

A fully-distributed proximal-point algorithm for Nash equilibrium seeking with linear convergence rate

Bianchi, Mattia; Belgioioso, Giuseppe; Grammatico, Sergio

DOI

[10.1109/CDC42340.2020.9304145](https://doi.org/10.1109/CDC42340.2020.9304145)

Publication date

2020

Document Version

Final published version

Published in

Proceedings of the 59th IEEE Conference on Decision and Control, CDC 2020

Citation (APA)

Bianchi, M., Belgioioso, G., & Grammatico, S. (2020). A fully-distributed proximal-point algorithm for Nash equilibrium seeking with linear convergence rate. In *Proceedings of the 59th IEEE Conference on Decision and Control, CDC 2020* (pp. 2303-2308). IEEE. <https://doi.org/10.1109/CDC42340.2020.9304145>

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.

A fully-distributed proximal-point algorithm for Nash equilibrium seeking with linear convergence rate

Mattia Bianchi, Giuseppe Belgioioso, and Sergio Grammatico

Abstract—We address the Nash equilibrium problem in a partial-decision information scenario, where each agent can only observe the actions of some neighbors, while its cost possibly depends on the strategies of other agents. Our main contribution is the design of a fully-distributed, single-layer, fixed-step algorithm, based on a proximal best-response augmented with consensus terms. To derive our algorithm, we follow an operator-theoretic approach. First, we recast the Nash equilibrium problem as that of finding a zero of a monotone operator. Then, we demonstrate that the resulting inclusion can be solved in a fully-distributed way via a proximal-point method, thanks to the use of a novel preconditioning matrix. Under strong monotonicity and Lipschitz continuity of the game mapping, we prove linear convergence of our algorithm to a Nash equilibrium. Furthermore, we show that our method outperforms the fastest known gradient-based schemes, both in terms of guaranteed convergence rate, via theoretical analysis, and in practice, via numerical simulations.

I. INTRODUCTION

Nash equilibrium (NE) problems have received increasing attention with the spreading of networked systems, due to the numerous engineering applications, including communication networks [1], formation control [2], charge scheduling of electric vehicles [3] and demand response in competitive markets [4]. These scenarios are characterized by the presence of multiple selfish decision-makers, or agents, that aim at optimizing their individual, yet inter-dependent, objective functions. From a game-theoretic perspective, one of the challenges is to assign to the agents behavioral rules that eventually ensure the attainment of a NE, a joint action from which no agent has interest 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, for example in the presence of a coordinator that broadcasts the data to the network [5]–[8]. However the existence of a central node with bidirectional communication with all the agents is impractical for many applications [9], [10]. One example is the Nash-Cournot competition model described in [11], 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. This motivates the development of fully-distributed algorithms, which allow to compute NEs relying on local data only.

M. Bianchi and S. Grammatico are with the Delft Center for Systems and Control (DCSC), TU Delft, The Netherlands. G. Belgioioso is with the Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH) Zürich, Switzerland. E-mail addresses: {m.bianchi, s.grammatico}@tudelft.nl, gbelgioioso@ethz.ch. This work was partially supported by NWO (project OMEGA, 613.001.702) and by the ERC (project COSMOS, 802348).

Two main approaches have been proposed, corresponding to two different information structures. For games where each agent can measure its own cost functions, pay-off based schemes were developed that do not require peer-to-peer communication [2], [12]. Instead, we consider the so-called *partial-decision information* scenario, where the agents hold an analytic expression of their own cost functions, but they are unable to evaluate the actual values, since they cannot access the strategies of all the competitors. To remedy the lack of knowledge, the agents engage in nonstrategic information exchange with some neighbors on a network; from the data received, they can estimate the strategies of all other agents, and eventually reconstruct the true values. This setup has only been introduced very recently. Most of the results available resort to (projected) gradient and consensus dynamics, both in continuous time [13], [14], and discrete time. For the discrete time case, early works [11], [15], focused on algorithms with vanishing step sizes, which typically result in slow convergence. More recently, fixed-step schemes were introduced in [16]–[18], building on a restricted monotonicity property, first revealed in [14]. The drawback is that, due to the partial-decision information assumption, small step sizes have to be chosen, affecting the speed of convergence. Of particular interest for this paper is the technique developed by Pavel in [18], that characterized the equilibria of a (generalized) game as the zeros of a monotone operator. The operator-theoretic approach is very elegant and convenient, since several splittings methods are already well established to solve monotone inclusions [19, §26]. For example, the authors of [20] adopted a preconditioned proximal-point algorithm (PPPA); yet, this results in a double-layer scheme, where the agents have to communicate multiple times to solve (inexactly) a subgame, at each step. Similarly, the proximal best-response dynamics proposed in [21] for stochastic games require an increasing number of data transmissions per iteration.

Contribution: Motivated by the above, in this paper we further exploit the restricted monotonicity property used in [16]–[18] to solve Nash equilibrium problems under partial-decision information. Specifically:

- We derive a simple, fully-distributed proximal-point algorithm (PPA), that is a proximal best-response augmented with consensual terms. Thanks to the use of a novel preconditioning matrix, our algorithm is single-layer, i.e., it requires only one communication per iteration. To the best of our knowledge, our PPPA is the first non-gradient based algorithm with this feature (§III-§IV);
- We prove global linear convergence of our algorithm to a

NE, under strong monotonicity and Lipschitz continuity of the game mapping, by providing a general result for the PPA of restricted strongly monotone operators (§IV);

- We demonstrate, both by comparing the theoretical convergence rates and in simulation, that our algorithm outperforms existing gradient-based dynamics, in terms of the number of iterations needed to converge (§V-§VI).

Basic notation: See [22].

