Distributed convergence to Nash equilibria by adversarial networks with directed topologies

Bahman Gharesifard Jorge Cortés

***Abstract*— This paper considers a class of strategic scenarios in which two cooperative groups of agents have opposing objectives with regards 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. We consider scenarios where the interaction topology within each cooperative network is given by a strongly connected and weight-balanced directed graph. We introduce a provably-correct distributed dynamics which converges to the set of Nash equilibria when the objective function is strictly concave-convex, differentiable, with globally Lipschitz gradient. The technical approach combines tools from algebraic graph theory, dynamical systems, convex analysis, and game theory.**

I. Introduction

The nature of interactions between individual agents in a variety of networked scenarions is strategic and not necessarily cooperative. Examples of strategic interactions occur in biological systems, e.g., selfishness and stealth in collective motion [1] and competitive interactions between cells and organs [2], cybersecurity [3], and collective bargaining and opinion dynamics in heterogeneous networks [4], [5], [6]. This paper considers a class of such strategic scenarios where two networks of agents, with directed topologies and opposing goals, are involved in a zero-sum game, where the objective function is a sum of concave-convex functions. Within each network, agents cooperate with their neighbors and have partial information about the state of the agents of the opposing network. Our goal is to design a continuous-time distributed dynamics that can be used by the networks to converge to the set of Nash equilibria. Specifically, we seek to generalize the results of [7] to allow for directed topologies.

Literature review: This work is related to the literature on zero-sum games and distributed optimization. The convergence of the continuous-time best-response dynamics for zero-sum games with concave-convex payoff functions is shown in [8]. The results can be extended to quasiconvex-quasiconcave payoff functions, as recently shown in [9]. Continuous-time gradient flow dynamics has also been used for finding Nash equilibria of zero-sum games [10], [11]. This dynamics may fail to converge for general concave-convex functions [12] but is convergent when both convexity and concavity assumptions are strict. This convergence result also holds true when the payoff function is linear in one argument and its Hessian is positive-definite in the other [11], [12]. It is also worth noting that finding the saddle point of function using (sub)gradient dynamics has also been studied in discrete time [11], [13], [14]. The distributed computation of Nash equilibria in noncooperative games has been investigated in different contexts, see for example [15], [16], [17].

Regarding the literature on distributed optimization, the design of distributed dynamics for optimization of a sum of convex functions has been studied intensively in recent years, see e.g. [18], [19], [20]. These are consensus-based dynamics, see [21], [22], [23], [24], and are typically designed in discrete time. Exceptions are the works [25], [26], [7] on continuous-time distributed optimization on undirected networks and [27] on directed networks.

Statement of contributions: The contributions of this paper are threefold. We start by formulating a distributed zero-sum game for two networks with directed topologies engaged in a strategic scenario. The networks' objectives are 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 network. We provide characterizations of the Nash equilibria of the game as saddle points of two newly-introduced functions that play a key role in the algorithm design. Secondly, we introduce a generalization of the saddle-point dynamics corresponding to these functions that also incorporates a design parameter. This strategy has a nice consensus plus gradient-based interpretation. Using the LaSalle Invariance Principle, we show that by appropriately choosing this parameter, the proposed dynamics asymptotically converges to the set of Nash equilibria for any pair of strongly connected weight-balanced adversarial networks and strictly concave-convex differentiable objective function with globally Lipschitz gradient. Interestingly, the interplay between the connectivity of the underlying networks and the Lipschitz constant of the gradient of the objective function plays a key role in determining the values of the design parameter. Finally, we provide a generalization to concave-convex functions of the known characterization of cocoercivity for concave functions, which plays a key role in our technical approach. The proofs are omitted for reasons of space and will appear elsewhere.

II. 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 $\imath : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$, $d_1, d_2 \in \mathbb{Z}_{\geq 1}$, $d_2 \geq d_1$, any natural inclusion which maps each vector in $\mathbb{R}^{d_1}$ to a vector in $\mathbb{R}^{d_2}$ by adding

Bahman Gharesifard and Jorge Cortés are with the Department of Mechanical and Aerospace Engineering, University of California, San Diego, {bgharesifard,cortes}@ucsd.edu.

zeros to the rest of its components. 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. A 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 [28]. 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 $x_2 \in \mathsf{X}_2$. A function $f : \mathbb{R}^d \to \mathbb{R}$ is *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 $|f(y) - f(z)| \leq C||y - z||$. For a differentiable function f, a point $x \in \mathbb{R}^d$ with $\nabla f(x) = 0$ is a *critical point* of f. A differentiable convex function f satisfies, for all $x, x' \in \mathbb{R}^d$, the *first-order condition* of convexity,

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

A. Stability analysis

Here, we recall some background on continuous-time dynamical systems following [29]. Consider a system on $\mathsf{X} \subset \mathbb{R}^d$ given by

$$\dot{x}(t) = \Psi(x(t)), \tag{2}$$

where $t \in \mathbb{R}_{\geq 0}$ and $\Psi : \mathsf{X} \subset \mathbb{R}^d \to \mathbb{R}^d$ is continuous. A solution to this dynamical system is a continuously differentiable curve $x : [0, T] \to \mathsf{X}$ which satisfies (2). The set of equilibria of (2) is denoted by $\mathrm{Eq}(\Psi) = \{x \in \mathsf{X} \mid \Psi(x) = 0\}$.

The LaSalle Invariance Principle for continuous-time systems is helpful to establish the asymptotic stability properties of systems of the form (2). A set $W \subset \mathsf{X}$ is *positively invariant* with respect to Ψ if each solution with initial condition in W remains in W for all subsequent times. The *Lie derivative* of a continuously differentiable function $V : \mathbb{R}^d \to \mathbb{R}$ along Ψ at $x \in \mathbb{R}^d$ is defined by $\mathcal{L}_\Psi V(x) = \nabla V(x) \cdot \Psi(x)$.

