Automatica 84 (2017) 48–55

Contents lists available at ScienceDirect

Automatica

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

Adaptation, coordination, and local interactions via distributed approachability ☆

Dario Bauso

Department of Automatic Control and Systems Engineering, The University of Sheffield, Mappin Street Sheffield, S1 3JD, UK
Dipartimento dell'Innovazione Industriale e Digitale, Università di Palermo, 90128 Palermo, Italy

ARTICLE INFO

Article history:
Received 27 August 2016
Received in revised form 6 January 2017
Accepted 20 April 2017

Keywords:
Repeated games
Approachability
Differential games
Robust control
Network flow

ABSTRACT

This paper investigates the relation between cooperation, competition, and local interactions in large distributed multi-agent systems. The main contribution is the game-theoretic problem formulation and solution approach based on the new framework of distributed approachability, and the study of the convergence properties of the resulting game model. Approachability theory is the theory of two-player repeated games with vector payoffs, and distributed approachability is here presented for the first time as an extension to the case where we have a team of agents cooperating against a team of adversaries under local information and interaction structure. The game model turns into a nonlinear differential inclusion, which after a proper design of the control and disturbance policies, presents a consensus term and an exogenous adversarial input. Local interactions enter into the model through a graph topology and the corresponding graph-Laplacian matrix. Given the above model, we turn the original questions on cooperation, competition, and local interactions, into convergence properties of the differential inclusion. In particular, we prove convergence and exponential convergence conditions around zero under general Markovian strategies. We illustrate our results in the case of decentralized organizations with multiple decision-makers.

1. Introduction

Cooperation, competition, and local interactions are three main co-existing elements in large distributed multi-agent systems with humans in the loop, see Fig. 1. The *state* of a decision-maker is captured by a time-varying abstract entity, which contains aggregate information on his past decisions and those of a subset of other decision-makers around him, as well as his cumulative or average payoff.

In abstract terms, *cooperation* refers to the capability of the decision-makers to make decisions to coordinate their states. The decision-makers try to reach consensus by exhibiting reciprocal attraction forces which may lead them to converge to a consensus equilibrium, see Olfati-Saber, Fax, and Murray (2007) and references therein.

By *competition* we refer to the capabilities of the decision-makers to let the collective state, a vector which involves the states of all the decision-makers, converge to a preassigned set or equilibrium point despite the presence of disturbances. A natural way to deal with such a scenario is via *approachability* theory, whose traditional formulation involves only two players, the decision-maker (player 1 or row player) and the adversarial disturbance (player 2 or column player) (Blackwell, 1956). The two players play repeatedly over time in a continuous- or discrete-time setting, and the outcome of the game at any time is a vector payoff. Both players try to influence the evolution of the average payoff. Existing results show that the approachability problem can be turned into a differential game in which the average payoff appears as the (collective) state of the game (Benaïm, Hofbauer, & Sorin, 2005, 2006; Soulaimani, Quincampoix, & Sorin, 2009). In particular player 1 plays to make the average payoff converge to a preassigned set, while player 2 tries to contrast him. Equivalence of Blackwell Approachability and No-Regret Learning is studied in Abernethy, Bartlett, and Hazan (2011). A dynamic programming approach to calculate approachable sets is presented in Kamble (2015). Approachability in Stackelberg Stochastic Games is investigated in Kalathil, Borkar, and Jain (2016). Convergence of the cumulative payoff rather than the average implies some variations of the conditions which are formalized in the context of *attainability* in two-player repeated games with vector payoffs, see e.g. Bauso, Solan, Lehrer, and Venel (2015) and Bauso (2016, Ch. 11). The distributed approachability problem that we formulate here assumes that player 1 is indeed made by a team of agents

☆ The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Andrey V. Savkin under the direction of Editor Ian R. Petersen.

E-mail address: d.bauso@sheffield.ac.uk.

http://dx.doi.org/10.1016/j.automatica.2017.06.017
0005-1098/

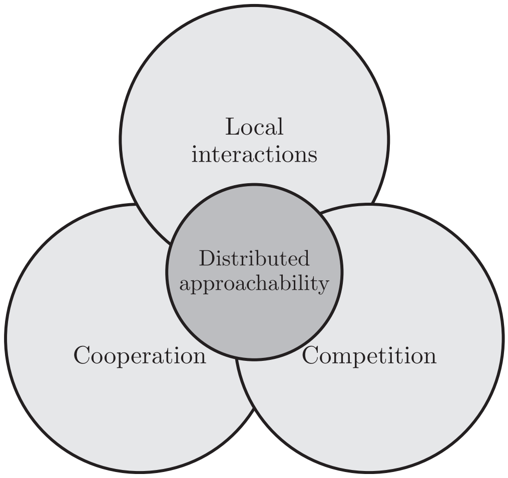

Fig. 1. Three dimensions of distributed decision making reframed within distributed approachability.

whose cooperation results in attraction forces against a team of adversarial disturbances, referred to as player 2, which exhibit external forces.