Operator-theoretic notation: For a function $\psi : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$, $\text{dom}(\psi) = \{x \in \mathbb{R}^n \mid \psi(x) < \infty\}$. The mapping $\iota_S : \mathbb{R}^n \rightarrow \{0, \infty\}$ denotes the indicator function for the set $S \subseteq \mathbb{R}^n$, i.e., $\iota_S(x) = 0$ if $x \in S$, ∞ otherwise. A set-valued mapping (or operator) $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is characterized by its graph $\text{gra}(\mathcal{F}) = \{(x, u) \mid u \in \mathcal{F}(x)\}$. $\text{dom}(\mathcal{F}) = \{x \in \mathbb{R}^n \mid \mathcal{F}(x) \neq \emptyset\}$, $\text{fix}(\mathcal{F}) = \{x \in \mathbb{R}^n \mid x \in \mathcal{F}(x)\}$ and $\text{zer}(\mathcal{F}) = \{x \in \mathbb{R}^n \mid 0 \in \mathcal{F}(x)\}$ denote the domain, set of fixed points and set of zeros, respectively. \mathcal{F}^{-1} denotes the inverse operator of \mathcal{F} , defined through its graph as $\text{gra}(\mathcal{F}^{-1}) = \{(u, x) \mid (x, u) \in \text{gra}(\mathcal{F})\}$. $\mathcal{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is (μ -strongly) monotone if $\langle u-v, x-y \rangle \geq 0$ ($\geq \mu \|x-y\|^2$), for all $(x, u), (y, v) \in \text{gra}(\mathcal{F})$, where $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. Id denotes the identity operator. For a function $\psi : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$, $\partial\psi : \text{dom}(\psi) \rightrightarrows \mathbb{R}^n$ denotes its subdifferential operator, defined as $\partial\psi(x) = \{v \in \mathbb{R}^n \mid \psi(z) \geq \psi(x) + \langle v, z-x \rangle \text{ for all } z \in \text{dom}(\psi)\}$. $N_S : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is the normal cone operator for the set $S \subseteq \mathbb{R}^n$, i.e., $N_S(x) = \emptyset$ if $x \notin S$, $\{v \in \mathbb{R}^n \mid \sup_{z \in S} \langle v, z-x \rangle \leq 0\}$ otherwise. If S is closed and convex, it holds that $\partial\iota_S = N_S$. $J_{\mathcal{F}} := (\text{Id} + \mathcal{F})^{-1}$ denotes the resolvent operator of \mathcal{F} .

II. MATHEMATICAL BACKGROUND

We consider a set of agents $\mathcal{I} := \{1, \dots, N\}$. Each agent $i \in \mathcal{I}$ shall choose its decision variable (i.e., strategy) 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 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 on the decision variables of the other agents $x_{-i} := \text{col}((x_j)_{j \in \mathcal{I} \setminus \{i\}})$. Then, the game is represented by the inter-dependent optimization problems:

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

Our goal here is to compute 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

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

The following assumptions are standard for NE problems, see, e.g., [17, Ass. 1], [14, Ass. 2].

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 continuously differentiable and convex for every x_{-i} . \square

Under Standing Assumption 1, a strategy x^* is a NE of the game in (1) if and only if it is a solution of the

variational inequality $\text{VI}(F, \Omega)$ ¹ [23, Prop. 1.4.2], where F is the *pseudo-gradient* mapping of the game:

$$F(x) := \text{col}((\nabla_{x_i} J_i(x_i, x_{-i}))_{i \in \mathcal{I}}). \quad (2)$$

Equivalently, x^* is a NE if and only if the following holds:

$$\mathbf{0}_n \in F(x^*) + N_{\Omega}(x^*), \quad (3)$$

A sufficient condition for the existence of a unique NE for the game in (1) is the strong monotonicity of the pseudo-gradient [23, Th. 2.3.3], as postulated next. This assumption is always used for NE seeking under partial-decision information with fixed step sizes, e.g., [18, Ass. 2], [13, Ass. 4], [17, Ass. 2].

Standing Assumption 2: The pseudo-gradient mapping F in (2) is μ -strongly monotone and θ_0 -Lipschitz continuous, for some $\mu, \theta_0 > 0$: for any $x, y \in \mathbb{R}^n$, $\langle x-y, F(x) - F(y) \rangle \geq \mu \|x-y\|^2$ and $\|F(x) - F(y)\| \leq \theta_0 \|x-y\|$. \square

III. DISTRIBUTED NASH EQUILIBRIUM SEEKING

In this section, we present an algorithm to seek a NE of the game in (1) in a fully-distributed way. Specifically, each agent i only knows 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 some neighbors over a communication network $\mathcal{G}(\mathcal{I}, \mathcal{E})$. The unordered pair (i, j) belongs to the set of edges \mathcal{E} if and only if agent i and j can mutually exchange information. We denote: $W \in \mathbb{R}^{N \times N}$ the mixing matrix of \mathcal{G} , with $w_{i,j} := [W]_{i,j}$ and $w_{i,j} > 0$ if $(i, j) \in \mathcal{E}$, $w_{i,j} = 0$ otherwise; $\mathcal{N}_i = \{j \mid (i, j) \in \mathcal{E}\}$ the set of neighbors of agent i .

Standing Assumption 3: The communication graph $\mathcal{G}(\mathcal{I}, \mathcal{E})$ is undirected and connected. \square

For ease of notation, we also assume that every node of the graph has a self-loop and that W is doubly stochastic; this condition is not strictly necessary and can be dropped (see §IV, Remark 4). Nonetheless, we note that such a mixing matrix can be distributedly generated on any undirected graph, e.g., by assigning Metropolis weights [22, §2].

Standing Assumption 4: The mixing matrix W satisfies the following conditions:

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

To cope with the lack of knowledge, the usual assumption for the partial-decision information scenario is that each agent keeps an estimate of all other agents' action [17], [16] [13]. We denote $\mathbf{x}_i = \text{col}((\mathbf{x}_{i,j})_{j \in \mathcal{I}}) \in \mathbb{R}^n$, where $\mathbf{x}_{i,i} := x_i$ and $\mathbf{x}_{i,j}$ is i 's estimate of agent j 's action, for all $j \neq i$; let also $\mathbf{x}_{j,-i} := \text{col}((\mathbf{x}_{j,\ell})_{\ell \in \mathcal{I} \setminus \{i\}})$. Our proposed fully-distributed NE seeking dynamics are summarized in Algorithm 1, where $\alpha > 0$ is a global constant parameter.

