Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Robust consensus in social networks and coalitional games ⋆

Dario Bauso * **Mark Cannon** ** **James Fleming** **

* *DICGIM, Università di Palermo, Viale delle Scienze I-90128, Palermo, Italy, dario.bauso@unipa.it.*
** *Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, UK, mark.cannon@eng.ox.ac.uk/james.fleming@worc.ox.ac.uk.*

Abstract: We study an n-player averaging process with dynamics subject to controls and adversarial disturbances. The model arises in two distinct application domains: i) coalitional games with transferable utilities (TU) and ii) opinion propagation. We study conditions under which the average allocations achieve robust consensus to some predefined target set.

Keywords: Game theory, networks, allocations, robust receding horizon control.

1. INTRODUCTION

We consider an n-player averaging process in which each player is described by a dynamic system with controlled and uncontrolled inputs, the latter being adversarial disturbances. Motivations for the dynamics can be found in two distinct application domains. Within the realm of allocation problems in coalitional games with Transferable Utilities (TU games) (Nedić and Bauso, 2013), the process describes a two-step distributed allocation process where, at each discrete time step, players first renegotiate their past allocations and second generate a new revenue and allocate it. On the other hand, the dynamics also arise naturally in opinion propagation, where the players adjust their opinions based on the inputs received from neighbors, and under the influence of persuaders.

The goal of our study is to provide a detailed analysis of the consensus dynamics under the umbrella of contractivity and invariance theory. In particular, we consider a predefined set and study convergence of the consensus value to such a set under a distributed receding horizon control law. This set can be thought of as (but it is not limited to) the core of the game for the allocation process. This is the set of imputations under which no coalition has a value greater than the sum of its members' payoffs. Therefore, no coalition has incentive to leave the grand coalition and receive a larger payoff. A similar problem is addressed by Bauso and Notarstefano (2012), which considers the case of allocations evolving according to doubly averaging dynamics (over time as well as space).

The main contribution of this paper is to introduce a distributed multi-stage receding horizon control strategy and to show that it ensures the existence of invariant and contractive sets for the collective dynamics (Theorem 1).

Related literature. Coalitional games with transferable utilities (TU) were first introduced by von Neumann and Morgenstern (1944). In this work a central issue is to determine whether the core is an "approachable" set, and which allocation processes can drive the "complaint vector" to that set. Approachability theory was developed by Blackwell (1956), and is captured in the well known Blackwell's Theorem. The geometric (approachability) principle that lies behind Blackwell's Theorem is among the fundamentals in allocation processes in coalitional games (Lehrer, 2002). The discrete-time dynamics analyzed in this paper follow the rules of typical consensus dynamics (see e.g. Nedić, Ozdaglar, and Parrilo, 2010, and references therein). Consensus dynamics also arise in the literature on agreement among multiple agents. An underlying communication graph for the agents and balancing weights are used to reach an agreement on a common decision variable in Nedić et al. (2010) and Ram et al. (2009) for distributed multi-agent optimization, and related problems of aggregation in multi-agent systems are studied in Shi and Hong (2009) and Liu et al. (2012).

The paper is organized as follows. In Section 2 we formulate the problem while Section 3 gives the main results. Section 4 provides numerical illustrations. Finally, in Section 5, provides concluding remarks and future directions.

Notation. We let x' denote the transpose of a vector x, and $\|x\|$ denote its Euclidean norm. An $n \times n$ matrix A is row-stochastic if the matrix has nonnegative entries a^i_j and $\sum_{j=1}^n a^i_j = 1$ for all $i = 1, \ldots, n$. For a matrix A, we use a^i_j or $[A]_{ij}$ to denote its ijth entry. A matrix A is doubly stochastic if both A and its transpose A' 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 work of D. Bauso was supported by the 2012 "Research Fellow" Program of the Dipartimento di Matematica, Università di Trento and by PRIN 20103S5RN3 "Robust decision making in markets and organizations, 2013-2016". D. Bauso is currently an academic visitor at the Department of Engineering Science, University of Oxford, UK. J. Fleming is supported by an EPSRC Doctoral Training Studentship.

2. DISTRIBUTED REWARD ALLOCATION

Every player in a set $N = \{1, \ldots, n\}$ is characterized by an average allocation vector $x_i(t) \in \mathbb{R}^n$. At every time he renegotiates with *neighbors* all past allocations and generates a new allocation vector $u_i(t)$. Then, the cumulative (over time) allocation $x_i(t)$ to player i evolves as follows

$$x_i(t+1) = \sum_{j=1}^{n} a^i_j(t) x_j(t) + u_i(t), \; t = 0, 1, \ldots \tag{1}$$

where $a^i = (a^i_1, \ldots, a^i_n)'$ is a vector of nonnegative weights consistent with the sparsity of the *communication graph* $\mathcal{G}(t) = (N, \mathcal{E}(t))$. A link $(j, i) \in \mathcal{E}(t)$ exists (and hence $a^i_j(t) \neq 0$) if player j is a neighbor of player i at time t, i.e. if player i renegotiates allocations with player j at time t. Figure 1 displays a possible communication graph.

