

Controlling network coordination games

Original

Controlling network coordination games / Durand, Stéphane; Como, Giacomo; Fagnani, Fabio. - ELETTRONICO. - 53:(2020), pp. 2862-2867. (Intervento presentato al convegno 21st IFAC World Congress (IFAC 2020) tenutosi a Berlin (Germany) nel July 12-17, 2020) [10.1016/j.ifacol.2020.12.957].

Availability:

This version is available at: 11583/2897499 since: 2021-04-28T19:19:04Z

Publisher:

Elsevier

Published

DOI:10.1016/j.ifacol.2020.12.957

Terms of use:

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

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2020. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.ifacol.2020.12.957>

(Article begins on next page)

Controlling network coordination games

Stéphane Durand* Giacomo Como* Fabio Fagnani*

* *Department of Mathematical Sciences “G.L. Lagrange”, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy*
(e-mail: {stephane.durand, giacomo.como, fabio.fagnani}@polito.it).

Abstract: We study a novel control problem in the context of network coordination games: the individuation of the smallest set of players capable of driving the system, globally, from one Nash equilibrium to another one. Our main contribution is the design of a randomized algorithm based on a time-reversible Markov chain with provable convergence guarantees.

Copyright © 2020 The Authors. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0>)

Keywords: Coordination games, Contagion, Control set, Resilience, Randomized search

1. INTRODUCTION

In a binary (0-1) network coordination game, where initially all agents are in the Nash equilibrium 0, what is the minimum number of agents that if forced to 1 will push the system to converge to the Nash equilibrium of all 1's under best response dynamics? This paper is devoted to the analysis of this problem and to the design of an algorithm for an efficient solution.

The considered problem can be framed in the more general setting of studying minimal strategies of intervention needed to drive a multi-agent system, governed by agents' myopic utility maximization, from a Nash equilibrium to a more desirable one. Typically, in game theory, interventions have been modeled as perturbations of the utility functions, e.g. taxes and prices in economic models or tolls in transportation systems. Here we take a different viewpoint: we select as small as possible a subset of nodes that if suitably controlled will lead the entire system to the desired configuration. The minimum cardinality of this set can also be interpreted as a measure of resilience of the system: the larger it is, the more difficult is for an external shock to destabilize it.

The problem of determining the best set of nodes to exert the most effective control in a networked system has recently appeared in other contexts: for instance, for linear opinion dynamics models, Acemoğlu et al. (2013), Yildiz et al. (2013), Vassio et al. (2014), and Como and Fagnani (2016) study the effect of stubborn nodes and the related problem of optimal selection of nodes to maximize influence on the rest of the network.

Binary coordination games have received a great attention in the recent years as one of the basic models for games with strategic complementarities, see Jackson and Zenou (2015). Its various applications include modeling of social and economic behaviors like the adoption of a new tech-

nology, the participation in an event or the participation to provide a public good effort.

This game is analyzed in detail in Morris (2000) where the key concept of cohesiveness of a set of players is introduced and then used in characterizing all NE's. The question of if an initial seed of influenced players (that maintain action 1 in all circumstances) is capable of propagating to the all network is also addressed in that paper including an equivalent characterization of this spreading phenomenon, expressed in terms of cohesiveness.

This contagion phenomenon is exactly what we want to analyze: subset of nodes from which propagation to the all network is successful are called sufficient control sets and our goal is to find such sets of minimum possible cardinality. However, the condition proposed in Morris (2000) is computationally quite demanding and cannot be used to directly solve our optimization problem. Indeed, even to determine if a single set is a sufficient control set, their approach requires a number of checks growing exponentially in the cardinality of the complement of any such set.

The dual problem of choosing a fixed number of players to have the maximum possible spread of the state 1 was studied in a seminal paper by Kempe et al. (2003). While their problem and ours are related, solving one does not provide a solution of the other. Another point worth stressing is that, in their setting, Kempe et al. (2003) consider agents equipped with random independent activation thresholds and take as the objective function the average size of the maximum spread. They prove that such function is sub-modular and then they design a greedy algorithm using this property. The randomness that they introduce is actually crucial in their approach, as the function considered would not be sub-modular for deterministic choices of thresholds. Different approaches to address related questions were considered, among other, in Rossi et al. (2019) in the large-scale limit.

In this paper we consider a scenario when all agents have a fixed threshold $1/2$, thus not covered in Kempe et al. (2003), and we design an iterative search randomized algorithm with provable properties of convergence towards

* Giacomo Como is also with the Department of Automatic Control, Lund University, Sweden. This work was partially supported by MIUR grant Dipartimenti di Eccellenza 2018–2022 [CUP: E11G18000350001], the Swedish Research Council, and by the Compagnia di San Paolo.

sufficient control sets of minimum cardinality. The core of the algorithm is a time-reversible Markov chains over the family of all sufficient control sets that starts with the full set, moves through all of them in an ergodic way, and concentrates on the smallest sets.