In steady state, agents should agree on their estimates, i.e., $\mathbf{x}_i = \mathbf{x}_j$ for all $i, j \in \mathcal{I}$. In fact, the updates of the estimates $\mathbf{x}_{i,-i}^k$ resembles a consensus protocol, and can be interpreted as the attempt of the agents to reach an agreement on the

¹Given a set $S \subseteq \mathbb{R}^m$ and a mapping $\psi : S \rightarrow \mathbb{R}^m$, the variational inequality $\text{VI}(\psi, S)$ is the problem of finding a vector $\omega^* \in S$ such that $\psi(\omega^*)^T(\omega - \omega^*) \geq 0$, for all $\omega \in S$.

Algorithm 1 Fully-distributed Nash equilibrium seeking via preconditioned proximal-point iteration

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

For all $k \geq 0$: Communication: The agents exchange the variables $\{x_i^k, \mathbf{x}_{i,-i}^k\}$ with their neighbors.
 Local variables update: each agent $i \in \mathcal{I}$ does:

$$\begin{aligned} \mathbf{x}_{i,-i}^{k+1} &= \frac{1}{2}(\mathbf{x}_{i,-i}^k + \sum_{j=1}^N w_{i,j} \mathbf{x}_{j,-i}^k) \\ x_i^{k+1} &= \operatorname{argmin}_{y \in \Omega_i} (J_i(y, \mathbf{x}_{i,-i}^{k+1}) + \frac{1}{2\alpha} \|y - x_i^k\|^2 + \frac{1}{2\alpha} \|y - \sum_{j=1}^N w_{i,j} x_{j,i}^k\|^2) \end{aligned}$$

time-varying quantity x . In turn, the strategy x_i of each agent is updated based on a proximal best-response, augmented with an extra disagreement penalization term. We remark that the agents evaluate their cost functions in their local estimates, not on the actual collective strategy.

Remark 1: The functions $J_i(\cdot, \mathbf{x}_{i,-i})$ are strongly convex, for all $\mathbf{x}_{i,-i}$, for all $i \in \mathcal{I}$, as a consequence of Standing Assumption 2. Therefore the argmin operator in Algorithm 1 is single-valued and the algorithm is always well-defined. \square

IV. CONVERGENCE ANALYSIS

In this section, we first derive Algorithm 1 as a PPPA. Then, we prove its convergence by leveraging a *restricted* monotonicity property, under which classical results for the PPA of monotone operators still hold.

Before proceeding, we need some definitions. We denote $\mathbf{x} = \operatorname{col}((x_i)_{i \in \mathcal{I}})$. Besides, let, for all $i \in \mathcal{I}$ [18, Eq.13-14],

$$\mathcal{R}_i := \begin{bmatrix} 0_{n_i \times n_{<i}} & I_{n_i} & \mathbf{0}_{n_i \times n_{>i}} \end{bmatrix}, \quad (4a)$$

$$\mathcal{S}_i := \begin{bmatrix} I_{n_{<i}} & \mathbf{0}_{n_{<i} \times n_i} & \mathbf{0}_{n_{<i} \times n_{>i}} \\ \mathbf{0}_{n_{>i} \times n_{<i}} & \mathbf{0}_{n_{>i} \times n_i} & I_{n_{>i}} \end{bmatrix}, \quad (4b)$$

where $n_{<i} := \sum_{j=1}^{i-1} n_j$, $n_{>i} := \sum_{j=i+1}^N n_j$. In simple terms, \mathcal{R}_i selects the i -th n_i dimensional component from an n -dimensional vector, while \mathcal{S}_i removes it. Thus, $\mathcal{R}_i \mathbf{x}_i = \mathbf{x}_{i,i} = x_i$ and $\mathcal{S}_i \mathbf{x}_i = \mathbf{x}_{i,-i}$. Let $\mathcal{R} := \operatorname{diag}((\mathcal{R}_i)_{i \in \mathcal{I}})$, $\mathcal{S} := \operatorname{diag}((\mathcal{S}_i)_{i \in \mathcal{I}})$. Hence, $\mathcal{R} \mathbf{x} = \mathbf{x}$ and $\mathcal{S} \mathbf{x} = \operatorname{col}((\mathbf{x}_{i,-i})_{i \in \mathcal{I}}) \in \mathbb{R}^{(N-1)n}$. Moreover, $\mathbf{x} = \mathcal{R}^\top \mathbf{x} + \mathcal{S}^\top \mathcal{S} \mathbf{x}$. We define the *extended pseudo-gradient* operator \mathbf{F} as

$$\mathbf{F}(\mathbf{x}) := \operatorname{col}((\nabla_{x_i} J_i(x_i, \mathbf{x}_{i,-i}))_{i \in \mathcal{I}}), \quad (5)$$

and the mappings

$$\mathbf{F}_a(\mathbf{x}) := \alpha \mathcal{R}^\top \mathbf{F}(\mathbf{x}) + (I_{Nn} - \mathbf{W}) \mathbf{x}, \quad (6)$$

$$\mathcal{A}(\mathbf{x}) := \mathbf{F}_a(\mathbf{x}) + \mathbf{N}_\Omega(\mathbf{x}), \quad (7)$$

where $\alpha > 0$ is a fixed parameter, $\mathbf{W} := W \otimes I_n$ and

$$\Omega := \{\mathbf{x} \in \mathbb{R}^{nN} \mid \mathcal{R} \mathbf{x} \in \Omega\}. \quad (8)$$

The following lemma relates the NE of the game in (1) to the operators \mathbf{F}_a and \mathcal{A} . The proof is analogous to [17, Prop. 1], and hence it is omitted.

