



On decentralized computation of the leader's strategy in bi-level games[☆]



Marko Maljkovic^{a,*}, Gustav Nilsson^b, Nikolas Geroliminis^a

^a The School of Architecture, Civil and Environmental Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland

^b Dipartimento di Ingegneria Industriale, Università di Trento, Trento, Italy

ARTICLE INFO

Article history:

Received 16 February 2024

Received in revised form 19 November 2024

Accepted 8 April 2025

ABSTRACT

Motivated by the omnipresence of hierarchical structures in many real-world applications, this study delves into the intricate realm of bi-level games, with a specific focus on exploring local Stackelberg equilibria as a solution concept. While existing literature offers various methods tailored to specific game structures featuring one leader and multiple followers, a comprehensive framework providing formal convergence guarantees appears to be lacking. Drawing inspiration from sensitivity results for nonlinear programs and guided by the imperative to maintain scalability and preserve agent privacy, we propose a decentralized approach based on the projected gradient descent with the Armijo stepsize rule. By the virtue of the Implicit Function Theorem, we establish convergence to a local Stackelberg equilibrium for a broad class of bi-level games. Moreover, for quadratic aggregative Stackelberg games, we also introduce a decentralized warm-start procedure based on the consensus alternating direction method of multipliers addressing the initialization issues reported in our previous work. Finally, we provide empirical validation through two case studies in smart mobility, showcasing the effectiveness of our general method in handling general convex constraints, and the effectiveness of its extension in tackling initialization issues.

© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In the realm of strategic decision-making, games with the inherent leader-follower structure have emerged as one of the fundamental frameworks to model the interplay between agents on multiple levels of hierarchy. Each lower-level agent typically competes to solve a personalized optimization problem shaped by the leader's decision variable and the actions of other followers. In turn, the leader seeks to minimize its objective while adhering to the equilibrium constraints imposed by the lower-level game, creating a hierarchical, bi-level, structure. Depending on the degree of interaction flexibility between the two levels, and with the pivotal works on games formalizing the concepts of Stackelberg (SG) (Von Stackelberg, 1952) and their broader format, Reverse Stackelberg games (RSG) (Ho, Luh, & Muralidharan, 1981; Ho, Luh,

& Olsder, 1982), various real-world problems in the domain of energy management (Aussel, Brotcorne, Lepaul, & von Niederhäusern, 2020; Motaleb, Siano, & Ghorbani, 2019), operational optimization (Hirose & Matsumura, 2019; Zardini, Lanzetti, Guerini, Frazzoli, & Dörfler, 2021) and transportation (Groot, Zaccour, & de Schutter, 2017; Maljkovic, Nilsson, & Geroliminis, 2023a; Maljkovic et al., 2023c) gained interest from the perspective of computing a no-regret solution for all participants.

In the context of pricing problems, it has been observed that for bi-level games with a certain lower-level structure, the concept of RSG can be effectively employed to incentivize the global optimum of the leader (Maljkovic et al., 2023a; Staňková, Olsder, & Bliemer, 2011). Specifically, formulating the leader's action, i.e., a dynamic pricing strategy, as a functional, rather than a real-valued vector, can help reshape the game among followers to produce an equilibrium that aligns with the minimizer of the leader's objective. This leads to a natural follow-up question: can a connection be established between these dynamic policies and the fixed, real-valued vector strategies commonly used in day-ahead markets and bidding mechanisms (Toubeau, Bottieau, De Grève, Vallée, & Bruninx, 2021; Wang et al., 2023)? In general, this is not straightforward, as limiting the leader's influence to merely parameterizing the lower-level game constrains their ability to reshape key mathematical properties of the lower-level problem such as convexity or smoothness. In fact,

[☆] This work was supported by the Swiss National Science Foundation under NCCR Automation Phase 2 project, fonds 565 776. This paper was recommended for publication in revised form by Associate Editor Behrouz Touri under the direction of Editor Christos G. Cassandras. The material in this paper was partially presented in (Maljkovic, Nilsson, & Geroliminis, 2023c), at the 22nd European Control Conference (ECC), June 13–16, 2023, Bucharest, Romania.

* Corresponding author.

E-mail addresses: marko.maljkovic@epfl.ch (M. Maljkovic), gustav.nilsson@unitn.it (G. Nilsson), nikolas.geroliminis@epfl.ch (N. Geroliminis).

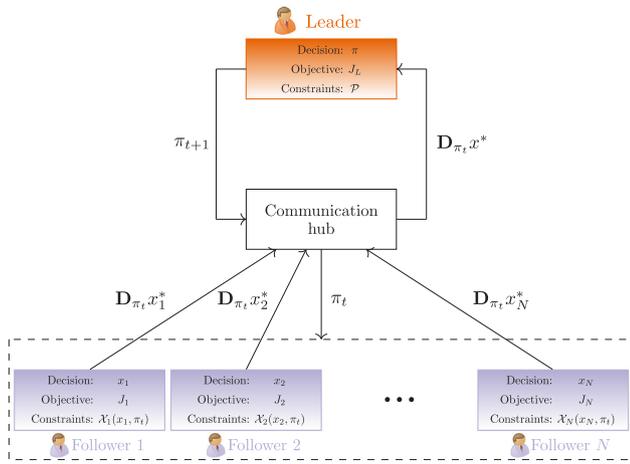


Fig. 1. Schematic sketch of the problem setup. Each of the N followers aims to optimize the personal objective J_i under the parametrized local constraints $x_i \in \mathcal{X}_i(x_i, \pi)$. The followers communicate with the leader through the communication hub that is used as a medium to collect the locally computed Jacobians D_{π_t, x_i^*} in every update step of the leader's action.