We conclude this introduction with a brief outline of the paper. In the final part of this section we report some basic notation used throughout the paper. Section 2 is dedicated to the formal introduction of the problem. The main technical parts are Sections 3 and 4. In Section 3, we introduce the important notion of monotone crusade (appeared for other purposes in Drakopoulos et al. (2014, 2016)) and we give an equivalent (but more operative) characterization of sufficient control sets. In Section 4 we introduce a family of reversible Markov chains whose invariant probability is proven to concentrate on the optimal sufficient control sets. Section 5 describes the algorithm, based on the Markov chains introduced in the previous section, and presents some simulation results. Finally, Section 6 points out to some directions for further studies.

1.1 Notation

Vectors are indicated in bold-face letters $\mathbf{x}, \mathbf{y}, \mathbf{z}$. For \mathbf{x}, \mathbf{y} , two vectors of the same dimension, the notation $\mathbf{x} \leq \mathbf{y}$ indicates that \mathbf{x} is lower or equal *component-wise* than \mathbf{y} . We define as usual the binary vectors δ_i : $(\delta_i)_i = 1$ and $(\delta_i)_j = 0$ for every $j \neq i$. If $S \subseteq \{1, \dots, n\}$, we put $\mathbf{1}_S = \sum_{i \in S} \delta_i$. Every $\mathbf{x} \in \{0, 1\}^n$ can be written as $\mathbf{x} = \mathbf{1}_S$ for some $S \subseteq \{1, \dots, n\}$. We call such a subset S the *support* of \mathbf{x} and we denote it $S_{\mathbf{x}}$. We use the notation $\mathbf{0}$ and $\mathbf{1}$ to denote, respectively, the vector of all 0 and the vector of all 1.

2. CONTROLLED MAJORITY DYNAMICS

We consider a set of players $\mathcal{V} = \{1, \dots, n\}$ connected by a simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the link set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is such that $(i, i) \notin \mathcal{E}$ for any i and $(i, j) \in \mathcal{E}$ if and only if $(j, i) \in \mathcal{E}$. We shall denote by $N_i = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\}$ the neighborhood of agent i in \mathcal{G} .

Let $\mathcal{A} = \{0, 1\}$ be the binary set of possible actions for all players. We denote by $\mathcal{X} = \mathcal{A}^n$ the set of strategy profiles and we define the *majority game* on \mathcal{G} as the game where each player $i \in \mathcal{V}$ has utility $u_i^c : \mathcal{X} \rightarrow \mathbb{R}$ given by

$$u_i^c(\mathbf{x}) = |\{j \in N_i \mid x_j = x_i\}|$$

In other words, the utility $u_i^c(\mathbf{x})$ of player i is given by the number of neighbors of i with which she is in agreement.

As usual in game theory, given a strategy profile $\mathbf{x} \in \mathcal{X}$ and a player i , we indicate with \mathbf{x}_{-i} the strategy profile restricted to all players but i and we consequently write $\mathbf{x} = (x_i, \mathbf{x}_{-i})$. The best response sets are then defined as

$$\mathcal{B}_i^c(\mathbf{x}) = \operatorname{argmax}_{\alpha \in \mathcal{A}} u_i^c(\alpha, \mathbf{x}_{-i}).$$

Using the notation

$$n_{i,\alpha}(\mathbf{x}) = |\{j \in N_i \mid x_j = \alpha\}|, \quad \alpha = 0, 1 \quad (1)$$

to indicate the number of neighbors of an agent i playing action α in the strategy profile \mathbf{x} , such best response sets can be more explicitly described as

$$\mathcal{B}_i^c(\mathbf{x}) = \begin{cases} \{0\}, & \text{if } n_{i,0}(\mathbf{x}) > n_{i,1}(\mathbf{x}) \\ \{0, 1\}, & \text{if } n_{i,0}(\mathbf{x}) = n_{i,1}(\mathbf{x}) \\ \{1\}, & \text{if } n_{i,0}(\mathbf{x}) < n_{i,1}(\mathbf{x}), \end{cases}$$

where

$$\mathcal{N} = \{\mathbf{x} \in \mathcal{X} : x_i \in \mathcal{B}_i^c(\mathbf{x}) \forall i \in \mathcal{V}\}$$

denotes the set of Nash equilibria. Notice that the set of Nash equilibria \mathcal{N} depends on the topology of the graph \mathcal{G} . However, for every graph \mathcal{G} , the constant strategy profiles and $\mathbf{1}$ are always Nash equilibria.

It is well known that the majority game is a potential game with potential function

$$\Phi_c(\mathbf{x}) = |\{(i, j) \in \mathcal{E}, x_j = x_i\}|. \quad (2)$$

This simply says that, for every two strategy profiles $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ such that $\mathbf{x}_{-i} = \mathbf{y}_{-i}$ for some player $i \in \mathcal{V}$,

$$\Phi_c(\mathbf{y}) - \Phi_c(\mathbf{x}) = u_i^c(\mathbf{y}) - u_i^c(\mathbf{x}). \quad (3)$$

(See Monderer and Shapley (1996).)

