Automatica 49 (2013) 1683–1692

Contents lists available at SciVerse ScienceDirect

Automatica

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

Distributed convergence to Nash equilibria in two-network zero-sum games ☆

B. Gharesifard [a,1], J. Cortés [b]

[a] Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, IL 61801, USA

[b] Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA 92093, USA

ARTICLE INFO

Article history:
Received 3 April 2012
Received in revised form
20 November 2012
Accepted 17 February 2013
Available online 6 April 2013

Keywords:
Adversarial networks
Distributed algorithms
Zero-sum game
Saddle-point dynamics
Nash equilibria

ABSTRACT

This paper considers a class of strategic scenarios in which two networks of agents have opposing objectives with regard to the optimization of a common objective function. In the resulting zero-sum game, individual agents collaborate with neighbors in their respective network and have only partial knowledge of the state of the agents in the other network. For the case when the interaction topology of each network is undirected, we synthesize a distributed saddle-point strategy and establish its convergence to the Nash equilibrium for the class of strictly concave–convex and locally Lipschitz objective functions. We also show that this dynamics does not converge in general if the topologies are directed. This justifies the introduction, in the directed case, of a generalization of this distributed dynamics which we show converges to the Nash equilibrium for the class of strictly concave–convex differentiable functions with globally Lipschitz gradients. The technical approach combines tools from algebraic graph theory, nonsmooth analysis, set-valued dynamical systems, and game theory.

1. Introduction

Recent years have seen an increasing interest on networked strategic scenarios where agents may cooperate or compete with each other towards the achievement of some objective, interact across different layers, have access to limited information, and are subject to evolving interaction topologies. This paper is a contribution to this body of work. Specifically, we consider a class of strategic scenarios in which two networks of agents are involved in a zero-sum game. We assume that the objective function can be decomposed as a sum of concave–convex functions and that the networks have opposing objectives regarding its optimization. Agents collaborate with the neighbors in their own network and have partial information about the state of the agents in the other network. Such scenarios are challenging because information is spread across the agents and possibly multiple layers, and networks, by themselves, are not the decision makers. Our aim is to design a distributed coordination algorithm that can be used by the agents to converge to the Nash equilibrium. Note that, for a 2-player zero-sum game of the type considered here, a pure Nash equilibrium corresponds to a saddle point of the objective function.

Literature review. Multiple scenarios involving networked systems and intelligent adversaries in sensor networks, filtering, finance, and wireless communications (Kim & Boyd, 2008; Wan & Lemmon, 2009) can be cast into the strategic framework described above. In such scenarios, the network objective arises as a result of the aggregation of agent-to-agent adversarial interactions regarding a common goal, and information is naturally distributed among the agents. The present work has connections with the literature on distributed optimization and zero-sum games. The distributed optimization of a sum of convex functions has been intensively studied in recent years; see e.g. Johansson, Rabi, and Johansson (2009), Nedic and Ozdaglar (2009a), Wan and Lemmon (2009) and Zhu and Martínez (2012). These works build on consensus-based dynamics (Bullo, Cortés, & Martínez, 2009; Mesbahi & Egerstedt, 2010; Olfati-Saber, Fax, & Murray, 2007; Ren & Beard, 2008) to find the solutions of the optimization problem in a variety of scenarios and are designed in discrete time. Exceptions include (Wang & Elia, 2010, 2011) on continuous-time distributed optimization on undirected networks and (Gharesifard & Cortés, 2012b) on directed networks.

Regarding zero-sum games, the works (Arrow, Hurwitz, & Uzawa, 1958; Maistroskii, 1977; Nedic & Ozdgalar, 2009b)

☆ This work was supported in part by Award FA9550-10-1-0499. This work was performed while B. Gharesifard was a postdoctoral researcher at the University of California, San Diego. The material in this paper was partially presented at the 2012 American Control Conference (ACC2012), June 27–29, 2012, Montréal, Canada, as Gharesifard and Cortés (2012a) and 51st IEEE Control and Decision Conference (CDC), December 10–13, 2012, Maui, Hawaii, USA, as Gharesifard and Cortés (in press). This paper was recommended for publication in revised form by Associate Editor Xiaobo Tan under the direction of Editor Miroslav Krstic.

E-mail addresses: bgharesi@illinois.edu (B. Gharesifard), cortes@ucsd.edu (J. Cortés).