Lemma 1: The following statements are equivalent:

- i) $\mathbf{x}^* = \mathbf{1}_N \otimes x^*$, with $x^* \in \Omega$ the NE of the game in (1);
- ii) \mathbf{x}^* solves $\operatorname{VI}(\mathbf{F}_a, \Omega)$;
- iii) $\mathbf{0}_{Nn} \in \mathcal{A}(\mathbf{x}^*)$. \square

A. Derivation of the algorithm

Lemma 1 is fundamental, because it provides a systematic way of deriving fully-distributed NE seeking algorithms, by applying standard solution methods for $\operatorname{VI}(\mathbf{F}_a, \Omega)$ (e.g., in [17], a projected gradient-method was developed) or operator splitting methods to compute a zero of the operator $\mathcal{A}(\mathbf{x})$ (e.g., [18] follows a similar approach for games with coupling constraints). Nonetheless, technical difficulties arise because of the partial-decision information assumption. Specifically, the operator $\mathcal{R}^\top \mathbf{F}$ is not monotone for most cases of interest, not even if strong monotonicity of the pseudo-gradient mapping F holds, i.e., Standing Assumption 2. Only when the estimates \mathbf{x} belong to the consensus subspace, i.e. $\mathbf{x} = \mathbf{1}_N \otimes x$, we have that $\mathbf{F}(\mathbf{x}) = F(x)$.

In fact, Algorithm 1 is an instance of (suitably preconditioned) PPA [19, Th. 23.41] to seek a zero of \mathcal{A} . We remark that many operator-theoretic properties are not guaranteed for the resolvent of a non-monotone operator $\mathcal{B} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$. For example, $\mathbf{J}_\mathcal{B} = (\operatorname{Id} + \mathcal{B})^{-1}$ may have a limited domain, or be set-valued. In this general case, we write the PPA as

$$\omega^{k+1} \in \mathbf{J}_\mathcal{B}(\omega^k), \quad (9)$$

that is well defined only if $\mathbf{J}_\mathcal{B}(\omega^k) \neq \emptyset$. Next, we show that Algorithm 1 is obtained by applying the iteration in (9) to the operator $\Phi^{-1} \mathcal{A}$, where

$$\Phi := I_{Nn} + \mathbf{W}, \quad (10)$$

is a symmetric nonsingular matrix, known as *preconditioning* matrix. We note that $\Phi \succ 0$, by Standing Assumption 4 and Gershgorin circle theorem, and that $\operatorname{zer}(\mathcal{A}) = \operatorname{zer}(\Phi^{-1} \mathcal{A})$.

Lemma 2: Algorithm 1 is equivalent to

$$\mathbf{x}^{k+1} \in \mathbf{J}_{\Phi^{-1} \mathcal{A}}(\mathbf{x}^k), \quad (11)$$

with \mathcal{A} as in (7), Φ as in (10). \square

Proof: By definition of inverse operator we have that

$$\begin{aligned} \mathbf{x}^{k+1} &\in (\operatorname{Id} + \Phi^{-1} \mathcal{A})^{-1} \mathbf{x}^k \\ &\Leftrightarrow \mathbf{0}_{Nn} \in \mathbf{x}^{k+1} + \Phi^{-1} \mathcal{A} \mathbf{x}^{k+1} - \mathbf{x}^k \\ &\Leftrightarrow \mathbf{0}_{Nn} \in \Phi(\mathbf{x}^{k+1} - \mathbf{x}^k) + \mathcal{A} \mathbf{x}^{k+1} \\ &\Leftrightarrow \mathbf{0}_{Nn} \in \mathbf{x}^{k+1} + \mathbf{W} \mathbf{x}^{k+1} - \mathbf{x}^k - \mathbf{W} \mathbf{x}^k + \mathbf{x}^{k+1} \\ &\quad - \mathbf{W} \mathbf{x}^{k+1} + \alpha \mathcal{R}^\top \mathbf{F}(\mathbf{x}^{k+1}) + \mathbf{N}_\Omega(\mathbf{x}^{k+1}). \quad (12) \end{aligned}$$

In turn, the last inclusion can be split in two components by left-multiplying both sides with \mathcal{R} and \mathcal{S} . Hence, by $\mathcal{S} \mathbf{N}_\Omega =$

$\mathbf{0}_{(N-1)n}$ and $\mathcal{S}\mathcal{R}^\top = \mathbf{0}_{(N-1)n \times n}$, (12) is equivalent to

$$\begin{cases} \mathbf{0}_{(N-1)n} \in \mathcal{S}(2\mathbf{x}^{k+1} - \mathbf{x}^k) - \mathcal{S}\mathbf{W}\mathbf{x}^k \\ \mathbf{0}_n \in 2\mathbf{x}^{k+1} - \mathbf{x}^k - \mathcal{R}\mathbf{W}\mathbf{x}^k + \mathbf{N}_\Omega(\mathbf{x}^{k+1}) \\ \quad \quad \quad + \alpha \mathbf{F}(\mathbf{x}^{k+1}, \mathcal{S}\mathbf{x}^{k+1}) \\ \mathbf{x}_{i,-i}^{k+1} = \frac{1}{2}(\mathbf{x}_{i,-i}^k + \sum_{j=1}^N w_{i,j} \mathbf{x}_{j,-i}^k) \\ \mathbf{0}_{n_i} \in \partial_{\mathbf{x}_i^{k+1}}(J_i(\mathbf{x}_i^{k+1}, \mathbf{x}_{i,-i}^{k+1}) + \frac{1}{2\alpha} \|\mathbf{x}_i^{k+1} - \mathbf{x}_i^k\|^2 \\ \quad + \frac{1}{2\alpha} \|\mathbf{x}_i^{k+1} - \sum_{j=1}^N w_{i,j} \mathbf{x}_{j,i}^k\|^2 \\ \quad + \iota_{\Omega_i}(\mathbf{x}_i^{k+1})). \end{cases} \stackrel{\Leftrightarrow}{\forall i \in \mathcal{I}}$$

The conclusion follows since the zeros of the subdifferential of a (strongly) convex function coincide with the minima (unique minimum) [19, Th. 16.3]. ■

