

Fully distributed Nash equilibrium seeking over time-varying communication networks with linear convergence rate

Bianchi, M.; Grammatico, S.

DOI

[10.1109/LCSYS.2020.3002734](https://doi.org/10.1109/LCSYS.2020.3002734)

Publication date

2021

Document Version

Final published version

Published in

IEEE Control Systems Letters

Citation (APA)

Bianchi, M., & Grammatico, S. (2021). Fully distributed Nash equilibrium seeking over time-varying communication networks with linear convergence rate. *IEEE Control Systems Letters*, 5(2), 499-504. <https://doi.org/10.1109/LCSYS.2020.3002734>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' – Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Fully Distributed Nash Equilibrium Seeking Over Time-Varying Communication Networks With Linear Convergence Rate

Mattia Bianchi^{ID} and Sergio Grammatico^{ID}, *Senior Member, IEEE*

Abstract—We design a distributed algorithm for learning Nash equilibria over time-varying communication networks in a partial-decision information scenario, where each agent can access its own cost function and local feasible set, but can only observe the actions of some neighbors. Our algorithm is based on projected pseudo-gradient dynamics, augmented with consensual terms. Under strong monotonicity and Lipschitz continuity of the game mapping, we provide a simple proof of linear convergence, based on a contractivity property of the iterates. Compared to similar solutions proposed in literature, we also allow for time-varying communication and derive tighter bounds on the step sizes that ensure convergence. In fact, in our numerical simulations, our algorithm outperforms the existing gradient-based methods, when the step sizes are set to their theoretical upper bounds. Finally, to relax the assumptions on the network structure, we propose a different pseudo-gradient algorithm, which is guaranteed to converge on time-varying balanced directed graphs.

Index Terms—Game theory, optimization algorithms, networked control systems.

I. INTRODUCTION

NASH equilibrium (NE) problems arise in several network systems, where multiple selfish decision-makers, or agents, aim at optimizing their individual, yet inter-dependent, objective functions. Engineering applications include communication networks [1], demand-side management in the smart grid [2], charging of electric vehicles [3] and demand response in competitive markets [4]. From a game-theoretic perspective, the challenge is to assign the agents behavioral rules that eventually ensure the attainment of a NE, a joint action from which no agent has an incentive to unilaterally deviate.

Literature review: Typically, NE seeking algorithms are designed under the assumption that each agent can access

the decisions of all the competitors [5], [6], [7]. This *full-decision information* hypothesis requires the presence of a coordinator, that broadcast the data to the network, and it is impractical for some applications [8], [9]. One example is the Nash-Cournot competition model described in [10], where the profit of each of a group of firms depends not only on its own production, but also on the whole amount of sales, a quantity not directly accessible by any of the firms. Therefore, in recent years, there has been an increased attention for fully distributed algorithms that allow to compute NEs relying on local information only. In this letter, we consider the so-called *partial-decision information* scenario, where the agents engage in nonstrategic information exchange with some neighbors on a network; based on the data received, they can estimate and eventually reconstruct the actions of all the competitors. This setup has only been introduced very recently. In particular, most of the results available resort to (projected) gradient and consensus dynamics, both in continuous time [11], [12], and discrete time. For the discrete time case, fixed-step algorithms were proposed in [13], [14], [15] (the latter for generalized games), all exploiting a certain restricted monotonicity property. Alternatively, the authors of [16] developed a gradient-play scheme by leveraging contractivity properties of doubly stochastic matrices. Nevertheless, in all these approaches theoretical guarantees are provided only for step sizes that are typically very small, affecting the speed of convergence. Furthermore, all the methods cited are designed for a time-invariant, undirected network. To the best of our knowledge, switching communication topologies have only been addressed with diminishing step sizes. For instance, the early work [10] considered aggregative games over time-varying jointly connected undirected graphs. This result was extended by the authors of [17] to games with coupling constraints. In [18], an asynchronous gossip algorithm was presented to seek a NE over directed graphs. The drawback is that vanishing steps typically result in slow convergence.

Contribution: Motivated by the above, in this letter we present the first fixed-step NE seeking algorithms for strongly monotone games over time-varying communication networks. Our novel contributions are summarized as follows:

- We propose a fully distributed projected gradient-play method, that is guaranteed to converge with linear rate when the network adjacency matrix is doubly stochastic.

Manuscript received March 15, 2020; revised May 14, 2020; accepted June 1, 2020. Date of publication June 16, 2020; date of current version July 3, 2020. This work was supported in part by Netherlands Organization for Scientific Research through Project OMEGA under Grant 613.001.702, and in part by the European Research Council through Project COSMOS under Grant 802348. Recommended by Senior Editor M. Guay. (Corresponding author: Mattia Bianchi.)

The authors are with the Delft Center for Systems and Control, TU Delft, 2613 GA Delft, The Netherlands (e-mail: m.bianchi@tudelft.nl; s.grammatico@tudelft.nl).

Digital Object Identifier 10.1109/LCSYS.2020.3002734

2475-1456 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.

With respect to [16], we consider a time-varying communication network and we allow for constrained action sets. Moreover, differently from the state of the art, we provide an upper bound on the step size that does not vanish as the number of agents increases (Section III);

- We show via numerical simulations that, even in the case of fixed networks, our algorithm outperforms the existing pseudo-gradient based dynamics, when the step sizes are set to their theoretical upper bounds (Section V);
- We prove that linear convergence to a NE on time-varying weight-balanced directed graphs can be achieved via a forward-backward algorithm [19, Sec. 12.7.2], which has been studied in [14], [15], but only for the special case of fixed undirected networks (Section IV).

