{x}_0}[T] \geq \mathbb{E}_{\delta^{(\mathcal{U})}}[T_Y] \geq \sum_{h=|\mathcal{U}|}^{n-1} \mathbb{E}[1/B_h].$$

Estimation (22) follows from the definition of $\eta(h)$. For the second inequality, consider the random time t_1 corresponding to the first jump of the process. From (21), using that $\boldsymbol{1}^\top \boldsymbol{u}$ is constant, we compute $\mathbb{E}[t_1] = 1/\boldsymbol{1}^\top \boldsymbol{u}$ and applying the previous part with initial condition $X(t_1) \leq \delta^{(\mathcal{U})}$ we obtain the claim. ∎

By comparing the lower bound in (23) with the corresponding upper bound in (13), we conclude that, for constant $\mathcal{U}$-controlled evolutionary dynamics, the control rate $\boldsymbol{1}^\top \boldsymbol{u}$ has a limited effect on the performance. In particular, $\mathbb{E}_{\boldsymbol{0}}[T]$ remains bounded away from 0 even in the limit case $\boldsymbol{1}^\top \boldsymbol{u} \to +\infty$, even though the expected control cost $\mathbb{E}_{\boldsymbol{0}}[J] = (\boldsymbol{1}^\top \boldsymbol{u})\mathbb{E}_{\boldsymbol{0}}[T]$ grows.

Note that (22) and (23) suggest that, together with the network structure, the support $\mathcal{U}$ of the control action may play a key role in achieving suitable spreading performance. In this direction, we propose another bound on the expected spreading time in which $\mathcal{U}$ has a central role.

Corollary 4: Let $(\mathcal{G}, \beta, U(t))$ be a $\mathcal{U}$-controlled evolutionary dynamics with initial condition $X(0) = \boldsymbol{0}$. Then,

$$\mathbb{E}_{\boldsymbol{0}}[T] \cdot \min_{\mathcal{R} \supseteq \mathcal{U}} \zeta(\mathcal{R}) \geq 1.$$

Proof: Using Lemmas 2 and 3, we bound the expected spreading time by considering a controlled evolutionary dynamics $(\mathcal{G}, 1, 0)$ with initial condition $Y(0) = \delta^{(\mathcal{R})}$ where $\mathcal{R} \supseteq \mathcal{U}$. Then, (20a) yields

$$\mathbb{E}_{\boldsymbol{0}}[T] \geq \mathbb{E}_{\delta^{(\mathcal{R})}}[T_Y] \geq \sum_{h=|\mathcal{R}|}^{n-1} \mathbb{E}[1/B_h] \geq 1/\zeta(\mathcal{R}),$$

where the last equality follows from the fact that $B_{|\mathcal{R}|} = \zeta(\mathcal{R})$. Since the above inequality holds true for every $\mathcal{R} \supseteq \mathcal{U}$, we obtain the claim. ∎

Remark 6: Denote now by N^c the r.v. counting the total number of activations due to exerting the control action, more formally in the notation above, $N^c = |\{h = 1, \ldots, n : t_h^c < t_h^s\}|$. Summing over h and averaging the second relation in (16b), we obtain that $\mathbb{E}[N^c] \leq \mathbb{E}[J]$. In other terms, the expected number of activations due to the control action is a lower bound on the control cost.

Remark 7: Lemmas 2 and 3 can be easily extended to the scenarios with heterogeneous evolutionary advantage and conflicts with no winners discussed in Remark 3. Hence, our fundamental performance limitations could be directly applied to these more general scenarios.

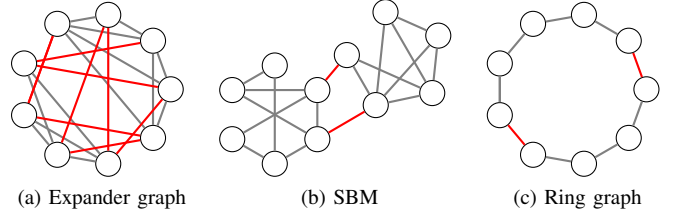(a) Expander graph   (b) SBM   (c) Ring graph

Fig. 2: Graph families analyzed in this paper. The red links represent $\phi(5)$, showing that expander graph are highly connected, while SBMs and rings have thigh bottlenecks.

IV. THREE KEY EXAMPLES

We analyze the behavior and performance of controlled evolutionary dynamics in some key examples of large-scale networks. Formally, we shall consider infinite families of graphs $(\mathcal{G}_n)_{n \in \mathbb{N}}$ parameterized by their order n and focus on three special cases: expanders, SBMs, and ring graphs (see Fig. 2 for an illustration). Expanders and ring graphs (formally defined in the following) constitute extreme cases and, as we will see, they represent benchmarks for topologies that are easy to control and hard to control, respectively. SBMs [25], instead, display an intermediate behavior: no fundamental limitation precludes fast spread, but constant control policies may fail in achieving it, whereas, as we shall see in Section V, feedback control policies may be effective.

In order to get a fair comparison between the three examples, we assume that all the links within each graph $\mathcal{G}_n$ have the same weight (i.e., $W_{ij} = w$, for all $\{i, j\} \in \mathcal{E}$). and that such value $w = \alpha/\Delta$ is inversely proportional to the maximum degree Δ in $\mathcal{G}_n$ as where $\alpha > 0$ is a constant independent from n. This prevents the sum of the weights of all links incident in a single node on a node from blowing up as n grows large.

