

Controlling Evolutionary Dynamics in Networks: A Case Study

Original

Controlling Evolutionary Dynamics in Networks: A Case Study / Zino, L., Como, G., Fagnani, F.. - ELETTRONICO. - 51:(2018), pp. 349-354. (7th IFAC workshop on distributed estimation and control in networked systems NecSys 2018 Groningen (NL) 27-28 Agosto 2018) [10.1016/j.ifacol.2018.12.060].

Availability:

This version is available at: 11583/2720821 since: 2025-03-04T21:23:59Z

Publisher:

Elsevier

Published

DOI:10.1016/j.ifacol.2018.12.060

Terms of use:

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

Publisher copyright

(Article begins on next page)

Controlling Evolutionary Dynamics in Networks: A Case Study

Lorenzo Zino ^{*,**} Giacomo Como ^{*,***} Fabio Fagnani ^{*}

^{*} *Department of Mathematical Sciences “G.L. Lagrange”, Politecnico di Torino, 10129 Torino, Italy*

(e-mail: {lorenzo.zino, giacomo.como, fabio.fagnani}@polito.it)

^{**} *Department of Mathematics “G. Peano”, Università di Torino, 10123 Torino, Italy. (e-mail: lorenzo.zino@unito.it)*

^{***} *Department of Automatic Control, Lund University, 22100 Lund, Sweden. (e-mail: giacomo.como@control.lth.se)*

Abstract: Due to their wide adaptability to different application fields spanning from opinion dynamics to biology, the analysis of evolutionary dynamics is a compelling problem in the science of networks and systems. In this paper, we deal with controlled evolutionary dynamics in networks. We discuss a novel approach to model these phenomena, which enables us to estimate the duration of the process depending on the network topology and on the control policy adopted. In a previous work, we have presented some preliminary results including a feedback control policy to speed up the dynamics. These encouraging results have pushed us toward deeper analysis of the problem. Here, we exhibit some critical issues concerning the feedback control policy originally proposed, which limit its applicability to real-world scenarios, and we address them by proposing a new improved control policy. Finally, using Monte Carlo simulations, we test the effectiveness of our approach to evolutionary dynamics and of the new control policy proposed here, against a real-world scenario, obtaining an extremely promising outcome for our future research.

Keywords: Multi-Agent Systems; Graph-based methods for networked control; Control over Networks; Feedback Control; Evolutionary Dynamics

1. INTRODUCTION

In the last few years, evolutionary dynamics have emerged as an effective modeling strategy to study and understand the effect of the introduction of a novel state in a network system and to predict its diffusion. We mention the works in Lieberman et al. (2005); Ohtsuki and Nowak (2006); Rychtář and Stadler (2008); Broom et al. (2011); Allen et al. (2017); Zino et al. (2017a). This paradigm yields very flexible models, which can be adapted to many different application fields, such as biology (e.g., to study the spread of a new species in a geographic region), opinion dynamics (e.g., to analyze the approval for a new political party), or diffusion of innovation (e.g., to predict the success of a new asset introduced in a population). However, few analytical results are available for these dynamics: most of the works in the literature are based on extensive Monte Carlo simulations, as in Lieberman et al. (2005); Broom et al. (2011); Allen et al. (2017), but for few analytical results for some very specific network topologies such as rings and small-world graphs, in Ohtsuki and Nowak (2006); Rychtář and Stadler (2008). In Zino et al. (2017a), we have proposed a change of perspective in the approach to model evolutionary dynamics, which has enabled us to perform analytical estimations of the time needed for the novel state to spread in a network, depending on its topological structure.

One of the most compelling objectives of the study of mathematical models is the development of effective control strategies to act on the evolution of a system and modify its outcome. Some paradigmatic examples come from the epidemiological field. The mathematical analysis of epidemics spreading in interconnected communities in Ganesh et al. (2005); Pastor-Satorras et al. (2015); Fagnani and Zino (2017), have allowed for the design of a bunch of control policies with relevant applications for the health system. See, e.g., Borgs et al. (2010); Drakopoulos et al. (2014); Nowzari et al. (2016). Another example lies in the modeling of opinion dynamics in Liggett (1985); Montanari and Saberi (2010); Young (2011). These studies have paved the way for several applications, such as the development of accurate strategies to maximize the influence of an opinion spreader in Kempe et al. (2003). Having considered these encouraging results obtained from the analysis of network dynamics, we set our main goal in the development of effective control strategies, which allows for speeding up the diffusion of the novel state over the network. Some preliminary results in this direction are available in Zino et al. (2017a), where a feedback control policy to solve this problem has been proposed.