we enter the domain of mathematical programs with complementarity constraints (MPCC) (Scheel & Scholtes, 2000), which often violate standard constraint qualifications and are therefore challenging to solve using conventional off-the-shelf solvers. If application permits centralized computation, then MPCCs can typically be tackled by iterative relaxations of the equilibrium constraints (Hoheisel, Kanzow, & Schwartz, 2011) or by recasting them into an instance of a mixed integer program (Fortuny-Amat & McCarl, 1981; Kleinert, Labbé, Ljubić, & Schmidt, 2021). However, the former approach is prone to numerical instabilities (Scheel & Scholtes, 2000), while the success of the latter heavily relies on the scalability of the underlying mixed-integer solver (Jara-Moroni, Pang, & Wächter, 2018). If local constraints are also involved, then the need to preserve privacy and ensure scalability often calls for customized decentralized approaches that leverage the specific communication infrastructure between agents (Fabiani, Tajeddini, Kebriaei, & Grammatico, 2022; Kebriaei & Iannelli, 2018).

Looking from a different angle, the sensitivity results for non-linear programs (Pedregosa, 2016a; Subotić, Hauswirth, & Dörfler, 2022; Wang, Xu, Perrault, Reiter, & Tambe, 2022), along with the success of gradient-descent-based methods in hyper-parameter optimization (Bengio, 2000; Pedregosa, 2016b; Rajeswaran, Finn, Kakade, & Levine, 2019), hint at the possibility of differentiating the Karush–Kuhn–Tucker (KKT) conditions of the followers' best-response optimization problems, in an attempt to estimate how the attained lower-level equilibrium reacts to a change in the leader's action. In Grazi, Franceschi, Pontil, and Salzo (2020), the sensitivity of the solution to a parametrized fixed-point equation representing the lower level of a bi-level optimization problem is thoroughly analyzed. From a game-theoretic perspective, this concept is appealing because Nash equilibria can often be computed in a similar way, given standard structural assumptions about the game (Paccagnan, Gentile, Parise, Kamgarpour, & Lygeros, 2019). To the best of our knowledge, Grontas et al. (2024) is the only work to explore this idea in the context of equilibrium seeking in specific classes of bi-level games.

Similar to Fabiani et al. (2022), Hu and Ralph (2007), Kulkarni and Shanbhag (2015), due to inherent non-convex nature of the bi-level problem arising from the equilibrium constraints, we also focus on computing local Stackelberg equilibria as the solution concept. In fact, this paper builds on our preliminary

work presented in Maljkovic et al. (2023c), but extends it to a general framework, depicted in Fig. 1, that can also account for any type of non-quadratic and non-aggregative games. Moreover, in contrast to existing methods (Fabiani et al., 2022; Grontas et al., 2024), we allow for any leader-parametrized form of the followers' constraint sets that satisfies simple regularity assumptions. This enables us to address scenarios commonly found in bi-level pricing problems (Golrezaei, Jaillet, Cheuk Nam Liang, & Mirrokni, 2023; Mayer & Steinhardt, 2016), which do not conform to the structural requirements of Fabiani et al. (2022), Grontas et al. (2024). For the proposed framework, we show how the computation of Jacobians describing the influence of the leader's strategy on the attained variational Nash Equilibrium (Maljkovic et al., 2023c) of the lower-level game can be performed locally by each of the followers. We start by briefly discussing the relationship between the RSG dynamic policies and the static ones used in the SG setup and continue by rigorously tackling the requirements of the Implicit Function Theorem (Dontchev & Rockafellar, 2009). Although these requirements are often assumed to be met in the existing literature, they can be violated even in games with relatively simple structures, as demonstrated in the case study of quadratic, aggregative games with polytopic constraints in Maljkovic et al. (2023c). In this context, we also address the reported initialization issue arising from the fact that in each iteration we differentiate the KKT conditions of an optimization problem equivalent to the standard best-response one. With that in mind, the main contributions of this paper can be summarized as follows:

- We propose a distributed, first-order-like, iterative method based on explicit fulfillment of the Implicit Function Theorem requirements. By ensuring a local improvement of the leader's objective at each iteration, we provide formal convergence guarantees for a broad class of bi-level games.
- For a subclass of quadratic, aggregative, Stackelberg games linearly parametrized by the decision variable of the leader, we propose a decentralized warm-start procedure based on the alternating direction method of multipliers (ADMM). In line with the distributed nature of the main algorithm, we compute a feasible leader's strategy that yields an interior point variational Nash Equilibrium of the lower-level game, i.e., renders no local inequality constraint active.
- For a subclass of quadratic, aggregative games with only equality constraints, we establish sufficient conditions for linear convergence and provide additional discussion points for other cases.
- We test the proposed method in an adaptation of the case study from the domain of smart mobility previously analyzed in Maljkovic et al. (2023c). Firstly, by introducing the so-called budget constraints parametrized by the leader's decision variable, we demonstrate the effectiveness of the main procedure for the bi-level games with general convex constraints. Then, by going back to the setup in Maljkovic et al. (2023c), we illustrate the effectiveness of the warm-start procedure in alleviating the initialization issues. Finally, we numerically evaluate the robustness of the proposed method in scenarios where the followers do not perfectly reach the Nash Equilibrium for a given value of the leader's decision variable. This demonstrates the algorithm's potential for future extension into an equilibrium-seeking mechanism.