Theorem 2.1: (LaSalle Invariance Principle): Let $W \subset \mathsf{X}$ be positively invariant under (2) and $V : \mathsf{X} \to \mathbb{R}$ a continuously differentiable function. Suppose the evolutions of (2) with initial conditions in W are bounded. Then any solution $x(t)$, $t \in \mathbb{R}_{\geq 0}$, starting in W converges to the largest positively invariant set M contained in $S_{\Psi,V} \cap W$, where $S_{\Psi,V} = \{x \in \mathsf{X} \mid \mathcal{L}_\Psi V(x) = 0\}$. When M is a finite collection of points, then the limit of each solution equals one of them.

B. Graph theory

We present some basic notions from algebraic graph theory following the exposition in [23]. A *directed graph*, or simply *digraph*, is a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a finite set called the vertex set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the edge set. A digraph is *undirected* if $(v, u) \in \mathcal{E}$ anytime $(u, v) \in \mathcal{E}$. We refer to an undirected digraph as a *graph*. A path is an ordered sequence of vertices such that any ordered pair of vertices appearing consecutively is an edge of the digraph. A digraph is *strongly connected* if there is a path between any pair of distinct vertices. For a graph, we refer to this notion simply as *connected*. A *weighted digraph* is a triplet $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathsf{A})$, where $(\mathcal{V}, \mathcal{E})$ is a digraph and $\mathsf{A} \in \mathbb{R}^{n\times n}_{\geq 0}$ is the *adjacency matrix* of $\mathcal{G}$, with the property that $a_{ij} > 0$ if $(v_i, v_j) \in \mathcal{E}$ and $a_{ij} = 0$, otherwise. The weighted out-degree and in-degree of v_i, $i \in \{1,\ldots,n\}$, are respectively, $d^{\mathsf{w}}_{\mathsf{out}}(v_i) = \sum_{j=1}^n a_{ij}$ and $d^{\mathsf{w}}_{\mathsf{in}}(v_i) = \sum_{j=1}^n a_{ji}$. The *weighted out-degree matrix* $\mathsf{D}_{\mathsf{out}}$ is the diagonal matrix defined by $(\mathsf{D}_{\mathsf{out}})_{ii} = d^{\mathsf{w}}_{\mathsf{out}}(i)$, for all $i \in \{1,\ldots,n\}$. The *Laplacian* matrix is $\mathsf{L} = \mathsf{D}_{\mathsf{out}} - \mathsf{A}$. Note that $\mathsf{L}\mathbf{1}_n = 0$. If $\mathcal{G}$ is strongly connected, then zero is a simple eigenvalue of L. $\mathcal{G}$ is undirected if $\mathsf{L} = \mathsf{L}^T$ and *weight-balanced* if $d^{\mathsf{w}}_{\mathsf{out}}(v) = d^{\mathsf{w}}_{\mathsf{in}}(v)$, for all $v \in \mathcal{V}$. Equivalently, $\mathcal{G}$ is weight-balanced if and only if $\mathbf{1}_n^T\mathsf{L} = 0$ if and only if $\mathsf{L} + \mathsf{L}^T$ is positive semidefinite. Furthermore, if $\mathcal{G}$ is weight-balanced and strongly connected, then zero is a simple eigenvalue of $\mathsf{L} + \mathsf{L}^T$. Note that any undirected graph is weight-balanced.

C. Zero-sum games

We recall some game theoretic notions from [30]. An n-player game is a triplet $\mathbf{G} = (P, \mathsf{X}, U)$, where P is the set of players, $n = |P| \in \mathbb{Z}_{\geq 2}$, $\mathsf{X} = \mathsf{X}_1 \times \ldots \times \mathsf{X}_n$, $\mathsf{X}_i \subset \mathbb{R}^{d_i}$ is the set of (pure) strategies of player $v_i \in P$, $d_i \in \mathbb{Z}_{\geq 1}$, and $U = (u_1,\ldots,u_n)$, where $u_i : \mathsf{X} \to \mathbb{R}$ is the payoff function of player v_i, $i \in \{1,\ldots,n\}$. The game $\mathbf{G}$ is called a *zero-sum game* if $\sum_{i=1}^n u_i(x) = 0$, for all $x \in \mathsf{X}$. If $x_i \in \mathsf{X}_i$, we denote by x_{-i} the strategy set of all players except v_i. An outcome $x^* \in \mathsf{X}$ is called a (pure) *Nash equilibrium* of $\mathbf{G}$ if for all $i \in \{1,\ldots,n\}$ and all $x_i \in \mathsf{X}_i$ we have

$$u_i(x_i^*, x_{-i}^*) \geq u_i(x_i, x_{-i}^*).$$

One can extend this notion to *mixed* Nash equilibria by assignment of probabilities to pure strategies [30]. In this paper, we focus on a particular class of two-players zero-sum games which have at least one pure Nash equilibrium. The following well-known Minmax Theorem [31] characterizes that the game $\mathbf{G} = (\{v_1, v_2\}, \mathsf{X}_1 \times \mathsf{X}_2, (u, -u))$ has a pure Nash equilibrium.

Theorem 2.2: (Minmax theorem): Let $\mathsf{X}_1 \subset \mathbb{R}^{d_1}$ and $\mathsf{X}_2 \subset \mathbb{R}^{d_2}$, $d_1, d_2 \in \mathbb{Z}_{\geq 1}$, be nonempty, closed, bounded, and convex. If $u : \mathsf{X}_1 \times \mathsf{X}_2 \to \mathbb{R}$ is continuous and the sets $\{x' \in \mathsf{X}_1 \mid u(x', y) \geq \alpha\}$ and $\{x' \in \mathsf{X}_2 \mid u(x, y') \leq \alpha\}$ are convex for all $x \in \mathsf{X}_1$, $y \in \mathsf{X}_2$, and $\alpha \in \mathbb{R}$, then