Remark 2: The preconditioning matrix Φ is designed to decouple the system of inclusion in (12) from the graph structure, i.e., to remove the term $\mathbf{W}\mathbf{x}^{k+1}$. This ensures that the resulting updates can be computed by the agents in a fully-distributed fashion. □

Remark 3: By Lemma 2 and the explicit expression of $J_{\Phi^{-1}\mathcal{A}}$ in Algorithm 1, we conclude that $\text{dom}(J_{\Phi^{-1}\mathcal{A}}) = \mathbb{R}^{Nn}$ and that $J_{\Phi^{-1}\mathcal{A}}$ is single-valued on its domain. □

B. Convergence analysis

Since the operator $\Phi^{-1}\mathcal{A}$ is not maximally monotone in general, the convergence of Algorithm 1 cannot be inferred by standard results for the PPA. In [17, Th. 5], the authors proposed an accelerated gradient NE seeking scheme, which achieves geometric convergence if the mapping F_a is strongly monotone. However, this is a limiting assumption, that can be guaranteed only for some classes of games (cf. [17, Rem. 3]). Instead, our analysis is based on a weaker condition, namely the *restricted* strong monotonicity of F_a only with respect to the NE [18], [16]. The main advantage is that, for suitable choices of the parameter α , restricted strong monotonicity (unlike strong monotonicity) of F_a holds for any game satisfying Standing Assumptions 1-4, as formalized in the next two statements.

Lemma 3 ([22, Lem. 1]): The mapping F in (5) is θ -Lipschitz continuous, for some $\theta \in [\mu, \theta_0]$: for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{Nn}$, $\|F(\mathbf{x}) - F(\mathbf{y})\| \leq \theta \|\mathbf{x} - \mathbf{y}\|$. □

Lemma 4 ([18, Lem. 3]): Let

$$M := \alpha \begin{bmatrix} \frac{\mu}{N} & -\frac{\theta_0 + \theta}{2\sqrt{N}} \\ -\frac{\theta_0 + \theta}{2\sqrt{N}} & \frac{\lambda_2(I - W)}{\alpha} - \theta \end{bmatrix}, \quad \begin{aligned} \alpha_{\max} &:= \frac{4\mu\lambda_2(I - W)}{(\theta_0 + \theta)^2 + 4\mu\theta}, \\ \rho_\alpha &:= \lambda_{\min}(M). \end{aligned} \quad (13)$$

For any $\alpha \in (0, \alpha_{\max})$, F_a is ρ_α -restricted strongly monotone with respect to the consensus subspace $\mathbf{E}_n := \{\mathbf{y} \in \mathbb{R}^{Nn} : \mathbf{y} = \mathbf{1}_N \otimes y, y \in \mathbb{R}^n\}$: for any $\mathbf{x} \in \mathbb{R}^{Nn}$ and any $\mathbf{y} \in \mathbf{E}_n$, it holds that $M \succ 0$ and also that

$$\langle \mathbf{x} - \mathbf{y}, F_a(\mathbf{x}) - F_a(\mathbf{y}) \rangle \geq \rho_\alpha \|\mathbf{x} - \mathbf{y}\|^2. \quad \square$$

Moreover, the operator $\Phi^{-1}\mathcal{A}$ retains this property, in the Hilbert space induced by the inner product $\langle \cdot, \cdot \rangle_\Phi$. Here, we denote, for a matrix $\Psi \succ 0$, by $\langle \cdot, \cdot \rangle_\Psi : (x, y) \mapsto \langle \Psi x, y \rangle$ and $\|\cdot\|_\Psi : x \mapsto \langle \Psi x, x \rangle$ the Ψ -weighted Euclidean inner product and norm, respectively.

Lemma 5: Let $\alpha_{\max}, \rho_\alpha$ be as in (13), $\alpha \in (0, \alpha_{\max})$. Then, $\Phi^{-1}\mathcal{A}$ is $\frac{\rho_\alpha}{\|\Phi\|}$ -restricted strongly monotone with respect to $\text{zer}(\mathcal{A})$ in the Φ -induced space: for all $(\mathbf{x}, \mathbf{u}), (\mathbf{y}, \mathbf{v}) \in \text{gra}(\Phi^{-1}\mathcal{A})$ such that $\mathbf{y} \in \text{zer}(\mathcal{A})$, it holds that

$$\langle \mathbf{u} - \mathbf{v}, \mathbf{x} - \mathbf{y} \rangle_\Phi \geq \frac{\rho_\alpha}{\|\Phi\|} \|\mathbf{x} - \mathbf{y}\|_\Phi^2. \quad \square$$

Towards our main result, we next prove the convergence of the iteration in (9) to an equilibrium, under restricted strong monotonicity of \mathcal{B} . The proof is based on the restricted contractivity of $J_{\mathcal{B}}$ with respect to its (unique) fixed point.

Theorem 1: Let $\mathcal{B} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$ be $\tilde{\mu}$ -restricted strongly monotone with respect to $\text{zer}(\mathcal{B})$, for some $\tilde{\mu} > 0$, in the space induced by some inner product $\langle \cdot, \cdot \rangle_\Psi$, i.e.,

$$\langle x - z, u - v \rangle_\Psi \geq \tilde{\mu} \|x - z\|_\Psi^2, \quad (14)$$

for any $(x, u), (z, v) \in \text{gra}(\mathcal{B})$ such that $z \in \text{zer}(\mathcal{B})$. Assume that $\text{zer}(\mathcal{B}) \neq \emptyset$ and that $\text{dom}(J_{\mathcal{B}}) = \mathbb{R}^m$. Then, for any $\omega^0 \in \mathbb{R}^m$, the proximal-point iteration in (9) converges with linear rate to the unique point $\{z\} = \text{zer}(\mathcal{B})$: for all $k \in \mathbb{N}$,

$$\|\omega^k - z\|_\Psi \leq \left(\frac{1}{1 + \tilde{\mu}}\right)^k \|\omega^0 - z\|_\Psi. \quad \square$$

We are now ready to show the main result of the paper, namely the linear convergence of Algorithm 1 to a NE.