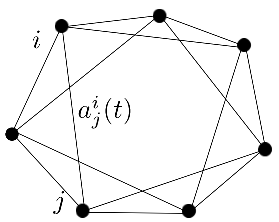

Fig. 1. Communication graph.

For each player $i \in N$, the input $u_i(\cdot)$ is the payoff of a repeated two-player game between player i (Player i_1) and an (external) adversary (Player i_2). Let S_1 and S_2 be the finite set of actions of players i_1 and i_2 respectively and let us denote the set of mixed action pairs by $\Delta(S_1) \times \Delta(S_2)$ (set of probability distributions on S_1 and S_2). For any pair of mixed strategies $(p(t), q(t)) \in \Delta(S_1) \times \Delta(S_2)$ for player i_1 and i_2 at time t, the expected payoff is

$$\begin{cases} u_i(t) = \displaystyle\sum_{j \in S_1, k \in S_2} p_j(t) \phi(j,k) q_k(t), \\ \displaystyle\sum_{j \in S_1} p_j(t) = 1, \; \sum_{k \in S_2} q_k(t) = 1, \;\; p_j, q_k \geq 0. \end{cases} \tag{2}$$

Essentially, in the above game $\phi(j, k) \in \mathbb{R}^n$ is the vector payoff when players i_1 and i_2 play pure strategies $j \in S_1$ and $k \in S_2$ respectively. Figure 2 illustrates the continuous action sets for the two players.

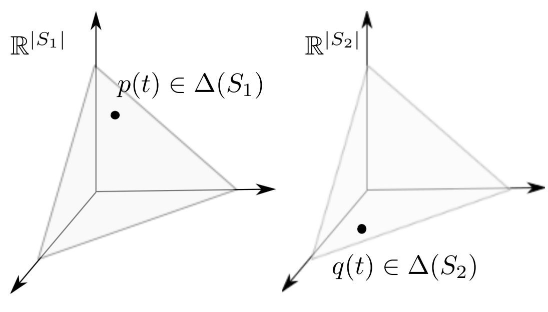

Fig. 2. Spaces of mixed strategies for the two players.

Let $X \subset \mathbb{R}^n$ be a closed convex target set (or a convex subset of a nonconvex target set), and consider the case where player i_1 seeks to drive the state $x_i(t)$ towards X, while player i_2 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 player i takes the form:

$$\begin{gathered} \min_{p(0)} \max_{q(0)} \cdots \min_{p(T-1)} \max_{q(T-1)} \sum_{t=0}^{T} |x_i(t)|^2_X \\ \left. \begin{aligned} & p(t) \in \Delta(S_1), \quad q(t) \in \Delta(S_2), \\ & x_i(t+1) = w_i(t) + u_i(t), \\ & u_i(t) = \sum_{j \in S_1, k \in S_2} p_j(t) \phi(j,k) q_k(t) \end{aligned} \right\} t = 0, \ldots, T-1 \end{gathered} \tag{3}$$

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

$$w_i(t) = \sum_{j=1}^{n} a^i_j(t) x_j(t). \tag{4}$$

Let $\xi(t) = \big(x_1(t), \ldots, x_n(t)\big)$ denote the collective state of all players in N at time t. We introduce a value function $V_{i,\tau}(\xi(t), t)$ representing the optimal cost over τ steps starting at $x_i(t)$, where $\tau = T - t$ for $t \in [0, T]$. Using dynamic programming and the Bellman principle, the value function must satisfy the following recursion

$$\begin{aligned} V_{i,\tau}(\xi(t), t) &= \min_{p(t) \in \Delta(S_1)} \max_{q(t) \in \Delta(S_2)} \Big\{ |x_i(t)|^2_X \\ &\qquad + V_{i,\tau-1}(\xi(t+1), t+1) \Big\} \\ &= |x_i(t)|^2_X + \min_{p(t) \in \Delta(S_1)} \max_{q(t) \in \Delta(S_2)} V_{i,\tau-1}(\xi(t+1), t+1) \end{aligned}$$

with final value $V_{i,0}(\xi(T), T) = |x_i(T)|^2_X$. The space average in (1) implies that the distance from X of the state x_i at any time more than one step into the future will depend on the current and future actions of players in N other than player i. To realise a distributed allocation strategy in which $u_i(t)$ is determined without collaboration or knowledge of the current or future actions of other players, we therefore incorporate maximization over $u_j(t)$, $j \in N$, $j \neq i$ when computing the worst case cost-to-go in the above problem formulation.