Basic notation: \mathbb{N} is the set of natural numbers, including 0. \mathbb{R} is the set of real numbers. $\mathbf{0}_n$ ($\mathbf{1}_n$) denotes the vector of dimension n with all elements equal to 0 (1); I_n the identity matrix of dimension n ; the subscripts might be omitted when there is no ambiguity. For a matrix $A \in \mathbb{R}^{m \times n}$, A^\top denotes its transpose; $[A]_{i,j}$ is the element on row i and column j , and the second subscript is omitted if $n = 1$; $\sigma_{\min}(A) = \sigma_1(A) \leq \dots \leq \sigma_n(A) =: \sigma_{\max}(A) = \|A\|$ denote its singular values. If $A \in \mathbb{R}^{n \times n}$, $\det(A)$ is its determinant; $A \succ 0$ stands for symmetric positive definite matrix; if A is symmetric, $\lambda_{\min}(A) = \lambda_1(A) \leq \dots \leq \lambda_n(A) =: \lambda_{\max}(A)$ denote its eigenvalues. \otimes denotes the Kronecker product. $\text{diag}(A_1, \dots, A_N)$ denotes the block diagonal matrix with A_1, \dots, A_N on its diagonal. Given N vectors x_1, \dots, x_N , $x := \text{col}(x_1, \dots, x_N) = [x_1^\top \dots x_N^\top]^\top$ and $x_{-i} = \text{col}(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N)$. $\|\cdot\|$ denotes the Euclidean vector norm. For a differentiable function $g: \mathbb{R}^n \rightarrow \mathbb{R}$, $\nabla_x g(x)$ denotes its gradient. A mapping $\mathcal{A}: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is ℓ -Lipschitz continuous if, for any $x, y \in \mathbb{R}^m$, $\|\mathcal{A}(x) - \mathcal{A}(y)\| \leq \ell \|x - y\|$. $\text{proj}_S: \mathbb{R}^n \rightarrow S$ denotes the Euclidean projection onto a closed convex set S . An operator $\mathcal{F}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is (μ) -strongly monotone if, for any $x, y \in \mathbb{R}^n$, $(\mathcal{F}(x) - \mathcal{F}(y))^\top (x - y) \geq 0$ ($\geq \mu \|x - y\|^2$). The variational inequality $\text{VI}(\mathcal{F}, S)$ is the problem of finding a vector $x^* \in S$ such that $\mathcal{F}(x^*)^\top (x - x^*) \geq 0$, for all $x \in S$.

II. MATHEMATICAL SETUP

We consider a set of agents $\mathcal{I} := \{1, \dots, N\}$, where each agent $i \in \mathcal{I}$ shall choose its action (i.e., decision variable) x_i from its local decision set $\Omega_i \subseteq \mathbb{R}^{n_i}$. Let $x = \text{col}((x_i)_{i \in \mathcal{I}}) \in \Omega$ denote the stacked vector of all the agents' decisions, $\Omega = \Omega_1 \times \dots \times \Omega_N \subseteq \mathbb{R}^n$ the overall action space and $n := \sum_{i=1}^N n_i$. The goal of each agent $i \in \mathcal{I}$ is to minimize its objective function $J_i(x_i, x_{-i})$, which depends on both the local variable x_i and the decision variables of the other agents $x_{-i} = \text{col}((x_j)_{j \in \mathcal{I} \setminus \{i\}})$. The game is then represented by the inter-dependent optimization problems:

$$\forall i \in \mathcal{I}: \quad \underset{y_i \in \Omega_i}{\text{argmin}} J_i(y_i, x_{-i}). \quad (1)$$

The technical problem we consider in this letter is the computation of a NE, as defined next.

Definition 1: A Nash equilibrium is a set of strategies $x^* = \text{col}((x_i^*)_{i \in \mathcal{I}}) \in \Omega$ such that, for all $i \in \mathcal{I}$:

$$J_i(x_i^*, x_{-i}^*) \leq \inf\{J_i(y_i, x_{-i}^*) | y_i \in \Omega_i\}.$$

The following regularity assumptions are common for NE problems, see, e.g., [15, Ass. 1], [14, Ass. 1].

Standing Assumption 1 (Regularity and Convexity): For each $i \in \mathcal{I}$, the set Ω_i is non-empty, closed and convex; J_i is continuous and the function $J_i(\cdot, x_{-i})$ is convex and continuously differentiable for every x_{-i} .

Under Standing Assumption 1, a joint action x^* is a NE of the game in (1) if and only if it solves the variational inequality $\text{VI}(\mathcal{F}, \Omega)$ [19, Prop. 1.4.2], or, equivalently, if and only if, for any $\alpha > 0$ [19, Prop. 1.5.8],

$$x^* = \text{proj}_\Omega(x^* - \alpha \mathcal{F}(x^*)), \quad (2)$$

where \mathcal{F} is the *pseudo-gradient* mapping of the game:

$$\mathcal{F}(x) := \text{col}((\nabla_{x_i} J_i(x_i, x_{-i}))_{i \in \mathcal{I}}). \quad (3)$$

Next, we postulate a sufficient condition for the existence of a unique NE, namely the strong monotonicity of the pseudo-gradient [19, Th. 2.3.3]. This assumption is always used for (G)NE seeking under partial-decision information with fixed step sizes, e.g., in [14, Ass. 2], [15, Ass. 3]. It implies strong convexity of the functions $J_i(\cdot, x_{-i})$ for every x_{-i} , but not necessarily (strong) convexity of J_i in the full argument.

Standing Assumption 2: The pseudo-gradient mapping in (3) is μ -strongly monotone and ℓ_0 -Lipschitz continuous, for some $\mu, \ell_0 > 0$.

In our setup, each agent i can only access its own cost function J_i and feasible set Ω_i . Moreover, agent i does not have full knowledge of x_{-i} , and only relies on the information exchanged locally with neighbors over a time-varying directed communication network $\mathcal{G}_k(\mathcal{I}, \mathcal{E}_k)$. The ordered pair (i, j) belongs to the set of edges, \mathcal{E}_k , if and only if agent i can receive information from agent j at time k . Let $W_k \in \mathbb{R}^{N \times N}$ denote the weighted adjacency matrix of \mathcal{G}_k , and $w_{i,j}^k := [W_k]_{i,j}$, with $w_{i,j}^k > 0$ if $(i, j) \in \mathcal{E}_k$, $w_{i,j}^k = 0$ otherwise; $D_k = \text{diag}((d_i^k)_{i \in \mathcal{I}})$ and $L_k = D_k - W_k$ the in-degree and Laplacian matrices of \mathcal{G}_k , with $d_i^k = \sum_{j=1}^N w_{i,j}^k$; $\mathcal{N}_i^k = \{j | (i, j) \in \mathcal{E}_k\}$ the set of in-neighbors of agent i .

Standing Assumption 3: For each $k \in \mathbb{N}$, the graph \mathcal{G}_k is strongly connected.

Assumption 1: For all $k \in \mathbb{N}$, the following hold:

- Self-loops:** $w_{i,i}^k > 0$ for all $i \in \mathcal{I}$;
- Double stochasticity:** $W_k \mathbf{1}_N = \mathbf{1}_N$, $\mathbf{1}_N^\top W_k = \mathbf{1}_N^\top$.