Theorem 2: Let $\alpha_{\max}, \rho_\alpha$ be as in (13), Φ as in (10), $\alpha \in (0, \alpha_{\max})$. For any initial condition, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 converges with linear rate to $\mathbf{x}^* = \mathbf{1}_n \otimes x^*$, where x^* is the unique NE of the game in (1): for all $k \in \mathbb{N}$,

$$\|\mathbf{x}^k - \mathbf{x}^*\| \leq \sqrt{\frac{\|\Phi\|}{\lambda_{\min}(\Phi)}} \left(\frac{1}{1 + \rho_\alpha / \|\Phi\|}\right)^k \|\mathbf{x}^0 - \mathbf{x}^*\|. \quad \square$$

Remark 4: If the mixing matrix W is not doubly stochastic (i.e., if the last condition in Standing Assumption 4 does not hold), an iteration analogous to Algorithm 1 can be derived by defining $\mathcal{A}(\mathbf{x}) = \alpha \mathcal{R}^\top F(\mathbf{x}) + (\mathbf{D} - \mathbf{W})\mathbf{x} + \mathbf{N}_\Omega(\mathbf{x})$ and $\Phi = \mathbf{D} + \mathbf{W}$, where $\mathbf{D} = \mathbf{D} \otimes \mathbf{I}_n$ and \mathbf{D} is the degree matrix of the communication graph, for which a convergence result analogous to Theorem 2 still holds. □

V. DISCUSSION ON THE CONVERGENCE RATE

In this section, we compare the convergence rate of Algorithm 1 with that of two gradient-based NE seeking schemes recently presented in [17]: GRANE [17, Alg. 1] and acc-GRANE [17, Alg. 2].

GRANE converges linearly under μ_{F_a} -restricted strongly monotonicity and θ_{F_a} -Lipschitz continuity of F_a in (6), with rate (in squared norm) $\mathcal{O}(\lambda_{\text{GRANE}}^k)$, where

$$\lambda_{\text{GRANE}} = 1 - \frac{1}{\theta_{F_a}^2 / \mu_{F_a}^2}. \quad (\text{GRANE})$$

Instead, acc-GRANE is only guaranteed to converge under the more restrictive assumption of (non-restricted) $\tilde{\mu}_{F_a}$ -strong monotonicity of the mapping F_a (which requires some additional conditions on the game mapping, see [17, Rem. 3]), where $\tilde{\mu}_{F_a} \leq \mu_{F_a}$, with squared norm rate $\mathcal{O}(\lambda_{\text{acc-GRANE}}^k)$, where

$$\lambda_{\text{acc-GRANE}} = 1 - \frac{1}{1 + \theta_{F_a} / \tilde{\mu}_{F_a}}. \quad (\text{acc-GRANE})$$

Finally, the convergence rate of Algorithm 1, in squared

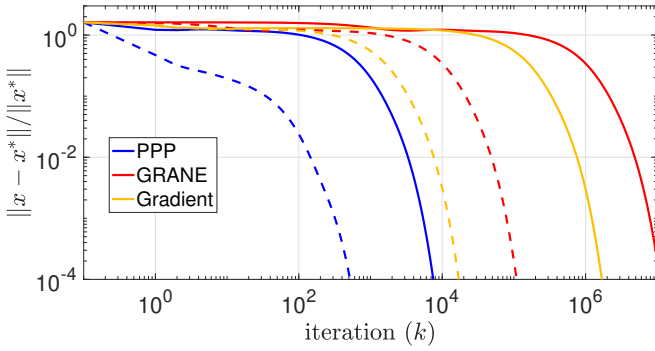


Fig. 1. Comparison of PPPA, gradient play and GRANE, with the theoretical step sizes that ensure convergence (solid lines) and with step sizes set 100 times larger than the theoretical upper bounds (dashed lines).

norm, is $\mathcal{O}(\lambda_{\text{PPP}}^k)$, where

$$\lambda_{\text{PPP}} = \left(\frac{1}{1 + \mu_{F_a} / \|\Phi\|} \right)^2. \quad (\text{PPP})$$

Since $\|\Phi\| = 2$ due to Standing Assumption 4, we get

$$\lambda_{\text{PPP}} = 1 + 1/(1 + 2/\mu_{F_a})^2 - 2/(1 + 2/\mu_{F_a}) \quad (15a)$$

$$\leq 1 - 1/(1 + 2/\mu_{F_a}), \quad (15b)$$

We note that $\lambda_2(I - W) \leq 2$, and that $\lambda_{\max}(I - W) \geq 1$, if the self-loop weights are chosen small enough. By the expression of α_{\max} in (13) and by picking $\alpha < \alpha_{\max}$ to ensure restricted strong monotonicity, it can be shown that αF is $\frac{\lambda_2(I - W)}{2}$ -Lipschitz continuous, and in turn that $3 \geq \lambda_{\max}(L) + \lambda_2(L)/2 \geq \theta_{F_a} \geq \lambda_{\max}(L) - \lambda_2(L)/2 \geq 1/2$.

We conclude that, when μ_{F_a} is small (which is typically the case), Algorithm 1 has a faster theoretical convergence rate than GRANE, by (15b). Moreover, if $\theta_{F_a} > 1$, then the guaranteed rate of Algorithm 1 is better than that of acc-GRANE by (15a), despite our PPPA is ensured to converge under milder conditions and requires only one communication per iteration instead of two.