The main contributions of this work, thus, consist in: *i*) pointing out some critical issues of the feedback control policy presented in our preliminary work; *ii*) proposing a novel feedback control policy that addresses these problems; and *iii*) testing the applicability of our model and

the effectiveness of the proposed feedback control policy against a real-world scenario for evolutionary dynamics.

The paper is organized as follows. In Section 2, we present the model for controlled evolutionary dynamics in network systems. Then, in Section 3, we discuss the control strategies for the process: first, we recall our preliminary results comparing the effectiveness of constant control policies and feedback ones, then we propose a novel feedback control policy that is capable of addressing the main critical issues of our preliminary proposal. Then, in Section 4, we present a case study based on a real-world parameter setting, in which the effectiveness of our novel feedback control policy is proved. Finally, Section 5 concludes the paper by discussing our results and presenting the future steps of our analysis.

1.1 Notation

We gather here some notation conventions used throughout this paper. The all-1 vector is denoted as $\mathbf{1}$ and the all-0 vector as $\mathbf{0}$. A vector of all-0 but a 1 in the i -th position is denoted by $\delta^{(i)}$. For a vector $\mathbf{x} \in \mathbb{R}^n$, \mathbf{x}^T indicates the transpose. \mathbb{R}^+ denotes the set of non-negative real numbers.

2. MODEL

The topological structure of the system is described as a graph, on whose nodes the novel state is introduced. Depending on the application field, the nodes may represent, e.g., geographic locations occupied by a species (in biological applications), or individuals having an opinion (in social science applications). For the sake of simplicity, in this paper we adopt the biological interpretation, that is the standard one used in the literature of evolutionary dynamics, e.g., in Lieberman et al. (2005). So, the network represents a geographic region and the novel state maps a mutant species introduced in the region. We assume that each node is fully occupied by just one species, either the native one or the mutants. Species occupation varies according to a spreading mechanism and an external control, detailed below.

- *The geographical graph.* We consider an undirected connected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, $W \geq 0$, entry-wise. The node set $\mathcal{V} = \{1, \dots, n\}$ represents the locations, link $\{i, j\} \in \mathcal{E}$ represents proximity of nodes i and j , and the weight W_{ij} measures the frequency of interactions between the pair of linked nodes. We suppose W to be symmetrical and $W_{ij} > 0 \iff \{i, j\} \in \mathcal{E}$.
- *The spreading mechanism.* We assume each undirected link $\{i, j\}$ to be equipped with an independent Poisson clock with rate W_{ij} , which models the times the two species in nodes i and j interact. When the clock associated with the link $\{i, j\}$ clicks, if the two species in i and j differ, then a conflict takes place and the winning species occupies both locations. Each conflict is won by mutants (each conflict independently of the others) with probability $\beta \in [0, 1]$.
- *The external control.* We fix an integrable non-negative function $u(t)$ and a target node $m(t) \in \mathcal{V}$, $t \in \mathbb{R}^+$. At time t , mutants are introduced in node $m(t)$ at rate $u(t)$.




Fig. 1. Transitions of the Markov process $X(t)$. Red nodes have state 1, white nodes have state 0.

Remark 1. The external control can be generalized by introducing mutants in more than one location at each time, as considered in Zino et al. (2017a). However, here we restrict our analysis to the case of a unique target node, both due to space constraints and because in real-world scenarios to control more than one node at a time may be excessively expensive and/or unfeasible. See Harris et al. (2011).