Remark 1: Assumption 1(i) is intended just to ease the notation. Instead, Assumption 1(ii) is stronger. It is typically used for networked problems on undirected symmetric graphs, e.g., in [10, Ass. 6], [17, Ass. 3], [16, Ass. 3], justified by the fact that it can be satisfied by assigning the following Metropolis weights to the communication:

$$\tilde{w}_{i,j}^k = \begin{cases} w_{i,j}^k / (\max\{d_i^k, d_j^k\} + 1) & \text{if } j \in \mathcal{N}_i^k \setminus \{i\}; \\ 0 & \text{if } j \notin \mathcal{N}_i^k; \\ 1 - \sum_{j \in \mathcal{N}_i^k \setminus \{i\}} \tilde{w}_{i,j}^k & \text{if } i = j. \end{cases}$$

In practice, to satisfy Assumption 1(ii) in case of symmetric communication, even under time-varying topology, it suffices for the agents to exchange their in-degree with their neighbors at every time step. Therefore, Standing Assumption 3 and Assumption 1 are easily fulfilled for undirected graphs connected at each step. For directed graphs, given any strongly connected topology, weights can be assigned such that the

resulting adjacency matrix (with self-loops) is doubly stochastic, via an iterative distributed process [20]. However, this can be impractical if the network is time-varying.

Under Assumption 1, it holds that $\sigma_{N-1}(W_k) < 1$, for all k , where $\sigma_{N-1}(W_k)$ denotes the second largest singular value of W_k . Moreover, for any $y \in \mathbb{R}^N$,

$$\|W_k(y - \mathbf{1}_N \bar{y})\| \leq \sigma_{N-1}(W_k) \|y - \mathbf{1}_N \bar{y}\|, \quad (4)$$

where $\bar{y} = \frac{1}{N} \mathbf{1}_N^\top y$ is the average of y . We will further assume that $\sigma_{N-1}(W_k)$ is bounded away from 1; this automatically holds if the networks \mathcal{G}_k are chosen among a finite family.

Assumption 2: There exists $\bar{\sigma} \in (0, 1)$ such that $\sigma_{N-1}(W_k) \leq \bar{\sigma}$, for all $k \in \mathbb{N}$.

III. DISTRIBUTED NASH EQUILIBRIUM SEEKING

In this section, we present a pseudo-gradient algorithm to seek a NE of the game (1) in a fully distributed way. To cope with partial-decision information, each agent keeps an estimate of all other agents' actions. Let $\mathbf{x}_i = \text{col}((\mathbf{x}_{i,j})_{j \in \mathcal{I}}) \in \mathbb{R}^{Nn}$, where $\mathbf{x}_{i,i} := x_i$ and $\mathbf{x}_{i,j}$ is agent i 's estimate of agent j 's action, for all $j \neq i$; also, $\mathbf{x}_{j,-i} = \text{col}((\mathbf{x}_{j,i})_{i \in \mathcal{I} \setminus \{j\}})$. The agents aim at asymptotically reconstructing the true value of the opponents' actions, based on the data received from their neighbors. The procedure is summarized in Algorithm 1. Each agent updates its estimates according to consensus dynamics, then its action via a gradient step. We remark that each agent computes the partial gradient of its cost in its local estimates \mathbf{x}_i , not on the actual joint action \mathbf{x} .

To write the algorithm in compact form, let $\mathbf{x} = \text{col}((\mathbf{x}_i)_{i \in \mathcal{I}})$; as in [15, eqs. 13-14], let, for all $i \in \mathcal{I}$,

$$\mathcal{R}_i := [\mathbf{0}_{n_i \times n_{<i}} \quad I_{n_i} \quad \mathbf{0}_{n_i \times n_{>i}}] \in \mathbb{R}^{n_i \times n}, \quad (5)$$

where $n_{<i} := \sum_{j=1}^{i-1} n_j$, $n_{>i} := \sum_{j=i+1}^N n_j$; let also $\mathcal{R} := \text{diag}((\mathcal{R}_i)_{i \in \mathcal{I}}) \in \mathbb{R}^{n \times n}$. In simple terms, \mathcal{R}_i selects the i -th n_i dimensional component from an n -dimensional vector. Thus, $\mathcal{R}_i \mathbf{x}_i = \mathbf{x}_{i,i} = x_i$, and $\mathbf{x} = \mathcal{R} \mathbf{x}$. We define the *extended pseudo-gradient* mapping \mathbf{F} as

$$\mathbf{F}(\mathbf{x}) := \text{col}((\nabla_{\mathbf{x}_i} J_i(x_i, \mathbf{x}_{i,-i}))_{i \in \mathcal{I}}). \quad (6)$$

Therefore, Algorithm 1 reads in compact form as:

$$\mathbf{x}^{k+1} = \text{proj}_{\Omega}(\mathbf{W}_k \mathbf{x}^k - \alpha \mathcal{R}^\top \mathbf{F}(\mathbf{W}_k \mathbf{x}^k)), \quad (7)$$

where $\Omega := \{\mathbf{x} \in \mathbb{R}^{Nn} | \mathcal{R} \mathbf{x} \in \Omega\}$ and $\mathbf{W}_k := W_k \otimes I_n$.

Lemma 1 [21, Lemma 3]: The mapping \mathbf{F} in (6) is ℓ -Lipschitz continuous, for some $\mu \leq \ell \leq \ell_0$.

Theorem 1: Let Assumptions 1-2 hold and let

$$M_\alpha = \begin{bmatrix} 1 - \frac{2\alpha\mu}{N} + \frac{\alpha^2\ell_0^2}{N} & \left(\frac{\alpha(\ell+\ell_0)+\alpha^2\ell_0\ell}{\sqrt{N}}\right)\bar{\sigma} \\ \left(\frac{\alpha(\ell+\ell_0)+\alpha^2\ell_0\ell}{\sqrt{N}}\right)\bar{\sigma} & (1 + 2\alpha\ell + \alpha^2\ell^2)\bar{\sigma}^2 \end{bmatrix}. \quad (8)$$

If the step size $\alpha > 0$ is chosen such that

$$\rho_\alpha := \lambda_{\max}(M_\alpha) = \|M_\alpha\| < 1, \quad (9)$$

then, for any initial condition, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 converges to $\mathbf{x}^* = \mathbf{1}_N \otimes x^*$, where x^* is the NE of the game in (1), with linear rate: for all $k \in \mathbb{N}$,

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq (\sqrt{\rho_\alpha})^k \|\mathbf{x}^0 - \mathbf{x}^*\|.$$

Algorithm 1 Fully Distributed NE Seeking

Initialization: for all $i \in \mathcal{I}$, set $x_i^0 \in \Omega_i$, $\mathbf{x}_{i,-i}^0 \in \mathbb{R}^{n-n_i}$.

