Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online
CONFERENCE PAPER ARCHIVE

IFAC-PapersOnLine 48-22 (2015) 076–081

Control of Stochastic Evolutionary Games on Networks ⋆

James R. Riehl * **Ming Cao** *

* *Faculty of Mathematics and Natural Sciences, ENTEG, University of Groningen, The Netherlands (e-mail: {j.r.riehl, m.cao}@rug.nl)*

Abstract: We investigate the control of stochastic evolutionary games on networks, in which each edge represents a two-player repeating game between neighboring agents. The games occur simultaneously at each time step, after which the agents can update their strategies based on local payoff and strategy information, while a subset of agents can be assigned strategies and thus serve as control inputs. We seek here the smallest set of control agents that will guarantee convergence of the network to a desired strategy state. After deriving an exact solution that is too computationally complex to be practical on large networks, we present a hierarchical approximation algorithm, which we show computes the optimal results for special cases of complete and ring networks, while simulations show that it yields near-optimal results on trees and arbitrary networks in a wide-range of cases, performing best on coordination games.

Keywords: networks, evolutionary games, optimization, control algorithms, stochastic systems

1. INTRODUCTION

The rapid growth in connectivity of our society and technology has led to ever-increasing complexity of the systems we rely upon, many of which can be characterized by agents making decisions across a network. These decisions are often based on local information and may be driven by competing objectives of the agents. Game theory has long been used to tackle these kinds of problems on a small scale and where the agents can be assumed to be perfectly rational, but for larger scale and more complex systems, these assumptions often no longer hold and the dynamics of the strategy choices are better modeled by *evolutionary game theory* (Sandholm (2010), Szabó and Fáth (2007)), where strategies propagate through the population based on the payoffs acquired by the agents. In biological terms, the mechanism for this propagation is an evolutionary *survival of the fittest* process in which the fitness is directly related to the payoffs of the game. The agents need not be simple organisms however; more complex systems such as human social networks and robotic networks can also fit well into an evolutionary game framework, but here the strategy propagation mechanism can be better thought of as a learning process or update rule. One of the well-established findings in the field is that evolutionary games can and often do lead to complex and undesired outcomes with respect to the population as a whole, such as in prisoner's dilemma games or *tragedy of the commons*, where selfishness tends to prevail over cooperation (Liebrand et al. (1986), Ostrom (2008)). There is consequently a strong incentive to understand the possibilities for influencing evolutionary games in order to catalyze better collective outcomes in populations.

Although a relatively recently emerging topic, researchers have attacked similar problems from several different and interesting angles. For example, in the setting of infinite and well-mixed populations, Kanazawa et al. (2009) showed how taxation and subsidies can promote the emergence of desired strategies in a routing game under the replicator dynamics, and Sandholm (2002) introduced pricing schemes to promote efficient choices on roadway networks, as well as in more general economic contexts (Sandholm (2007)). For more complex networks under best-response dynamics, Balcan et al. (2014) showed that broadcasted information can guarantee convergence to an equilibrium that is within a given factor of the social optimum. Also, Cheng et al. (2015) presented a framework for studying the control of networked evolutionary games using large-scale logical dynamic networks to model transitions between all possible strategy states. They used this framework to derive equivalent conditions for reachability and consensus of strategies on a network given a particular set of control agents. However, the optimal control of evolutionary games on networks remains a challenging open problem and is the primary focus of this paper, which extends our earlier work in Riehl and Cao (2014a) and Riehl and Cao (2014b).

After defining a general evolutionary game framework in Section 2, we formulate a minimum agent control problem in Section 3. Although we derive an exact solution algorithm in Section 4, the high computational complexity makes it practical only for small networks. Faced with this obstacle, we proceed by designing an algorithm in Section 5 that uses a hierarchical approach to approximate the solution. We show in Section 6 that the resulting approximation is exact for classes of games on complete and ring networks. Moreover in Section 7, simulations show that the results are quite accurate on trees as well as a broad class of geometric random networks. We summarize

⋆ This work was supported in part by the European Research Council (ERCStG-307207).

10.1016/j.ifacol.2015.10.310

our contributions and discuss important directions for future work in Section 8.

2. EVOLUTIONARY GAME FRAMEWORK

In this section we define the evolutionary game framework, which consists of a network, payoff matrix, and strategy update dynamics.