In the three examples, we consider controlled evolutionary dynamics with initial condition $X(0) = \boldsymbol{0}$. Before presenting the examples, we note that the trivial bound $\eta(h) \leq \alpha h$ substituted in (22) yields, for every $\mathcal{U}$-controlled evolutionary dynamics, the following bound on the expected spreading time

$$\mathbb{E}_{\boldsymbol{0}}[T] \geq \sum_{h=|\mathcal{U}|}^{n-1} \frac{1}{\alpha h} \geq \frac{1}{\alpha} \log \frac{n}{|\mathcal{U}|}. \tag{24}$$

Hence, if $\mathcal{U}$ is assumed to be constant in n, the expected spreading time grows at least *logarithmically* in n. In view of this observation, we say that a graph family $\mathcal{G}_n$ exhibits *fast* spread if there exists a constant $K > 0$ such that $\mathbb{E}_{\boldsymbol{0}}[T] \leq K \ln n$ for every n in $\mathbb{N}$, whereas it exhibits *slow* spread if the above condition is not satisfied, that is, spreading time grows more than logaritmically in the network size. Note that, as a consequence of Remarks 3 and 7, the presence of heterogeneity in the evolutionary advantage β (but always greater than $1/2$) or of conflicts with no winners have no impact on whether a graph family exhibits fast or slow spread, since they may only yield multiplicative constants, which do not change the scaling factor with respect to the network size.

Example 1 (Expander graphs): A family of graphs $\mathcal{G}_n$ is referred to as *expander* if there exists a constant $\gamma > 0$ such
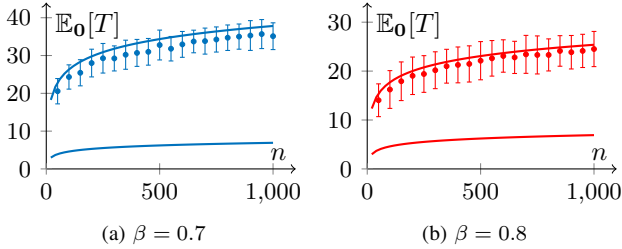
(a) $\beta = 0.7$ (b) $\beta = 0.8$

Fig. 3: Monte Carlo estimation (200 simulations) of the expected spreading time $\mathbb{E}_{\mathbf{0}}[T]$ on complete graphs for different values of n ($\alpha = 1$), with 90% confidence intervals. The two solid curves are the theoretical bounds from (24) and (26).

that, for every n in $\mathbb{N}$, the minimum conductance profile of the graph $\mathcal{G}_n$ satisfies

$$\phi(h) \geq \gamma \cdot \min\{h, n-h\}, \qquad \forall 0 \leq h \leq n. \tag{25}$$

i.e., if all subsets of nodes have weighted boundary at least proportional to their cardinality. Important examples of expander families are complete graphs, as well as, with high probability, Erdős-Rényi (illustrated in Fig. 2a), small-world models, and many scale-free graphs, which are often used to represent well connected real-world systems and social networks.

We now show that expander graph families exhibit fast spread. In fact, applying (25) to the right-hand side of (13) in Corollary 1, we obtain

$$\mathbb{E}_{\mathbf{0}}[T] \leq \frac{\beta}{(2\beta - 1)\mathbf{1}^\top \boldsymbol{u}} + \frac{2\ln(n/2) + 2}{\gamma(2\beta - 1)}. \tag{26}$$

Figure 3 reports Monte Carlo computations of the expected spreading time for complete graphs, together with our analytical bounds. The bound in (26) appears to be tight for the complete graph. In fact, for complete graphs the inequality (4) holds true as an equality when the function $f(h)$ coincides with the minimum conductance profile (since all subsets with h nodes have the same weighted boundary). This suggests that the spreading time can be expected to be close to the upper bound (12) from Theorem 1 when one can determine a function $f(h)$ so that the inequality (4) is close to an equality.

Example 2 (SBMs): SBMs are made of dense (expander) subgraphs linked among each other by few connections (see Fig. 2b). For the sake of simplicity, we consider the case of two subgraphs (termed communities), but our results can be easily generalized. We consider the following implementation. Fixed $0 < c \leq 1/2$, we partition the nodes into two disjoint sets $\mathcal{V} = \mathcal{V}^1 \cup \mathcal{V}^2$, with $n_1 = \lfloor cn \rfloor$ and $n_2 = \lceil (1-c)n \rceil$ nodes, respectively. Nodes in each subset are linked as in Erdős-Rényi random graph, where each link is present with probability $p \in (0, 1]$, independent of the others [31]. On top of this, L links positioned uniformly at random connect nodes belonging to the two different subgraphs.

We now show that SBMs exhibit slow spread, unless the control is exerted on both communities.[1] Assume that the

[1]Since SBMs are random graphs, results will be provided with high probability (w.h.p.), i.e., with probability converging to 1 as n grows.

control is exerted only in $\mathcal{V}^1$, if at all. Then we fix $\mathcal{U} \subseteq \mathcal{V}^1$ and consider any $\mathcal{U}$-controlled evolutionary dynamics $(\mathcal{G}, \beta, U(t))$ on $\mathcal{G}$. Considering the standing assumptions that all nonzero weights are equal to $w = \alpha/\Delta$ and that the maximal degree satisfies the inequality $\Delta \geq n_2 p \geq (1-c)np$ (w.h.p.), we have that $\zeta(\mathcal{V}^1) \leq \alpha L/(1-c)np$. Hence, Corollary 4 yields

$$\mathbb{E}_{\mathbf{0}}[T] \geq \frac{(1-c)np}{\alpha L}, \tag{27}$$

that is, slow spread. We demonstrate that such an estimation is asymptotically order-tight for constant control policies $U(t) = \boldsymbol{u}$ by leveraging Corollary 1. To this aim, we estimate the minimum conductance profile as follows. First, note that (even though SMBs are no expanders) each of the two subgraphs is expander with $\gamma = cnpw/2$ [2]. Using the trivial bound $\Delta \leq n$, we get $\gamma \geq c\alpha p/2$. Given an integer $0 < h < n$, any subset $\mathcal{W} \subseteq \mathcal{V}$ with $|\mathcal{W}| = h$ can be written as $\mathcal{W} = \mathcal{W}^1 \cup \mathcal{W}^2$, with $\mathcal{W}^i \subseteq \mathcal{V}^i$, $|\mathcal{W}^i| = h_i$, and $h_1 + h_2 = h$. Therefore,

$$\phi(h) \geq \frac{c\alpha p}{2} \cdot \begin{cases} \min\{h, n_1 - h\}, & h \leq n_1, \\ \min\{h - n_1, n_2 - h\}, & n_1 < h \leq n_2, \\ \min\{h - n_2, n - h\}, & n_2 < h < n. \end{cases} \tag{28}$$