1 Tel.: +1 858 366 5163; fax: +1 858 822 3107.

0005-1098/$ – see front matter
http://dx.doi.org/10.1016/j.automatica.2013.02.062

study the convergence of discrete-time subgradient dynamics to a saddle point. Continuous-time best-response dynamics for zero-sum games converges to the set of Nash equilibria for both convex–concave (Hofbauer & Sorin, 2006) and quasiconvex–quasiconcave (Barron, Goebel, & Jensen, 2010) functions. Under strict convexity–concavity assumptions, continuous-time subgradient flow dynamics converges to a saddle point (Arrow, Hurwitz, & Uzawa, 1951; Arrow et al., 1958). Asymptotic convergence is also guaranteed when the Hessian of the objective function is positive definite in one argument and the function is linear in the other (Arrow et al., 1958; Feijer & Paganini, 2010). The saddle-point and learning dynamics reviewed above are not directly applicable to the problem considered here because the information about the game is distributed across the networked scenario and is not centrally available anywhere. The distributed computation of Nash equilibria in noncooperative games, where all players are adversarial, is a challenging problem. The algorithm in Li and Başar (1987) relies on all-to-all communication and does not require players to know each other's payoff functions (which must be strongly convex). In Frihauf, Krstic, and Başar (2012) and Stankovic, Johansson, and Stipanovic (2012), players are unaware of their own payoff functions but have access to the payoff value of an action once it has been executed. These works design distributed strategies based on extremum seeking techniques to seek the set of Nash equilibria.

Statement of contributions. We introduce the problem of distributed convergence to Nash equilibria for two networks engaged in a strategic scenario. The networks aim to either maximize or minimize a common objective function which can be written as a sum of concave–convex functions. Individual agents collaborate with neighbors in their respective network and have partial knowledge of the state of the agents in the other one. Our first contribution is the introduction of an aggregate objective function for each network which depends on the interaction topology through its Laplacian and the characterization of a family of points with a saddle property for the pair of functions. We show the correspondence between these points and the Nash equilibria of the overall game. When the graphs describing the interaction topologies within each network are undirected, the gradients of these aggregate objective functions are distributed. Building on this observation, our second contribution is the synthesis of a consensus-based saddle-point strategy for adversarial networks with undirected topologies. We show that the proposed dynamics is guaranteed to asymptotically converge to the Nash equilibrium for the class of strictly concave–convex and locally Lipschitz objective functions. Our third contribution focuses on the directed case. We show that the transcription of the saddle-point dynamics to directed topologies fails to converge in general. This leads us to propose a generalization of the dynamics, for strongly connected weight-balanced topologies, that incorporates a design parameter. We show that, by appropriately choosing this parameter, the new dynamics asymptotically converges to the Nash equilibrium for the class of strictly concave–convex differentiable objective functions with globally Lipschitz gradients. The technical approach employs notions and results from algebraic graph theory, nonsmooth and convex analysis, set-valued dynamical systems, and game theory. As an intermediate result in our proof strategy for the directed case, we provide a generalization of the known characterization of cocoercivity of concave functions to concave–convex functions.

The results of this paper can be understood as a generalization to competing networks of the results we obtained in Gharesifard and Cortés (2012b) for distributed optimization. This generalization is nontrivial because the payoff functions associated with individual agents now also depend on information obtained from the opposing network. This feature gives rise to a hierarchy of saddle-point dynamics whose analysis is technically challenging and requires, among other things, a reformulation of the problem as a constrained zero-sum game, a careful understanding of the coupling between the dynamics of both networks, and the generalization of the notion of cocoercivity to concave–convex functions.

Organization. Section 2 contains preliminaries on nonsmooth analysis, set-valued dynamical systems, graph theory, and game theory. In Section 3, we introduce the zero-sum game for two adversarial networks involved in a strategic scenario and introduce two novel aggregate objective functions. Section 4 presents our algorithm design and analysis for distributed convergence to Nash equilibrium when the network topologies are undirected. Section 5 presents our treatment for the directed case. Section 6 gathers our conclusions and ideas for future work. Appendix contains the generalization to concave–convex functions of the characterization of cocoercivity of concave functions.