$$\max_x \min_y u(x, y) = \min_y \max_x u(x, y).$$

III. Problem statement

Consider two networks Σ_1 and Σ_2 composed of agents $\{v_1,\ldots,v_{n_1}\}$ and agents $\{w_1,\ldots,w_{n_2}\}$, respectively. Throughout this paper, Σ_1 and Σ_2 are either connected undirected graphs, c.f. Section IV, or strongly connected weight-balanced digraphs, c.f. Section V. Since the latter case includes the first one, throughout this section, we assume the latter. The state of Σ_1, denoted by x_1, belongs to $\mathsf{X}_1 \subset \mathbb{R}^{d_1}$, $d_1 \in \mathbb{Z}_{\geq 1}$. Likewise, the state of Σ_2, denoted by x_2, belongs to $\mathsf{X}_2 \subset \mathbb{R}^{d_2}$, $d_2 \in \mathbb{Z}_{\geq 1}$. In this paper, we do not get into the details of what these states represent (as a particular case, the network state could correspond to the collection of the states of agents in it). In addition, each agent v_i in Σ_1 has an estimate $x_1^i \in \mathbb{R}^{d_1}$ of what the network state is, which may differ from the actual value x_1. Similarly, each agent w_j in Σ_2 has an estimate $x_2^j \in \mathbb{R}^{d_2}$ of what the network state is. Within each network, neighboring agents can share their estimates. Networks can also obtain information about each other. This is modeled by means of a bipartite directed graph Σ_{eng}, called *engagement* graph, with disjoint vertex sets $\{v_1, \dots, v_{n_1}\}$ and $\{w_1, \dots, w_{n_2}\}$, where every agent has at least one out-neighbor. According to this model, an agent in Σ_1 obtains information from its out-neighbors in Σ_{eng} about their estimates of the state of Σ_2, and vice versa.

For each $i \in \{1, \dots, n_1\}$, let $f_1^i : \mathsf{X}_1 \times \mathsf{X}_2 \to \mathbb{R}$ be a locally Lipschitz concave-convex function only available to agent $v_i \in \Sigma_1$. Similarly, let $f_2^j : \mathsf{X}_1 \times \mathsf{X}_2 \to \mathbb{R}$ be a locally Lipschitz concave-convex function only available to agent $w_j \in \Sigma_2$, $j \in \{1, \dots, n_2\}$. The networks Σ_1 and Σ_2 are engaged in a zero-sum game with payoff function $U : \mathsf{X}_1 \times \mathsf{X}_2 \to \mathbb{R}$

$$U(x_1, x_2) = \sum_{i=1}^{n_1} f_1^i(x_1, x_2) = \sum_{j=1}^{n_2} f_2^j(x_1, x_2), \qquad (3)$$

where Σ_1 wishes to maximize U, while Σ_2 wishes to minimize it. The objective of the networks is therefore to settle upon a Nash equilibrium, i.e., to solve the following maxmin problem

$$\max_{x_1 \in \mathsf{X}_1} \min_{x_2 \in \mathsf{X}_2} U(x_1, x_2). \qquad (4)$$

We refer to the this zero-sum game as the 2*-network zero-sum game* and denote it by $\mathbf{G}_{\text{adv-net}} = (\Sigma_1, \Sigma_2, \Sigma_{\text{eng}}, U)$. We assume that $\mathsf{X}_1 \subset \mathbb{R}^{d_1}$ and $\mathsf{X}_2 \subset \mathbb{R}^{d_2}$ are compact convex. For convenience, let $\boldsymbol{x}_1 = (x_1^1, \dots, x_1^{n_1})^T$ and $\boldsymbol{x}_2 = (x_2^1, \dots, x_2^{n_2})^T$ denote vector of agent estimates about the state of the respective networks.

Remark 3.1: (Applications to distributed problems in the presence of adversaries): Multiple scenarios involving networked systems and intelligent adversaries in sensor networks, filtering, finance, and communications [32], [33] can be cast into the strategic framework described above. Here we present a class of examples from communications inspired by [34, Section 5.5.3]. Consider n Gaussian communication channels, each with signal power $p_i \in \mathbb{R}_{\geq 0}$ and noise power $\eta_i \in \mathbb{R}_{\geq 0}$, for $i \in \{1, \dots, n\}$. The capacity of each channel is proportional to $\log(1+\beta p_i/(\sigma_i+\eta_i))$, where $\beta \in \mathbb{R}_{>0}$ and $\sigma_i > 0$ is the receiver noise. Note that capacity is concave in p_i and convex in η_i. Both signal and noise powers must satisfy a budget constraint, i.e., $\sum_{i=1}^n p_i = P$ and $\sum_{i=1}^n \eta_i = C$, for some given $P, C \in \mathbb{R}_{>0}$. Two networks of n agents are involved in this scenario, one, Σ_1, selecting signal powers to maximize capacity, the other one, Σ_2, selecting noise powers to minimize it. The network Σ_1 has decided that m_1 channels will have signal power x_1, while $n-1-m_1$ will have signal power x_2. The remaining nth channel has its power determined to satisfy the budget constraint, i.e., $P - m_1 x_1 - (n-1-m_1)x_2$. Likewise, the network Σ_2 does something similar with m_2 channels with noise power y_1, $n-1-m_2$ channels with noise power y_2, and one last channel with noise power $C - m_2 y_1 - (n-1-m_2)y_2$. Each network is aware of the partition made by the other one. The individual objective function of the two agents (one from Σ_1, the other from Σ_2) making decisions on the power levels of the ith channel is the channel capacity itself. For $i \in \{1, \dots, n-1\}$, this takes the form

$$f^i(x, y) = \log\Big(1 + \frac{\beta x_a}{\sigma_i + y_b}\Big),$$

for some $a, b \in \{1, 2\}$. Here $x = (x_1, x_2)$ and $y = (y_1, y_2)$. For $i = n$, it takes instead the form

$$f^n(x, y) = \log\Big(1 + \frac{\beta(P - m_1 x_1 - (n-1-m_1)x_2)}{\sigma_n + C - m_2 y_1 - (n-1-m_2)y_2}\Big).$$

Note that $\sum_{i=1}^n f^i(x, y)$ is the total capacity of the n communication channels. •

A. Reformulation of the 2-network zero-sum game

In this section, we describe how agents in each network use the information obtained from their neighbors to compute the value of their own objective functions. Based on these estimates, we introduce a reformulation of the $\mathbf{G}_{\text{adv-net}} = (\Sigma_1, \Sigma_2, \Sigma_{\text{eng}}, U)$ which is instrumental for establishing some of our results.