The paper is outlined as follows: the rest of this section is devoted to introducing some basic notation. In Section 2, we introduce the general bi-level setup, discuss the connection between the Stackelberg and Reverse Stackelberg games, postulate

the main standing assumptions, and formally introduce the problem. In Section 3, we revise and generalize the decentralized method for computing the local Stackelberg equilibrium previously outlined in Maljkovic et al. (2023c). Section 4 then focuses on the subclass of quadratic, aggregative Stackelberg games and presents the proposed warm-start procedure. Finally, we conclude the paper with Sections 5 and 6 where we present the numerical examples and propose some ideas for future research.

Notation: Let \mathbb{R} denote the set of real numbers, \mathbb{R}_+ the set of non-negative reals, and \mathbb{Z}_+ the set of non-negative integers. Let $\mathbf{0}_m$ and $\mathbf{1}_m$ denote the all zero and all one vectors of length m respectively, and \mathbb{I}_m the identity matrix of size $m \times m$. For a finite set \mathcal{A} we let $|\mathcal{A}|$ be the cardinality of \mathcal{A} . Furthermore, if \mathcal{A} is a finite set of vectors x_i , we let $x := \text{col}(x_i)_{i \in \mathcal{A}}$ be their concatenation. For $x \in \mathbb{R}^n$, we let $\text{Dg}(x) \in \mathbb{R}^{n \times n}$ denote a diagonal matrix whose elements on the diagonal correspond to vector x . If \mathcal{A} is a set of k matrices $A_i \in \mathbb{R}^{m \times n}$, then $\text{blkdiag}\{A_i\}_{i \in \mathcal{A}} \in \mathbb{R}^{km \times kn}$ denotes the corresponding block-diagonal matrix and $\text{vstack}(A_i)_{i \in \mathcal{A}} \in \mathbb{R}^{km \times n}$ denotes their vertical concatenation. For $A \in \mathbb{R}^{n \times n}$, $A \succ 0$ ($\succeq 0$) is equivalent to $x^T A x > 0$ (≥ 0) for all $x \in \mathbb{R}^{n \times n}$. We let $A \otimes B$ denote the Kronecker product between two matrices. For a differentiable function $f(x) : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we let $\mathbf{D}_x f \in \mathbb{R}^{m \times n}$ denote the Jacobian matrix of f defined as $(\mathbf{D}_x f)_{ij} := \frac{\partial f_i}{\partial x_j}$. If $f(x)$ is a real-valued function, i.e., $m = 1$, we adopt $\nabla_x f := \mathbf{D}_x f \in \mathbb{R}^n$. Finally, for a set-valued mapping $\mathcal{F} : \mathcal{Y} \rightrightarrows \mathcal{X}$, $\text{gph}(\mathcal{F}) := \{(y, x) \in \mathcal{Y} \times \mathcal{X} \mid x \in \mathcal{F}(y)\}$ denotes its graph.

2. Theoretical preliminaries

Throughout the paper, we consider a bi-level game with a set of $N + 1$ agents $\bar{\mathcal{I}} = \mathcal{I} \cup \{L\}$, where L represents the leading agent and each $i \in \mathcal{I}$ represents one of the N followers. In this setup, the leading agent is the first one to choose an action π from its feasible set \mathcal{P} , to which all N followers will respond at once with a personal decision vector x_i from their feasible set $\mathcal{X}_i(\pi)$ that is in their best interest. If $m_f \in \mathbb{N}$ represents the dimension of the follower's decision space, we assume $\mathcal{X}_i(\pi)$ in the form of $\mathcal{X}_i(\pi) := \{x_i \in \mathbb{R}^{m_f} \mid \mathbf{g}_i^{\text{inq}}(x_i, \pi) \leq \mathbf{0}_{m_{\text{inq},i}} \wedge \mathbf{g}_i^{\text{eq}}(x_i, \pi) = \mathbf{0}_{m_{\text{eq},i}}\}$, where $m_{\text{inq},i}, m_{\text{eq},i} \in \mathbb{N}$ denote the number of inequality and equality constraints encompassed in $\mathcal{X}_i(\pi)$. If $m_L \in \mathbb{N}$ represents the dimension of the leader's action, the nature of the leader's strategy can lead to two types of games in general:

- **Reverse Stackelberg Games (RSG):** where the leader's strategy $\pi \in \mathcal{P}$ is a map $\pi : \mathbb{R}^{m_f} \rightarrow \mathbb{R}^{m_L}$;
- **Stackelberg Games (SG):** where the leader's strategy is a fixed real vector, $\pi \in \mathcal{P} \subseteq \mathbb{R}^{m_L}$.