2.1 Network and single-game payoffs

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote an undirected network consisting of an agent set $\mathcal{V} = \{1, \dots, n\}$ and an edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$, where each edge represents a 2-player symmetric game between neighbors.[1] The agents choose strategies from a binary set $\mathcal{S} := \{A, B\}$ and receive payoffs upon completion of each game according to the matrix

$$M = \begin{array}{c c} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} a & b \\ c & d \end{pmatrix} \end{array}. \quad (1)$$

We assume that at each time step, players use a single strategy against all opponents, and thus the games occur simultaneously. We denote the strategy state by $x(t) = [x_1(t), \dots, x_n(t)]^T$, where $x_i(t) \in \mathcal{S}$ is the strategy of agent i at time t. Total payoffs for each agent are given by

$$y_i(t) = w_i \sum_{j \in \mathcal{N}_i} M_{x_i(t), x_j(t)}, \quad (2)$$

where $\mathcal{N}_i := \{j \in \mathcal{V} : \{i, j\} \in \mathcal{E}\}$ is the neighbor set of agent i and the most common values for the weights w_i are 1 for cumulative payoffs and $\frac{1}{|\mathcal{N}_i|}$ for averaged payoffs. The total payoffs are collected into the vector $y(t) = [y_1(t), \dots, y_n(t)]^T$.

2.2 Strategy update dynamics

A fundamental concept behind evolutionary games is that better performing strategies are adopted more often, meaning that rather than rationally choosing best-response strategies, players imitate strategies in their neighborhood that result in higher payoffs. We capture this dynamic with a strategy update rule that is a function of the strategies and payoffs of neighboring agents:

$$x_i(t+1) = f\left(\{x_j(t), y_j(t) : j \in \mathcal{N}_i \cup \{i\}\}\right). \quad (3)$$

The only restrictions we make on the update rule is that it is *payoff monotone*, *i.e.* players only switch to strategies with which at least one agent in the neighborhood achieves a greater payoff (Szabó and Fáth (2007)), and *persistent*, meaning that if there exists a better performing strategy in an agent's neighborhood, then the agent will switch to that strategy with a probability that is lower bounded by $\epsilon > 0$.

2.3 Example: proportional imitation

One example of such dynamics is the *proportional imitation* rule, in which each agent chooses a neighbor randomly and if this neighbor received a higher payoff in the previous round by using a different strategy, then the agent will switch with a probability proportional to the payoff difference. This is a widely studied model that has some nice properties, in particular that the strategy distributions in well-mixed populations using proportional imitation is approximated by the replicator dynamics (Schlag (1998)). The proportional imitation rule can be expressed as follows:

$$p\big(x_i(t+1) = x_j(t)\big) := \left[\frac{\lambda}{|\mathcal{N}_i|}\big(y_j(t) - y_i(t)\big)\right]_0^1 \quad (4)$$

for each agent $i \in \mathcal{V}$ where $j \in \mathcal{N}_i$ is a uniformly randomly chosen neighbor, $\lambda > 0$ is an arbitrary rate constant, and the notation $[z]_0^1$ indicates $\max(0, \min(1, z))$. This update rule is clearly payoff monotone and it is also persistent with a lower bound ϵ that can be computed from λ, the maximum degree of the network, and the smallest positive payoff difference, which exists since there are only a finite number of possible agent payoffs.

2.4 Dynamics with control agents

We add to the general framework a set of *control agents* $\mathcal{L} \subseteq \mathcal{V}$ whose strategy can be externally manipulated, either through direct control or by adding neighbors with very high payoffs. Although the most general formulation would allow dynamic control sequences for these agents, we restrict the inputs here to fixed strategies and focus initially on the simpler case of the problem. This results in the following controlled dynamics:

$$x_i(t+1) = f_c\left(\{x_j(t), y_j(t) : j \in \mathcal{N}_i \cup \{i\}\}, \mathcal{L}\right),$$

which are an extension of (3) with a special case for the control agents.

$$x_i(t+1) = \begin{cases} A, & i \in \mathcal{L} \\ f(\cdot), & \text{otherwise} \end{cases} \quad (5)$$

The combination of a network, payoff matrix, and update rule forms what we call a *network game* $\Gamma := (\mathcal{G}, M, f)$.

3. PROBLEM FORMULATION

Now that we have a general dynamic evolutionary game model with control inputs, we are interested in how one can influence the network through efficient use of these inputs in order to achieve some desired outcome of strategies. In this work, we focus on achieving uniform adoption of strategy A and pose what we call the *Minimum Agent Consensus Control* (MACC) problem.

Problem 1. (**MACC**). Given a network game Γ and initial strategy state $x(0)$, find the smallest set of control agents $\mathcal{L}$ such that $x_i(t) \to A$ for each agent $i \in \mathcal{V}$.

We say that $x_i(t)$ converges *almost surely* to X if $\lim_{t\to\infty}[p(x(t) = X)] = 1$, and indicate this with the shorthand notation $x_i(t) \to X$.

Remark 1. Since we are concerned with optimality, we seek solutions corresponding to given initial strategy states. Although one could modify the proposed approach to compute a set of control agents that would work for any initial condition, due to the complexity of the underlying network, this would almost certainly lead to very conservative results.

[1] We do not require connectivity of the network.