The quadruple $(\mathcal{G}, \beta, u(t), m(t))$ defines a *controlled evolutionary system*. To it, we now associate a n -dimensional Markov process $X(t)$ that describes the evolution of the system. Details on Markov processes can be found in Levin et al. (2009). Component $X_i(t) \in \{0, 1\}$ represents whether location i at time t is fully occupied by the native species ($X_i(t) = 0$) or by the mutants ($X_i(t) = 1$). Time $t = 0$ is the moment mutants are started to be introduced in the system, so the initial configuration is $X(0) = \mathbf{0}$. Formally, $X(t)$ is a non homogeneous Markov jump process on the configuration space $\{0, 1\}^n$. The only transitions that can take place from a generic state $X(t) = \mathbf{x} = (x_1, \dots, x_n)$ are the ones to states that differ from \mathbf{x} in a single entry. Their rates are, for $i \in \mathcal{V}$:

$$\begin{aligned} \lambda_i^+(\mathbf{x}) &= \begin{cases} \beta(1-x_i)(W\mathbf{x})_i + (1-x_i)u(t) & \text{if } i = m(t) \\ \beta(1-x_i)(W\mathbf{x})_i & \text{if } i \neq m(t) \end{cases} \\ \lambda_i^-(\mathbf{x}) &= (1-\beta)x_i [W(\mathbf{1}-\mathbf{x})]_i, \end{aligned} \quad (1)$$

where $\lambda_i^\pm(\mathbf{x})$ denotes the transition rate from state \mathbf{x} to state $\mathbf{x} \pm \delta^{(i)}$, as shown in Fig. 1.

In the following of this work, two reasonable assumptions are made. First, in order to force the presence of an external control, we assume that

$$X(t) = \mathbf{0} \implies u(t) > 0. \quad (2)$$

Second, we consider a mutant exhibiting some evolutionary advantages, modeled by $\beta > 1/2$.

From (1) and (2), it is straightforward that $\mathbf{1}$, the all-mutants configuration, is the only absorbing state of the process and that, in the long run, mutants will almost surely occupy the whole network. Hence, what is of interest in the applications is the analysis transient behavior, captured by the following two quantities: the *expected diffusion time*, defined as

$$\tau = \mathbb{E} [\inf \{t \in \mathbb{R}^+ : X(t) = \mathbf{1}\}], \quad (3)$$

and the *expected control cost*, defined as

$$J = \mathbb{E} \left[\int_0^\infty u(t) dt \right], \quad (4)$$

with the understanding that, once the process is absorbed, then $u(t)$ is set to 0. The effectiveness of various control policies will be compared through the analysis of these two quantities.

In order to tackle the analysis of this evolutionary dynamics, we introduce two stochastic processes, being one-dimensional observables on the system:

$$Z(t) := \mathbf{1}^T X(t), \quad (5)$$

which counts the number of locations occupied by mutants, and

$$B(t) := X(t)^T W (\mathbf{1} - X(t)), \quad (6)$$

which counts the total weight of the links in the boundary between locations occupied by the native species and mutants.

The estimation of τ and J is a non-trivial problem. However, a simple heuristic which can give an interesting intuition on the main quantities that influence τ can be found in the analysis of the N-Intertwined Mean Field Approximation (NIMFA) of the evolutionary process, which considers a continuous-state relaxation of the system, as detailed in Van Mieghem (2011).

Remark 2. Let the state of each node assume a continuous value $x_i \in [0, 1]$, which represents the probability that the node is occupied by mutants. Then, according to the NIMFA, the evolution of the vector \mathbf{x} is governed by the following ODE, obtained following the procedure in Kurtz (1981):

$$\begin{aligned} \dot{x} &= \lambda^+(x) - \lambda^-(x) \\ &= \beta \text{diag}(\mathbf{1} - x) W x + (1 - x_{m(t)}) u(t) \delta^{(m)} \\ &\quad - (1 - \beta) \text{diag}(x) W (\mathbf{1} - x). \end{aligned} \quad (7)$$

Similar to the stochastic processes, we can define $z(t) = \mathbf{1}^T x(t)$ and $b(t) = x(t)^T W (\mathbf{1} - x(t))$. Then, from (7), $z(t)$ is the solution of the Cauchy problem

$$\begin{cases} \dot{z} = (2\beta - 1)b(t) + (1 - x_{m(t)})u(t) \\ z(0) = 0. \end{cases} \quad (8)$$

Since $\beta > 1/2$, from (8) we deduce that the velocity of the process is proportional to *i*) the magnitude of the boundary between the native species and mutants, i.e., $b(t)$; and *ii*) the control rate in nodes occupied by the native species.