The receding horizon implementation of the optimal strategy for player i defines $p(t)$, and hence $u_i(t)$ in (2), as the minimizing argument for the T-stage problem with optimal value function $V_{i,T}(\xi(\cdot), \cdot)$. The stability of a receding horizon control law can be ensured (Mayne et al., 2000) by imposing a terminal constraint such as $x_i(T) \in X \; \forall i \in N$. However this would require collaboration or knowledge of the actions of other players in N, and instead we therefore impose the local constraint $|x_i(t+1)|_X \leq |w_i(t)|_X$. Since this constraint involves only $u_i(t)$ and $w_i(t)$, it does not depend on bounds on the allocations of other players which could lead to conservative terminal constraints. We show that it nonetheless results in a global contractivity property. The modified problem formulation can now be restated as

$$\begin{aligned} V_{i,\tau}(\xi(t), t) &= |x_i(t)|^2_X \\ &\quad + \min_{p(t) \in \Delta(S_1)} \max_{\substack{q(t) \in \Delta(S_2) \\ u_j(t), j \neq i, j \in N}} V_{i,\tau-1}(\xi(t+1), t+1) \\ &\qquad \text{subject to } |x_i(t+1)|_X \leq |w_i(t)|_X. \end{aligned} \tag{5}$$

Our goal is to study contractivity and invariance of sets for the collective dynamics (1)-(2). In particular, we consider the collective value function $\sum_{i=1}^{n} V_{i,T}(\xi(t), t)$ assuming each player $i \in N$ employs a T-stage receding horizon strategy with the optimal cost $V_{i,T}(\xi(\cdot), \cdot)$ defined in (5).

2.1 Motivations

Coalitional games. The set X introduced above can be thought of as the core of a coalitional game with Transferable Utilities (TU game).

A coalitional TU game is defined by a pair $\langle N, \eta \rangle$, where $N = \{1, \ldots, n\}$ is a set of players and $\eta : 2^N \to \mathbb{R}$ a function defined for each coalition $S \subseteq N$ $(S \in 2^N)$. The function η determines the value $\eta(S)$ assigned to each coalition $S \subset N$, with $\eta(\emptyset) = 0$. We let η_S be the value $\eta(S)$ of the characteristic function η associated with a nonempty coalition $S \subseteq N$. Given a TU game $\langle N, \eta \rangle$, let $C(\eta)$ be the core of the game,

$$C(\eta) = \left\{ x \in \mathbb{R}^n \;\Big|\; \sum_{j \in N} [x]_j = \eta_N, \right.$$
$$\left. \sum_{j \in S} [x]_j \geq \eta_S \text{ for all nonempty } S \subset N \right\}.$$

Essentially, the core of the game is the set of all allocations that make the grand coalition stable with respect to all subcoalitions. Condition $\sum_{j \in N} [x]_j = \eta_N$ is also called an efficiency condition. Condition $\sum_{j \in S} [x]_j \geq \eta_S$ for all nonempty $S \subset N$ is referred to as "stability with respect to subcoalitions", since it guarantees that the total amount given to the members of a coalition exceeds the value of the coalition itself. Also, the averaging process in (4) can be justified as *inequity aversion* on the part of the players.

Social networks. Opinion dynamics has attracted the attention of many scientists over the the past few years. The propagation of the opinions describe the time evolution of the beliefs of a large population of agents as a result of repeated interactions among the agents, in many cases over a social network (see e.g. Castellano et al., 2009, Sect. III, and Acemoğlu and Ozdaglar, 2011). In continuous opinion dynamics models, beliefs or opinions are represented by scalars or vectors, evolving according to some averaging process. The latter consists in each opinion moving towards a convex combination of (a subset of) other agents' current beliefs, thus modeling the attractive nature of social influence. There are many models that, under the assumption that the underlying social network is connected, prove that the agents' opinions reach consensus asymptotically. Some exceptions can be found in the models by Krause (2000) where the authors introduce *homophily* in the form of "bounded confidence", by that meaning that the agents are not influenced by far beliefs. A similar behavior can be found also in models with competing stubborn agents (Acemoğlu et al., 2013), the latter being agents that do not change their opinions but try to influence the others' opinions. Such stubborn agents might represent leaders, political parties or media sources. For instance, Como and Fagnani (2011) provide scaling limits showing that if the agents' population is homogeneous, then, as the population size grows large, the empirical belief distribution converges towards the solution of a certain deterministic mean-field differential equation in the space of probability measures. Such results are in the spirit of the propagation of chaos (Sznitman, 1991) in interacting particle systems.

2.2 Main assumptions

Following (Nedić et al., 2010) (see also Nedić and Bauso, 2013) we make the following assumptions on the information structure. Let $A(t)$ be the weight matrix with entries $a_j^i(t)$.