In any case, we refer to the phase of choosing the optimal $x_i \in \mathcal{X}_i(\pi)$ as the *Lower-level game* and the process of choosing the optimal leader's strategy knowing that the followers will play a best-response as the *Upper-level game*. Furthermore, we define the joint strategy of all followers as $x := \text{col}(x_i)_{i \in \mathcal{I}} \in \mathcal{X}(\pi)$ and for every $i \in \mathcal{I}$, we define $x_{-i} := \text{col}(x_j)_{j \in \mathcal{I} \setminus i} \in \mathcal{X}_{-i}(\pi)$, such that $\mathcal{X}(\pi) := \prod_{i \in \mathcal{I}} \mathcal{X}_i(\pi)$ and $\mathcal{X}_{-i}(\pi) := \prod_{j \in \mathcal{I} \setminus i} \mathcal{X}_j(\pi)$.

2.1. Lower-level game

The followers choose their strategies in an attempt to minimize personal objectives $J_i(x_i, x_{-i}, \pi)$ by optimally responding to other players' strategies. In addition to defining the leader's influence on the functional form of the followers' objectives, adopting specific general forms for the constraint sets $\mathcal{X}_i(\pi)$ allows us to capture different paradigms studied in the literature. For instance, in Stackelberg games, we can structurally distinguish between two types of constraint sets $\mathcal{X}_i(\pi)$. The first type

includes inequality constraints parametrized by π that take a polytopic form, i.e., $\mathbf{g}_i^{\text{inq}}(x_i, \pi) = W_i x_i + W_i^T \pi + b_i$, where W_i , W_i^T , and b_i have appropriate dimensions. This structure is more prevalent in the literature, as it often models restricted access to shared resources by the followers (Grontas et al., 2024; Nabetani, Tseng, & Fukushima, 2011). The second type involves bi-linear constraints, i.e., $\mathbf{g}_i^{\text{inq}}(x_i, \pi) = x_i^T W_i \pi + b_i$, and is commonly used in bi-level pricing problems (Golrezaei et al., 2023; Mayer & Steinhardt, 2016) to represent players' limited personal budgets. Notably, much of the existing literature (Fabiani et al., 2022; Grontas et al., 2024) relies on polytopic constraint sets, making their proposed methods less applicable to problems involving budget constraints.

In any case, in this paper we keep a general, π -induced, lower-level game $G^0(\mathcal{I}; \pi)$ in the form of N coupled optimization problems, i.e., $G^0(\mathcal{I}; \pi) := \{G_i^0(\pi, x_{-i}) \mid i \in \mathcal{I}\}$, where each follower aims to solve:

$$G_i^0(\pi, x_{-i}) := \left\{ \begin{array}{l} \min_{x_i \in \mathbb{R}^{m_f}} J_i(x_i, x_{-i}, \pi) \\ \text{s.t. } x_i \in \mathcal{X}_i(\pi) \end{array} \right\}. \quad (1)$$

For such a game, a viable no-regret solution for all players is given by the concept of a Nash equilibrium, which is formally introduced in Definition 1.

Definition 1 (Nash Equilibrium). For any leader's strategy $\pi \in \mathcal{P}$, a joint strategy $x^* \in \mathcal{X}$ is a Nash Equilibrium (NE) of the game $G^0(\mathcal{I}; \pi)$ if, for all $i \in \mathcal{I}$ and all $x_i \in \mathcal{X}_i(\pi)$, it holds that $J_i(x_i^*, x_{-i}^*, \pi) \leq J_i(x_i, x_{-i}^*, \pi)$.

For a particular $\pi \in \mathcal{P}$, it is rarely possible to find a closed-form characterization of the full set of NE in a general setup. Therefore, we postulate standard assumptions about the structure of $G^0(\mathcal{I}; \pi)$ that allow us to focus on the variational Nash Equilibria (v-NE) as a solution concept of the lower-level game.

Standing Assumption 1. For every $i \in \mathcal{I}$ and any $\pi \in \mathcal{P}$, $x_{-i} \in \mathcal{X}_{-i}$, the cost $J_i(x_i, x_{-i}, \pi)$ is convex and continuously differentiable in x_i . Moreover, it is continuous in $x \in \mathcal{X}(\pi)$ and the sets $\mathcal{X}_i(\pi)$ are nonempty, compact, convex and satisfy Slater's constraint qualification.