For $h \in \{n_1, n_2\}$, the previous bound reduces to the trivial inequality $\phi(h) \geq 0$. However, the presence of L links between the two subgraphs ensures that $\phi(h) \geq L\alpha/n$, for $h \in \{n_1, n_2\}$. Combining this with (28), using bounds on the harmonic series, and applying Corollary 1, we finally obtain

$$\mathbb{E}_{\mathbf{0}}[T] \leq \frac{1}{2\beta - 1}\left(\frac{2}{L\alpha}n + \frac{12}{c\alpha p}\ln\frac{en}{2}\right) + \frac{\beta}{(2\beta - 1)\mathbf{1}^\top \boldsymbol{u}}. \tag{29}$$

This, together with (27), shows that the expected spreading time grows linearly with the network size n, as confirmed by the simulations (blue curves) in Fig. 4.

Note that, if we allow the control nodes in both communities, we instead obtain a logarithmic growth lower bound on $\mathbb{E}_{\mathbf{0}}[T]$. These considerations extend to more general SBMs composed of many dense subgraphs. They lead to the conclusion that to obtain a fast spread in these graphs using constant control policies or, more generally, control policies with constant support $\mathcal{U}$, the set $\mathcal{U}$ must necessarily have a nonempty intersection with all the communities. This is a drawback in practical applications as, not only it may require the cardinality of $\mathcal{U}$ to be large, but also needs precise a-priori information on the network topology to suitably position the control nodes.

Example 3 (Ring graphs): We now consider the family $\mathcal{G}_n$ of undirected ring graphs, whereby, for every n in $\mathbb{N}$, $\mathcal{G}_n$ has node set $\mathcal{V}_n = \{1, \ldots, n\}$ and every node i is connected to nodes $i - 1$ and $i + 1$ (modulo n) by a link of weight $w = \alpha/2$ (see Fig. 2c). First, we observe that rings are instances of poorly connected graphs and are not expander, since $\phi_n(h) = \alpha$ for all $0 \leq h \leq n$. We bound the performance comparing our process $X(t)$ with the Markov process $Y(t)$ associated with the controlled evolutionary dynamics $(\mathcal{G}, 1, U(t))$ with $Y(0) = \mathbf{0}$, since Lemma 2 ensures that the expected spreading time and cost verify $\mathbb{E}_{\mathbf{0}}[T_X] \geq \mathbb{E}_{\mathbf{0}}[J_Y]$ and $\mathbb{E}_{\mathbf{0}}[J_X] \geq \mathbb{E}_{\mathbf{0}}[J_Y]$, where the subscripts denote the process.

We bound the performance of $Y(t)$ applying Corollary 2. To this aim, we start with two simple remarks. First, for every $\mathcal{W} \subseteq \mathcal{V}$, we have $\zeta(\mathcal{W}) \leq \alpha|\mathcal{W}|$. Second, the spreading mechanism cannot increase the boundary $B(t)$: it decreases by α, if the two neighbors of the node that changes its state are both state-1 nodes, or, otherwise, it remains the same. Moreover, considering that every jump due to exerting a control action can increase the boundary by α, we have, for every h, $B_h \leq \alpha N^{\mathrm{c}}$. where N^{c} is the total number of control activations. Consider now the event $E := \{N^{\mathrm{c}} \leq 2\mathbb{E}[N^{\mathrm{c}}]\}$. Using the inequality above and Markov inequality we obtain

$$
\begin{aligned}
\mathbb{E}\left[\frac{1 - \mathbb{E}[J_h \mid \sigma_{h-1}]}{B_{h-1}}\right] &\geq \mathbb{E}\left[\frac{1 - \mathbb{E}[J_h \mid \sigma_{h-1}]}{\alpha N^{\mathrm{c}}}\mathbb{1}_E\right] \\
&\geq \frac{\mathbb{E}[1 - \mathbb{E}[J_h \mid \sigma_{h-1}]]}{2\alpha\mathbb{E}[N^{\mathrm{c}}]}\mathbb{P}[E] \\
&\geq \frac{1 - \mathbb{E}[\mathbb{E}[J_h \mid \sigma_{h-1}]]}{4\alpha\mathbb{E}[N^{\mathrm{c}}]}.
\end{aligned}
$$

From Corollary 2 and Remark 6, we finally obtain

$$
\mathbb{E}_{\mathbf{0}}[T_X] \geq \mathbb{E}_{\mathbf{0}}[T_Y] \geq \frac{n - \mathbb{E}_{\mathbf{0}}[J_Y]}{4\mathbb{E}_{\mathbf{0}}[J_Y]} \geq \frac{1}{4\mathbb{E}_{\mathbf{0}}[J_X]}n - \frac{1}{4}.
$$

This allows us to conclude that the family of ring graphs exhibits slow spread, unless adopting a control policy whose cost $\mathbb{E}_{\mathbf{0}}[J]$ blows up with the network size.