Assumption 1. The matrix $A(t)$ is doubly stochastic with positive diagonal. Furthermore, there exists a scalar $\alpha > 0$ such that $a_j^i(t) \geq \alpha$ whenever $a_j^i(t) > 0$.

At any time, the instantaneous graph $\mathcal{G}(t)$ need not be connected. However, for the proper behavior of the process, the union of the graphs $\mathcal{G}(t)$ over a period of time is assumed to be connected.

Assumption 2. There exists an integer $Q \geq 1$ such that the graph $\left(N, \bigcup_{\tau=tQ}^{(t+1)Q-1} \mathcal{E}(\tau)\right)$ is strongly connected for every $t \geq 0$.

For simplicity the one-shot vector-payoff game (S_1, S_2, x_i) is denoted by G.

Let $\lambda \in \mathbb{R}^n$. Denote by $\langle \lambda, G \rangle$ the zero-sum one-shot game whose set of players and their action sets are as in the game G, and the payoff that player i_2 pays to player i_1 is $\lambda' \phi(j, k)$ for every $(j, k) \in S_1 \times S_2$.

The resulting zero-sum game is described by the matrix
$$\Phi_\lambda = [\lambda' \phi(j, k)]_{j \in S_1, k \in S_2}.$$
As a zero-sum one-shot game, the game $\langle \lambda, G \rangle$ has a value v_λ, where
$$v_\lambda := \min_{p \in \Delta(S_1)} \max_{q \in \Delta(S_2)} p' \Phi_\lambda q = \max_{q \in \Delta(S_2)} \min_{p \in \Delta(S_1)} p' \Phi_\lambda q.$$

Following Blackwell (1956), we assume next that the value of the projected game is always negative.

Assumption 3. The payoff $\phi(j, k)$ is bounded, and
(a). $v_\lambda < 0$, for all $\lambda \in \mathbb{R}^n$;
(b). for all $q \in \Delta(S_2)$ and $\epsilon > 0$, there exists $p \in \Delta(S_1)$ such that $\left\| \sum_{j \in S_1, k \in S_2} p_j \phi(j, k) q_j \right\| \leq \epsilon$.

Condition (a) of Assumption 3 is among the foundations of approachability theory as it guarantees that the average vector payoff of a two-player repeated game converges almost surely to X (see e.g. Blackwell, 1956, and also Cesa-Bianchi and Lugosi, 2006, chapter 7). For future purposes we introduce the maximal value of the projected game, which is given by
$$\tilde{v} := \max_{\lambda \in \mathbb{R}^n, \|\lambda\|=1} v_\lambda \tag{6}$$
We show below that $\tilde{v}$ is related to the rate of convergence to the target set. Condition (b) implies that the norm of $u_i(t)$ in (2) can be made arbitrarily small, so that player i_1 effectively has the right of veto over any strategy of their opponent i_2. This condition ensures that any target set X is controlled invariant, as we show below.

3. MAIN RESULT

The main result of this paper establishes contractivity and invariance for the collective dynamics (1-2) under the multi-stage receding horizon strategy defined in (5). Before stating the theorem we introduce three lemmas. The first of these establishes that the space averaging process in (1) reduces the total distance (i.e. the sum of distances) of the states from the set X.

Lemma 1. Let Assumption 1 hold. Then the total distance from X decreases when replacing the states $x_i(t)$ by their space averages $w_i(t)$, i.e.,

$$\sum_{i=1}^{n} |w_i(t)|_X^2 \leq \sum_{i=1}^{n} |x_i(t)|_X^2.$$

Proof. Given in the appendix.

As a preliminary step to the next result, observe that, from the definition of $|\cdot|_X$ and from (1) and (3), we can write

$$\begin{aligned} |x_i(t+1)|_X^2 &= \|x_i(t+1) - P_X[x_i(t+1)]\|^2 \\ &\leq \|x_i(t+1) - P_X[w_i(t)]\|^2 \\ &= \|w_i(t) + u_i(t) - P_X[w_i(t)]\|^2 \\ &= \|w_i(t) - P_X[w_i(t)]\|^2 + \|u_i(t)\|^2 \\ &\quad + 2(w_i(t) - P_X[w_i(t)])' u_i(t). \end{aligned} \tag{7}$$

The following lemma states that, under the approachability assumption, there necessarily exists an input $u_i(t)$ given by (2) that places the successor state $x_i(t+1)$ closer to X than the space average $w_i(t)$.

Lemma 2. Let Assumption 3 hold. Then for each $i \in N$ there exists $u_i(t)$ satisfying (2) and

$$\begin{cases} |x_i(t+1)|_X < |w_i(t)|_X, & \text{if } |w_i(t)|_X > 0 \\ |x_i(t+1)|_X = 0, & \text{if } |w_i(t)|_X = 0 \end{cases}$$

Proof. Given in the appendix.