Strictly speaking, under Standing Assumption 1, for every $i \in \mathcal{I}$, the Nash Equilibrium strategy $x_i^* \in \mathcal{X}_i(\pi)$ is the solution of the best-response optimization problem (1) for x_{-i}^* , i.e., $G_i^0(\pi, x_{-i}^*)$. The optimality of $x_i^* \in \mathcal{X}_i(\pi)$ is guaranteed if and only if x_i^* solves the KKT system of equations $l_i(z_i, \pi \mid x_{-i}^*) = 0$, where the vector mapping l_i is defined for every $z_i = (x_i, \lambda_i, v_i)$ as

$$l_i(z_i, \pi \mid x_{-i}^*) := \begin{bmatrix} \nabla_{x_i} \mathcal{L}_i(z_i, \pi) \\ \text{Dg}(\lambda_i) \mathbf{g}_i^{\text{inq}}(x_i, \pi) \\ \mathbf{g}_i^{\text{eq}}(x_i, \pi) \end{bmatrix}, \quad (2)$$

with the Lagrangian $\mathcal{L}_i(z_i, \pi) = J_i(x_i, x_{-i}^*, \pi) + \lambda_i^T \mathbf{g}_i^{\text{inq}}(x_i, \pi) + v_i^T \mathbf{g}_i^{\text{eq}}(x_i, \pi)$, and $\lambda_i \in \mathbb{R}_+^{m_{\text{inq},i}}$ and $v_i \in \mathbb{R}^{m_{\text{eq},i}}$ representing the dual variables associated with the inequality and equality constraints. Under Standing Assumption 1, if some $\hat{z}_i = (\hat{x}_i, \hat{\lambda}_i, \hat{v}_i)$, with feasible \hat{x}_i and $\hat{\lambda}_i$, satisfies $l_i(\hat{z}_i, \pi \mid x_{-i}^*) = 0$, then \hat{z}_i is the optimizer of (1). On the other hand, based on Facchinei and Pang, Prop.1.4.2, Standing Assumption 1 also ensures that a joint strategy $x \in \mathcal{X}(\pi)$ is a NE if and only if it solves a variational inequality problem, hence providing a closed-form description of the lower-level game's solution set. Namely, if $F(x, \pi) := \text{col}(\nabla_{x_i} J_i(x_i, x_{-i}, \pi))_{i \in \mathcal{I}}$ denotes the pseudo-gradient of $G^0(\mathcal{I}; \pi)$, we can adopt the following assumption about the followers.

Standing Assumption 2. For any $\pi \in \mathcal{P}$, the agents $i \in \mathcal{I}$ play a joint strategy $x \in \mathcal{N}_0(\pi)$, where $\mathcal{N}_0(\pi)$ is the set of all v-NE of the game $G^0(\mathcal{I}; \pi)$, given by $\mathcal{N}_0(\pi) := \{x \in \mathcal{X}(\pi) \mid (y - x)^T F(x, \pi) \geq 0, \forall y \in \mathcal{X}(\pi)\}$.

2.2. Upper-level game

On the upper level, finding the optimal strategy $\pi \in \mathcal{P}$ imposes solving a minimization problem of the leader's objective $J_L : \mathbb{R}^{N_{mf}} \times \mathcal{P} \rightarrow \mathbb{R}$. Instances of both SG and RSG can be compactly written as:

$$G_1 := \left\{ \begin{array}{l} \min_{\pi \in \mathcal{P}} J_L(x^*, \pi) \\ \text{s.t. } (x^*, \pi) \in \text{gph}(\mathcal{N}_0) \cap (\mathbb{R}^{mf} \times \mathcal{P}) \end{array} \right\}. \quad (3)$$

In general, the optimal π in G_1 is a possibly non-unique solution (Maljkovic et al., 2023c, Corr. 1) to a non-convex problem that requires the ability to understand the leader's influence on the position of the lower-level game's v-NE. Moreover, depending on the properties of $F(x, \pi)$, the lower-level game could admit multiple NE for a particular parametrization. In this study, we restrict ourselves to cases where $G^0(\mathcal{I}; \pi)$ admits a unique NE for any $\pi \in \mathcal{P}$. Hence, we state the following assumption, common in existing literature (Bianchi, Belgioioso, & Grammatico, 2022; Paccagnan, Gentile, Parise, Kamgarpour, & Lygeros, 2016), that ensures the existence and uniqueness of the lower-level game's v-NE (Facchinei & Pang, Th. 2.3.3).

Standing Assumption 3. For any $\pi \in \mathcal{P}$, the pseudo-gradient $F(x, \pi)$ is strongly monotone in $x \in \mathcal{X}(\pi)$.

Concerning the nature of the leader's decision variable, it is evident that SG represents a distinct instance of RSG, wherein the leader's strategy assumes a constant function. This limits the flexibility to incentivize a certain NE of the lower-level game, as the leader's strategies in the form of feedback policies offer a means to directly shape the functional form of the followers' optimization problems. Moreover, with many real-world applications requiring different notions of fairness, solving a SG can be considered arguably more challenging. To elucidate this contrast, we look at the following example for a specific class of games, referred to as quadratic aggregative games in Maljkovic et al. (2023a), Paccagnan, Kamgarpour, and Lygeros (2016).

Definition 2 (Quadratic Aggregative Games). Let the leader's objective be $J_L = \frac{1}{2}\sigma(x)^T P_L \sigma(x) + q_L^T \sigma(x)$, where $\sigma(x) = \sum_{i \in \mathcal{I}} x_i$. Moreover, let the lower-level game $G^0(\mathcal{I}; \pi)$ be defined by

$$J_i(x_i, x_{-i}, \pi) = \frac{1}{2}x_i^T P_i x_i + x_i^T Q_i \sigma(x_{-i}) + r_i^T x_i + x_i^T S_i \pi_i,$$

where $\sigma(x_{-i}) = \sigma(x) - x_i$, $\pi_i \in \mathbb{R}^{mf}$, $\pi = \text{col}(\pi_i)_{i \in \mathcal{I}}$, the matrices P_L, P_i, Q_i, S_i and vectors q_L, r_i are all real valued, $P_L \succ 0$, for every $i \in \mathcal{I}$, $P_i, S_i \succ 0$, and the Standing assumptions 1, 2 and 3 all hold.