V. FEEDBACK CONTROL POLICIES

The examples discussed in the previous section revealed different behavior of the controlled evolutionary dynamics, depending on the network structure. Specifically, we characterized families of network topologies that are *easy-to-control* (e.g., expander graphs), where simple constant control policies can guarantee fast spread, and others that are *hard-to-control* (e.g., rings) where fast spread is not possible under any control policy. Furthermore, we found that there are networks (e.g., SBMs) that belong to neither of these two classes, for which the way the control is exerted plays a key role in determining the performance. In such scenarios, the use of feedback control policies may allow one to achieve fast spread with reasonable control efforts, while simple open-loop control policies fail.

We observe that, in the presence of evolutionary advantage for the novel states ($\beta > 1/2$), the evolutionary dynamics naturally induces a positive feedback loop enhancing the spreading process. However, the presence of bottlenecks in the graph topology (as in the example of SBMs) may slow down such a diffusion process. Hence, the objective of our control policies will be to enforce such a positive feedback whenever the process enters in a bottleneck. Technically, our policies rely on the following two considerations: i) it is of no use to exert control on state-1 nodes: avoiding this yields $C(t) = \mathbf{1}^\top U(t)$; and ii) in order to apply Theorem 1, the feedback control law $U(t) = \nu(X(t))$ must enforce a relation like (4). This suggests to consider feedback laws whose instantaneous rate $\nu(X(t))$ only depends on $X(t)$ through $A(t)$ and $B(t)$. Moreover, (4) suggests that a valuable control policy should compensate for the boundary when $B(t)$ becomes too small, in order to keep $B(t) + C(t)$ always sufficiently large.

Based on these considerations, we design a feedback control policy that is determined by two ingredients: a *target function* $\iota : \{0,1\}^n \to \{1, \dots, n\}$, which selects the node in which the control is exerted (as a function of the current configuration); and a *rate function* $\mu : \{0, \dots, n\} \times \mathbb{R}_+ \to \mathbb{R}_+$, which determines the control rate (as a function of the number of 1's in the current configuration, that is, $A(t)$, and of the boundary $B(t)$). We assume the target function to be any function $\iota(\boldsymbol{x})$ such that $x_{\iota(\boldsymbol{x})} = 0$ if $\boldsymbol{x} \neq \mathbf{1}$ (i.e., we always choose to control a node that is currently in state 0). The rate function is assumed to have the following form. We fix a parameter $K > 0$ and we put

$$
\mu(a, b) = \begin{cases} K - b & \text{if } a < n \text{ and } b < K, \\ 0 & \text{otherwise.} \end{cases} \tag{30}
$$

Finally, the feedback control law $\nu(X(t))$ is given by:

$$
\nu_i(X(t)) = \begin{cases} \mu(A(t), B(t)) & \text{if } i = \iota(X(t)), \\ 0 & \text{else.} \end{cases} \tag{31}
$$

Briefly, in the selected node $i = \iota(X(t))$ we exert a control action with rate μ as in (30), while in all other nodes no control is exerted.

To analyze the expected spreading time under this feedback control policy, it is useful to consider the following *floor* version of the conductance profile.

Definition 2: For $K \geq 0$, the *K-floor conductance profile* of an undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ of order $|\mathcal{V}| = n$ is the function $\phi_K : \{1, \dots, n-1\} \to \mathbb{R}_+$ defined as

$$
\phi_K(a) := \max\{\phi(a), K\} = \max\left\{\min_{\mathcal{W} \subset \mathcal{V}, |\mathcal{W}| = a} \zeta(\mathcal{W}), K\right\}.
$$

From Theorem 1 we establish the following upper bounds on the expected spreading time and on the expected cost of the feedback control policy in (31).

Corollary 5: Consider the feedback controlled evolutionary dynamics $(\mathcal{G}, \beta, \nu)$ with initial condition $X(0) = \mathbf{0}$, where ν follows (31). Then, the expected spreading time verifies

$$
\mathbb{E}_{\mathbf{0}}[T] \leq \frac{\beta}{(2\beta - 1)K} + \frac{1}{2\beta - 1}\sum_{a=1}^{n-1}\frac{1}{\phi_K(a)}, \tag{32}
$$

and the expected control cost is bounded as

$$
\mathbb{E}_{\mathbf{0}}[J] \leq \frac{\beta}{2\beta - 1} + \frac{|\{1 \leq a < n : \phi(a) < K\}|}{2\beta - 1}. \tag{33}
$$

Proof: From (30) it follows that, for all t such that $A(t) < n$, $C(t) = \max\{K - B(t), 0\}$. For such values of t we thus have $B(t) + C(t) = \max\{B(t), K\} \geq \phi_K(A(t))$. The upper bound (32) follows from Theorem 1.

The expected control cost is estimated as follows. First, note that $C(t) \leq K$ for every t. Moreover, if a is such that $\phi(a) >$

K, then $B(t) > K$ and $C(t) = 0$. Using the bound on the time spent by the process $A(t)$ in each state from (5), we conclude

$$
\begin{aligned}
\mathbb{E}_{\mathbf{0}}[J] \leq \quad & K\Big(\mathbb{E}[T_0 | X(0) = \mathbf{0}] \\
& + \sum_{a=1}^{n-1} \mathbb{E}[T_a | X(0) = \mathbf{0}] \mathbb{1}_{[0,K]}(\phi(a)) \Big) \\
= \quad & K\Big(\frac{\beta}{2\beta - 1} \frac{1}{\phi_K(0)} + \frac{1}{2\beta - 1} \sum_{a=1}^{n-1} \frac{\mathbb{1}_{[0,K]}(\phi(a))}{\phi_K(a)} \Big) \\
= \quad & \frac{\beta}{2\beta - 1} + \frac{|\{1 \leq a < n : \phi(a) < K\}|}{2\beta - 1},
\end{aligned}
$$

thus proving (33). ∎

We observe that the two bounds in Corollary 5 depend on the control policy ν through the parameter K and on the network structure through the minimum conductance profile ϕ. Due to the monotonicity of ϕ_K, the bound in (32) is nonincreasing in K, while (33) is nondecreasing in K. This establishes a trade-off behavior between faster spread and higher control cost, paving the way for the formalization of optimization problems to choose the value K best compromising fast spread (large K) and affordable cost (small K).