The (asynchronous) best response dynamics is a discrete time Markov chain X_t on the strategy profile space \mathcal{X} where, at every time t , a player i is chosen uniformly at random and she modifies her action choosing an element uniformly at random within $\mathcal{B}_i^c((X_t)_{-i})$. Denote with $P_{\mathbf{x},\mathbf{y}}$ the transition matrix (on $\mathcal{X} \times \mathcal{X}$) of the Markov X_t and note that, given $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} = (\alpha, \mathbf{x}_{-i})$, it holds that

$$P_{\mathbf{x},\mathbf{y}} > 0 \Leftrightarrow \Phi_c(\alpha, \mathbf{x}_{-i}) \geq \Phi_c(\mathbf{x}) \quad (4)$$

From the fact that the potential is not decreasing along the trajectories of X_t , it follows the classical result that, with probability 1, X_t converges in finite time to the set \mathcal{N} of Nash equilibria.

The question we pose is: what is the minimal number of agents that if forced to 1 will ensure almost surely that the best response dynamics reach the state $\mathbf{1}$.

Given a subset $\mathcal{C} \subseteq \mathcal{V}$, we indicate with $X_t^{\mathcal{C}}$ the Markov chain where only the agents in $\mathcal{V} \setminus \mathcal{C}$ update their action according to the best response rule defined above, while agents in \mathcal{C} maintain action 1. This new Markov takes values in the subset of strategy profiles

$$\mathcal{X}^{(\mathcal{C})} = \{x \in \mathcal{X} \mid x_i = 1 \forall i \in \mathcal{C}\}$$

This restricted game remains a potential game. This new dynamics will converge too to its set of Nash equilibria

$$\mathcal{N}^{\mathcal{C}} = \mathcal{N} \cap \mathcal{X}^{(\mathcal{C})}.$$

The following is our main object of study in this paper.

Definition 1. (Sufficient control set). A subset of players $\mathcal{C} \subseteq \mathcal{V}$ is a *sufficient control set* if

$$\mathbb{P}(\exists t : X_t^{\mathcal{C}} = \mathbf{1}, \mid X_0^{\mathcal{C}} = x_0) = 1 \quad (5)$$

for every initial strategy profile $x_0 \in \mathcal{X}^{(\mathcal{C})}$. A sufficient control set is *minimal* if none of its strict subsets is a sufficient control set. A sufficient control set is *optimal* if no sufficient control set has strictly smaller cardinality.

Our objective is to find optimal sufficient control sets. To give a more intuitive idea of what control sets resemble, a few illustrative examples are displayed in Figures 1–3.

3. MONOTONE CRUSADES AND VALID CONTROL SETS

In this section, we introduce the concept of monotone crusade, that will play a crucial role in our theory.




Fig. 1. Example: On trees, the set of the leaves (in green) is always a valid control set. In red you can see another valid control set, of size 2 and optimal for this particular tree




Fig. 2. An example of a clique, the sets of size $\lfloor \frac{k}{2} \rfloor - 1$ are exactly the minimal control sets and also the optimal control sets




Fig. 3. If all nodes have degree at most 2, then choosing one node per connected component gives a control set

These are particular sequences of strategy profiles, that start from a given configuration and end at $\mathbf{1}$, such that every step increases the number of players using strategy 1 without decreasing the potential. More formally:

Definition 2. (Adapted Monotone Crusade). Let $\mathcal{C} \subseteq \mathcal{V}$. A *monotone crusade from \mathcal{C}* is a sequence of vectors $\mathbf{x}^k \in \mathcal{X}$, for $k = 0, \dots, m$ such that

- (i) $\mathbf{x}^0 = \mathbf{1}_{\mathcal{C}}, \mathbf{x}^m = \mathbf{1}$
- (ii) for every $k = 1, \dots, m - 1$ there exists $i_k \in \mathcal{V} \setminus \mathcal{C}$ such that $\mathbf{x}^{k+1} = \mathbf{x}^k + \delta_{i_k}$

Moreover, if a function $V : \mathcal{X} \rightarrow \mathbb{R}$ satisfies

- (iii) $V(\mathbf{x}^{k+1}) \geq V(\mathbf{x}^k)$ for $k = 0, \dots, m - 1$,

then the sequence \mathbf{x}^k is called a *V-adapted monotone crusade from \mathcal{C}* .

A few comments on the above definition are in order:

Remark 1. All nodes i_1, \dots, i_m appearing in (2) are necessarily distinct otherwise the condition $\mathbf{x}^k \in \{0, 1\}^n$ for all k would be violated. Indeed, we must have that $\mathcal{V} \setminus \mathcal{C} = \{i_1, \dots, i_m\}$ and thus $m = |\mathcal{V} \setminus \mathcal{C}|$. This allows for a monotone crusade from \mathcal{C} to be equivalently characterized by the sequence of nodes (i_k) , the induced order on the nodes or by the sequence of increasing support sets (S_k) defined by $S_k = S_{\mathbf{x}^k}$, for $k = 0, \dots, m$ having the property that $S_0 = \mathcal{C}$ and $S_m = \mathcal{V}$.

