Automatica 90 (2018) 204–211

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Consensus in opinion dynamics as a repeated game ☆

Dario Bauso [a,b,*], Mark Cannon [c]

[a] Department of Automatic Control and Systems Engineering, The University of Sheffield, Mappin Street Sheffield, S1 3JD, United Kingdom
[b] Dipartimento di Ingegneria Chimica, Gestionale, Informatica, Meccanica, Università di Palermo, V.le delle Scienze, 90128 Palermo, Italy
[c] Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, United Kingdom

ARTICLE INFO

Article history:
Received 10 December 2014
Received in revised form 29 August 2017
Accepted 30 November 2017
Available online 15 February 2018

Keywords:
Game theory
Networks
Allocations
Robust receding horizon control

ABSTRACT

We study an n-agent averaging process with dynamics subject to controls and adversarial disturbances. The model arises in multi-population opinion dynamics with macroscopic and microscopic intertwined dynamics. The averaging process describes the influence from neighbouring populations, whereas the input term indicates how the distribution of opinions in the population changes as a result of dynamical evolutions at a microscopic level (individuals' changing opinions). The input term is obtained as the vector payoff of a two player repeated game. We study conditions under which the agents achieve robust consensus to some predefined target set. Such conditions build upon the approachability principle in repeated games with vector payoffs.

1. Introduction

We consider an n-agent averaging process in which each agent is described by a dynamic system with controlled and uncontrolled inputs, the latter being adversarial disturbances.

We specialize the model to multi-population opinion dynamics. The averaging process describes the influence from neighbouring populations, whereas the input term indicates how the distribution of opinions in the population changes as a result of dynamical evolutions at a microscopic level (individuals' changing opinions). The input term is obtained as the vector payoff of a two player repeated game (Bauso, Lehrer, Solan, & Venel, 2015; Blackwell, 1956; Lehrer, 2002). Motivations for the dynamics can be found in coalitional games with Transferable Utilities (TU games) (von Neumann & Morgenstern, 1944), bargaining (Bauso & Notarstefano, 2012; Nedić & Bauso, 2013), consensus (Liu, Xie, & Wang, 2012; Nedić, Ozdaglar, & Parrilo, 2010; Shi & Hong, 2009; Sundhar Ram, Nedić, & Veeravalli, 2009), opinion dynamics (Acemoğlu, Como, Fagnani, & Ozdaglar, 2013; Acemoğlu & Ozdaglar, 2011; Aeyels & De Smet, 2008; Banerjee, 1992; Blondel, Hendrickx, & Tsitsiklis, 2010; Castellano, Fortunato, & Loreto, 2009; Como & Fagnani, 2011; Hegselmann & Krause, 2002; Krause, 2000; Pluchino, Latora, & Rapisarda, 2006; Sznitman, 1991) and in multi-population games with macroscopic and microscopic dynamics.

The main contribution of this paper is to introduce a distributed multi-stage receding horizon control strategy that ensures the existence of invariant and contractive sets for the collective dynamics, and which can be used to enforce convergence of consensus to a specified set.

This paper improves (Bauso, Cannon, & Fleming, 2014) as it links the model to opinion dynamics and multi-population games, it includes exponential stability and identifies regions of attraction that are dependent and independent of the horizon, and it provides new numerical results. An alternative way to deal with the problem is to include a deterministic adversarial disturbance in the spirit of set inclusion theory (Hofbauer, Benaïm, & Sorin, 2005, 2006).

The paper is organized as follows. In Section 2 we formulate the problem. In Section 3 we discuss motivations. Section 4 gives the main control theoretic results. Numerical illustrations are presented in Section 5, and concluding remarks are provided in Section 6.

Notation. We denote the Euclidean norm of a vector x as $\|x\|$, and we use a^i_j or $[A]_{ij}$ to denote the ijth entry of a matrix A. We say that $A \in \mathbb{R}^{n\times n}$ is row-stochastic if $a^i_j \geq 0$ for all $i, j \in \{1, \ldots, n\}$ and $\sum_{j=1}^n a^i_j = 1$ for all $i \in \{1, \ldots, n\}$. Matrix A is doubly stochastic if both A and its transpose $A^\top$ are row-stochastic. We use $|S|$ for the cardinality of a given finite set S. We write $P_X[x]$ to denote the projection of a vector x on a set X, and we write $|x|_X$ for the distance from x to X, i.e., $P_X[x] = \arg\min_{y\in X} \|x - y\|$ and $|x|_X = \|x - P_X[x]\|$, respectively.