From this heuristic, we can hypothesize that a technique to achieve fast diffusion should: *i*) compensate for the slowdowns of the process when the total weight of the links on the boundary $B(t)$ is small (i.e., in correspondence to bottlenecks of the graph), and *ii*) avoid wasting energy introducing mutants in nodes with state 1 and when the process is already evolving fast.

3. CONSTANT VS. FEEDBACK CONTROL POLICIES

In Zino et al. (2017a), we have proposed two families of control policies:

- *Constant control policies.* We fix $m(t) = m, \forall t \in \mathbb{R}^+$, and given $u \in \mathbb{R}^+$,

$$u(t) = \begin{cases} u & \text{if } X_m(t) = 0 \\ 0 & \text{if } X_m(t) = 1. \end{cases} \quad (9)$$

- *Feedback control policies.* We assume $m(t)$ to be moved in such a way that $X_{m(t)} = 0$, and $u(t)$ to be a feedback control of the state of the system. We set




Fig. 2. Examples of geographical graphs (a) easy to control even with constant policies, (b) easy to control only with feedback policies, and (c) hard to control with any policy.

$$u(t) = \begin{cases} u_0 & \text{if } Z(t) = 0 \\ 0 & \text{if } Z(t) = n \\ B(t)\tilde{u}(Z(t)) & \text{else,} \end{cases} \quad (10)$$

where $\tilde{u}(z)$ is the actual control function that depends only on the number of nodes occupied by mutants $Z(t)$.

Remark 3. The rationale for the feedback control policy (10) lies in the observations stated in Remark 2. In fact, the main idea is to activate the control (i.e., to set $\tilde{u}(z) > 0$) only for those values of $Z(t)$ that may correspond to bottlenecks of the graph.

These two specific families of control policies allow for an analytical estimation of τ and J . These estimations have been performed in Zino et al. (2017a) (detailed proofs are available in Zino et al. (2017b)) and yield to classify geographical graphs, depending on their topology, into three classes: *i*) *easy to control even with constant control policies* (e.g., expander structures such as complete or Erdos-Renyi random graphs), for which a constant control policy guarantees fast diffusion of the mutants; *ii*) *easy to control only with feedback control policies* (e.g., barbell graphs), for which constant control policies fail to achieve fast diffusion, but feedback control policies succeed; and *iii*) *hard to control with any control policy* (e.g., rings), for which mutants need a long time to diffuse under any control policy. Fig. 2 depicts a remarkable example for each one of the three classes.

The main strength of the proposed feedback control policy is in the fact that it allows for speeding up the diffusion process, relying on just two observables: $Z(t)$ and $B(t)$, and under the only assumption that it is possible to move the target node to a generic location occupied by the native species. We remark that no optimization is needed in the choice of such target node, which is usually a computationally hard problem, see Kempe et al. (2003).

However, when considering real-world applications of our proposed control policy, three problems may arise, restricting its usability. First, the detection of the bottlenecks of a graph is a NP-complete problem, which implies that a heavy computational effort is required for their exact computation in large-scale systems. Second, the control policy in (10) is very sensible to small errors in the data. For instance, if a node is added to the graph, then the control policy could fail in activating in correspondence to the bottlenecks. This also implies that the bottlenecks of the graph should be detected exactly (solving the NP-complete problem), since small errors in their estima-

tion could hamper the effectiveness of the control policy. Finally, the feedback control policy originally proposed could waste energy, inserting mutants when the process is already evolving fast.