Remark 2. We can also define a decreasing version of the monotone crusade where $\mathbf{x}^0 = \mathbf{1}, \mathbf{x}^m = \mathbf{1}_{\mathcal{C}}$ and where, for every $k, \mathbf{x}^{k+1} = \mathbf{x}^k - \delta_{i_k}$. This will be called a *decreasing crusade to \mathcal{C}* (*V-adapted* if (iii) in Definition 2 is satisfied).

Definition 3. (Valid control set). A set \mathcal{C} is *V-valid* if there exists a *V-adapted monotone crusade from \mathcal{C}* .

The main goal of the rest of this section is to show that the class of Φ_c -valid control sets (we recall that Φ_c is the potential of the majority game) coincides with the class of sufficient control sets defined in (Definition 1).

The following property is instrumental to our results.

Lemma 1. (Monotonicity of Coordination Game). For every two strategy profiles $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ and every player $i \in \mathcal{V}$, the following holds true:

- (1) if $\mathbf{x} \leq \mathbf{y}$ and $\Phi_c(\mathbf{1}, \mathbf{x}_{-i}) \geq \Phi_c(\mathbf{x})$ then $\Phi_c(\mathbf{1}, \mathbf{y}_{-i}) \geq \Phi_c(\mathbf{y})$;
- (2) if $\mathbf{x} \geq \mathbf{y}$ and $\Phi_c(\mathbf{0}, \mathbf{x}_{-i}) \geq \Phi_c(\mathbf{x})$ then $\Phi_c(\mathbf{0}, \mathbf{y}_{-i}) \geq \Phi_c(\mathbf{y})$.

Proof We only prove the first assertion, the second can be obtained by exchanging the role of 0 and 1.

If $y_i = 1$, we have that $(\mathbf{1}, \mathbf{y}_{-i}) = \mathbf{y}$ and there is nothing to prove. If $x_i = 1$, then the inequality $\mathbf{x} \leq \mathbf{y}$ ensure that $y_i = 1$ and we are thus in the previous case.

We now consider the case when both x_i and y_i have value 0.

Note that, for any configuration \mathbf{z} if \mathbf{z} is such that $z_i = 0$, the variation of the potential when player i changes her action from 0 to 1 can be expressed as

$$\Phi_c(\mathbf{1}, \mathbf{z}_{-i}) - \Phi_c(\mathbf{z}) = n_{i,1}(\mathbf{z}) - n_{i,0}(\mathbf{z})$$

where, we recall, $n_{i,1}(\mathbf{z})$ and $n_{i,0}(\mathbf{z})$ are the number of neighbors of i whose action is, respectively, 1 and 0.

As $\mathbf{x} \leq \mathbf{y}$, we have that $n_{i,0}(\mathbf{y}) \leq n_{i,0}(\mathbf{x})$ and $n_{i,1}(\mathbf{y}) \geq n_{i,1}(\mathbf{x})$. Hence,

$$\begin{aligned} \Phi_c(\mathbf{1}, \mathbf{y}_{-i}) - \Phi_c(\mathbf{y}) &= n_{i,1}(\mathbf{y}) - n_{i,0}(\mathbf{y}) \\ &\geq n_{i,1}(\mathbf{x}) - n_{i,0}(\mathbf{x}) = \Phi_c(\mathbf{1}, \mathbf{x}_{-i}) - \Phi_c(\mathbf{x}) \end{aligned}$$

This yields the thesis.

Proposition 1. (Monotonicity for the Inclusion). A superset of a Φ_c -valid control set is a Φ_c -valid control set.

Proof Assume that \mathcal{C} is a Φ_c -valid control set and let $\mathcal{C}' \supseteq \mathcal{C}$. Let \mathbf{x}^k be a Φ_c -adapted monotone crusade from \mathcal{C} with associated sequence of points (i_k) for $k = 1, \dots, m = n - |\mathcal{C}|$ such that $\mathbf{x}^{k+1} - \mathbf{x}^k = \delta_{i_k}$ for each k . Consider the subsequence of points $i_{k_1}, i_{k_2}, \dots, i_{k_{m'}}$ that are in $\mathcal{V} \setminus \mathcal{C}'$ and put $\mathbf{y}^h = \max\{\mathbf{1}_{\mathcal{C}'}, \mathbf{x}^{k_h}\}$. By construction, we have that $\mathbf{y}^h \geq \mathbf{x}^{k_{h+1}-1}$ and thus, by Lemma 1 and the fact that \mathbf{x}^k is a Φ_c -adapted monotone crusade from \mathcal{C} , we have that $\Phi_c(\mathbf{y}^h) \leq \Phi_c(\mathbf{y}^{h+1})$.