Lemma 2 shows that the constraint incorporated in the definition of the receding horizon strategy in (5) is feasible for all collective states $\xi(t)$. The next result provides upper and lower bounds on the collective value function $\sum_{i=1}^{n} V_{i,T}$ in terms of the sum of the distances of each invidual player's state from X.

Lemma 3. The value functions $V_{i,T}(\xi, \cdot)$, $i \in N$ satisfy, for all $\xi = (x_1, \ldots, x_n) \in \mathbb{R}^n \times \cdots \times \mathbb{R}^n$,

$$\sum_{i=1}^{n} |x_i|_X^2 \leq \sum_{i=1}^{n} V_{i,T}(\xi, \cdot) \leq (T+1) \sum_{i=1}^{n} |x_i|_X^2. \tag{8}$$

Proof. Given in the appendix.

For $r > 0$, let $\Phi(r)$ denote the set

$$\Phi(r) = \Big\{ (x_1, \ldots, x_n) \in \mathbb{R}^n \times \cdots \times \mathbb{R}^n \;\Big|\; \sum_{i=1}^{n} |x_i|_X^2 \leq r^2 \Big\}$$

and, for any positive integer T, define r_T by

$$r_T = \max \Big\{ r \;\Big|\; \sum_{i=1}^{n} |\hat{x}_i(T)|_X^2 = 0 \text{ for all } \xi(0) \in \Phi(r) \Big\}$$

where $\hat{x}_i(t)$ for $t \in [0, T]$ denotes the evolution of (1)-(2) under the min-max strategy with value function $V_{i,T-t}\big((\hat{x}_1(t), \ldots, \hat{x}_n(t)), t\big)$ in (5) and with $\hat{x}_i(0) = x_i(0)$.

We are now ready to state the main result.

Theorem 1. **(Contractivity and invariance)** Let Assumptions 1-3 hold and let $\xi(0) \in \Phi(r)$ for some $r > 0$. Then under the receding horizon strategy with optimal cost $V_{i,T}(\xi(t), t)$ for all $i \in N$ and all $t = 0, 1, \ldots$ we have

$$\xi(t) \in \Phi(r) \quad \forall t > 0, \tag{9}$$

and if $r \leq r_T$, then $\xi(t)$ converges exponentially to X:

$$\sum_{i=1}^{n} |x_i(t)|_X^2 \leq \Big(\frac{T}{T+1} \Big)^t \sum_{i=1}^{n} V_{i,T}(\xi(0), 0) \quad \forall t > 0. \tag{10}$$

Proof. Given in the appendix.

Discussion Theorem 1 demonstrates that there exist both invariant and contractive sets for the collective dynamics. In particular, the receding horizon policy renders the set $\Phi(r)$ invariant irrespective of the choice of the horizon T. Essentially this is due to the imposition of the constraint $|x_i(t+1)|_X \leq |w_i(t)|_X$ in the definition of the receding horizon policy, which, as a result of the space averaging process in (1), enforces invariance globally without requiring knowledge of other players' current actions.

The convergence rate of the collective state to X in (10) decreases as T increases. Essentially this is a result of the looseness of the upper bound in Lemma 3, according to which the rate of decrease of the accumulated cost represents a smaller fraction of the bound on $V_{i,T}(\xi(t), t)$ as T increases. By increasing T however, it is possible to reach the target set X in T steps from a larger set of initial conditions $\xi(0)$; increasing T thus increases the applicability of the contractivity result in (10) by allowing for larger values of r_T.

4. SIMULATION EXAMPLE

ν	x_{min}	x_{max}	dt	$std(m_0)$	T	$\overline{m}_0$
10^3	0	100	0.01	{8,10,15}	40	50

Table 1. Simulation parameters.

We provide here numerical studies showing contractivity and invariance in opinion propagation. The algorithm used to perform the simulations is illustrated below. We consider a number of players $\nu = 10^3$ and a discretized set of states $\mathcal{X} = \{x_{min}, x_{min}+1, \ldots, x_{max}\}$ where $x_{min} = 0$ (minimum state) and $x_{max} = 100$ (maximum state). The simulation parameters are listed in Table 1. We assume that the step size for the simulation is $dt = 0.01$. The horizon length, which corresponds to the number of iterations, is chosen as $T = 40$. This length of horizon is sufficient to highlight the convergence properties of the population opinions.

For a specific choice of weights a^i, $i = 1, \ldots, n$, the dynamics (1) have the form

$$\begin{cases} x_i(t+1) = \text{round}\ (\beta(\bar{m}(t) - x_i(t)) + u_i(t)), \\ x_0 \in \{x_{min}, x_{min}+1, \ldots, x_{max}\}. \end{cases} \tag{11}$$

where we have replaced $\sum_{j=1}^{n} a_j^i(t) x_j(t)$ by

$$\bar{m}(t) := \frac{1}{n} \sum_{j=1}^{n} x_j(t).$$

The above dynamics represents the case of a fully connected network where every agent's opinion is subject to an attractive force from the average $\bar{m}(t)$.