Each agent in Σ_1 has a locally Lipschitz, concave-convex function $\tilde{f}_1^i : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2 n_2} \to \mathbb{R}$ with the properties:

- **(Extension of own payoff function):** for any $x_1 \in \mathbb{R}^{d_1}$, $x_2 \in \mathbb{R}^{d_2}$,
$$\tilde{f}_1^i(x_1, \mathbf{1}_{n_2} \otimes x_2) = f_1^i(x_1, x_2). \qquad (5a)$$
- **(Distributed over Σ_{eng}):** there exists $\mathfrak{f}_1^i : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2 |\mathcal{N}^{\text{in}}_{\Sigma_{\text{eng}}}(v_i)|} \to \mathbb{R}$ such that, for any $x_1 \in \mathbb{R}^{d_1}$ $\boldsymbol{x}_2 \in \mathbb{R}^{d_2 n_2}$,
$$\tilde{f}_1^i(x_1, \boldsymbol{x}_2) = \mathfrak{f}_1^i(x_1, \pi_1^i(\boldsymbol{x}_2)), \qquad (5b)$$
with $\pi_1^i : \mathbb{R}^{d_2 n_2} \to \mathbb{R}^{d_2 |\mathcal{N}^{\text{out}}_{\Sigma_{\text{eng}}}(v_i)|}$ the projection of $\boldsymbol{x}_2$ to the values received by v_i from its out-neighbors in Σ_{eng}.

Each agent in Σ_2 has a function $\tilde{f}_2^j : \mathbb{R}^{d_1 n_1} \times \mathbb{R}^{d_2} \to \mathbb{R}$ with similar properties. The collective payoff functions of the two networks are

$$\tilde{U}_1(\boldsymbol{x}_1, \boldsymbol{x}_2) = \sum_{i=1}^{n_1} \tilde{f}_1^i(x_1^i, \boldsymbol{x}_2), \qquad (6a)$$

$$\tilde{U}_2(\boldsymbol{x}_1, \boldsymbol{x}_2) = \sum_{j=1}^{n_2} \tilde{f}_2^j(\boldsymbol{x}_1, x_2^j). \qquad (6b)$$

In general, the functions $\tilde{U}_1$ and $\tilde{U}_2$ need not be the same. However, $\tilde{U}_1(\mathbf{1}_{n_1} \otimes x_1, \mathbf{1}_{n_1} \otimes x_2) = \tilde{U}_2(\mathbf{1}_{n_1} \otimes x_1, \mathbf{1}_{n_1} \otimes x_2)$, for any $x_1 \in \mathbb{R}^{d_1}$, $x_2 \in \mathbb{R}^{d_2}$. When both functions coincide, the next result shows that the original game can be lifted to a (constrained) zero-sum game.

Lemma 3.2: (Reformulation of the 2-network zero-sum game): Assume that the individual payoff functions $\{\tilde{f}_1^i\}_{i=1}^{n_1}$, $\{\tilde{f}_2^j\}_{j=1}^{n_2}$ satisfying (5) are such that the network payoff functions defined in (6) satisfy $\tilde{U}_1 = \tilde{U}_2$, and let $\tilde{U}$ denote this common function. Then, the problem (4) on $\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ is equivalent to the following problem on $\mathbb{R}^{n_1 d_1} \times \mathbb{R}^{n_2 d_2}$,

$$\max_{\boldsymbol{x}_1 \in \mathsf{X}_1^{n_1}} \min_{\boldsymbol{x}_2 \in \mathsf{X}_2^{n_2}} \tilde{U}(\boldsymbol{x}_1, \boldsymbol{x}_2),$$
$$\text{subject to} \quad \mathbf{L}_1 \boldsymbol{x}_1 = \mathbf{0}_{n_1 d_1}, \quad \mathbf{L}_2 \boldsymbol{x}_2 = \mathbf{0}_{n_2 d_2}, \qquad (7)$$

with $\mathbf{L}_\ell = \mathsf{L}_\ell \otimes \mathsf{I}_{d_\ell}$ and L_ℓ the Laplacian of Σ_ℓ, $\ell \in \{1,2\}$. We denote by $\tilde{\mathbf{G}}_{\text{adv-net}} = (\Sigma_1, \Sigma_2, \Sigma_{\text{eng}}, \tilde{U})$ the constrained zero-sum game defined by (7) and refer to this situation by saying that $\mathbf{G}_{\text{adv-net}}$ can be lifted to $\tilde{\mathbf{G}}_{\text{adv-net}}$. Our objective is to design a coordination algorithm that is implementable with the information that agents in Σ_1 and Σ_2 possess and leads them to find a Nash equilibrium of $\tilde{\mathbf{G}}_{\text{adv-net}}$, which corresponds to a Nash equilibrium of $\mathbf{G}_{\text{adv-net}}$ by Lemma 3.2. Achieving this goal, however, is nontrivial because individual agents, not networks themselves, are the decision makers. From the point of view of agents in each network, the objective is to agree on the states of both their own network and the other network, and that the resulting states correspond to a Nash equilibrium of $\mathbf{G}_{\text{adv-net}}$.

We finish this section by presenting a characterization of the Nash equilibria of $\tilde{\mathbf{G}}_{\text{adv-net}}$, instrumental for proving some of our upcoming results.