Iterate until convergence: for all $i \in \mathcal{I}$,

Distributed averaging: $\hat{\mathbf{x}}_i^k = \sum_{j=1}^N w_{i,j}^k \mathbf{x}_j^k$

Local variables update: $x_i^{k+1} = \text{proj}_{\Omega_i}(\hat{\mathbf{x}}_{i,i}^k - \alpha \nabla_{x_i} J_i(\hat{\mathbf{x}}_i^k))$

$$\mathbf{x}_{i,-i}^{k+1} = \hat{\mathbf{x}}_{i,-i}^k.$$

Lemma 2: The condition in (9) holds if $\alpha > 0$ and

$$\alpha < \frac{\bar{\sigma}}{3\ell_0} \quad (10a)$$

$$\alpha < \frac{2\mu}{\ell_0^2} \quad (10b)$$

$$0 < 2\mu(1 - \bar{\sigma}^2) - \alpha(\bar{\sigma}^2(2\ell_0\ell + \ell^2 + 4\mu\ell + 2\ell_0^2) - \ell_0^2) - \alpha^2(\ell_0\ell^2 + \mu\ell^2 + 2\ell_0^2\ell)2\bar{\sigma}^2 - \alpha^3 2\ell_0^2\ell^2\bar{\sigma}^2. \quad (10c)$$

Proof: The condition in (10a) implies that $M_\alpha > 0$ (by diagonal dominance and positivity of the diagonal elements, as can be checked by recalling that $\ell \leq \ell_0$, $\mu \leq \ell_0$, $N \geq 2$, $\bar{\sigma} < 1$). The inequalities in (10b)-(10c) are the Sylvester's criterion for the matrix $I_2 - M_\alpha$: they impose that $[I_2 - M_\alpha]_{1,1} > 0$ (10b) and $\det(I_2 - M_\alpha) > 0$ (10c), hence $I_2 - M_\alpha > 0$. Altogether, this implies $\|M_\alpha\| < 1$. ■

Remark 2: The conditions in (10) always hold for α small enough, since, in the monomial inequality (10c), the constant term is $2\mu(1 - \bar{\sigma}^2) > 0$. While explicit solutions are known for cubic equations, we prefer the compact representation in (10c). The bounds in (10) are not tight, and in practice better bounds on the step size α are obtained by simply checking the Euclidean norm of the 2×2 matrix M_α in (8). Instead, the key observation is that the conditions in (10) do not depend on the number of agents: given the parameters $\bar{\sigma}$, μ , ℓ_0 and ℓ , a constant α that ensures convergence can be chosen independently of N . On the contrary, the rate $\sqrt{\rho_\alpha}$ does depend on N and, in fact, it approaches 1 as N grows unbounded (analogously to the results in [13], [14], [16]).

Remark 3: Compared to [16, Algorithm 7] (or [14, Algorithm 1]), in Algorithm 1 the agents *first* exchange information with their neighbors, and *then* evaluate their gradient term, resulting in better bounds on the step size α . Moreover, differently from [16, Th. 1], Theorem 1 provides a contractivity property for the iterates in (7) that holds at *each* step. This has beneficial consequences in terms of robustness, see Remark 6.

A. Technical Discussion

In Algorithm 1, the partial gradients $\nabla_{\mathbf{x}_i} J_i$ are evaluated on the local estimates $\mathbf{x}_{i,-i}$, not on the actual strategies \mathbf{x}_{-i} . Only if the estimates of all the agents coincide with the actual value, i.e., $\mathbf{x} = \mathbf{1}_N \otimes x$, we have that $\mathbf{F}(\mathbf{x}) = F(x)$. As a consequence, the mapping $\mathcal{R}^\top \mathbf{F}$ is not necessarily monotone, not even under strong monotonicity of the game mapping. Indeed, the loss of monotonicity is the main technical difficulty arising from the partial-decision information setup. Some works [12], [13], [14], [15], [21] deal with this issue by leveraging a restricted strong monotonicity property, which can be ensured, by opportunistically choosing the parameter γ , for the augmented mapping $\mathbf{F}_a(\mathbf{x}) := \gamma \mathcal{R}^\top \mathbf{F}(\mathbf{x}) + \mathbf{L} \mathbf{x}$, where $\mathbf{L} = L \otimes I_n$ and L is the

Laplacian of a fixed undirected connected network. Since the unique solution of the VI(F_a, Ω) is $\mathbf{x}^* = \mathbf{1}_N \otimes x^*$, with x^* the unique NE of the game in (1) [14, Prop. 1], one can design NE seeking algorithms via standard solution methods for variational inequalities (or the corresponding monotone inclusions, [15]). For instance, in [14], a forward-backward algorithm [19, 12.4.2] is proposed to solve VI(F_a, Ω), resulting in the algorithm

$$\mathbf{x}^{k+1} = \text{proj}_{\Omega}(\mathbf{x}^k - \tau(F_a(\mathbf{x}))). \quad (11)$$

We also recover this iteration when considering [15, Algorithm 1] in the absence of coupling constraints. However, exploiting the monotonicity of F_a results in conservative upper bounds on the parameters τ and γ , and hence in slow convergence (see Sections IV-V). More recently, the authors of [16] studied the convergence of (11) based on contractivity of the iterates, in the case of a fixed undirected network with doubly stochastic adjacency matrix W , unconstrained action sets (i.e., $\Omega = \mathbb{R}^n$), and by fixing $\tau = 1$, which results in the algorithm:

$$\mathbf{x}^{k+1} = (W \otimes I_N)\mathbf{x} - \alpha \mathcal{R}^\top F(\mathbf{x}^k). \quad (12)$$

Nonetheless, the upper bound on α provided in [16, Th. 1] is decreasing to zero when the number of agents N grows unbounded (in contrast with that in Theorem 1, see Lemma 2).

IV. BALANCED DIRECTED GRAPHS

In this section, we relax Assumption 1 to the following.

Assumption 3: For all $k \in \mathbb{N}$, the communication graph \mathcal{G}_k is weight balanced: $(\mathbf{1}_N^\top W_k)^\top = W_k \mathbf{1}_N$.

For weight-balanced digraphs, in-degree and out-degree of each node coincide. Therefore, the matrix $\tilde{L}_k := (L_k + L_k^\top)/2 = D_k - (W_k + W_k^\top)/2$ is itself the symmetric Laplacian of an undirected graph. Besides, such a graph is connected by Standing Assumption 3; hence \tilde{L}_k has a simple eigenvalue in 0, and the others are positive, i.e., $\lambda_2(\tilde{L}_k) > 0$.

Assumption 4: There exist $\bar{\sigma}, \bar{\lambda} > 0$ such that $\sigma_{\max}(L_k) \leq \bar{\sigma}$ and $\lambda_2(\tilde{L}_k) \geq \bar{\lambda}$, for all $k \in \mathbb{N}$.

Remark 4: Assumptions 2 and 4 always hold if the networks switch among a finite family. Yet, $\bar{\sigma}$, $\bar{\sigma}$ and $\bar{\lambda}$ are global parameters, that could be difficult to compute in a distributed way; upper/lower bounds might be available for special classes of networks, e.g., unweighted graphs.