Therefore, we propose here a novel feedback control policy capable of addressing these three issues. Specifically, considering the importance of the observable $B(t)$ that has been highlighted in Remark 2, we seek to exploit its knowledge in a more thoughtful way. The core idea is to use this observable in order to activate the control only in the presence of slowdowns of the spreading mechanisms, corresponding to bottlenecks of the topology. This allows, on the one hand, to improve the robustness of the control technique, that is not prone anymore to small errors in the data, and, on the other hand, to avoid useless waste of energy when the evolutionary advantage of the mutants guarantees the speed of the process to be sufficiently large. Moreover, in this new control policy, there is no need for actually detecting the bottlenecks of the network, since they are automatically identified through the observable $B(t)$ when the process slows down. Hence, we should intuitively act on the system only when $B(t)$ is small. To this aim, fixed a positive parameter $C > 0$, we let

$$u(t) = \begin{cases} C - B(t) & \text{if } Z(t) \neq n, B(t) < C \\ 0 & \text{else,} \end{cases} \quad (11)$$

where $C > 0$ guarantees assumption (2) to be verified.

Remark 4. The new control policy in (11) is able to address the three issues posed by the original proposal. In fact, the explicit use of the observable $B(t)$ enables us to immediately recognize when the process enters in a bottleneck and promptly compensates for the slowdown, without any useless waste of energy when the process is already evolving fast.

We observe that, in order to tune the parameter C , one has to estimate the magnitude of the bottlenecks of the geographical graph to decide when to activate the control to compensate for the slowdowns. However, the exact detection of these bottlenecks is not required by our new control policy: one can estimate them using fast algorithms and define the parameter C accordingly, making our policy feasible for the application to large-scale systems.

The theoretical analysis of our novel feedback control policy (11) in terms of the computation of an upper-bound on the expected diffusion time τ and on its expected cost J will be part of our future research. Instead, we devote the remaining of this paper to show the potentiality of our mathematical model on a realistic case study. Specifically, we aim to numerically prove the effectiveness of our new control technique by applying it in a parameters setting based on a real-world scenario.

4. CASE STUDY

In this section, we present a case study to show the potentiality our feedback control policy to speed up evolutionary dynamics in network systems. Inspired by a current hot topic in epidemic control, we consider a possible strategy to control the Zika outbreak in Rwanda (details about the outbreak can be found in CDC Centers for Disease Control and Prevention (2018a)) by substituting *Aedes aegypti*



Fig. 3. Topological graph of the considered geographic area. Links connect locations within 11.7 km.

mosquitoes with genetically modified organisms (GMOs) that are similar to the disease-spreading mosquitoes but do not transmit the Zika virus. Similar control strategies that involve the use of GMOs have been proposed and adopted in trials and experiments to fight other mosquitoes-borne diseases such as dengue fever in Central and South America. More details can be found in Harris et al. (2011); Carvalho et al. (2015).

4.1 Parameter Settings

The geographical network is constructed as follows. We consider a data set of 1621 locations in Rwanda with their GPS coordinates from National Imagery and Mapping Agency (2018). Two nodes are connected if and only if the mosquitoes in the two locations can contact. Hence, we establish a threshold corresponding to 11.7 km, that is the maximum distance traveled by mosquitoes to lay their eggs, according to Bogojević et al. (2007): locations within this distance are connected. The so obtained network is represented in Fig. 3. Since the duration of the life cycle of an *Aedes aegypti* lasts in average 10 days, according to CDC Centers for Disease Control and Prevention (2018b), we set the activation rate of each undirected link $(i, j) \in \mathcal{E}$ equal to $W_{ij} = 0.1$. Finally, we give the GMOs a little evolutionary advantage. We model it by setting $\beta = 0.53$. The parameters used in the case study are summarized in Table 1.

4.2 Simulations

We have performed 200 Monte Carlo simulations of the evolutionary dynamics on the network for each control policy (constant control policy and the new feedback one). Simulations are generated according to a Gillespie algorithm, following Gillespie (1976). In order to compare the two policies, we set the specific parameters of the two control policies in such a way that the average costs coincide.

Table 1. Parameters of the Rwanda case study

Parameter	Meaning	Value
n	Number of locations	1621
W_{ij}	Activation rate	0.1
β	Evolutionary advantage	0.53
u	Control rate (constant)	2
C	Control parameter (feedback)	1.5
t	Time unit	day




Fig. 4. Diffusion time and control cost of the 200 simulations performed, both under the constant control policy (blue squares) and under the new feedback control policy (red circles).