VI. NUMERICAL EXAMPLE: A CONNECTIVITY PROBLEM

We consider the connectivity problem described in [2]. Some mobile sensor devices have to coordinate their actions via wireless communication, to perform some task, e.g., exploration or surveillance. Mathematically, the sensors (agents) aim at autonomously finding the positions which minimize some global cost. This can be robustly achieved by designing individual cost functions for the agents, such that the Nash equilibrium of the resulting game coincide with an optimum of the global objective [2]. Specifically, each agent i of a group $\mathcal{I} = \{1, \dots, N\}$ is a mobile sensor moving on a plane, designed to achieve some private primary objective related to its Cartesian position $x_i \in \mathbb{R}^2$, provided that overall connectivity is preserved over the network. This is represented by the cost functions $J_i(x_i, x_{-i}) = q_i x_i^\top x_i + r_i x_i + \sum_{j=1}^N m_{i,j} \|x_i - x_j\|^2$, where $q_i > 0$, r_i and $m_{i,j} \geq 0$ are local parameters, for all $i, j \in \mathcal{I}$. The agents cannot measure the positions of the other sensors, but communicate with some neighbors over a (randomly generated) communication network. We set $N = 10$, $m_{i,j} = 1/N \forall i, j \in \mathcal{I}$;




Fig. 2. Comparison of PPPA, GRANE, and acc-GRANE, with theoretical step sizes and compact action sets.

we pick randomly with uniform distribution q_i in $[1, 2]$ and r_i in $[-2, 2]$. Because of the quadratic structure of the game and the choice of the parameters, all of our assumptions are satisfied. We compare the performance of Algorithm 1 with that of some gradient-based NE seeking algorithms proposed in the literature, for random initial conditions.

Unconstrained action sets: We compare Algorithm 1 with GRANE [17, Alg. 1] and the gradient play in [24, Eq. 7]. We set $\alpha \approx 10^{-2}$, which satisfies the condition in Theorem 1; for the other algorithms we choose the best step sizes with theoretical convergence guarantees. Figure 1 (solid lines) shows that both gradient algorithms are outperformed by far by our PPPA. Indeed, both GRANE and the gradient play are converging very slowly to the NE, mostly due to the small step sizes employed. However, our numerical experience suggests that the theoretical bounds for the parameters are very conservative. Hence, we repeat the experiment by taking all the step sizes 100 times bigger than their theoretical upper bounds (dashed lines). The convergence appears faster for all the algorithms, but Algorithm 1 is still orders of magnitude better than the gradient-based schemes.

Constrained action sets: We assume each coordinate of the position of each sensor to be constrained in the interval $[0.1, 0.5]$. We test Algorithm 1 against GRANE [17, Alg. 1], the inexact ADMM algorithm in [16, Alg. 1], and acc-GRANE [17, Alg. 2] (the latter is guaranteed to converge only under *non-restricted* strong monotonicity of the mapping F_a in (6): we check numerically that this condition holds for our case study). For all the algorithms, we select the best step sizes with theoretical guarantees. The results are illustrated in Figure 2.

Remark 5: Our PPPA requires each agent to solve a strongly convex optimization problem at each iteration. While this can be efficiently done via iterative algorithms, it might be more demanding than a projected pseudo-gradient step, which in general requires to solve a strongly convex *quadratic* optimization problem. Nonetheless, our simulations show that our method can drastically reduce the number of iterations, and thus of data transmissions, needed to converge. This is often advantageous, even at the price of an increased local computational effort, since communication between agents is typically expensive in terms of both time and energy consumption (especially in wireless systems). \square

VII. CONCLUSION AND OUTLOOK

Nash equilibrium problems under partial-decision information can be solved via a fully-distributed preconditioned proximal-point algorithm, under strong monotonicity and Lipschitz continuity of the game mapping. Our algorithm has proven much faster than the existing gradient-based methods, at least in our numerical experience. The extension of our results to games with coupling constraints or played on time-varying communication networks is left as future research.

APPENDIX

1) *Proof of Lemma 5:* Let $(\mathbf{x}, \mathbf{u}), (\mathbf{y}, \mathbf{v}) \in \text{gra}(\Phi^{-1}\mathcal{A})$ such that $\mathbf{y} \in \text{zer}(\mathcal{A}) = \text{zer}(\Phi^{-1}\mathcal{A})$. By definition, $(\mathbf{x}, \Phi\mathbf{u}), (\mathbf{y}, \Phi\mathbf{v}) \in \text{gra}(\mathcal{A})$. Besides, the operator \mathcal{A} in (7) is ρ_α -restricted strongly monotone with respect to the consensus space \mathbf{E}_n , by Lemma 4 and monotonicity of the normal cone [19, Th. 20.25]. Since $\mathbf{y} \in \mathbf{E}_n$ by Lemma 1, we can write:

$$\begin{aligned} \langle \mathbf{u} - \mathbf{v}, \mathbf{x} - \mathbf{y} \rangle_\Phi &= \langle \Phi\mathbf{u} - \Phi\mathbf{v}, \mathbf{x} - \mathbf{y} \rangle \\ &\geq \rho_\alpha \|\mathbf{x} - \mathbf{y}\|^2 \geq \frac{\rho_\alpha}{\|\Phi\|} \|\mathbf{x} - \mathbf{y}\|_\Phi^2. \quad \blacksquare \end{aligned}$$

2) *Proof of Theorem 1:* We start by noting that $\text{fix}(\mathbf{J}_B) = \text{zer}(B)$, since, for any $z \in \mathbb{R}^m$, $\mathbf{0}_m \in \mathcal{B}(z) \Leftrightarrow z \in z + \mathcal{B}(z) \Leftrightarrow (\text{Id} + \mathcal{B})^{-1}(z) \ni z$. Let $x \in \mathbb{R}^m$, $y \in \mathbf{J}_B(x) = (\text{Id} + \mathcal{B})^{-1}(x)$, $z \in \text{zer}(\mathcal{B})$. By definition of inverse operator, $x - y \in \mathcal{B}(y)$; therefore, by taking $v = \mathbf{0}_m$ in (14), we have

$$\begin{aligned} \langle y - z, x - y - (z - z) \rangle_\Psi \\ = -\|y - z\|_\Psi^2 + \langle y - z, x - z \rangle_\Psi \geq \tilde{\mu} \|y - z\|_\Psi^2. \end{aligned} \quad (16)$$

In turn, by the Cauchy-Schwarz inequality we obtain

$$\|y - z\|_\Psi \|x - z\|_\Psi \geq (1 + \tilde{\mu}) \|y - z\|_\Psi^2. \quad (17)$$