Example 4 (Fast spread on SBMs): We consider the feedback control policy in (31) on the implementation of SBMs introduced in Example 2, choosing $K < c\alpha p/2$. From (28), we have that $\phi(a) \geq K$ for $a \in \{1, \ldots, n-1\} \setminus \{n_1, n_2\}$, w.h.p. as n grows large. Therefore, $\phi_K(a) = \max\{K, \phi(a)\}$. The application of (32) yields

$$
\mathbb{E}_{\mathbf{0}}[T] \leq \frac{\beta}{(2\beta - 1)K} + \frac{1}{2\beta - 1}\Big(\frac{2}{K} + \frac{12}{c\alpha p}\Big(\ln\Big(\frac{n}{2}\Big) + 1 \Big) \Big),
\tag{34}
$$

where the algebraic computations follow the ones carried on in Section 2. Using (33), we bound $\mathbb{E}_{\mathbf{0}}[J] \leq \frac{2+\beta}{2\beta-1}$. Hence, with a bounded control cost, we achieve an expected spreading time that grows logarithmically in n, that is, fast spread. Figure 4 compares our feedback control policy with constant policies, highlighting the improvement in performance of the former with respect to the latter. We notice that a larger improvement is observed for highly connected communities (i.e., for large values of p). We remark that the control strategy is designed without knowledge of the exact partition of nodes into communities, and even of their size. This provides robustness in real-world situations, where the network is known with uncertainty.

VI. DISCUSSION AND CONCLUSION

We have proposed a continuous-time stochastic dynamical model that allows for an analytical treatment of evolutionary dynamics and incorporates exogenous control input mechanisms. The main contributions of this paper are i) the formulation and a rigorous study of a link-based spreading process incorporating an exogenous control and leading to a nonhomogenous Markov process; ii) the derivation of a set of explicit estimations of the spreading time and the control cost for specific control policies; and iii) the design and analysis of an effective low complexity feedback control policy.

Through our analysis, we characterized three classes of network topologies depending on their controllability. *Topologies*
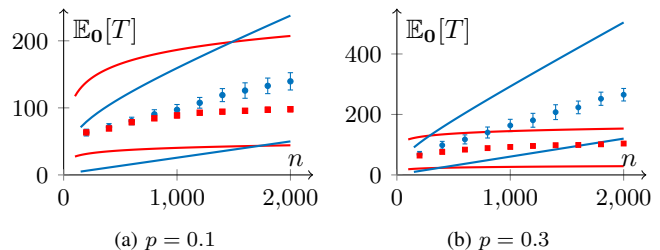

Fig. 4: Monte Carlo estimation (200 runs) and 90% confidence intervals of the expected spreading time $\mathbb{E}_{\mathbf{0}}[T]$ on SBMs with $\alpha = 1$, $\beta = 0.8$, $c = 0.4$, and $L = 5$, feedback control (31) with $K = 1/4$ (red squares) and constant control (blue circles), for two values of p. Red solid curves represent the bounds (22) and (34). Blue solid curves represent the bounds (27) and (29).

easy to control even with constant control policies, which include highly connected structures such as expander graphs. For these topologies, little effort is required to achieve fast spread (expected spreading time logarithmic in the network size), for every constant control policy. *Topologies easy to control only with feedback control policies*, of which stochastic block model are a representative example. For these structures, no fundamental limit precludes fast spread, but constant control policies fail to achieve it. A major contribution of this paper consists in the development of a feedback control policy that remarkably improves the speed of the process, achieving fast spread. *Topologies hard to control under any control policy*, of which the ring graph is a representative element. For these poorly connected topologies, any feasible control policy yields slow spread (spreading time linear in the network size).

The generality of Theorem 1 and the effectiveness of the feedback control policy proposed pave the way for the design of optimal control policies for real-world applications (e.g. introduction of genetically modified mosquitoes [21]). When more information on the network structure is available, we believe that such information might be exploited to design and study targeted introduction policies. For instance, inspired by targeted vaccination strategies [6], the introduction of the novel states might be prioritized in nodes with high network centrality. Other avenues of future research include the use of the technical tools developed in this paper to study control policies on other dynamical processes on networks, such as opinion dynamics, diffusion of information, and epidemics.

REFERENCES

[1] L. Zino, G. Como, and F. Fagnani, "Fast diffusion of a mutant in controlled evolutionary dynamics," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11 908–11 913, 2017, 20th IFAC World Congress.

[2] A. Ganesh, L. Massoulié, and D. Towsley, "The effect of network topology on the spread of epidemics," *Proceedings 24th IEEE INFOCOM*, vol. 2, pp. 1455–1466, 2005.

[3] M. Draief, "Epidemic processes on complex networks: The effect of topology on the spread of epidemics," *Physica A*, vol. 363, no. 1, pp. 120–131, 2006.

[4] W. Mei, S. Mohagheghi, S. Zampieri, and F. Bullo, "On the dynamics of deterministic epidemic propagation over networks," *Annual Reviews in Control*, vol. 44, pp. 116–128, 2017.