We assume that the set of initial states are obtained from an initial Gaussian distribution m_0 with mean $\overline{m}_0$ equal to

	$\beta\,[10^{-2}]$	$\tilde{v}$
Case I	0.0015	−0.08
Case II	15	−0.08

Table 2. Varying simulation parameters with different regimes.

Algorithm

Input: Set of parameters as in Table 1
Output: Distribution function m_t, mean $\overline{m}_t$ and standard deviation $std(m_t)$.
Initialize. Generate $x_i(0)$ from Gaussian distribution with mean $\bar{m}_0$ and standard dev. $std(m_0)$,
for time $t = 0, 1, \ldots, T-1$ **do**
  if $t > 0$, **then** compute m_t, $\overline{m}_t$, and $std(m_t)$,
  end if
  for player $i = 1, 2, \ldots, n$ **do**
    compute $x_i(t+1)$ by solving (3),
  end for
end for;
STOP

50. For the three simulation examples we set the standard deviation $std(m_0)$ equal to 8, 10, and 15, respectively.

For Case I, Figure 3 shows on the left the time plots of the microscopic evolution of each agent's opinion in the three examples, $std(m_0) = 8$ (top), $std(m_0) = 10$ (middle), $std(m_0) = 15$ (bottom). Figure 3 (right) displays the time plot $\overline{m}(t)$ (solid line and y-axis labeling on the left) and the evolution of the standard deviation $std(m(t))$ (dashed line and y-axis labeling on the right). Note that both the mean distribution $\overline{m}(t)$ and the standard deviation $std(m_t)$ converge to zero at approximately $t = 30$, which means that all the agents' opinions have reached ε-consensus around zero.

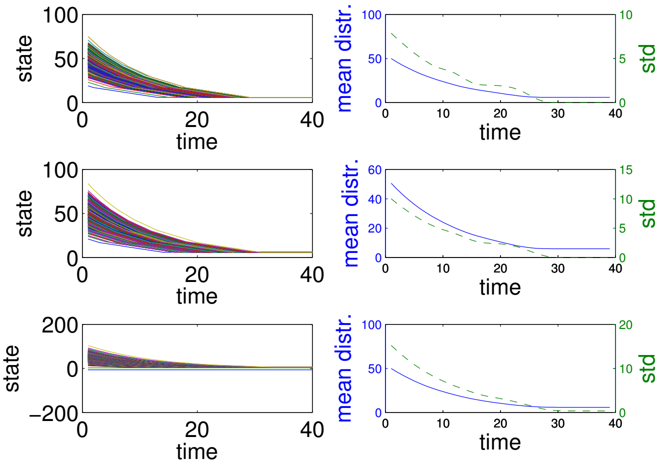

Fig. 3. Microscopic time plot (left) and time plot of mean distribution and standard deviation (right) for Case I.

For Case II we increase β, which means that the attraction force among the opinions is stronger. Figure 4 (left) from top to bottom shows that the stream of opinions reach consensus while the consensus value approaches zero. This behavior reflects also in a rapid decrease of the standard deviation $std(m_t)$ which reaches zero around $t = 10$.

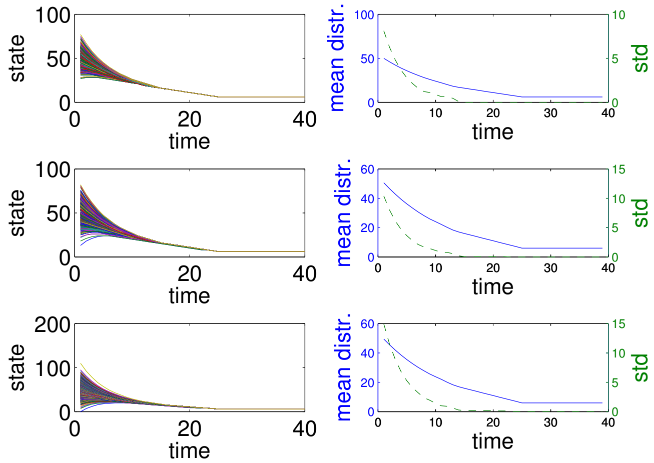

Fig. 4. Microscopic time plot (left) and time plot of mean distribution and standard deviation (right) for Case II.

5. CONCLUSIONS

We have analyzed convergence conditions of a distributed allocation process arising in the context of TU games. Future directions include the extension of our results to population games with mean-field interactions, and averaging algorithms driven by Brownian motions.

REFERENCES

D. Acemoğlu and A. Ozdaglar. Opinion dynamics and learning in social networks. *International Review of Economics*, 1(1):3–49, 2011.