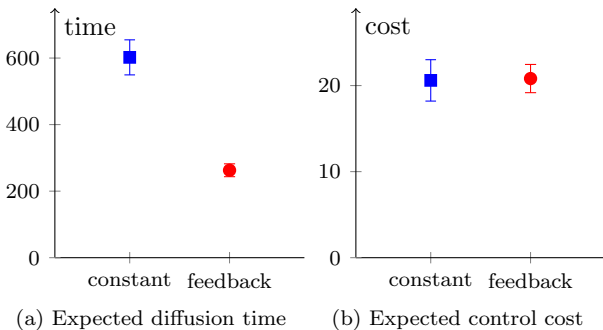


Fig. 5. Monte Carlo estimation (200 simulations) of the expected diffusion time τ , in (a), and of the expected control cost, in (b), with 95% confidence intervals for the constant (blue squares) and the new feedback control policy (red circles).

To this aim, we fix $u = 2$ (for the constant control) and $C = 1.5$ (for the feedback one). Our simulations, reported in Fig. 4, allow for estimating the magnitude of the improvement gained by adopting the feedback control policy: at the same cost, the diffusion time reduces in average by more than 56% (such an improvement is guaranteed by a strong statistical significance of the result, given by a p -value $p < 0.001$), as depicted in Fig. 5. This comparison also shows that the outcomes of feedback control policy are less variable than the ones of constant control policy. Hence, the feedback control policy we proposed seems to outperform the constant one, not only in average, but also considering worst case scenarios.

5. CONCLUSION

In this paper we have studied controlled evolutionary dynamics, where a novel state is introduced in a network and its diffusion is studied. First, we have presented our framework, recalling some preliminary results for these dynamics from Zino et al. (2017a), in which the effectiveness of feedback control policies to speed up the diffusion process had been proved. Then, we have presented the main contributions of this paper, that are *i*) the discussion of some concerning problems related with the feedback control policy originally developed in Zino et al. (2017a)

that limit its applicability, and the proposal of a novel feedback control policy capable to address and overcome these issues; and *ii*) the analysis of a real-world biological evolutionary system as a case study. The goal of the presentation and the analysis of this case study is twofold: first, we want to show the applicability of our model to real-world scenarios, far beyond the toy examples considered in our preliminary work; and second, we aim to prove the validity and the effectiveness of our proposed feedback control policy.

Given the promising results obtained in the case study presented in Section 4, in our future research we are planning to tackle the analysis of the proposed control policy in order to derive accurate bounds on the expected diffusion time and the expected control cost and to investigate its robustness. On the other hand, we are aiming to establish fundamental limits on the diffusion time achievable under any control policy. By combining these two future results, the effectiveness of our new control policy will be analytically evaluated, paving the way for further researches seeking for an optimal strategy within the family of feedback control policies we proposed. Finally, from an application perspective, we aim to study the feasibility of the proposed feedback control policy in real-world scenarios.

FUNDING STATEMENT

This work was partly supported by the Italian Ministry of Education, University, and Research (MIUR) through the DISMA – Dipartimento di Eccellenza 2018-2022 project, the Swedish Research Council through Project Research Grant 2015-04066, and by the Compagnia di San Paolo through a PoliTo Starting Grant.

REFERENCES

- Allen, B., Lippner, G., Chen, Y.T., Fotouhi, B., Momeni, N., Yau, S.T., and Nowak, M.A. (2017). Evolutionary dynamics on any population structure. *Nature*, 544(7649), 227–230.
- Bogojević, M.S., Hengl, T., and Merdić, E. (2007). Spatiotemporal monitoring of floodwater mosquito dispersal in Osijek, Croatia. *Journal of the American Mosquito Control Association*, 23(2), 99–108.
- Borgs, C., Chayes, J., Ganesh, A., and Saberi, A. (2010). How to distribute antidote to control epidemics. *Random Structures and Algorithms*, 37(2), 204–222.
- Broom, M., Rychtář, J., and Stadler, B.T. (2011). Evolutionary dynamics on graphs - the effect of graph structure and initial placement on mutant spread. *Journal of Statistical Theory and Practice*, 5(3), 369–381.
- Carvalho, D.O., McKemey, A.R., Garziera, L., Lacroix, R., Donnelly, C.A., Alphey, L., Malavasi, A., and Capurro, M.L. (2015). Suppression of a field population of *Aedes aegypti* in Brazil by sustained release of transgenic male mosquitoes. *PLoS Negl. Trop. Dis.*, 9(7), 1–15.
- CDC Centers for Disease Control and Prevention (2018a). Health Information for Travelers to Rwanda. <https://wwwnc.cdc.gov/travel/destinations/traveler/none/rwanda>. [Online; accessed 11-Jul-2018].
- CDC Centers for Disease Control and Prevention (2018b). Mosquito Life-Cycle. https://www.cdc.gov/dengue/entomologyecology/m_lifecycle.html. [Online; accessed 11-Jul-2018].