Remark 3. The full set is always a Φ_c -valid control set

The following result clarifies the connection between valid control sets for the majority game and sufficient control sets introduced in the previous section.

Theorem 1. (Characterization). A subset $\mathcal{C} \subseteq \mathcal{V}$ is a sufficient control set if and only if it is a Φ_c -valid control set.

Proof We first show that a sufficient control set is Φ_c -valid. If \mathcal{C} is a sufficient control set, there exists a sequence of vectors $\mathbf{y}^0, \dots, \mathbf{y}^T \in \mathcal{X}^{(\mathcal{C})}$ such that $\mathbf{y}^0 = \mathbf{1}_{\mathcal{C}}$ and $\mathbf{y}^T = \mathbf{1}$ that the best response dynamics follows with positive probability. This is equivalent to saying, using

the definition of best response dynamics (see in particular property (4)), that

- (1) $\mathbf{y}^{k+1} - \mathbf{y}^k = \pm \delta_{i_k}$ for all $k = 0, \dots, T - 1$;
- (2) $\Phi_c(\mathbf{y}^0) \leq \dots \leq \Phi_c(\mathbf{y}^T)$.

For every $i \notin \mathcal{C}$, define

$$k(i) = \min\{k = 1, \dots, T \mid \mathbf{y}^k - \mathbf{y}^{k-1} = \delta_i\}$$

that is the first time when agent i change her action to 1 in the sequence \mathbf{y}^t . Order now the agents in $\mathcal{V} \setminus \mathcal{C}$ as i_1, \dots, i_m in such a way that $k_{i_1} < k_{i_2} < \dots < k_{i_m}$. Consider the increasing monotone crusade \mathbf{x}^h associated with the sequence of points (i_h) , namely,

$$\mathbf{x}^h = \mathbf{1}_C + \sum_{h' \leq h} \delta_{i_{h'}}$$

and notice that $\mathbf{x}^{h-1} \geq \mathbf{y}^{k(i_h)-1}$. Since $\Phi_c(\mathbf{y}^{k(i_h)-1}) \leq \Phi_c(\mathbf{y}^{k(i_h)})$, it follows from Lemma 1 that $\Phi_c(\mathbf{x}^{h-1}) \leq \Phi_c(\mathbf{x}^h)$. This tells us that \mathbf{x}^h is a Φ_c -adapted monotone crusade from \mathcal{C} and, thus, \mathcal{C} a Φ_c -valid control set.

We now show that a Φ_c -valid control set \mathcal{C} is sufficient. We fix any initial condition $\mathbf{x}_0 \in \mathcal{X}^{(C)}$ and we put $\mathcal{C}' = S_{\mathbf{x}_0}$. \mathcal{C}' is a superset of \mathcal{C} and, on the basis of Proposition 1, \mathcal{C}' is also a Φ_c -valid control set. Let \mathbf{x}^k be a corresponding Φ_c -adapted monotone crusade from \mathcal{C}' . By the properties of adapted monotone crusades (properties 2. and 3. in Definition 2) and the characterization (4) of the transition matrix of the best response dynamics X_t , starting from \mathbf{x}_0 , the Markov X_t will follow such a sequence with positive probability. Thus, from any initial condition, X_t will reach $\mathbf{1}$ with positive probability. The standard result on Markov that if there is a state such that for any other state there exists a path to it with nonzero probability, then the state will be visited with probability one, then yields the claim.

4. MARKOV CHAINS AND BACKWARD SEARCH ALGORITHMS

The characterization of sufficient control sets through the concept of monotone crusades suggest the possibility that such sets can be found starting from the strategy profile $\mathbf{1}$, iteratively replacing 1's with 0's in the attempt to follow backwards a monotone crusade. To this aim we now introduce a family of Markov chains Z_t^ϵ on the binary space \mathcal{X} , parameterized by $\epsilon \in [0, 1]$ that will be the core part of our algorithms.

Transitions of Z_t^ϵ are described as follows:

At every discrete time, a node uniformly at random i is activated. If her neighbors with current action 1 ($n_{i,1}$) are strictly less than her neighbors with current action 0 ($n_{i,0}$), she stays still. Otherwise, if her action is 1 it changes to 0 with probability 1, while if her action is 0, she changes to 1 with probability ϵ .

The only non zero non trivial transition probabilities of Z_t^ϵ are the following. Given $\mathbf{x} \in \mathcal{X}$,

$$\begin{aligned} x_i = 1, n_{i,1}(\mathbf{x}) \geq n_{i,0}(\mathbf{x}) &\Rightarrow P_{\mathbf{x},(0,\mathbf{x}_{-i})}^\epsilon = 1/n \\ x_i = 0, n_{i,1}(\mathbf{x}) \geq n_{i,0}(\mathbf{x}) &\Rightarrow P_{\mathbf{x},(1,\mathbf{x}_{-i})}^\epsilon = \epsilon/n \end{aligned} \quad (6)$$

In the case when $\epsilon = 0$, only transitions from 1 to 0 are allowed. In this case, the Markov has absorbing points. The relation of these points with sufficient control sets is studied in the next result.