If one regards the game in Definition 2 as a Reverse Stackelberg game, it suffices to choose the leader's strategy as a mapping of the form

$$\pi_i(x_i, x_{-i}) = S_i^{-1} \left[\frac{1}{2} \bar{P}_i x_i + \bar{Q}_i \sigma(x_{-i}) + \bar{r}_i \right], \quad (4)$$

where $\bar{P}_i = P_L - P_i$, $\bar{Q}_i = P_L - Q_i$ and $\bar{r}_i = q_L - r_i$, so that the leader's objective and the objectives of the followers satisfy $J_L(x_i, x_{-i}) - J_L(\tilde{x}_i, x_{-i}) = J_i(x_i, x_{-i}) - J_i(\tilde{x}_i, x_{-i})$ for any fixed $x_{-i} \in \mathcal{X}_{-i}$ and any two $x_i, \tilde{x}_i \in \mathcal{X}_i$. This implies that due to (4), the leader's objective J_L becomes the exact potential (Monderer & Shapley, 1996) of the lower-level game by definition. Consequently, this guarantees that the minimizer of J_L aligns with the v-NE of $G^0(\mathcal{I}; \pi)$, which can, in this case, be computed using a decentralized, iterative, fixed-point method as in Maljkovic et al. (2023a, 2023c), Paccagnan, Kamgarpour, and Lygeros (2016). Conversely, if one regards the game in Definition 2 as an instance of Stackelberg games, such manipulation is no longer possible. If

x^R represents the NE obtained when applying (4) in the RSG setup, one might naively try to plug back x^R into (4) to obtain a static pricing vector $\pi_i^R = \pi_i(x_i^R, x_{-i}^R)$ and use it in the setup of a SG. However, in a general case, this does not yield a viable solution, as illustrated in the following proposition.

Proposition 1. Let a bi-level game be defined as in Definition 2 such that $\mathcal{X}_i(\pi) = \mathbb{R}^{mf}$ for all $i \in \mathcal{I}$. Moreover, let the mapping $\pi = \text{col}(\pi_i)_{i \in \mathcal{I}}$ be given by $\pi_i : \mathbb{R}^{N_{mf}} \rightarrow \mathbb{R}^{mf}$ and let (4) yield $x^R \in \mathcal{X}(\pi)$. If $\pi_i^R = \pi_i(x_i^R, x_{-i}^R)$, then utilizing π_i^R in a SG gives rise to a NE equal to x^R if and only if $P_i x_i^R = P_i x_i^R$.

Proof. It suffices to look at the KKT systems $l_i(z_i, \pi_i | x_{-i}^R) = 0$, given by (2), for the SG and RSG scenarios. Namely, after applying (4), the derivative of the i th follower's Lagrangian evaluated at x_i^R satisfies $P_L x_i^R + P_i \sigma(x_{-i}^R) + q_L = 0$. On the other hand, if π_i^R is applied in the context of SG, x_i^R will remain the NE of $G^0(\mathcal{I}; \pi)$ if and only if $\frac{1}{2}(P_L + P_i)x_i^R + P_i \sigma(x_{-i}^R) + q_L = 0$ holds. Hence, x^R remains the NE of $G^0(\mathcal{I}; \pi)$ if and only if $\frac{1}{2}P_L x_i^R = \frac{1}{2}P_i x_i^R$, which does not always hold. \square

Therefore, even the Stackelberg games that exhibit favorable mathematical properties such as the one in Definition 2 pose significant difficulties in computing the leader's strategy. As previously mentioned in the introduction, if the nature of the application allows centralized computation, one can formulate an MPCC that can be recast into an instance of mixed-integer linear or quadratic problems (Kleinert et al., 2021) using the big-M reformulation (Fortuny-Amat & McCarl, 1981) as demonstrated in Fochesato, Cenedese, and Lygeros (2022), Hobbs (2001), Maljkovic, Nilsson, and Geroliminis (2023b). Apart from suffering from poor scalability (Jara-Moroni et al., 2018), such computation could breach the privacy of the lower-level agents in many real-world applications, particularly in terms of sharing information about personal constraint sets \mathcal{X}_i .

With that in mind, the focus of this paper is entirely redirected towards the decentralized computation of the leader's strategy in Stackelberg games. Before formally introducing the problem setup, we highlight that the proposed modeling framework is particularly well-suited for applications such as day-ahead markets, bidding mechanisms, transportation pricing, and energy markets, where upper-level updates typically occur over extended time intervals. For systems with rapidly changing dynamics, special care needs to be taken to ensure that the lower-level equilibrium is attained before each upper-level update.