2. Preliminaries

We start with some notational conventions. Let $\mathbb{R}$, $\mathbb{R}_{\geq 0}$, $\mathbb{Z}$, $\mathbb{Z}_{\geq 1}$ denote the set of real, nonnegative real, integer, and positive integer numbers, respectively. We denote by $\|\cdot\|$ the Euclidean norm on $\mathbb{R}^d$, $d \in \mathbb{Z}_{\geq 1}$, and also use the short-hand notation $\mathbf{1}_d = (1, \ldots, 1)^T$ and $\mathbf{0}_d = (0, \ldots, 0)^T \in \mathbb{R}^d$. We let I_d denote the identity matrix in $\mathbb{R}^{d\times d}$. For matrices $A \in \mathbb{R}^{d_1\times d_2}$ and $B \in \mathbb{R}^{e_1\times e_2}$, $d_1, d_2, e_1, e_2 \in \mathbb{Z}_{\geq 1}$, we let $A \otimes B$ denote their Kronecker product. The function $f : \mathsf{X}_1 \times \mathsf{X}_2 \to \mathbb{R}$, with $\mathsf{X}_1 \subset \mathbb{R}^{d_1}$, $\mathsf{X}_2 \subset \mathbb{R}^{d_2}$ closed and convex, is *concave–convex* if it is concave in its first argument and convex in the second one (Rockafellar, 1997). A point $(x_1^*, x_2^*) \in \mathsf{X}_1 \times \mathsf{X}_2$ is a *saddle point* of f if $f(x_1, x_2^*) \leq f(x_1^*, x_2^*) \leq f(x_1^*, x_2)$ for all $x_1 \in \mathsf{X}_1$ and $x_2 \in \mathsf{X}_2$. Finally, a set-valued map $f : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ takes elements of $\mathbb{R}^d$ to subsets of $\mathbb{R}^d$.

2.1. Nonsmooth analysis

We recall some notions from nonsmooth analysis (Clarke, 1983). A function $f : \mathbb{R}^d \to \mathbb{R}$ is *locally Lipschitz* at $x \in \mathbb{R}^d$ if there exists a neighborhood $\mathcal{U}$ of x and $C_x \in \mathbb{R}_{\geq 0}$ such that $|f(y)-f(z)| \leq C_x\|y-z\|$, for $y, z \in \mathcal{U}$. f is locally Lipschitz on $\mathbb{R}^d$ if it is locally Lipschitz at x for all $x \in \mathbb{R}^d$ and *globally Lipschitz* on $\mathbb{R}^d$ if for all $y, z \in \mathbb{R}^d$, there exists $C \in \mathbb{R}_{\geq 0}$ such that $|f(y)-f(z)| \leq C\|y-z\|$. Locally Lipschitz functions are differentiable almost everywhere. The *generalized gradient* of f is

$$\partial f(x) = \mathrm{co}\Big\{\lim_{k\to\infty} \nabla f(x_k) \mid x_k \to x,\ x_k \notin \Omega_f \cup S\Big\},$$

where Ω_f is the set of points where f fails to be differentiable and S is any set of measure zero.

Lemma 2.1 (*Continuity of the Generalized Gradient Map*)**.** *Let $f : \mathbb{R}^d \to \mathbb{R}$ be a locally Lipschitz function at $x \in \mathbb{R}^d$. Then the set-valued map $\partial f : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ is upper semicontinuous and locally bounded at $x \in \mathbb{R}^d$ and moreover, $\partial f(x)$ is nonempty, compact, and convex.*

For $f : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ and $z \in \mathbb{R}^d$, we let $\partial_x f(x, z)$ denote the generalized gradient of $x \mapsto f(x, z)$. Similarly, for $x \in \mathbb{R}^d$, we let $\partial_z f(x, z)$ denote the generalized gradient of $z \mapsto f(x, z)$. A point $x \in \mathbb{R}^d$ with $\mathbf{0} \in \partial f(x)$ is a *critical point* of f. A function $f : \mathbb{R}^d \to \mathbb{R}$ is *regular* at $x \in \mathbb{R}$ if for all $v \in \mathbb{R}^d$, the right directional derivative of f, in the direction of v, exists at x and coincides with the generalized directional derivative of f at x in the direction of v. We refer the reader to Clarke (1983) for definitions of these notions. A convex and locally Lipschitz function at x is regular (Clarke, 1983, Proposition 2.3.6). The notion of regularity plays an important role when considering sums of Lipschitz functions.

Lemma 2.2 (*Finite Sum of Locally Lipschitz Functions*)**.** *Let $\{f^i\}_{i=1}^n$ be locally Lipschitz at $x \in \mathbb{R}^d$. Then $\partial(\sum_{i=1}^n f^i)(x) \subseteq \sum_{i=1}^n \partial f^i(x)$, and equality holds if f^i is regular for $i \in \{1, \ldots, n\}$.*

A locally Lipschitz and convex function f satisfies, for all $x, x' \in \mathbb{R}^d$ and $\xi \in \partial f(x)$, the *first-order condition* of convexity:

$$f(x') - f(x) \geq \xi \cdot (x' - x). \tag{1}$$