To seek a NE over switching balanced digraphs, we propose the iteration in Algorithm 2. In compact form, it reads as

$$\mathbf{x}^{k+1} = \text{proj}_{\Omega}(\mathbf{x}^k - \tau(\gamma \mathcal{R}^\top F(\mathbf{x}^k) + L_k \mathbf{x}^k)) \quad (13)$$

where $L_k = L_k \otimes I_n$. Clearly, (13) is the same scheme of (11), just adapted to take the switching topology into account. In fact, the proof of convergence of Algorithm 2 is based on a restricted strong monotonicity property of the operator

$$F_a^k(\mathbf{x}) := \gamma \mathcal{R}^\top F(\mathbf{x}) + L_k \mathbf{x}, \quad (14)$$

that still holds for balanced directed graphs, as we show next.

Theorem 2: Let Assumptions 3-4 hold, and let

$$\begin{aligned} M &:= \gamma \left[-\frac{\frac{\mu}{N}}{2\sqrt{N}} \quad \frac{-\frac{\ell_0 + \ell}{2\sqrt{N}}}{\gamma - \theta} \right], \quad \gamma_{\max} := \frac{4\mu\bar{\lambda}}{(\ell_0 + \ell)^2 + 4\mu\bar{\theta}}, \\ \bar{\mu} &:= \lambda_{\min}(M), \quad \bar{\ell} := \ell + \bar{\sigma}, \quad \tau_{\max} := 2\bar{\mu}/\bar{\ell}^2, \end{aligned} \quad (15)$$

Algorithm 2 Fully Distributed NE Seeking

Initialization: for all $i \in \mathcal{I}$, set $x_i^0 \in \Omega_i$, $x_{i,-i}^0 \in \mathbb{R}^{n-n_i}$.

Iterate until convergence: for all $i \in \mathcal{I}$,

$$\begin{aligned} \hat{\mathbf{x}}_i^k &= \sum_{j=1}^N w_{i,j}^k (\mathbf{x}_i^k - \mathbf{x}_j^k) \\ x_i^{k+1} &= \text{proj}_{\Omega_i}(\mathbf{x}_i^k - \tau(\gamma \nabla_{x_i} J_i(\mathbf{x}_i^k) + \hat{\mathbf{x}}_{i,i}^k)) \\ \mathbf{x}_{i,-i}^{k+1} &= \hat{\mathbf{x}}_{i,-i}^k. \end{aligned}$$

$$\rho_{\gamma, \tau} := 1 - 2\tau\bar{\mu} + \tau^2\bar{\ell}^2.$$

If $\gamma \in (0, \gamma_{\max})$, then $M \succ 0$ and, for any $\tau \in (0, \tau_{\max})$, for any initial condition, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 2 converges to $\mathbf{x}^* = \mathbf{1}_N \otimes x^*$, where x^* is the unique NE of the game in (1), with linear rate: for all $k \in \mathbb{N}$,

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq (\sqrt{\rho_{\gamma, \tau}})^k \|\mathbf{x}^0 - \mathbf{x}^*\|.$$

Remark 5: Differently from the bound α_{\max} in (8), τ_{\max} in (15) vanishes as N grows (fixed the other parameters), as $\bar{\mu}$ decreases to 0 (by continuity of the eigenvalues).

Remark 6: Based on Theorems 1, 2, it can be proven that the discrete-time systems (7), (13) are input-to-state-stable (ISS) with respect to additive disturbances, with ISS-Lyapunov function $\|\mathbf{x} - \mathbf{x}^*\|^2$. By Lipschitz continuity of the updates, this implies ISS for noise both on the communication and in the evaluation of the partial gradients.

V. NUMERICAL EXAMPLE: A NASH-COURNOT GAME

We consider the Nash-Cournot game in [15, Sec. 6]. N firms produce a commodity that is sold to m markets. Each firm $i \in \mathcal{I} = \{1, \dots, N\}$ can only participate in $n_i \leq m$ of the markets; its action $x_i \in \mathbb{R}^{n_i}$ is the vector of quantities of product to be sent to these n_i markets, bounded by the local constraints $\mathbf{0}_{n_i} \leq x_i \leq X_i$. Let $A_i \in \mathbb{R}^{m \times n_i}$ be the matrix that specifies which markets firm i participates in. Specifically, $[A_i]_{k,j} = 1$ if $[x_i]_j$ is the amount of product sent to the k -th market by agent i , $[A_i]_{k,j} = 0$ otherwise, for all $k = 1, \dots, m$, $j = 1, \dots, n_i$. Let $A := [A_1 \dots A_N]$; then $Ax = \sum_{i=1}^N A_i x_i \in \mathbb{R}^m$ are the quantities of total product delivered to each market. Firm i aims at maximizing its profit, i.e., minimizing the cost function $J_i(x_i, x_{-i}) = c_i(x_i) - p(Ax)^\top A_i x_i$. Here, $c_i(x_i) = x_i^\top Q_i x_i + q_i^\top x_i$ is firm i 's production cost, with $Q_i \in \mathbb{R}^{n_i \times n_i}$, $Q_i \succ 0$, $q_i \in \mathbb{R}^{n_i}$. Instead, $p: \mathbb{R}^m \rightarrow \mathbb{R}^m$ associates to each market a price that depends on the amount of product delivered to that market. Specifically, the price for the market k , for $k = 1, \dots, m$, is $[p(Ax)]_k = \bar{P}_k - \chi_k [Ax]_k$, where \bar{P}_k , $\chi_k > 0$. We set $N = 20$, $m = 7$. The market structure is as in [15, Fig. 1], that defines which firms are allowed to participate in which markets. Therefore, $x \in \mathbb{R}^n$, with $n = 32$. We select randomly with uniform distribution r_k in $[1, 2]$, Q_i diagonal with diagonal elements in $[14, 16]$, q_i in $[1, 2]$, \bar{P}_k in $[10, 20]$, χ_k in $[1, 3]$, X_i in $[5, 10]$, for all $i \in \mathcal{I}$, $k = 1, \dots, m$. The resulting setup satisfies Standing Assumptions 1-2 [15, Sec. 6]. The firms cannot access the production of all the competitors, but can communicate with some neighbors on a network.

We first consider the case of a fixed, undirected graph, under Assumption 1. Algorithm 2 in this case reduces to [14, Algorithm 1]. We compare Algorithms 1-2 with the




Fig. 1. Distance from the NE for different pseudo-gradient NE seeking methods, with step sizes that guarantee convergence.