[5] F. Fagnani and L. Zino, "Time to extinction for the SIS epidemic model: New bounds on the tail probabilities," *IEEE Transactions on Network Science and Engineering*, vol. 6, no. 1, pp. 74–81, 2019.

[6] F. Chung, P. Horn, and A. Tsiatas, "Distributing antidote using pagerank vectors," *Internet Mathematics*, vol. 6, no. 2, pp. 237–254, 2009.

[7] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, "How to distribute antidote to control epidemics," *Random Structures and Algorithms*, vol. 37, no. 2, pp. 204–222, 2010.

[8] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, "An efficient curing policy for epidemics on graphs," *IEEE Transactions on Network Science and Engineering*, vol. 1, no. 2, pp. 67–75, 2014.

[9] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Analysis and control of epidemics: a survey of spreading processes on complex networks," *IEEE Control Systems Magazine*, vol. 36, no. 1, pp. 26–46, 2016.

[10] T. M. Ligget, *Interacting particle systems*. New York, NY, USA: Springer-Verlag, 1985.

[11] A. Montanari and A. Saberi, "The spread of innovations in social networks," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 107, no. 47, pp. 20 196–20 201, 2010.

[12] H. Peyton Young, "The dynamics of social innovation," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, no. 4, pp. 21 285–21 291, 2011.

[13] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proceedings of the 9th ACM SIGKDD*, Aug. 2003, pp. 137–146.

[14] G. Como, S. Durand, and F. Fagnani, "Optimal targeting in super-modular games," *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6366–6380, 2022.

[15] J. R. Riehl and M. Cao, "Towards optimal control of evolutionary games on networks," *IEEE Transactions on Automatic Control*, vol. 62, no. 1, pp. 458–462, 2017.

[16] E. Lieberman, C. Hauert, and M. A. Nowak, "Evolutionary dynamics on graphs," *Nature*, vol. 433, 2005.

[17] J. Rychtář and B. Stadler, "Evolutionary dynamics on small-world networks," *International Journal of Computational and Mathematical Sciences*, vol. 2, no. 1, pp. 1–4, 2008.

[18] H. Ohtsuki and M. A. Nowak, "Evolutionary games on cycles," *Proceedings Biological Sciences*, vol. 273, no. 1598, pp. 2249–2256, 2006.

[19] M. Broom, J. Rychtář, and B. T. Stadler, "Evolutionary dynamics on graphs - the effect of graph structure and initial placement on mutant spread," *Journal of Statistical Theory and Practice*, vol. 5, no. 3, pp. 369–381, 2011.

[20] B. Allen, G. Lippner, Y.-T. Chen, B. Fotouhi, N. Momeni, S.-T. Yau, and M. A. Nowak, "Evolutionary dynamics on any population structure," *Nature*, vol. 544, no. 7649, pp. 227–230, 2017.

[21] L. Zino, G. Como, and F. Fagnani, "Controlling evolutionary dynamics in networks: A case study," *IFAC-PapersOnLine*, vol. 51, no. 23, pp. 349–354, 2018, 7th IFAC Workshop on Distributed Estimation and Control in Networked Systems NECSYS 2018.

[22] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. Providence, RI, USA: American Mathematical Society, 2009.

[23] M. Mihail, "Conductance and convergence of markov chains-a combinatorial treatment of expanders," in *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, Oct 1989, pp. 526–531.

[24] G. Roberts and J. Rosenthal, "Quantitative bounds for convergence rates of continuous time markov processes," *Electronic Journal of Probability*, vol. 1, 1996.

[25] S. Fortunato and D. Hric, "Community detection in networks: A user guide," *Physics Reports*, vol. 659, pp. 1–44, 2016.

[26] A. Buchman *et al.*, "Broad dengue neutralization in mosquitoes expressing an engineered antibody," *PLoS Pathogens*, vol. 16, no. 1, pp. 1–20, 2020.

[27] I. Ferreira, "The probability of survival for the biased voter model in a random environment," *Stochastic Processes and Their Applications*, vol. 34, no. 1, pp. 25–38, 1990.

[28] D. Acemoğlu, G. Como, F. Fagnani, and A. Ozdaglar, "Opinion fluctuations and disagreement in social networks," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 1–27, 2013.

[29] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. part i," *Annual Reviews in Control*, vol. 43, pp. 65–79, 2017.

[30] S. Goel, R. Montenegro, and P. Tetali, "Mixing time bounds via the spectral profile," *Electronic Journal of Probability*, vol. 11, pp. 1–26, 2006.

[31] B. Bollobas, *Random Graphs*. Cambridge, UK: Cambridge University Press, 2001.

APPENDIX A
PROOF OF LEMMA 1

For a given time $t_0 \in \mathbb{R}_+$, let X_k, A_k, B_k, C_k be the discrete-time processes of $X(t)$, $A(t)$, $B(t)$ $C(t)$, respectively, starting from time t_0^+, with $X_0 = X(t_0) = \boldsymbol{x}$. Clearly, $\mathbb{P}[A(t) \geq a, \, t \geq t_0 \,|\, X(t_0) = \boldsymbol{x}] = \mathbb{P}[A_k \geq a, \, k \geq 0 \,|\, X_0 = \boldsymbol{x}]$. It follows from (3) that the increase and decrease transition probabilities of A_k conditioned to X_k at time k, satisfy $p_k^-(a|\boldsymbol{x}) \leq 1 - \beta$ and $p_k^+(a|\boldsymbol{x}) \geq \beta$, for $a \notin \{0, n\}$, while $p_k^+(0|0) = 1$ for all k, and $a = n$ is an absorbing state.