☆ The material in this paper was partially presented at the 19th World Congress of the International Federation of Automatic Control (IFAC 2014), August 24–August 29, 2014, Cape Town, South Africa. This paper was recommended for publication in revised form by Associate Editor Hideaki Ishii under the direction of Editor Ian R. Petersen.

* Corresponding author at: Department of Automatic Control and Systems Engineering, The University of Sheffield, Mappin Street Sheffield, S1 3JD, United Kingdom.

E-mail addresses: d.bauso@sheffield.ac.uk (D. Bauso), mark.cannon@eng.ox.ac.uk (M. Cannon).

https://doi.org/10.1016/j.automatica.2017.12.062
0005-1098/

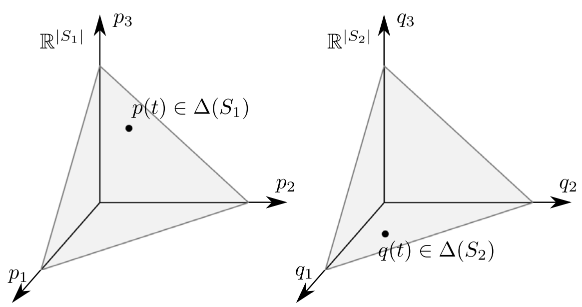

Fig. 1. Spaces of mixed strategies for the two players.

2. Model and problem set-up

For each i in a set $N = \{1, \dots, n\}$, agent i is characterized by a state $x_i(t) \in \mathbb{R}^{\tilde{n}}$. At every time t this state evolves according to a distributed averaging process representing the interaction of the agent with its neighbours, and under the influence of an input variable $u_i(t)$.

Formally, the state $x_i(t)$ of agent i evolves as follows:

$$x_i(t+1) = \sum_{j=1}^{n} a_j^i(t) x_j(t) + u_i(t), \; t = 0, 1, \dots \tag{1}$$

where $a^i = (a_1^i, \dots, a_n^i) \in \mathbb{R}^n$ is a vector of nonnegative scalar weights relating to the *communication graph* $\mathcal{G}(t) = \big(N, \mathcal{E}(t)\big)$. A link $(j, i) \in \mathcal{E}(t)$ exists (and hence $a_j^i(t) \neq 0$) if agent j is a neighbour of agent i at time t.

For each agent $i \in N$, the input $u_i(\cdot)$ is the payoff of a repeated two-player game between player i (Player A) and an (external) adversary (Player B). Let S_A and S_B be the finite sets of actions of players A and B, respectively, and let us denote the set of mixed action pairs by $\Delta(S_A) \times \Delta(S_B)$ (set of probability distributions on S_A and S_B). For any pair of mixed strategies $(p(t), q(t)) \in \Delta(S_A) \times \Delta(S_B)$ for player A and B at time t, the expected payoff is

$$\begin{cases} u_i(t) = \displaystyle\sum_{j \in S_A, k \in S_B} p_j^i(t) \phi(j, k) q_k^i(t), \\ \displaystyle\sum_{j \in S_A} p_j^i(t) = 1, \;\; \sum_{k \in S_B} q_k^i(t) = 1, \;\; p_j^i, q_k^i \geq 0. \end{cases} \tag{2}$$

Essentially, in the above game $\phi(j, k) \in \mathbb{R}^{\tilde{n}}$ is the vector payoff when players A and B play pure strategies $j \in S_A$ and $k \in S_B$, respectively. Fig. 1 illustrates the continuous action sets for the two players, for the case that $S_A = \{1, 2, 3\}$ and $S_B = \{1, 2, 3\}$.

Let $X \subset \mathbb{R}^{\tilde{n}}$ be a closed convex target set, and assume that player A seeks to drive the state $x_i(t)$ to X, while player B tries to push the state far from it. The resulting strategy can be formulated as the solution of a robust optimization problem, with one player minimizing and the other maximizing the distance of the state from X.