Fig. 2. Distance from the NE for Algorithm 1, with step size $\alpha = 2 \cdot 10^{-3}$ (upper bound in Theorem 1), and the method in [16, Algorithm 1], with step size $\alpha = 4 \cdot 10^{-6}$ (upper bound in [16, Th. 1]). Algorithm 1 converges much faster, thanks to the larger step size. The scheme in [16, Algorithm 1] still converges if we set $\alpha = 2 \cdot 10^{-3}$ (dashed line, not supported theoretically).




Fig. 3. Comparison of Algorithms 1 and 2, on a time-varying graph, for 20, 50 or 100 agents, with the step sizes set to their theoretical upper bounds.




Fig. 4. Distance from the NE for Algorithm 2, on a time-varying digraph. Since the networks are sparse, Theorem 2 ensures convergence only for small step sizes ($\gamma = 5 \cdot 10^{-4}$, $\tau = 3 \cdot 10^{-4}$), and convergence is slow (solid line). However, the bounds are conservative: the iteration still converges with τ 1000 times larger than the theoretical value (dashed line).

inexact ADMM in [13] and the accelerated gradient method in [14], for the step sizes that ensure convergence. Specifically, we set α as in Theorem 1 for Algorithm 1. The convergence of all the other Algorithms is based on the monotonicity of F_a in (14); hence we set γ as in Theorem 2. Instead of using the conservative bounds in (15) for the parameters, $\bar{\mu}$ and $\bar{\ell}$, we obtain a better result by computing the values numerically. F_a is (non-restricted) strongly monotone for our parameters, hence also the convergence result for [14, Algorithm 2] holds. Figure 1 shows that Algorithm 1 outperforms all the other

methods (we also note that the accelerated gradient in [14, Algorithm 2] requires two projections and two communications per iterations). As a numerical example, we also compare Algorithm 1 with the scheme in (12) by removing the local constraints, in Figure 2.

For the case of doubly stochastic time-varying networks, we randomly generate 5 connected graphs and for each iteration we pick one with uniform distribution. In Figure 3, we compare the performance of Algorithms 1-2, for step sizes set to their upper bounds as in Theorems 1-2. Since the theoretical convergence rate in Theorems 1-2 worsens as the number of agents grows, to show how the performance is affected in practice, we repeat the experiment for different values of N and random market structures (Figure 3).

Finally, in Figure 4, we test Algorithm 2 with communication topology chosen at each step with uniform distribution between two unweighted balanced directed graphs: the directed ring, where each agent i can send information to the agent $i+1$ (with the convention $N+1 \equiv 1$), and a graph where agent i is also allowed to transmit to agent $i+2$, for all $i \in \mathcal{I}$.

VI. CONCLUSION

Nash equilibrium problems on time-varying graphs can be solved with linear rate via fixed-step pseudo-gradient algorithms, if the network is connected at every iteration and the game mapping is Lipschitz continuous and strongly monotone. Our algorithm proved much faster than the existing gradient-based methods, when the step sizes satisfy their theoretical upper bounds. The extension to games with coupling constraints is left as future research. It would be also valuable to relax our uniform connectedness assumption, i.e., allowing for jointly strongly connected directed graphs.

APPENDIX

A. Proof of Theorem 1

We define the estimate consensus subspace $E := \{y \in \mathbb{R}^{Nn} | y = \mathbf{1}_N \otimes y, y \in \mathbb{R}^n\}$ and its orthogonal complement $E_\perp = \{y \in \mathbb{R}^{Nn} | (\mathbf{1}_N \otimes I_n)^\top y = \mathbf{0}_n\}$. Thus, any vector $x \in \mathbb{R}^{Nn}$ can be written as $x = x_\parallel + x_\perp$, where $x_\parallel = \text{proj}_E(x) = \frac{1}{N}(\mathbf{1}_N \mathbf{1}_N^\top \otimes I_n)x$, $x_\perp = \text{proj}_{E_\perp}(x)$, and $x_\parallel^\top x_\perp = 0$. Also, we use the shorthand notation Fx in place of $F(x)$ and $F(x)$. We recast the iteration in (7) as

$$x^{k+1} = \text{proj}_\Omega(\hat{x}^k - \alpha \mathcal{R}^\top F \hat{x}^k), \quad \hat{x}^k = W_k x^k. \quad (16)$$

Let x^* be the unique NE of the game in (1), and $x^* = \mathbf{1}_N \otimes x^*$. We recall that $x^* = \text{proj}_\Omega(x^* - \alpha F x^*)$ by (2), and then $x^* = \text{proj}_\Omega(x^* - \alpha \mathcal{R}^\top F x^*)$. Moreover, $W_k x^* = (W_k \otimes I_n)(\mathbf{1}_N \otimes x^*) = \mathbf{1}_N \otimes x^* = x^*$; hence x^* is a fixed point for (16). Let $x^k = x \in \mathbb{R}^{Nn}$ and $\hat{x} = W_k x = \hat{x}_\parallel + \hat{x}_\perp = \mathbf{1}_N \otimes \hat{x}_\parallel + \hat{x}_\perp \in \mathbb{R}^{Nn}$. Thus, it holds that

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &= \|\text{proj}_\Omega(\hat{x} - \alpha \mathcal{R}^\top F \hat{x}) - \text{proj}_\Omega(x^* - \alpha \mathcal{R}^\top F x^*)\|^2 \\ &\leq \|(\hat{x} - \alpha \mathcal{R}^\top F \hat{x}) - (x^* - \alpha \mathcal{R}^\top F x^*)\|^2 \\ &= \|\hat{x}_\parallel + \hat{x}_\perp - x^* + \alpha \mathcal{R}^\top (-F \hat{x} + F x^* + F \hat{x}_\parallel - F \hat{x}_\perp)\|^2 \\ &= \|\hat{x}_\parallel - x^*\|^2 + \|\hat{x}_\perp\|^2 \\ &\quad + \alpha^2 \|\mathcal{R}^\top (F \hat{x} - F \hat{x}_\parallel + F \hat{x}_\perp - F x^*)\|^2 \\ &\quad - 2\alpha (\hat{x}_\parallel - x^*)^\top \mathcal{R}^\top (F \hat{x} - F \hat{x}_\parallel) \end{aligned}$$