Consider now a discrete-time birth-and-death chain $\tilde{A}_k$ with state space $\{0, \ldots, n\}$ and transition probabilities $p^+(0) = 1$, $p^-(0) = p^-(n) = p^+(0) = 0$, $p^+(a) = \beta$, and $p^-(a) = 1 - \beta$, for $a = \{1, \ldots, n-1\}$, with $\tilde{A}_0 = a$. A standard argument allows us to couple the two processes A_k and $\tilde{A}_k$ in such a way that $A_k \geq \tilde{A}_k$ for every k, yielding $\mathbb{P}[A_k \geq a, \, k \geq 0 \,|\, X_0 = \boldsymbol{x}] \geq \mathbb{P}[\tilde{A}_k \geq a, \, k \geq 0]$. On the other hand, a direct computation for the birth-and-death chain $\tilde{A}_k$ implies that

$$\mathbb{P}[\tilde{A}_k \geq a, \, \forall k] = \frac{1 - \left(\frac{1-\beta}{\beta}\right)}{1 - \left(\frac{1-\beta}{\beta}\right)^{n-a+1}} \geq \frac{2\beta - 1}{\beta},$$

which yields the claim, since t_0 can be chosen arbitrarily. ∎

APPENDIX B
PROOF OF LEMMA 2

First, we consider $\beta = \gamma$ and $\boldsymbol{x}_0 \leq \boldsymbol{y}_0$. We define the coupled process $Z(t) = (X(t), Y(t))$ on the state space $\{0, 1\}^n \times \{0, 1\}^n$, with initial condition $Z(0) = (\boldsymbol{x}_0, \boldsymbol{y}_0)$, associated with the same graph $\mathcal{G}$. The coupling mechanism is the following. Each link $\{i, j\}$ is equipped with an independent rate-W_{ij} Poisson clock. When the clock associated with link $\{i, j\}$ ticks, the spreading mechanism of a controlled evolutionary dynamics acts on that link for both $X(t)$ and $Y(t)$ and, if a conflict occurs in both processes, then the outcome is the same. Each node i is equipped with a Poisson clock with rate $U_i(t)$ associated with the exogenous control in node i. When the clock associated with node i ticks, then both X_i and Y_i are set to 1. The marginals of $X(t)$ and $Y(t)$ coincide with the distribution of a controlled evolutionary dynamics $(\mathcal{G}, \beta, U(t))$ with initial condition $X(0) = \boldsymbol{x}_0$ and $Y(0) = \boldsymbol{y}_0$, respectively.

We show that, under this coupling, $Y(t) \geq X(t)$, for every $t \geq 0$. Since $\boldsymbol{x}_0 \leq \boldsymbol{y}_0$, the inequality holds for $t = 0$. Then, we prove that any transition of $Z(t)$ preserves the inequality. Assume that a transition occurs at time t. If the transition is triggered by the spreading mechanism on link $\{i, j\}$, it means that a conflict occurs. Since the inequality holds before the transition, only three scenarios are possible:

a) $X_i(t^-) = Y_i(t^-) = 0$, $X_j(t^-) = 0$, and $Y_j(t^-) = 1$;
b) $X_i(t^-) = Y_i(t^-) = 0$, and $X_j(t^-) = Y_j(t^-) = 1$;
c) $X_i(t^-) = 0$, $Y_i(t^-) = 1$, and $X_j(t^-) = Y_j(t^-) = 1$.

In all these scenarios, it is straightforward to verify that, after the transition, the inequality $Y(t^+) \geq X(t^+)$ holds, irrespective of the state that wins the conflict. If instead the transition is triggered by the control mechanism on node i, then it yields $X_i(t^+) = Y_i(t^+) = 1$, preserving the inequality.

The proof for $\beta < \gamma$ and $\boldsymbol{x}_0 = \boldsymbol{y}_0$ follows a similar argument. We define the coupled process $Z(t) = (X(t), Y(t))$ in which each link $\{i, j\}$ is equipped with an independent Poisson clock with rate W_{ij}. When the clock associated with link $\{i, j\}$ ticks, the spreading mechanism acts on that link for both $X(t)$ and $Y(t)$. If a conflict occurs in only one of the two processes, then it is solved as in a standard controlled evolutionary dynamics with probability for the novel state to win the conflict equal to β for $X(t)$ and γ for $Y(t)$, respectively. If the conflict occurs in both processes, then with probability β the novel state wins in both $X(t)$ and $Y(t)$, with probability $\gamma - \beta$ it wins only in $Y(t)$, and with probability $1 - \gamma$ the novel state loses in both $X(t)$ and $Y(t)$. Each node i is given an nonhomogeneous Poisson clock with rate $U_i(t)$. When the clock associated with node i ticks, then both X_i and Y_i turn to 1. We immediately deduce that the two marginals $X(t)$ and $Y(t)$ are controlled evolutionary dynamics $(\mathcal{G}, \beta, U(t))$ and $(\mathcal{G}, \gamma, U(t))$, respectively, with the same initial condition $X(0) = Y(0) = \boldsymbol{x}_0 = \boldsymbol{y}_0$.

Under this coupling, the inequality $Y(t) \geq X(t)$ holds true for every $t \geq 0$. In fact, at $t = 0$ it is verified, since $X(0) = Y(0)$. Then, the analysis of all possible transitions and their effect on the inequality is performed similar to above and is omitted due to space constraints. Finally, the case $\beta < \gamma$ and $\boldsymbol{x}_0 < \boldsymbol{y}_0$ is obtained by combining the two couplings above.