Let us set $x = w^k$ and $y = w^{k+1}$. If $y \neq z$, by dividing both sides of (17) by $\|y - z\|_\Psi$, we obtain for the iteration in (9):

$$\|w^{k+1} - z\|_\Psi \leq \left(\frac{1}{1 + \tilde{\mu}} \right) \|w^k - z\|_\Psi;$$

this trivially holds also if $y = z$. By recursion, we conclude linear convergence of $(w^k)_{k \in \mathbb{N}}$ to z . As z is an arbitrary point in $\text{zer}(\mathcal{B})$, it also follows that $\text{zer}(\mathcal{B})$ is a singleton. \blacksquare

3) *Proof of Theorem 2:* Algorithm 1 can be written as the PPA in (11) by Lemma 2. Since $\text{zer}(\mathcal{A}) = \text{zer}(\Phi^{-1}\mathcal{A}) = \{\mathbf{x}^*\}$ by Lemma 1, $\text{dom}(\mathbf{J}_{\Phi^{-1}\mathcal{A}}) = \mathbb{R}^{Nn}$ by Remark 3 and in view of Lemma 5, we can apply Theorem 1 to conclude that

$$\begin{aligned} \sqrt{\lambda_{\min}(\Phi)} \|\mathbf{x}^k - \mathbf{x}^*\| &\leq \|\mathbf{x}^k - \mathbf{x}^*\|_\Phi \\ &\leq \left(\frac{1}{1 + \rho_\alpha / \|\Phi\|} \right)^k \|\mathbf{x}^0 - \mathbf{x}^*\|_\Phi \\ &\leq \sqrt{\|\Phi\|} \left(\frac{1}{1 + \rho_\alpha / \|\Phi\|} \right)^k \|\mathbf{x}^0 - \mathbf{x}^*\|. \quad \blacksquare \end{aligned}$$

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 University Press, 2009, p. 443–493.
- [2] M. S. Stankovic, K. H. Johansson, and D. M. Stipanovic, "Distributed seeking of Nash equilibria with applications to mobile sensor networks," *IEEE Transactions on Automatic Control*, vol. 57, no. 4, pp. 904–919, 2012.
- [3] S. Grammatico, "Dynamic control of agents playing aggregative games with coupling constraints," *IEEE Transactions on Automatic Control*, vol. 62, no. 9, pp. 4537–4548, 2017.
- [4] N. Li, L. Chen, and M. A. Dahleh, "Demand response using linear supply function bidding," *IEEE Transactions on Smart Grid*, vol. 6, no. 4, pp. 1827–1838, 2015.
- [5] C. Yu, M. van der Schaar, and A. H. Sayed, "Distributed learning for stochastic generalized Nash equilibrium problems," *IEEE Transactions on Signal Processing*, vol. 65, no. 15, pp. 3893–3908, 2017.
- [6] G. Belgioioso and S. Grammatico, "Projected-gradient algorithms for generalized equilibrium seeking in aggregative games are preconditioned forward-backward methods," in *2018 European Control Conference*, 2018, pp. 2188–2193.
- [7] J. S. Shamma and G. Arslan, "Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria," *IEEE Transactions on Automatic Control*, vol. 50, no. 3, pp. 312–327, 2005.
- [8] P. Yi and L. Pavel, "An operator splitting approach for distributed generalized Nash equilibria computation," *Automatica*, vol. 102, pp. 111–121, 2019.
- [9] 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.
- [10] K. Bimpikis, S. Ehsani, and R. Ilkiliç, "Cournot competition in networked markets," in *15th ACM Conference on Economics and Computation*, ser. EC '14. ACM, 2014, p. 733.
- [11] J. Koshal, A. Nedić, and U. V. Shanbhag, "Distributed algorithms for aggregative games on graphs," *Operations Research*, vol. 64, pp. 680–704, 2016.
- [12] P. Frihauf, M. Krstic, and T. Basar, "Nash equilibrium seeking in noncooperative games," *IEEE Transactions on Automatic Control*, vol. 57, no. 5, pp. 1192–1207, 2012.
- [13] C. De Persis and S. Grammatico, "Distributed averaging integral Nash equilibrium seeking on networks," *Automatica*, vol. 110, p. 108548, 2019.
- [14] D. Gadjov and L. Pavel, "A passivity-based approach to Nash equilibrium seeking over networks," *IEEE Transactions on Automatic Control*, vol. 64, no. 3, pp. 1077–1092, 2019.
- [15] F. Salehisadaghiani and L. Pavel, "Distributed Nash equilibrium seeking: A gossip-based algorithm," *Automatica*, vol. 72, pp. 209–216, 2016.
- [16] 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, 2019.
- [17] T. Tatarenko, W. Shi, and A. Nedić, "Geometric convergence of gradient play algorithms for distributed Nash equilibrium seeking," *arXiv preprint arXiv:1809.07383*, 2018.
- [18] L. Pavel, "Distributed GNE seeking under partial-decision information over networks via a doubly-augmented operator splitting approach," *IEEE Transactions on Automatic Control*, vol. 65, no. 4, pp. 1584–1597, 2020.
- [19] H. H. Bauschke and P. L. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2017, vol. 2011.
- [20] P. Yi and L. Pavel, "Distributed generalized Nash equilibria computation of monotone games via double-layer preconditioned proximal-point algorithms," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 1, pp. 299–311, 2019.
- [21] J. Lei and U. V. Shanbhag, "Linearly convergent variable sample-size schemes for stochastic Nash games: Best-response schemes and distributed gradient-response schemes," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 3547–3552.
- [22] M. Bianchi and S. Grammatico, "Fully distributed Nash equilibrium seeking over time-varying communication networks with linear convergence rate," *IEEE Control Systems Letters*, vol. 5, no. 2, pp. 499–504, 2021.
- [23] F. Facchinei and J. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer, 2007.
- [24] T. Tatarenko and A. Nedić, "Geometric convergence of distributed gradient play in games with unconstrained action sets," *arXiv preprint arXiv:1907.07144*, 2019.