$$\begin{aligned}
& -2\alpha(\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*)^{\top} \mathcal{R}^{\top} (\mathbf{F}\hat{\mathbf{x}}_{\parallel} - \mathbf{F}\mathbf{x}^*) \\
& -2\alpha\hat{\mathbf{x}}_{\perp}^{\top} \mathcal{R}^{\top} (\mathbf{F}\hat{\mathbf{x}}_{\perp} - \mathbf{F}\mathbf{x}^*) \\
& -2\alpha\hat{\mathbf{x}}_{\perp}^{\top} \mathcal{R}^{\top} (\mathbf{F}\hat{\mathbf{x}}_{\parallel} - \mathbf{F}\mathbf{x}^*) \quad (17) \\
& \leq \|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|^2 + \|\hat{\mathbf{x}}_{\perp}\|^2 + \alpha^2(\ell^2\|\hat{\mathbf{x}}_{\perp}\|^2 + \frac{\ell_0^2}{N}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|^2 \\
& + \frac{2\ell_0\ell}{\sqrt{N}}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|\|\hat{\mathbf{x}}_{\perp}\|) + \frac{2\alpha\ell}{\sqrt{N}}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|\|\hat{\mathbf{x}}_{\perp}\| \\
& - \frac{2\alpha\mu}{N}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|^2 + 2\alpha\ell\|\hat{\mathbf{x}}_{\perp}\|^2 + \frac{2\alpha\ell_0}{\sqrt{N}}\|\hat{\mathbf{x}}_{\perp}\|\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|,
\end{aligned}$$

where the first inequality follows by nonexpansiveness of the projection ([22, Prop. 4.16]), and to bound the addends in (17) we used, in the order:

- 3rd term: $\|\mathcal{R}\| = 1$, Lipschitz continuity of \mathbf{F} , and $\|\mathbf{F}\hat{\mathbf{x}}_{\parallel} - \mathbf{F}\mathbf{x}^*\| = \|\mathbf{F}\hat{\mathbf{x}}_{\parallel} - \mathbf{F}\mathbf{x}^*\| \leq \ell_0\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\| = \frac{\ell_0}{\sqrt{N}}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|$;
- 4th term: $\|\mathcal{R}(\mathbf{1} \otimes (\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*))\| = \|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\| = \frac{1}{\sqrt{N}}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|$;
- 5th term: $(\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*)^{\top} \mathcal{R}^{\top} (\mathbf{F}\hat{\mathbf{x}}_{\parallel} - \mathbf{F}\mathbf{x}^*) = (\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*)^{\top} (\mathbf{F}\hat{\mathbf{x}}_{\parallel} - \mathbf{F}\mathbf{x}^*) \geq \mu\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|^2 = \frac{1}{N}\|\hat{\mathbf{x}}_{\parallel} - \mathbf{x}^*\|^2$;
- 6th term: Lipschitz continuity of \mathbf{F} ;
- 7th term: $\|\mathbf{F}\hat{\mathbf{x}}_{\perp} - \mathbf{F}\mathbf{x}^*\| \leq \frac{\ell_0}{\sqrt{N}}\|\hat{\mathbf{x}}_{\perp} - \mathbf{x}^*\|$ as above.

Besides, for every $\mathbf{x} = \mathbf{x}_{\parallel} + \mathbf{x}_{\perp} \in \mathbb{R}^{Nn}$ and for all $k \in \mathbb{N}$, it holds that $\hat{\mathbf{x}} = \mathbf{W}_k \mathbf{x} = \mathbf{x}_{\parallel} + \mathbf{W}_k \mathbf{x}_{\perp}$, where $\mathbf{W}_k \mathbf{x}_{\perp} \in \mathbf{E}_{\perp}$, by doubly stochasticity of \mathbf{W}_k , and $\|\hat{\mathbf{x}}_{\perp}\| = \|\mathbf{W}_k \mathbf{x}_{\perp}\| \leq \bar{\sigma}\|\mathbf{x}_{\perp}\|$ by (4) and properties of the Kronecker product. Therefore we can finally write, for all $k \in \mathbb{N}$, for all $\mathbf{x}^k \in \mathbb{R}^{Nn}$,

$$\begin{aligned}
\|\mathbf{x}^{k+1} - \mathbf{x}^*\|^2 & \leq \begin{bmatrix} \|\mathbf{x}_{\parallel}^k - \mathbf{x}^*\| \\ \|\mathbf{x}_{\perp}^k\| \end{bmatrix}^{\top} M_{\alpha} \begin{bmatrix} \|\mathbf{x}_{\parallel}^k - \mathbf{x}^*\| \\ \|\mathbf{x}_{\perp}^k\| \end{bmatrix} \\
& \leq \lambda_{\max}(M_{\alpha})(\|\mathbf{x}_{\parallel}^k - \mathbf{x}^*\|^2 + \|\mathbf{x}_{\perp}^k\|^2) \\
& = \lambda_{\max}(M_{\alpha})\|\mathbf{x}^k - \mathbf{x}^*\|^2.
\end{aligned}$$

B. Proof of Theorem 2

Let \mathbf{x}^* be the unique NE of the game in (1), and $\mathbf{x}^* = \mathbf{1}_N \otimes \mathbf{x}^*$. We recall that the null space $\text{null}(\mathbf{L}_k) = \mathbf{E} = \{\mathbf{y} \in \mathbb{R}^{Nn} | \mathbf{y} = \mathbf{1}_N \otimes \mathbf{y}, \mathbf{y} \in \mathbb{R}^n\}$ by Standing Assumption 3. Therefore, $\mathbf{L}_k \mathbf{x}^* = \mathbf{0}_N$ and \mathbf{x}^* is a fixed point of the iteration in (13) by (2). With \mathbf{F}_a^k as in (14), for all $k \in \mathbb{N}$, for any $\mathbf{x} \in \mathbb{R}^{Nn}$, it holds that $(\mathbf{x} - \mathbf{x}^*)^{\top} (\mathbf{F}_a^k \mathbf{x} - \mathbf{F}_a^k \mathbf{x}^*) = (\mathbf{x} - \mathbf{x}^*)^{\top} \gamma \mathcal{R}^{\top} (\mathbf{F}\mathbf{x} - \mathbf{F}\mathbf{x}^*) + (\mathbf{x} - \mathbf{x}^*)^{\top} \mathbf{L}_k (\mathbf{x} - \mathbf{x}^*) = (\mathbf{x} - \mathbf{x}^*)^{\top} \gamma \mathcal{R}^{\top} (\mathbf{F}\mathbf{x} - \mathbf{F}\mathbf{x}^*) + (\mathbf{x} - \mathbf{x}^*)^{\top} \tilde{\mathbf{L}}_k (\mathbf{x} - \mathbf{x}^*)$, where $\tilde{\mathbf{L}}_k = (\mathbf{L}_k + \mathbf{L}_k^{\top})/2 = (\mathbf{L}_k + \mathbf{L}_k^{\top}) \otimes \mathbf{I}_n/2 = \tilde{\mathbf{L}}_k \otimes \mathbf{I}_n$, and $\tilde{\mathbf{L}}_k$ is the Laplacian of a connected graph (see Section IV) and $\lambda_2(\tilde{\mathbf{L}}_k) > \bar{\lambda}$ by Assumption 4. Therefore we can apply [15, Lemma 3] to conclude that $(\mathbf{x} - \mathbf{x}^*)^{\top} (\mathbf{F}_a^k \mathbf{x} - \mathbf{F}_a^k \mathbf{x}^*) \geq \bar{\mu}\|\mathbf{x} - \mathbf{x}^*\|^2$, with $\bar{\mu} > 0$ as in (15). Also, \mathbf{F}_a^k is Lipschitz continuous with constant $\bar{\ell} = \ell + \bar{\sigma}$, $\bar{\sigma}$ as in Assumption 4. Therefore we have