Proposition 3.3: (Characterization of the Nash equilibria of $\tilde{\mathbf{G}}_{\text{adv-net}}$): For Σ_1, Σ_2 strongly connected and weight-balanced, define F_1 and F_2 by

$$F_1(\boldsymbol{x}_1, \boldsymbol{z}_1, \boldsymbol{x}_2) = -\tilde{U}(\boldsymbol{x}_1, \boldsymbol{x}_2) + \boldsymbol{x}_1^T \mathbf{L}_1 \boldsymbol{z}_1 + \frac{1}{2} \boldsymbol{x}_1^T \mathbf{L}_1 \boldsymbol{x}_1,$$

$$F_2(\boldsymbol{x}_2, \boldsymbol{z}_2, \boldsymbol{x}_1) = \tilde{U}(\boldsymbol{x}_1, \boldsymbol{x}_2) + \boldsymbol{x}_2^T \mathbf{L}_2 \boldsymbol{z}_2 + \frac{1}{2} \boldsymbol{x}_2^T \mathbf{L}_2 \boldsymbol{x}_2.$$

Then, F_1 and F_2 are convex in their first argument, linear in their second one, and concave in their third one. Moreover, assume $(\boldsymbol{x}_1^*, \boldsymbol{z}_1^*, \boldsymbol{x}_2^*, \boldsymbol{z}_2^*)$ satisfies the following *saddle property* for (F_1, F_2): $(\boldsymbol{x}_1^*, \boldsymbol{z}_1^*)$ is a saddle point of $(\boldsymbol{x}_1, \boldsymbol{z}_1) \mapsto F_1(\boldsymbol{x}_1, \boldsymbol{z}_1, \boldsymbol{x}_2^*)$ and $(\boldsymbol{x}_2^*, \boldsymbol{z}_2^*)$ is a saddle point of $(\boldsymbol{x}_2, \boldsymbol{z}_2) \mapsto F_2(\boldsymbol{x}_2, \boldsymbol{z}_2, \boldsymbol{x}_1^*)$. Then,

(i) $(\boldsymbol{x}_1^*, \boldsymbol{z}_1^* + \mathbf{1}_{n_1} \otimes a_1, \boldsymbol{x}_2^*, \boldsymbol{z}_2^* + \mathbf{1}_{n_2} \otimes a_2)$ satisfies the saddle property for (F_1, F_2) for any $a_1 \in \mathbb{R}^{d_1}$, $a_2 \in \mathbb{R}^{d_2}$, and

(ii) $(\boldsymbol{x}_1^*, \boldsymbol{x}_2^*)$ is a Nash equilibrium of $\tilde{\mathbf{G}}_{\text{adv-net}}$.

Furthermore,

(iii) if $(\boldsymbol{x}_1^*, \boldsymbol{x}_2^*)$ is a Nash equilibrium of $\tilde{\mathbf{G}}_{\text{adv-net}}$ then there exists $\boldsymbol{z}_1^*, \boldsymbol{z}_2^*$ such that $(\boldsymbol{x}_1^*, \boldsymbol{z}_1^*, \boldsymbol{x}_2^*, \boldsymbol{z}_2^*)$ satisfies the saddle property for (F_1, F_2).

IV. Distributed Nash seeking dynamics for undirected graphs

Here, we review following [7] a dynamics which solves (7) when Σ_1 and Σ_2 are undirected. In this scenario, the gradients of F_1 and F_2 are, respectively, distributed over Σ_1 and Σ_2. By Proposition 3.3, it is natural to consider the saddle-point dynamics for F_1 and F_2 to solve (4), i.e.,

$$\dot{\boldsymbol{x}}_1 + \mathbf{L}_1 \boldsymbol{x}_1 + \mathbf{L}_1 \boldsymbol{z}_1 = \nabla_{\boldsymbol{x}_1} \tilde{U}(\boldsymbol{x}_1, \boldsymbol{x}_2), \tag{8a}$$

$$\dot{\boldsymbol{z}}_1 = \mathbf{L}_1 \boldsymbol{x}_1, \tag{8b}$$

$$\dot{\boldsymbol{x}}_2 + \mathbf{L}_2 \boldsymbol{x}_2 + \mathbf{L}_2 \boldsymbol{z}_2 = -\nabla_{\boldsymbol{x}_2} \tilde{U}(\boldsymbol{x}_1, \boldsymbol{x}_2), \tag{8c}$$

$$\dot{\boldsymbol{z}}_2 = \mathbf{L}_2 \boldsymbol{x}_2, \tag{8d}$$

where $\boldsymbol{x}_j, \boldsymbol{z}_j \in \mathbb{R}^{n_j d_j}$, $j \in \{1,2\}$. The following result establishes the convergence properties of this dynamics.

Theorem 4.1: (Asymptotic convergence of the undirected distributed Nash seeking dynamics): Consider the zero-sum game $\tilde{\mathbf{G}}_{\text{adv-net}} = (\Sigma_1, \Sigma_2, \Sigma_{\text{eng}}, \tilde{U})$, where

(i) Σ_1 and Σ_2 are connected and undirected,

(ii) $\tilde{U} : \mathsf{X}_1^{n_1} \times \mathsf{X}_2^{n_2} \to \mathbb{R}$, X_1 and X_2 compact convex subsets of, respectively, $\mathbb{R}^{d_1}$ and $\mathbb{R}^{d_2}$, is a differentiable strictly concave-convex function, distributed over Σ_{eng} and also Σ_1 and Σ_2 in the sense of (6).

Then the projection onto the first and third components of the solutions to (8) asymptotically converges to the solution of (7).

It is worth mentioning that this result, in fact, also holds true when Σ_1 and Σ_2 are undirected and U is the sum of locally Lipschitz concave-convex functions, see [7].

V. Distributed Nash seeking dynamics for directed graphs

In this section, we introduce a continuous-time Nash seeking dynamics implementable over strongly connected and weight-balanced directed topologies. This dynamics is distributed over each individual network and can find the Nash equilibria of the zero-sum game, provided that the payoff function is differentiable, strictly concave-convex, with globally Lipschitz gradient. This result generalizes the Nash seeking saddle-point dynamics of (8) to directed topologies.