We denote

$$\begin{aligned} \mathcal{Z} &= \{x \in \mathcal{X} \mid \mathbb{P}(\exists t_0 : Z_{t_0}^0 = x \mid Z_0^0 = \mathbf{1}) > 0\} \\ \mathcal{Z}^\infty &= \{x \in \mathcal{X} \mid \mathbb{P}(\exists t_0 : Z_t^0 = x \forall t \geq t_0 \mid Z_0^0 = \mathbf{1}) > 0\} \end{aligned}$$

the sets of reachable and absorbing state of the chain Z^0 .

Theorem 2. The following facts hold:

- (1) \mathcal{C} is a sufficient control set if and only if $\mathbf{1}_C \in \mathcal{Z}$;
- (2) if \mathcal{C} is a minimal sufficient control set, then $\mathbf{1}_C \in \mathcal{Z}^\infty$.

Proof

(1): By definition, if $\mathbf{x} = \mathbf{1}_C \in \mathcal{Z}$, there exists a sequence of strategy profiles \mathbf{y}^k , for $k = 0, \dots, m$ such that $\mathbf{y}^0 = \mathbf{1}$ and $\mathbf{y}^m = \mathbf{1}_C$ satisfying the properties

- (1) $\mathbf{y}^k - \mathbf{y}^{k+1} = \delta_{i_k}$ for all $k = 0, \dots, m - 1$;
- (2) $\Phi_c(\mathbf{y}^0) \geq \dots \geq \Phi_c(\mathbf{y}^T)$.

Then $x^k = \mathbf{y}^{m-k}$ is a Φ_c -adapted monotone crusade from \mathcal{C} and this yields that \mathcal{C} is a Φ_c -valid control set and thus also a sufficient control set by virtue of Theorem 1. Inverting this argument we prove the other implication.

(2): If \mathcal{C} is a minimal sufficient control set, we know from (1) that $\mathbf{1}_C \in \mathcal{Z}$. If, by contradiction, $\mathbf{1}_C \notin \mathcal{Z}^\infty$, then, from $\mathbf{x} = \mathbf{1}_C$, the Markov Z_t^0 could reach, in one step, a different state $\mathbf{x}' = \mathbf{1}_{C'}$ with $C' \subsetneq C$. This contradicts minimality.

Theorem 2 allows to reformulate the optimization problem as follows:

$$\min_{\mathbf{x} \in \mathcal{Z}} \|\mathbf{x}\|_1 \quad (7)$$

where $\|\mathbf{x}\|_1 = \sum_i x_i$. Optimal sufficient control sets \mathcal{C} are those for which $\mathbf{1}_C$ solves (7).

The Markov Z_t^0 is naturally related to the minority game whose definition we briefly recall thus: Given an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we define the minority game on \mathcal{G} as the binary game where each player $i \in \mathcal{V}$ has utility $u_i^a : \mathcal{X} \rightarrow \mathbb{R}$ given by

$$u_i^a(\mathbf{x}) = |\{j \in N_i \mid x_j \neq x_i\}|$$

that is simply the number of neighbors with which i is in disagreement. We denote by \mathcal{N}_a the set of Nash equilibria of the minority game. This game is potential with a potential Φ_a that is just the opposite than the potential of the majority game:

$$\Phi_a(\mathbf{x}) = -\Phi_c(\mathbf{x}) \quad (8)$$

The following property clarifies the relation of the minority game with our problem.

Proposition 2. $\mathcal{N}_a \subseteq \mathcal{Z}$: Nash equilibria of the minority game are valid control sets.

Proof Let $\mathbf{x} \in \mathcal{N}_a$ and let \mathbf{y}^k , for $k = 1, \dots, m$ be any decreasing crusade from $\mathbf{1}$ to \mathbf{x} . By construction, $\mathbf{x} \leq \mathbf{y}^k \leq \mathbf{1}$ for all k . Consider the sequence of nodes (i_k) such that $\mathbf{y}^{k-1} - \mathbf{y}^k = \delta_{i_k}$. We have that $\mathbf{x}_{i_k} = 0$ for every k . For all k , since \mathbf{x} is a Nash equilibrium then 0 is in i_k 's minority best response, thus $\Phi_a(\mathbf{x}) \geq \Phi_a(\mathbf{1}, \mathbf{x}_{-i_k})$, or equivalently, $\Phi_c(\mathbf{x}) \leq \Phi_c(\mathbf{1}, \mathbf{x}_{-i_k})$. By Lemma 1, it follows that $\Phi_c(\mathbf{y}^k) \leq \Phi_c(\mathbf{1}, \mathbf{y}_{-i_k}^k) = \Phi_c(\mathbf{y}^{k-1})$. Therefore $\mathbf{x}^h = \mathbf{y}^{m-h}$ is a Φ_c -adapted monotone crusade from $\mathcal{C} = S_{\mathbf{x}}$. By virtue of Theorem 2, we have that $\mathbf{x} \in \mathcal{Z}$.