Existence of a coupling $Z(t) = (X(t), Y(t))$ such that $Y(t) \geq X(t)$ for every $t \geq 0$ implies the stochastic domination $Y(t) \succeq X(t)$ [3], yielding $\mathbb{E}_{\boldsymbol{y}_0}[T_Y] \leq \mathbb{E}_{\boldsymbol{x}_0}[T_X]$. Finally, Since $U(t) \geq 0$, we conclude that

$$\mathbb{E}_{\boldsymbol{x}_0}[J_X] = \int_0^{\mathbb{E}_{\boldsymbol{x}_0}[T_X]} U(t)\mathrm{d}t \geq \int_0^{\mathbb{E}_{\boldsymbol{y}_0}[T_Y]} U(t)\mathrm{d}t = \mathbb{E}_{\boldsymbol{y}_0}[J_X],$$

thus completing the proof. ∎

APPENDIX C
PROOF OF LEMMA 3

We define a coupled process $Z(t) = (X(t), Y(t))$ on the state space $\{0, 1\}^n \times \{0, 1\}^n$, with initial condition $Z(0) = (\boldsymbol{x}_0, \boldsymbol{y}_0)$. Here, the marginal distributions of $X(t)$ and $Y(t)$ are controlled evolutionary dynamics $(\mathcal{G}, 1, U(t))$ and $(\mathcal{G}, 1, 0)$, respectively, associated with a the same graph $\mathcal{G}$. The coupling mechanism is the following. Each link $\{i, j\}$ of the graph $\mathcal{G}$ is equipped with an independent rate-W_{ij} Poisson clock. When the clock associated with link $\{i, j\}$ ticks, the spreading mechanism acts on that link for both $X(t)$ and $Y(t)$, with $\beta = 1$. Each node i is given a rate-$U_i(t)$ nonhomogeneous Poisson clock, associated with the exogenous control in node i. When the clock associated with node i ticks, the state $X_i(t)$ turns to 1. We immediately deduce that the two marginals of $X(t)$ and $Y(t)$ are controlled evolutionary dynamics $(\mathcal{G}, 1, U(t))$ and $(\mathcal{G}, 1, 0)$ with the desired initial conditions, respectively.

Now, we prove that, under this coupling, $Y(t) \geq X(t)$, for every t. At $t = 0$, this is verified by assumption. We now show that any transition of $Z(t)$ preserves the inequality. Assume that a transition occurs at time t. If the transition is triggered by the spreading mechanism on link $\{i, j\}$, the same argument used in the proof of Lemma 2 yields $Y(t^+) \geq X(t^+)$. If the

transition is triggered by the control mechanism on node i, we observe that necessarily $Y_i(t^-) = 1$ (since we can only control nodes in which $Y_i(0) = 1$ and, being $\beta = 1$, if $Y_i(0) = 1$ then $Y_i(t) = 1$, $\forall t \geq 0$). Hence, the transition yields $X_i(t^+) = Y_i(t^+) = 1$, and the inequality is preserved. ∎

Lorenzo Zino (M'21) is an Assistant Professor with the Department of Electronics and Telecommunications at Politecnico di Torino (Turin, Italy), since 2022. He received the BS in Applied Mathematics and MS in Mathematical Engineering (summa cum laude) from from Politecnico di Torino, Torino, Italy, in 2012 and 2014, respectively, and the PhD in Pure and Applied Mathematics (with honors) from Politecnico di Torino and Università di Torino (joint doctorate program), in 2018. He was a Research Fellow at Politecnico di Torino (2018–19) and University of Groningen (2019–22), and a Visiting Research Assistant at New York University, Tandon School of Engineering (2017–18 and 2019). His current research interests include modeling, analysis, and control of dynamics over complex networks, applied probability, network analysis, and game theory. Since 2021, he is an Associate Editor of the *Journal of Computational Science*.

Giacomo Como (M'12) is a Professor at the Department of Mathematical Sciences, Politecnico di Torino, Italy. He is also a Senior Lecturer at the Automatic Control Department, Lund University, Sweden. He received the B.Sc., M.S., and Ph.D. degrees in Applied Mathematics from Politecnico di Torino, in 2002, 2004, and 2008, respectively. He was a Visiting Assistant in Research at Yale University in 2006–2007 and a Postdoctoral Associate at the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, from 2008 to 2011. Prof. Como currently serves as Senior Editor of the *IEEE Transactions on Control of Network Systems*, as Associate Editor of *Automatica* and as chair of the IEEE-CSS Technical Committee on Networks and Communications. He served as Associate Editor for the *IEEE Transactions on Network Science and Engineering* (2015-2021). He was the IPC chair of the IFAC Workshop Nec-Sys'15 and a semiplenary speaker at the International Symposium MTNS'16. He is recipient of the 2015 George S. Axelby Outstanding Paper Award. His research interests are in dynamics, information, and control in network systems with applications to cyber-physical systems, infrastructure networks, and social and economic networks.

Fabio Fagnani received the Laurea degree in Mathematics from the University of Pisa and the Scuola Normale Superiore, Pisa, Italy, in 1986. He received the PhD degree in Mathematics from the University of Groningen, Groningen, The Netherlands, in 1991. From 1991 to 1998, he was an Assistant Professor of Mathematical Analysis at the Scuola Normale Superiore. In 1997, he was a Visiting Professor at the Massachusetts Institute of Technology, Cambridge MA. Since 1998, he has been with the Politecnico of Torino, where he is currently (since 2002) a Full Professor of Mathematical Analysis. From 2006 to 2012, he has acted as Coordinator of the PhD program in Mathematics for Engineering Sciences at Politecnico di Torino. From June 2012 to September 2019, he served as the Head of the Department of Mathematical Sciences, Politecnico di Torino. His current research topics are on cooperative algorithms and dynamical systems over graphs, inferential distributed algorithms, and opinion dynamics. He is an Associate Editor of the *IEEE Transactions on Automatic Control* and served in the same role for the *IEEE Transactions on Network Science and Engineering* and *IEEE Transactions on Control of Network Systems*.