We start by modifying the dynamics of (8) as

$$\dot{\boldsymbol{x}}_1 + \alpha \mathbf{L}_1 \boldsymbol{x}_1 + \mathbf{L}_1 \boldsymbol{z}_1 = \nabla \tilde{U}_{\boldsymbol{x}_1}(\boldsymbol{x}_1, \boldsymbol{x}_2), \tag{9a}$$

$$\dot{\boldsymbol{z}}_1 = \mathbf{L}_1 \boldsymbol{x}_1, \tag{9b}$$

$$\dot{\boldsymbol{x}}_2 + \alpha \mathbf{L} \boldsymbol{x}_2 + \mathbf{L} \boldsymbol{z}_2 = -\nabla \tilde{U}_{\boldsymbol{x}_2}(\boldsymbol{x}_1, \boldsymbol{x}_2), \tag{9c}$$

$$\dot{\boldsymbol{z}}_2 = \mathbf{L}_2 \boldsymbol{x}_2, \tag{9d}$$

where $\alpha \in \mathbb{R}_{>0}$ is a design parameter and the payoff function is differentiable with globally Lipschitz gradient. The reason behind including the parameter α in the dynamics is that (8) may fail to converge when transcribed to directed graphs, for the same reason that the continuous-time saddle-point distributed optimization dynamics may fail on undirected graphs, see [27].

Next, we show that a suitable choice of this design parameter, makes this dynamics convergent.

Theorem 5.1: (Asymptotic convergence of the directed distributed Nash seeking dynamics): Consider the zero-sum game $\tilde{\mathbf{G}}_{\text{adv-net}} = (\Sigma_1, \Sigma_2, \Sigma_{\text{eng}}, \tilde{U})$, where

(i) Σ_1 and Σ_2 are strongly connected and weight-balanced,

(ii) $\tilde{U} : \mathsf{X}_1^{n_1} \times \mathsf{X}_2^{n_2} \to \mathbb{R}$, X_1 and X_2 compact convex subsets of, respectively, $\mathbb{R}^{d_1}$ and $\mathbb{R}^{d_2}$, is a differentiable strictly concave-convex function with globally Lipschitz gradient, distributed over Σ_{eng} and also Σ_1 and Σ_2 in the sense of (6).

Let $h : \mathbb{R}_{>0} \to \mathbb{R}$ be defined by

$$h(r) = \frac{1}{2} \Lambda_*^{\min} \left(-\frac{r^4 + 3r^2 + 2}{r} + \sqrt{\left(\frac{r^4 + 3r^2 + 2}{r} \right)^2 - 4} \right) + \frac{K r^2}{(1 + r^2)}, \tag{10}$$

$\Lambda_*^{\min} = \min_{j=1,2} \{ \Lambda_*(\mathsf{L}_j + \mathsf{L}_j^T) \}$, where $\Lambda_*(\cdot)$ denotes the smallest non-zero eigenvalue and $K \in \mathbb{R}_{>0}$ is the Lipschitz constant for the gradient of $\tilde{U}$. Then there exists $\beta^* \in \mathbb{R}_{>0}$ with $h_j(\beta^*) = 0$, $j \in \{1, 2\}$, such that for all $0 < \beta < \beta^*$, the projection onto the first and third components of the solutions of (9) with $\alpha = \frac{\beta^2+2}{\beta}$ asymptotically converges to the solution of (7).

Remark 5.2: (Comparison with the best-response dynamics): The advantage of using the gradient flow is that it avoids the cumbersome computation of the best-response map. This, however, does not come for free. There are concave-convex functions for which the (distributed) gradient flow dynamics, unlike the best-response dynamics, fails to converge to the saddle point, see [12] for an example. •

Remark 5.3: (Scenarios with more than two adversarial networks): It is known that there are continuous-time zero-sum games with three players and strictly concave-convex payoff functions, for which even the best-response dynamics fails to converge, see [9]. This leaves little hope for extensions of Theorems 4.1 and 5.1 to N-network zero-sum games, with $N \in \mathbb{Z}_{\geq 3}$. •

We finish this section with an example.

Example 5.4: (Distributed adversarial selection of signal and noise power via (9)*):* Recall the communication scenario described in Remark 3.1. Consider 5 channels, $\{v_1, v_2, v_3, v_4, v_5\}$, for which the network Σ_1 has decided that $\{v_1, v_3\}$ have signal power x_1 and $\{v_2, v_4\}$ have signal power x_2. Channel v_5 has its signal power determined to satisfy the budget constraint $P \in \mathbb{R}_{>0}$, i.e., $P - 2x_1 - 2x_2$. Similarly, the network Σ_2 has decided that v_1 has noise power y_1, $\{v_2, v_3, v_4\}$ have noise power y_2, and v_5 has noise power $C - y_1 - 3y_2$ to meet the budget constraint $C \in \mathbb{R}_{>0}$. We let $\boldsymbol{x} = (x^1, x^2, x^3, x^4, x^5)$ and $\boldsymbol{y} = (y^1, y^2, y^3, y^4, y^5)$, where $x^i = (x_1^i, x_2^i) \in [0, P]^2$ and $y^i = (y_1^i, y_2^i) \in [0, C]^2$, for each $i \in \{1, \dots, 5\}$.

The networks Σ_1 and Σ_2, which are weight-balanced and strongly connected, and the engagement topology Σ_{eng} are shown in Figure 1. Note that, according to this topology, each

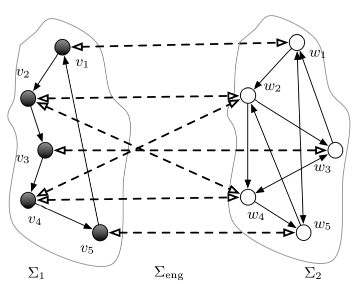

Fig. 1. The Σ_1, Σ_2 and Σ_{eng} for the case study of Example 5.4 are shown. Edges which correspond to Σ_{eng} are dashed.

agent can observe the power employed by its adversary in its channel and, additionally, the agents in channel 2 can obtain information about the estimates of the opponent in channel 4 and vice versa. The payoff functions of the agents are given in Remark 3.1, where we take $\sigma_i = \sigma_1$, for $i \in \{1, 3, 5\}$, and $\sigma_i = \sigma_2$, for $i \in \{2, 4\}$, with $\sigma_1, \sigma_2 \in \mathbb{R}_{>0}$.