We have the following simple but not obvious consequence.

Corollary 1. (Existence). For any graph, there exists a sufficient control set whose size is less than or equal to half the total number of nodes.

Proof Since the minority game is a potential game, it admits at least one Nash equilibrium $\mathbf{x} \in \mathcal{N}_a$. By symmetry $\tilde{\mathbf{x}} = \mathbf{1} - \mathbf{x}$ is also a Nash equilibrium. Proposition 2 guarantees that they are both sufficient control sets and one of the two necessarily contains not more than half the nodes in the network.

Notice that the reciprocal of the second part of Theorem 2 and the reciprocal of Proposition 2 are not true in general, as shown by the following examples.




Fig. 4. The red set above is a minimal control set with respect to the inclusion, yet the rightmost node is not in her best response, making it not a Nash equilibrium.




Fig. 5. The red set above is optimal, but not a Nash equilibrium.




Fig. 6. Two sufficient control sets (red nodes) on the same graph, both in \mathcal{Z}^∞ only the one on the right is minimal.

The above considerations show that we cannot directly use the dynamics Z_t^0 as an algorithm to find optimal sufficient control sets as its absorbing state \mathcal{Z}^∞ are not minimal in general. To overcome this difficulty we will instead use the Markov Z_t^ϵ with $\epsilon > 0$. The presence of transitions in the opposite directions allows the algorithm to make a full exploration of the set \mathcal{Z} .

Theorem 3. Let $\epsilon > 0$. The following facts hold:

- (1) Z_t^ϵ is ergodic inside the set of states \mathcal{Z} ;
- (2) Z_t^ϵ is time-reversible and its unique invariant probability is given by $\mu^\epsilon(\mathbf{x}) := K\epsilon^{||\mathbf{x}||_1}$ where $K > 0$ is the normalization constant;
- (3) As ϵ tends to 0, μ^ϵ converges to a probability measure μ concentrated on the subset $\operatorname{argmin}_{\mathbf{x} \in \mathcal{Z}} ||\mathbf{x}||_1$.

We do not present a proof of Theorem 3 here due to space limitations and since Theorem 3 is a special case of a more general result proved in Durand et al. (2020) for supermodular games.

By virtue of the reformulation (7), we have that, for small ϵ and sufficiently large t , the Markov Z_t^ϵ will spend most of the time in strategy profiles whose support are optimal sufficient control sets. This observation is at the base of a practical algorithm described in the next section.

5. IMPLEMENTATION AND SIMULATIONS

We have implemented an iterative algorithm based on the Markov chain Z_t^ϵ studied in the previous section. For the sake of increasing the speed of convergence, we actually considered a modification of the Markov chain Z_t^ϵ with all trivial self-loop transitions removed. This induces a little modification in the invariant probability of the Markov chain, but does not affect the minimal set that is the output of the algorithm. Also, our algorithm keeps track of the best strategy profile (the smallest $||x||_1$) found so far. This algorithms is written in Algorithm 1.

We have applied the algorithm to random realizations of Erdős-Rényi graphs with different number of nodes n and probability $p = 1/2$. For every value of n , we ran 500 executions on 20 randomly generated graphs. The algorithm is stopped after $100n$ iterations, using for epsilon the constant value $\epsilon = 0.2$.

As a point of comparison, we computed a benchmark optimum consisting of the exhaustive optimum for small graphs, and a much longer execution on bigger graphs. Figure 7 shows the average values of the size of the sufficient control sets computed by the algorithm. We compare it with the benchmark optimum and also with the result obtained looking at the very last step of the algorithm. This plot shows a remarkable performance of the algorithm that in linear time gets quite close to the optimum. It also shows that the Markov, though fluctuating, as $\epsilon > 0$, still remains close to the optimal strategy profiles. The important question of how to tune the parameter ϵ for optimize performance has not been addressed here.

Finally, notice how optimal sufficient control sets are scaling linearly with respect to the size of the graph. This suggests that Erdős-Rényi random graphs are somewhat hard to control in this sense. In other terms, they show resilience to this type of external actions.

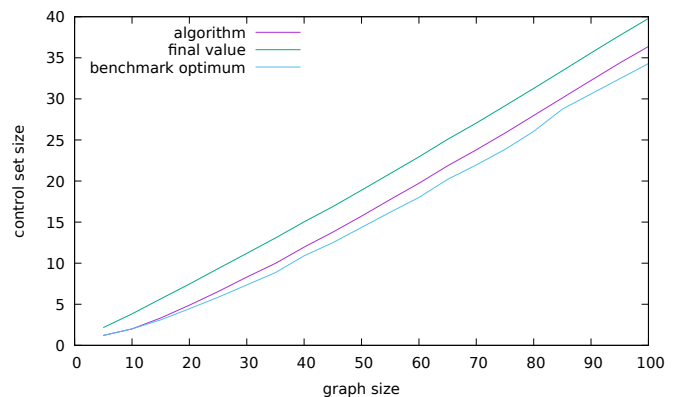


Fig. 7. In blue, the benchmark optimum. In purple, the minimal control set encountered in $100n$ steps. In green, the size of the control set at the end.