In compact form the problem with finite horizon $[0, T]$ to be solved by agent i takes the form:

$$\min_{p^i(0)} \max_{q^i(0)} \cdots \min_{p^i(T-1)} \max_{q^i(T-1)} \sum_{t=0}^{T} |x_i(t)|_X^2$$
$$\left. \begin{array}{l} p^i(t) \in \Delta(S_A), \;\; q^i(t) \in \Delta(S_B), \\ x_i(t+1) = y_i(t) + u_i(t), \\ u_i(t) = \displaystyle\sum_{j \in S_A, k \in S_B} p_j^i(t) \phi(j, k) q_k^i(t) \end{array} \right\} t = 0, \dots, T-1 \tag{3}$$

where $y_i(t)$ is the *space average* defined as

$$y_i(t) = \sum_{j=1}^{n} a_j^i(t) x_j(t). \tag{4}$$

Through the above problem we can study contractivity and invariance of sets for the collective dynamics (1)–(2). In the following we simplify notation and drop the dependence on i of p and q.

3. Multi-population opinion dynamics

A simple model of opinion dynamics is derived from a classical model of consensus dynamics that also arises in the Kuramoto oscillator model (Pluchino et al., 2006). In this perspective, the dynamic model (1) appears as a discrete-time model of a consensus problem (Olfati-Saber, Fax, & Murray, 2007), in which the coupling term accounts for emulation (an individual's opinion is influenced by those of its neighbours), and which includes an additional input term (the natural opinion changing rate). In addition, the target set X can be used to enforce consensus. For instance we can set $X := \{x\} \subset \mathbb{R}^{\tilde{n}}$, in which case $\lim_{t \to \infty} x_i(t) = x$ for all $i \in N$ also implies that $\lim_{t \to \infty} x_i(t) - x_j(t) = 0$ for all $i, j \in N$. Note that this notion of consensus may be in general different from the consensus studied in distributed algorithms (Olfati-Saber et al., 2007).

In the following we consider n distinct populations of agents interacting according to a predefined topology. Let the collective state be $\xi(t) = \big(x_1(t), \dots, x_n(t)\big)$, which we now see as a collection of n macro-states. For each population, and at every time $t \in [0, T]$, a probability distribution function $x_i(t)$, $i \in N$, describes the probability distribution of agents over a discrete set of micro-states. In other words, consider a finite discrete space of micro-states $\{1, \dots, \tilde{n}\}$, and let a probability distribution function be given, $m_i : \{1, \dots, \tilde{n}\} \times [0, +\infty) \to [0, 1], (j, t) \mapsto m_i(j, t)$, which satisfies $\sum_{j \in \{1, \dots, \tilde{n}\}} m_i(j, t) = 1$ for every t. Now, let us collect all distribution values $m_i(j, t), j \in \{1, \dots, \tilde{n}\}$ in the macro-state vector of population i, namely:

$$x_i(t) := \big(m_i(1, t), m_i(2, t) \dots, m_i(\tilde{n}, t)\big) \in [0, 1]^{\tilde{n}}.$$

Thus, the averaging term in (1) describes the influence from neighbour populations.

As for the input term, consider, from a microscopic perspective, the case that the political opinions in a single population are distributed between two states, vote *left* and vote *right*, and such a distribution is subject to transitions from one state to the other. This is represented by the network depicted in Fig. 2 where nodes 1 and 2 correspond to the two states. Two persuaders, one of which is the controller (player A), the other the disturbance (player B), can influence the transitions described by the controlled flows $\hat{v}_j$, $j = 1, \dots, 4$ and disturbance parameters $\hat{w}_k$, $k = 3, 4$. In particular, player A can influence all the transitions, while player B has influence only on the transitions from node 2.

More generally, the terms $\hat{v}_j$ and $\hat{w}_k$ determine the transition rates between state 1 (vote *left*) and state 2 (vote *right*). In other words, a political campaign can make voters change their political opinion, and the controlled transition rates $\hat{v}_j$, $j = 2, 4$ represent the rates of change from one state to the other as a consequence of such a deliberate action. The parameters $\hat{w}_k$ modulate these flows and are representatives of unpredicted or uncontrolled events that can influence voters' opinions.

In this case $\tilde{n} = 2$ and the evolution of the distribution is given by

$$x_i(t+1) = \Big(I + B\tilde{B}^\top \big(\hat{v}_2(t), \hat{v}_4(t)\big) + D\tilde{D}^\top \big(\hat{v}_4(t), \hat{w}_4(t)\big)\Big) x_i(t) \tag{5}$$