- Drakopoulos, K., Ozdaglar, A., and Tsitsiklis, J.N. (2014). An efficient curing policy for epidemics on graphs. *IEEE Transactions on Network Science and Engineering*, 1(2), 67–75.
- Fagnani, F. and Zino, L. (2017). Time to extinction for the SIS epidemic model: new bounds on the tail probabilities. *IEEE Transactions on Network Science and Engineering*. Published online.
- Ganesh, A., Massoulié, L., and Towsley, D. (2005). The effect of network topology on the spread of epidemics. *Proceedings - IEEE INFOCOM*, 2, 1455–1466.
- Gillespie, D.T. (1976). A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.*, 22, 403–434.
- Harris, A.F., Nimmo, D., McKemey, A.R., Kelly, N., Scaife, S., Donnelly, C.A., Beech, C., Petrie, W.D., and Alphey, L. (2011). Field performance of engineered male mosquitoes. *Nature Biotechnology*, 29(11), 1034–1037.
- Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '03*, 137. ACM Press, New York, New York, USA.
- Kurtz, T.G. (1981). *Approximation of Population Processes*, volume 36. SIAM, Philadelphia, PA.
- Levin, D.A., Peres, Y., and Wilmer, E.L. (2009). *Markov chains and mixing times*. American Mathematical Society, Providence, RI, USA.
- Lieberman, E., Hauert, C., and Nowak, M.A. (2005). Evolutionary dynamics on graphs. *Nature*, 433.
- Liggett, T.M. (1985). *Interacting particle systems*. Springer-Verlag, New York, NY, USA.
- Montanari, A. and Saberi, A. (2010). The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47), 20196–20201.
- National Imagery and Mapping Agency (2018). 1621 locations in Rwanda. <http://www.math.uwaterloo.ca/tsp/world/rw1621.tsp>. [Online; accessed 11-Jul-2018].
- Nowzari, C., Preciado, V.M., and Pappas, G.J. (2016). Analysis and control of epidemics: a survey of spreading processes on complex networks. *IEEE Control Systems*, 36(1), 26–46.
- Ohtsuki, H. and Nowak, M.A. (2006). Evolutionary games on cycles. *Proceedings of the Royal Society B-Biological Sciences*, 273(1598), 2249–2256.
- Pastor-Satorras, R., Castellano, C., Van Mieghem, P., and Vespignani, A. (2015). Epidemic processes in complex networks. *Reviews of Modern Physics*, 87(3), 925–979.
- Rychtář, J. and Stadler, B. (2008). Evolutionary dynamics on small-world networks. *International Journal of Computational and Mathematical Sciences*, 2(1), 1–4.
- Van Mieghem, P. (2011). The n-intertwined SIS epidemic network model. *Computing*, 93(2), 147–169.
- Young, H.P. (2011). The dynamics of social innovation. *Proceedings of the National Academy of Sciences*, 108(4), 21285–21291.
- Zino, L., Como, G., and Fagnani, F. (2017a). Fast Diffusion of a Mutant in Controlled Evolutionary Dynamics. *IFAC-PapersOnLine*, 50(1), 11908–11913.
- Zino, L., Como, G., and Fagnani, F. (2017b). Fast Diffusion of a Mutant in Controlled Evolutionary Dynamics. <http://calvino.polito.it/~zino/ifacconf.pdf>. [Online; accessed 11-Jul-2018].