6. CONCLUSION

Graph(Adjacency list of sets \mathbf{A} :
 $i \leftrightarrow j \Leftrightarrow i \in A_j \wedge j \in A_i$;
 List of degree-threshold \mathbf{L} : max number of 0 among
 the neighbor for the strategy 1 to be the best response;
 0-partial degree list \mathbf{D} : D_i is the number of i 's
 neighbor with strategy 0, initialized to $\forall i, D_i^0 \leftarrow 0$;
 Strategy \mathbf{X} , initialized to $\forall i, X_i^0 \leftarrow 1$;
 number $n1$ of 1, initialized to $n1 \leftarrow |V|$)

Input: graph's \mathbf{A} and \mathbf{L} , trajectory length T

Output: a Valid Control Set

$m \leftarrow |V|$, $S \leftarrow \{0, 1 \dots, |V|\}$;

Initialize $\mathbf{D}, \mathbf{X}, n1$;

for $t = 0$ **to** T **do**

 // weights

forall $i \in V$ **do**

if $D_i \leq T_i$ **then**

if $X_i = 1$ **then**

 | $w_i \leftarrow 1$

else

 | $w_i \leftarrow \epsilon$

end

else

 | $w_i \leftarrow 0$

end

end

 // random value, conditioned on the step
 having an effect. The sum is never 0 as
 the last node chosen still verifies the
 inequality

 choose j with probability $\frac{w_j}{\sum w_i}$;

 // effect of the step

if $X_j = 0$ **then**

forall $k \in A_j$ **do**

 | $D_k \leftarrow D_k - 1$

end

$X_j \leftarrow 1$;

$n1 \leftarrow n1 + 1$;

else

forall $k \in A_j$ **do**

 | $D_k \leftarrow D_k + 1$

end

$X_j \leftarrow 0$;

$n1 \leftarrow n1 - 1$;

end

 // comparison

if $m < n1$ **then**

$m \leftarrow n1$;

$CS \leftarrow \emptyset$;

forall $i \in V$ **do**

if $X_i = 1$ **then**

 | $CS = CS \cup \{i\}$

end

end

end

return CS ;

end

Algorithm 1: Pseudocode for the algorithm employed
 in the simulations.

We have formulated the problem of finding, in a network coordination game, the minimum number of players to be controlled in order to drive the system from one Nash equilibrium to another one. To the scope, we have designed a low complexity randomized algorithm and proven its convergence properties. We have finally carried on some numerical simulations corroborating the results.

Many challenging issues naturally arise from our analysis and simulations. Erdős-Rényi graphs exhibited optimal sufficient control sets growing linearly in the number of nodes. It would be of interest to prove this analytically, as well to extend the analysis to other graph families, connecting resilience properties to the topological structure.

The problem studied in this paper is an instance of a more general problem of studying the effect of control actions in games. Future research will analyze similar control problems for general games with strategic complements as well strategic substitutes. A first step in this direction can be found in Durand et al. (2020).

REFERENCES

- Acemođlu, D., Como, G., Fagnani, F., and Ozdaglar, A. (2013). Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research*, 38(1), 1–27.
- Como, G. and Fagnani, F. (2016). From local averaging to emergent global behaviors: The fundamental role of network interconnections. *Systems and Control Letters*, 95, 70–76.
- 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.
- Drakopoulos, K., Ozdaglar, A., and Tsitsiklis, J.N. (2016). When is a network epidemic hard to eliminate? *Mathematics of Operations Research*, 42(1-14).
- Durand, S., Como, G., and Fagnani, F. (2020). Optimal targeting in supermodular games. <https://staff.polito.it/giacomo.como/material/TS.pdf>.
- Jackson, M. and Zenou, Y. (2015). Games on networks. *Handbook of Game Theory with Economic Applications*.
- Kempe, D., Kleinberg, J., and Tardos, E. (2003). Maximizing the spread of influence through a social network. In *Proceedings of SIGKDD'03*, 137–146.
- Monderer, D. and Shapley, L.S. (1996). Potential games. *Games and Economic Behavior*.
- Morris, S. (2000). Contagion. *The Review of Economic Studies*.
- Rossi, W., Como, G., and Fagnani, F. (2019). Threshold models of cascades in large-scale networks. *IEEE Transactions on Network Science and Engineering*, 6(2), 158–172.
- Vassio, L., Fagnani, F., Frasca, P., and Ozdaglar, A. (2014). Message passing optimization of harmonic influence centrality. *IEEE Transactions on Control of Network Systems*, 1(1), 109–120.
- Yildiz, E., Ozdaglar, A., Acemođlu, D., Saberi, A., and Scaglione, A. (2013). Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*, 1(4), 19–30.