This example fits into the approach described in Section III-A by considering the following extended payoff functions:

$$
\begin{aligned}
\tilde{f}_1^1(x^1, \boldsymbol{y}) &= \log(1 + \frac{\beta x_1^1}{\sigma_1 + y_1^1}), \\
\tilde{f}_1^2(x^2, \boldsymbol{y}) &= \frac{1}{3}\log(1 + \frac{\beta x_2^2}{\sigma_2 + y_2^4}) + \frac{2}{3}\log(1 + \frac{\beta x_2^2}{\sigma_2 + y_2^2}), \\
\tilde{f}_1^3(x^3, \boldsymbol{y}) &= \log(1 + \frac{\beta x_1^3}{\sigma_1 + y_2^3}), \\
\tilde{f}_1^4(x^4, \boldsymbol{y}) &= \frac{1}{3}\log(1 + \frac{\beta x_2^4}{\sigma_2 + y_2^2}) + \frac{2}{3}\log(1 + \frac{\beta x_2^4}{\sigma_2 + y_2^4}), \\
\tilde{f}_1^5(x^5, \boldsymbol{y}) &= \log\Big(1 + \frac{\beta(P - 2x_1^5 - 2x_2^5)}{\sigma_1 + C - y_1^5 - 3y_2^5}\Big), \\
\tilde{f}_2^1(\boldsymbol{x}, y^1) &= \tilde{f}_1^1(x^1, \boldsymbol{y}), \quad \tilde{f}_2^3(\boldsymbol{x}, y^3) = \tilde{f}_1^3(x^3, \boldsymbol{y}), \\
\tilde{f}_2^2(\boldsymbol{x}, y^2) &= \frac{2}{3}\log(1 + \frac{\beta x_2^2}{\sigma_2 + y_2^2}) + \frac{1}{3}\log(1 + \frac{\beta x_2^4}{\sigma_2 + y_2^2}), \\
\tilde{f}_2^4(\boldsymbol{x}, y^4) &= \frac{1}{3}\log(1 + \frac{\beta x_2^2}{\sigma_2 + y_2^4}) + \frac{2}{3}\log(1 + \frac{\beta x_2^4}{\sigma_2 + y_2^4}), \\
\tilde{f}_2^5(\boldsymbol{x}, y^5) &= \tilde{f}_1^5(x^5, \boldsymbol{y}).
\end{aligned}
$$

Note that these functions are strictly concave and thus the zero-sum game defined has a unique saddle point on the set $[0, P]^2 \times [0, C]^2$. These functions satisfy (5) and $\tilde{U}_1 = \tilde{U}_2$. Figure 2 shows the convergence of the dynamics (9) to the Nash equilibrium of the resulting 2-network zero-sum game. •

VI. Conclusions and future work

We have considered a class of strategic scenarios in which two networks of agents are involved in a zero-sum game. The networks' objectives are to either maximize or minimize a common objective function. Individual agents collaborate with neighbors in their respective network and have partial knowledge of the state of the agents in the other network. Specifically, we have considered directed networks where information flows unidirectionally. We have introduced the directed distributed Nash-seeking dynamics and shown that, for appropriate parameter choices, this dynamics is guaranteed to converge to the Nash equilibrium for strictly concave-convex and differentiable objective functions with globally Lipschitz gradients. Future work will include relaxing the assumptions on the problem data under which convergence is guaranteed, including the smoothness, strict concavity-convexity properties, and sum decomposition of the objective function, and exploring the application of our results to various areas, including competitive social networks, collective bargaining, and collaborative pursuit-evasion.

Acknowledgments

This work was supported in part by Award FA9550-10-1-0499.

References

[1] R. Stocker and W. M. Durham, "Tumbling or stealth," *Science*, vol. 325, pp. 400–402, 2009.

[2] L. A. Johnson, "Competition interactions between cells: death, growth, and geography," *Science*, vol. 324, pp. 1679–1682, 2009.

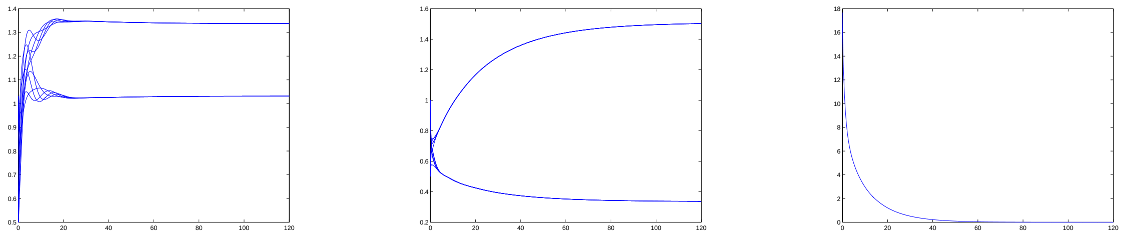

(a) (b) (c)

Fig. 2. Execution of (9) over the networked strategic scenario described in Example 5.4, with $\beta = 8$, $\sigma_1 = 1$, $\sigma_2 = 4$, $P = 6$, and $C = 4$. (a) and (b) show the evolution of the agent's estimates of the state of networks Σ_1 and Σ_2, respectively, and (c) shows the value of the Lyapunov function. Here, $\alpha = 3$ in (9) and initially, $\boldsymbol{x}^0 = ((1, 0.5), (0.5, 1), (0.5, 0.5), (0.5, 1), (0.5, 1))^T$, $\boldsymbol{z}_1^0 = \mathbf{0}_{10}$, $\boldsymbol{y}^0 = ((1, 0.5), (0.5, 1), (0.5, 1), (0.5, 0.5), (1, 0.5))^T$ and $\boldsymbol{z}_2^0 = \mathbf{0}_{10}$. The equilibrium $(\boldsymbol{x}^*, \boldsymbol{z}_1^*, \boldsymbol{y}^*, \boldsymbol{z}_2^*)$ is $\boldsymbol{x}^* = (1.3371, 1.0315)^T \otimes \mathbf{1}_5$, $\boldsymbol{y}^* = (1.5027, 0.3366)^T \otimes \mathbf{1}_5$, $\boldsymbol{z}_1^* = (0.7508, 0.5084, 0.1447, 0.5084, 0.1447, -0.1271, -0.5201, -0.1271, -0.5201, -0.7626)^T$ and $\boldsymbol{z}_2^* = (0.1079, -0.0987, -0.0002, 0.2237, 0.0358, 0.2875, -0.0360, 0.0087, -0.1076, -0.4213)$.