2.3. Problem formulation – Stackelberg games

Owing to the inherently non-convex nature of the hierarchical problem arising from the equilibrium constraints, we concentrate on instances of (3) that are guaranteed to admit a solution. Based on (Luo, Pang, & Ralph, 1996, Th.1.4.1), this is ensured by posing the following mild assumptions concerning the leader's objective.

Standing Assumption 4. The leader's objective function $J_L : \mathbb{R}^{N_{mf}} \times \mathcal{P} \rightarrow \mathbb{R}$ is coercive and continuously differentiable in both x^* and π .

We can now focus on finding the leader's strategy $\pi \in \mathcal{P}$ corresponding to the concept of local Stackelberg equilibria (l-SE) previously studied in Fabiani et al. (2022), Hu and Ralph (2007), Kulkarni and Shanbhag (2015), that is guaranteed to exist due to Standing Assumption 4. For the sake of completeness, we reintroduce this concept formally in Definition 3.

Definition 3 (Local Stackelberg Equilibrium). Let G_1 be a Stackelberg game as in (3). A pair $(\hat{x}^*, \hat{\pi}) \in \text{gph}(\mathcal{N}_0) \cap (\mathbb{R}^{mf} \times \mathcal{P})$ is a

local Stackelberg equilibrium of G_1 if there exist open neighborhoods $\Omega_{\hat{x}^*}$ and $\Omega_{\hat{\pi}}$ of \hat{x}^* and $\hat{\pi}$ respectively, such that

$$J_L(\hat{x}^*, \hat{\pi}) \leq \inf_{(x^*, \pi) \in \text{gph}(\mathcal{N}_0) \cap \Omega} J_L(x^*, \pi), \quad (5)$$

where $\Omega := \Omega_{\hat{x}^*} \times (\mathcal{P} \cap \Omega_{\hat{\pi}})$.

Interestingly, restricting ourselves to the framework of I-SE implies that finding the leader's strategy in (3) reduces to finding the local optimum of J_L as a function of π . Namely, under [Standing Assumption 1](#) and [3](#), for any $\pi \in \mathcal{P}$, we have that $|\mathcal{N}_0(\pi)| = 1$. This means that to find the I-SE, we need to find $\hat{\pi}$ and its neighborhood $\Omega_{\hat{\pi}}$, since the condition (5) will always be fulfilled for $\hat{x}^* = \mathcal{N}_0(\hat{\pi})$ and the open ball of radius R given by $\Omega_{\hat{x}^*} := \{x \in \mathcal{X} \mid \|x - \hat{x}^*\| < R\}$, where $R > \max_{x \in \hat{\mathcal{N}}(\Omega_{\hat{\pi}})} \|x - \hat{x}^*\|$ and $\hat{\mathcal{N}}(\Omega_{\hat{\pi}}) = \bigcup_{\pi \in \Omega_{\hat{\pi}}} \mathcal{N}_0(\pi)$.

To summarize, in the following sections we will focus on designing an iterative, decentralized, gradient descent-based algorithm that leverages the guarantees provided by the Implicit Function Theorem ([Dontchev & Rockafellar, 2009](#)) concerning the continuous differentiability of $J_L(x^*(\pi), \pi)$ at the current π value. However, before delving deeper into the details, we establish the regularity of the leader's optimization problem through [Standing Assumption 5](#).

Standing Assumption 5. *The leader's constraint set $\mathcal{P} \subseteq \mathbb{R}^{m_L}$ is nonempty, compact and convex, and for every $i \in \mathcal{I}$, each element of $\mathbf{g}_i^{\text{inc}}(x_i, \pi)$ and $\mathbf{g}_i^{\text{eq}}(x_i, \pi)$ is continuously differentiable in both x_i and π . For every $x_{-i}^* \in \mathcal{X}_{-i}(\pi)$, every component of the derivative of the Lagrangian associated with the KKT system $l_i(z_i, \pi \mid x_{-i}^*) = 0$, i.e., $\nabla_{x_i} \mathcal{L}(z_i, \pi)$, is continuously differentiable at both x_i and π .*

3. Decentralized computation of the local Stackelberg equilibrium

To tackle the problem of computing the local Stackelberg equilibrium, the initial step involves introducing the idea of Projected Gradient descent incorporating the Armijo rule. This concept, along with the Implicit Function Theorem, will form the foundation of our method.

3.1. Projected gradient descent with Armijo rule

We start by adopting the projected gradient descent method with 'Armijo step-size rule along the projection arc' explored in [Bertsekas \(1999\)](#). To update π at iteration $t \in \mathbb{N}$, we first define the mapping $\pi^+ : \mathcal{P} \times \mathbb{R}_+ \rightarrow \mathcal{P}$ as

$$\pi^+(\pi_t, s) := \Pi_{\mathcal{P}} \left[\pi_t - s \frac{dJ^L(x_{\pi_t}^*, \pi)}{d\pi} \Big|_{\pi=\pi_t} \right],$$

where $\Pi_{\mathcal{P}}$ is the projection operator on the leader's constraint set for some particular step size $s \in \mathbb{R}_+$ and $x_{\pi_t}^*$ emphasizes the dependence of the Nash Equilibrium on π . Let β, \bar{s} and δ be fixed scalars such that $\beta, \delta \in (0, 1)$ and $\bar{s} > 0$. Moreover, let $l_t \in \mathbb{Z}_{\geq 0}$ be the smallest non-negative integer such that for $s_t = \beta^{l_t} \bar{s}$