$$\begin{aligned}
& \|\mathbf{x}^{k+1} - \mathbf{x}^*\|^2 \\
& = \|\text{proj}_{\Omega}(\mathbf{x}^k - \tau \mathbf{F}_a^k(\mathbf{x}^k)) - \text{proj}_{\Omega}(\mathbf{x}^* - \tau \mathbf{F}_a^k \mathbf{x}^*)\|^2 \\
& \leq \|(\mathbf{x}^k - \tau \mathbf{F}_a^k \mathbf{x}^k) - (\mathbf{x}^* - \tau \mathbf{F}_a^k \mathbf{x}^*)\|^2 \\
& = \|\mathbf{x}^k - \mathbf{x}^*\|^2 - 2\tau(\mathbf{x}^k - \mathbf{x}^*)^{\top} (\mathbf{F}_a^k \mathbf{x}^k - \mathbf{F}_a^k \mathbf{x}^*) \\
& \quad + \tau^2 \|\mathbf{F}_a^k \mathbf{x}^k - \mathbf{F}_a^k \mathbf{x}^*\|^2 \\
& \leq (1 - 2\tau\bar{\mu} + \tau^2(\bar{\ell} + \bar{\sigma})^2)\|\mathbf{x}^k - \mathbf{x}^*\|^2 = \rho_{\gamma, \tau}\|\mathbf{x}^k - \mathbf{x}^*\|^2,
\end{aligned}$$

where in the first inequality we used [22, Prop. 4.16], and $\rho_{\gamma, \tau} \in (0, 1)$ if τ is chosen as in Theorem 2.

REFERENCES

- [1] F. Facchinei and J. Pang, "Nash equilibria: The variational approach," in *Convex Optimization in Signal Processing and Communications*, D. P. Palomar and Y. C. Eldar, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2009, pp. 443–493.
- [2] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications," *IEEE Signal Process. Mag.*, vol. 29, no. 5, pp. 86–105, Sep. 2012.
- [3] S. Grammatico, "Dynamic control of agents playing aggregative games with coupling constraints," *IEEE Trans. Autom. Control*, vol. 62, no. 9, pp. 4537–4548, Sep. 2017.
- [4] N. Li, L. Chen, and M. A. Dahleh, "Demand response using linear supply function bidding," *IEEE Trans. Smart Grid*, vol. 6, no. 4, pp. 1827–1838, Jul. 2015.
- [5] C. Yu, M. van der Schaar, and A. H. Sayed, "Distributed learning for stochastic generalized Nash equilibrium problems," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 3893–3908, Aug. 2017.
- [6] G. Belgioioso and S. Grammatico, "Projected-gradient algorithms for generalized equilibrium seeking in aggregative games are preconditioned forward-backward methods," in *Proc. Eur. Control Conf.*, 2018, pp. 2188–2193.
- [7] J. S. Shamma and G. Arslan, "Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria," *IEEE Trans. Autom. Control*, vol. 50, no. 3, pp. 312–327, Mar. 2005.
- [8] J. Ghaderi and R. Srikant, "Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate," *Automatica*, vol. 50, no. 12, pp. 3209–3215, 2014.
- [9] K. Bimpikis, S. Ehsani, and R. Ilkiliç, "Cournot competition in networked markets," in *Proc. 15th ACM Conf. Econ. Comput. (EC)*, 2014, p. 733.
- [10] J. Koshal, A. Nedić, and U. V. Shanbhag, "Distributed algorithms for aggregative games on graphs," *Oper. Res.*, vol. 64, no. 3, pp. 680–704, 2016.
- [11] M. Ye and G. Hu, "Distributed Nash equilibrium seeking by a consensus based approach," *IEEE Trans. Autom. Control*, vol. 62, no. 9, pp. 4811–4818, Sep. 2017.
- [12] D. Gadjov and L. Pavel, "A passivity-based approach to Nash equilibrium seeking over networks," *IEEE Trans. Autom. Control*, vol. 64, no. 3, pp. 1077–1092, Mar. 2019.
- [13] F. Salehisadaghiani, W. Shi, and L. Pavel, "Distributed Nash equilibrium seeking under partial-decision information via the alternating direction method of multipliers," *Automatica*, vol. 103, pp. 27–35, May 2019.
- [14] T. Tatarenko, W. Shi, and A. Nedić. (2018). *Geometric Convergence of Gradient Play Algorithms for Distributed Nash Equilibrium Seeking*. [Online]. Available: <https://arxiv.org/abs/1809.07383>
- [15] L. Pavel, "Distributed GNE seeking under partial-decision information over networks via a doubly-augmented operator splitting approach," *IEEE Trans. Autom. Control*, vol. 65, no. 4, pp. 1584–1597, Apr. 2020.
- [16] T. Tatarenko and A. Nedić. (2019). *Geometric Convergence of Distributed Gradient Play in Games With Unconstrained Action Sets*. [Online]. Available: <https://arxiv.org/abs/1907.07144>
- [17] G. Belgioioso, A. Nedić, and S. Grammatico. (2019). *Distributed Generalized Nash Equilibrium Seeking in Aggregative Games on Time-Varying Networks*. [Online]. Available: <https://arxiv.org/abs/1907.00191>
- [18] F. Salehisadaghiani and L. Pavel, "Nash equilibrium seeking with non-doubly stochastic communication weight matrix," *EAI Endorsed Trans. Collab. Comput.*, vol. 4, no. 13, pp. 1–11, 2019.
- [19] F. Facchinei and J. Pang, *Finite-Dimensional Variational Inequalities and Complementarity Problems*. New York, NY, USA: Springer, 2007.
- [20] B. Gharesifard and J. Cortés, "Distributed strategies for generating weight-balanced and doubly stochastic digraphs," *Eur. J. Control*, vol. 18, no. 6, pp. 539–557, 2012.
- [21] M. Bianchi and S. Grammatico. (2020). *A Continuous-Time Distributed Generalized Nash Equilibrium Seeking Algorithm Over Networks for Double-Integrator Agents*. [Online]. Available: <https://arxiv.org/abs/1910.11608>
- [22] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, vol. 2011. Cham, Switzerland: Springer, 2017.