D. Acemoğlu, G. Como, F. Fagnani, and A. Ozdaglar. Opinion fluctuations and disagreement in social networks. *Math. of Operation Research*, 38(1):1–27, 2013.

D. Bauso and G. Notarstefano. Distributed n-player approachability via time and space average consensus. In *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 198–203, Santa Barbara, CA, USA, Sept 2012.

D. Blackwell. An analog of the minmax theorem for vector payoffs. *Pacific J. Math.*, 6:1–8, 1956.

C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Rev. Mod. Phys.*, 81:591–646, 2009.

N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

G. Como and F. Fagnani. Scaling limits for continuous opinion dynamics systems. *The Annals of Applied Probability*, 21(4):1537–1567, 2011.

U. Krause. A discrete nonlinear and non-autonomous model of consensus formation. In *Communications in Difference Equations, S. Elaydi, G. Ladas, J. Popenda, and J. Rakowski editors, Gordon and Breach, Amsterdam*, pages 227–236, 2000.

E. Lehrer. Allocation processes in cooperative games. *International J. of Game Theory*, 31:341–351, 2002.

H. Liu, G. Xie, and L. Wang. Necessary and sufficient conditions for containment control of networked multi-agent systems. *Automatica*, 48(7):1415–1422, 2012.

D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.

A. Nedić and D. Bauso. Dynamic coalitional tu games: Distributed bargaining among players' neighbors. *IEEE Trans. Autom. Control*, 58(6):1363–1376, 2013.

A. Nedić, A. Ozdaglar, and P.A. Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Trans. Autom. Control*, 55(4):922–938, 2010.

S. Sundhar Ram, A. Nedić, and V.V. Veeravalli. Incremental stochastic subgradient algorithms for convex optimization. *SIAM Journal on Optimization*, 20(2): 691–717, 2009.

G. Shi and Y. Hong. Global target aggregation and state agreement of nonlinear multi-agent systems with switching topologies. *Automatica*, 45(5):1165–1175, 2009.

A.-S. Sznitman. Topics in propagation of chaos. *Springer Lecture Notes in Mathematics*, 1464:165–251, 1991.

J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton Univ. Press, 1944.

APPENDIX

Proof of Lemma 1. By convexity of the distance function $|\cdot|_X$ and from (4) we have $|w_i(t)|_X \leq \sum_{j=1}^n a_j^i(t)|x_j(t)|_X$. Hence convexity of $(\cdot)^2$ implies

$$|w_i(t)|_X^2 \leq \sum_{j=1}^n a_j^i(t)|x_j(t)|_X^2,$$

Summing both sides over $i = 1, \ldots, n$ we obtain

$$\begin{aligned}\sum_{i=1}^n |w_i(t)|_X^2 &\leq \sum_{i=1}^n \sum_{j=1}^n a_j^i(t)|x_j(t)|_X^2 \\ &= \sum_{j=1}^n \Big(\sum_{i=1}^n a_j^i(t)\Big)|x_j(t)|_X^2 = \sum_{j=1}^n |x_j(t)|_X^2,\end{aligned}$$

where the last equality follows from the stochasticity of $A(t)$ in Assumption 1. □

Proof of Lemma 2. Rearranging equation (7) we obtain

$$\begin{aligned}&|x_i(t+1)|_X^2 - |w_i(t)|_X^2 \\ &\qquad \leq \|u_i(t)\|^2 + 2\big(w_i(t) - P_X[w_i(t)]\big)'u_i(t). \quad (12)\end{aligned}$$

Now Assumption 3 implies that for any $w_i(t) \in \mathbb{R}^n$, there exists a mixed strategy $p(t) \in \Delta(S_1)$ for Player i_1 such that, for any mixed strategy $q(t) \in \Delta(S_2)$ of Player i_2, $u_i(t) = \sum_{j\in S_1}\sum_{k\in S_2} p_j(t)\phi(j,k)q_k(t)$ satisfies

$$\big(w_i(t) - P_X[w_i(t)]\big)'u_i(t) < \tilde{v}\,|w_i(t)|_X < 0$$

whenever $|w_i(t)|_X > 0$. Furthermore by Assumption 3, $p(t) \in \Delta(S_1)$ can be chosen so that $\|u_i(t)\|$ is arbitrarily small, and from the convexity of $\Delta(S_1)$ it follows that

$$\big(w_i(t) - P_X[w_i(t)]\big)'u_i(t) < \tilde{v}\,|w_i(t)|_X\|u_i(t)\|/L < 0$$

for some $L > 0$. From (12) we therefore obtain

$$|x_i(t+1)|_X^2 - |w_i(t)|_X^2 < \|u_i(t)\|\Big(\|u_i(t)\| - 2\frac{|\tilde{v}|}{L}|w_i(t)|_X\Big).$$

This bound is negative if $\|u_i(t)\| < 2\frac{|\tilde{v}|}{L}|w_i(t)|_X$, whereas $|x_i(t+1)|_X = 0$ if $|w_i(t)|_X = 0$ and $u_i(t) = 0$. □

Proof of Lemma 3. The lower bound in (8) follows directly from the fact that $V_{i,T-1}(\xi,\cdot) \geq 0$ for any horizon $T \geq 1$ and all $\xi \in \mathbb{R}^n \times \cdots \times \mathbb{R}^n$.