[3] S. Roy, C. Ellis, S. Shiva, D. Dasgupta, V. Shandilya, and Q. Wu, "A survey of game theory as applied to network security," in *Proceedings of the 43rd International Conference on System Sciences*, (Koloa, Kauai, Hawaii), Jan. 2010.

[4] M. Broecheler, O. Shakarian, and V. S. Subrahmanian, "A scalable framework for modeling competitive diffusion in social networks," in *IEEE International Conf. on Social Computing*, (Minneapolis, MN), pp. 295 – 302, Aug. 2010.

[5] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis)information in social networks," *Games and Economic Behavior*, vol. 7, no. 2, 2010. 194-227.

[6] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar, "Opinion fluctuations and disagreement in social networks," *Mathematics of Operation Research*, Sept. 2010. Submitted. Available at `http://arxiv.org/abs/1009.2653`.

[7] B. Gharesifard and J. Cortés, "Distributed convergence to Nash equilibria by adversarial networks with undirected topologies," in *American Control Conference*, (Montréal, Canada), pp. 5881–5886, 2012.

[8] J. Hofbauer and S. Sorin, "Best response dynamics for continuous zero-sum games," *Discrete and Continuous Dynamical Systems Ser. B*, vol. 6, no. 1, pp. 215–224, 2006.

[9] E. N. Barron, R. Goebel, and R. R. Jensen, "Best response dynamics for continuous games," *Proceeding of the American Mathematical Society*, vol. 138, no. 3, pp. 1069–1083, 2010.

[10] K. Arrow, L. Hurwitz, and H. Uzawa, *A Gradient Method for Approximating Saddle Points and Constrained Maxima*. United States Army Air Forces: Rand Corporation, 1951.

[11] K. Arrow, L. Hurwitz, and H. Uzawa, *Studies in Linear and Non-Linear Programming*. Stanford, California: Stanford University Press, 1958.

[12] D. Feijer and F. Paganini, "Stability of primal-dual gradient dynamics and applications to network optimization," *Automatica*, vol. 46, pp. 1974–1981, 2010.

[13] D. Maistroskii, "Gradient methods for finding saddle points," *Matekon*, vol. 13, pp. 3–22, 1977.

[14] A. Nedic and A. Ozdgalar, "Subgradient methods for saddle-point problems," *Journal of Optimization Theory & Applications*, vol. 142, no. 1, pp. 205–228, 2009.

[15] S. Li and T. Başar, "Distributed algorithms for the computation of noncooperative equilibria," *Automatica*, vol. 23, no. 4, pp. 523–533, 1987.

[16] P. Frihauf, M. Krstić, and T. Başar, "Nash equilibrium seeking in noncooperative games," *IEEE Transactions on Automatic Control*, 2012. To appear.

[17] M. S. Stankovic, K. H. Johansson, and D. M. Stipanovic, "Distributed seeking of Nash equilibria in mobile sensor networks," in *IEEE Conf. on Decision and Control*, (Atlanta, GA), pp. 5598–5603, 2010.

[18] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.

[19] A. Nedic, A. Ozdaglar, and P. A. Parrilo, "Constrained consensus and optimization in multi-agent networks," *IEEE Transactions on Automatic Control*, vol. 55, no. 4, pp. 922–938, 2010.

[20] B. Johansson, M. Rabi, and M. Johansson, "A randomized incremental subgradient method for distributed optimization in networked systems," *SIAM Journal on Control and Optimization*, vol. 20, no. 3, pp. 1157–1170, 2009.

[21] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[22] W. Ren and R. W. Beard, *Distributed Consensus in Multi-vehicle Cooperative Control*. Communications and Control Engineering, Springer, 2008.

[23] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Applied Mathematics Series, Princeton University Press, 2009. Electronically available at http://coordinationbook.info.

[24] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Applied Mathematics Series, Princeton University Press, 2010.

[25] J. Wang and N. Elia, "Control approach to distributed optimization," in *Allerton Conf. on Communications, Control and Computing*, (Monticello, IL), pp. 557–561, Oct. 2010.

[26] J. Wang and N. Elia, "A control perspective for centralized and distributed convex optimization," in *IEEE Conf. on Decision and Control*, (Orlando, Florida), pp. 3800–3805, 2011.

[27] B. Gharesifard and J. Cortés, "Distributed continuous-time convex optimization on weight-balanced digraphs," *IEEE Transactions on Automatic Control*, 2012. Conditionally accepted.

[28] R. T. Rockafellar, *Convex Analysis*. Princeton Landmarks in Mathematics and Physics, Princeton, NJ: Princeton University Press, 1997. Reprint of 1970 edition.

[29] H. K. Khalil, *Nonlinear Systems*. Prentice Hall, 3 ed., 2002.

[30] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. SIAM, 2 ed., 1999.

[31] E. A. Ok, *Real Analysis with Economic Applications*. Princeton, NJ: Princeton University Press, 2007.

[32] J. S. Kim and S. Boyd, "A minimax theorem with applications to machine learning, signal processing, and finance," *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1344–1367, 2008.

[33] P. Wan and M. D. Lemmon, "Event-triggered distributed optimization in sensor networks," in *Symposium on Information Processing of Sensor Networks*, (San Francisco, CA), pp. 49–60, 2009.

[34] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.