Local Coordination captures the idea that the decision-makers have (i) local information, namely they know only some state components, and (ii) local influence, namely their decisions influence only some state components. To model local coordination we refer to the concept of distributed Markovian or state-feedback control policies. Back to the approachability interpretation, the state of the decision-maker is the subset of payoff components he can monitor and control. As it will be clear later on, the term *distributed approachability* is here used to address such a concept. This term has already appeared in Bauso and Notarstefano (2015) in the context of coalitional games.

Contribution. As main contribution this paper builds a mathematical model involving each of the above dimensions: cooperation, competition, and local interactions. To capture competition the model takes the form of a distributed approachability problem, thus departing in an original way from the traditional two-player approachability formulation. A further contribution is in that the model links in an original way to a stylized model in the literature of decentralized organizations thus contributing to the cross-fertilization of engineering and social science.

Building on existing results (Benaïm et al., 2005, 2006; Soulaimani et al., 2009), which show that an approachability problem can be turned into a differential game, the game is ultimately transformed into a nonlinear differential inclusion describing the continuous-time evolution of the cumulative or average payoff. Here, the distributed control involves the mixed actions of all the decision-makers (player 1) and the distributed disturbance is the mixed actions of all adversaries (player 2). The decision-makers coordinate to drive the vector payoff to a preassigned set against the actions of the adversaries. Nonlinearity is due to bounds on controls and disturbances. Given such a system, we look at equilibrium points, which represent conditions under which the attraction forces counterbalance the external ones.

We show that cooperation results in a consensus term in the differential inclusion which describes the attraction forces. Under such forces the states of the decision-makers tend to get closer one to each other.

Competition takes the form of an exogenous signal. In other words, the adversary tries to attract the local states by exhibiting some centrifugal force.

Local interaction enters into the model through a graph topology. We study the influence of such topology both on the stationary solution and on the transient dynamics. The graph topology appears in the consensus term, through the graph-Laplacian matrix.

Given the above model, we can turn the original questions on cooperation, competition, and local interactions, into convergence properties of the differential inclusion. In particular, we prove convergence and exponential convergence conditions around zero under general Markovian strategies using approachability theorem by Blackwell. We observe that when we use distributed Markovian strategies, we obtain a robust consensus dynamics and for such a dynamics we study the corresponding convergence properties.

The main assumption is in the form of set inclusion, and represents properties of the action sets of the game. This assumption is borrowed from the literature on robust control of network systems (Blanchini, Miani, & Ukovich, 2000; Blanchini, Rinaldi, & Ukovich, 1997).

To place the contribution of this paper in proper context, we illustrate our results in the case of decentralized organizations with multiple decision-makers that must perform n specialized tasks (Dessein & Santos, 2006). The decision-makers, each one associated to a single task, choose the levels of adaptation and coordination. A higher level of adaptation implies that the workers show higher flexibility to adapt their tasks. A higher level of coordination entails an increase in the communication between workers. The performance of the organization depends on: (i) how well each task is adapted to specific market conditions, operational conditions, and consumers' needs and (ii) how well all tasks are coordinated with each other.

This paper is organized as follows. In Section 2 we introduce approachability and distributed approachability. In Section 3 we turn the game into a dynamical system. In Section 4 we provide the main results on convergence and exponential convergence. In Section 5 we discuss the results in the context of decentralized organizations. In Section 6 we provide conclusions.

2. Distributed approachability

In this section, we first introduce the traditional approachability setting involving two players and a continuous-time repeated game with vector payoffs. Then, we formulate the problem at hand in the form of a distributed approachability problem with a team of decision-makers playing against a team of adversaries.

2.1. Approachability

The traditional approachability setting involves a two-player repeated game with vector payoffs, which we refer to as Γ. The set of players is $N = \{1, 2\}$, and the finite set of actions of each player i is A_i. The instantaneous payoff is given by a biaffine function $g : A_1 \times A_2 \to \mathbb{R}^m$, where m is a natural number.

We extend g to the set of mixed actions pairs, $\Delta(A_1) \times \Delta(A_2)$, in a bilinear fashion. The one-shot vector-payoff game $(\Delta(A_1), \Delta(A_2), u)$ has compact convex action sets and is denoted by G.

The game Γ is played in continuous-time over the time interval $[0, \infty)$. We assume that the players use non-anticipative behavior strategies, according to the definition provided below.

Denote by $\mathcal{C}_i$ the set of all *actions* of player i, that is, the set of all measurable functions from the time space, $[0, \infty)$, to player i's mixed actions. That is,

$$\mathcal{C}_i := \{a_i : [0, \infty) \to \Delta(A_i),\ a_i \text{ is measurable}\}.$$

Definition 2.1. A function $\sigma_i = \sigma_i[\cdot] : \mathcal{C}_{-i} \to \mathcal{C}_i$ is a *non-anticipative behavior strategy* for player i, if

$$a_{-i}(s) = a'_{-i}(s) \quad \forall s \in [0, t]$$
$$\implies \quad \sigma_i[a_{-i}](s) = \sigma_i[a'_{-i}](s) \quad \forall s \in [0, t].$$