To prove the upper bound in (8), consider first the case of $T = 1$. The definition of $V_{i,0}(\xi,\cdot) = |x_i(\cdot)|_X^2$ and the constraint $|x_i(t+1)|_X \leq |w_i(t)|_X$ (which is necessarily feasible by Lemma 2) imply, for all $i \in N$, that

$$\begin{aligned}V_{i,1}(\xi(t),t) &= |x_i(t)|_X^2 + \min_{p\in\Delta(S_1)}\max_{q\in\Delta(S_2)} |x_i(t+1)|_X^2 \\ &\leq |x_i(t)|_X^2 + |w_i(t)|_X^2.\end{aligned}$$

Summing over $i \in N$ and using Lemma 1, we obtain

$$\sum_{i=1}^n V_{i,1}(\xi(t),t) \leq 2\sum_{i=1}^n |x_i(t)|_X^2. \quad (13)$$

Consider next the case of $T > 1$. If the upper bound in (8) holds for a horizon of $T-1$, then since the constraints $|x_i(t+1)|_X \leq |w_i(t)|_X$ for all $i \in N$ imply (by Lemma 1) that $\sum_{i=1}^n |x_i(t+1)|_X^2 \leq \sum_{i=1}^n |x_i(t)|_X^2$, we have

$$\begin{aligned}\sum_{i=1}^n V_{i,T}(\xi(t),t) &\leq \sum_{i=1}^n |x_i(t)|_X^2 \\ &\quad + \sum_{i=1}^n \min_{p(t)\in\Delta(S_1)}\max_{q(t)\in\Delta(S_2)} T|x_i(t+1)|_X^2 \\ &\leq (T+1)\sum_{i=1}^n |x_i(t)|_X^2. \quad (14)\end{aligned}$$

The upper bound in (8) follows by induction using (13) and (14). □

Proof of Theorem 1. The positive invariance of $\Phi(r)$ for any $r > 0$ in (9) is a consequence of the constraint $|x_i(t+1)|_X \leq |w_i(t)|_X$ that is imposed on the optimal receding horizon policy in (5). Specifically, by Lemma 2 this constraint is necessarily feasible, and Lemma 1 implies

$$\sum_{i=1}^n |x_i(t+1)|_X^2 \leq \sum_{i=1}^n |w_i(t)|_X^2 \leq \sum_{i=1}^n |x_i(t)|_X^2.$$

The contractivity and exponential convergence in (10) results from $\xi(0) \in \Phi(r_T)$, since, from the definition of r_T, the terminal state of (1)-(2) under the min-max optimal sequence $\{(p(0),q(0)),\ldots,(p(T-1),q(T-1))\}$ for (5) satisfies $|\hat{x}_i(T)|_X = 0$, from which it follows that

$$V_{i,T-1}\big(\xi(0),0\big) = V_{i,T}\big(\xi(0),0\big).$$

Furthermore (9) implies $\xi(t) \in \Phi(r_T)\ \forall t \geq 0$, and therefore

$$\begin{aligned}V_{i,T}(\xi(t),t) &= |x_i(t)|_X^2 \\ &\quad + \min_{p(t)\in\Delta(S_1)}\ \max_{\substack{q(t)\in\Delta(S_2)\\ u_j(t), j\neq i, j\in N}} V_{i,T-1}(\xi(t+1),t+1) \\ &\leq |x_i(t)|_X^2 + V_{i,T-1}(\xi(t+1),t+1) \\ &= |x_i(t)|_X^2 + V_{i,T}(\xi(t+1),t+1)\end{aligned}$$

for all $t \geq 0$. Summing this inequality over $i \in N$ and using the upper bound of Lemma 3 gives

$$\begin{aligned}\sum_{i=1}^n \big[V_{i,T}(\xi(t+1),t+1) - V_{i,T}(\xi(t),t)\big] &\leq -\sum_{i=1}^n |x_i(t)|_X^2 \\ &\leq -\frac{1}{T+1}\sum_{i=1}^n V_{i,T}(\xi(t),t)),\end{aligned}$$

and hence $\sum_{i=1}^n V_{i,T}(\xi(t),t) \leq \left(\frac{T}{T+1}\right)^t \sum_{i=1}^n V_{i,T}(\xi(0),0)$. Applying the lower bound of Lemma 3 therefore yields (10). □