$$\begin{aligned} J^L(x_{\pi_t}^*, \pi_t) - J^L(x_{\pi^+(\pi_t, s_t)}^*, \pi^+(\pi_t, s_t)) &\geq \\ &\geq \delta \left(\frac{dJ^L(x_{\pi_t}^*, \pi)}{d\pi} \Big|_{\pi=\pi_t} \right)^T (\pi_t - \pi^+(\pi_t, s_t)) \end{aligned} \quad (6)$$

holds. Then, the leader's strategy is updated as

$$\pi_{t+1} = \pi^+(\pi_t, s_t). \quad (7)$$

Under [Standing Assumption 5](#), to observe that l_t is well defined, i.e., a stepsize s_t will be found after a finite number of trials based on the test given by (6), it suffices to invoke the following adaptation of [Bertsekas \(1999, Prop. 2.3.3\)](#).

Lemma 1 (*Proposition 2.3.3 of [Bertsekas \(1999\)](#)*). *Let the set \mathcal{P} satisfy [Standing Assumptions 4](#) and [5](#), $J_L(x_{\pi}^*, \pi)$ be continuously differentiable on \mathcal{P} and $\delta \in (0, 1)$. Then, for every $\pi \in \mathcal{P}$, there exists $s_{\pi} > 0$ such that $J_L(x_{\pi}^*, \pi) - J_L(x_{\pi^+(\pi, s)}^*, \pi^+(\pi, s)) \geq \delta \nabla_{\pi} J_L(x_{\pi}^*, \pi)^T (\pi - \pi^+(\pi, s))$ holds for every $s \in [0, s_{\pi}]$.*

Therefore, the complexity of each update step boils down to ensuring that $J_L(x_{\pi}^*, \pi)$ is continuously differentiable, i.e., showing that the gradient of the leader's objective with respect to the current strategy given by

$$\frac{dJ_L(x_{\pi}^*, \pi)}{d\pi} = \frac{\partial J_L(x_{\pi}^*, \pi)}{\partial \pi} + \mathbf{D}_{\pi}^T x_{\pi}^* \frac{\partial J_L(x_{\pi}^*, \pi)}{\partial x_{\pi}^*}, \quad (8)$$

is well-defined. In that case, if $\bar{s} > s_{\pi_t}$ we have $l_t = 0$, otherwise the testing procedure (6) terminates after $l_t = \lceil \log_{\beta}(\frac{\bar{s}}{s_{\pi_t}}) \rceil$ iterations.

The challenging aspect of computing (8) stems from the requirement to compute the Jacobian $\mathbf{D}_{\pi} x_{\pi}^*$, i.e., from having to estimate how the NE of $G^0(\mathcal{I}; \pi)$ reacts to variations in π . This is particularly challenging as in general there exists no closed-form functional description of the connection between π and the obtained NE x_{π}^* . Therefore, we aim to achieve this by virtue of the Implicit Function Theorem. Namely, with the constraint sets $\mathcal{X}_i(\pi)$ being local, and knowing that

$$\mathbf{D}_{\pi}^T x_{\pi}^* \frac{\partial J_L(x_{\pi}^*, \pi)}{\partial x_{\pi}^*} = \sum_{i \in \mathcal{I}} \mathbf{D}_{\pi}^T x_{\pi, i}^* \frac{\partial J_L(x_{\pi}^*, \pi)}{\partial x_{\pi, i}^*},$$

we can compute $\mathbf{D}_{\pi} x_{\pi}^*$ in a distributed manner such that each follower remains in charge of only computing the personal Jacobian $\mathbf{D}_{\pi} x_{\pi, i}^*$. The Jacobians are then communicated to the leader as illustrated in [Fig. 1](#), who, in return, calculates

$$\frac{dJ_L(x_{\pi}^*, \pi)}{d\pi} = \frac{\partial J_L(x_{\pi}^*, \pi)}{\partial \pi} + \sum_{i \in \mathcal{I}} \mathbf{D}_{\pi}^T x_{\pi, i}^* \frac{\partial J_L(x_{\pi}^*, \pi)}{\partial x_{\pi, i}^*}$$

before updating its decision via (7). To obtain individual $\mathbf{D}_{\pi} x_{\pi, i}^*$, we leverage the fact that the computed lower-level NE has to solve the best-response optimization problem of the corresponding follower. Namely, to tackle the requirements of the Implicit Function Theorem, for every $i \in \mathcal{I}$, we formulate an optimization problem equivalent to (1) and directly apply the theorem to the problem's KKT mapping $l_i(z_i, \pi \mid x_{-i}^*)$. To ease the notation in the following sections, we will suppress the subscript denoting dependence on π when it is clear from the context and refer to the Jacobian of follower $i \in \mathcal{I}$ as $\mathbf{D}_{\pi} x_i^*$.

3.2. Differentiating the KKT conditions

For a given $\pi \in \mathcal{P}$ and the corresponding unique v-NE $x^* \in \mathcal{X}(\pi)$ of the lower-level game, the Implicit Function Theorem allows us to locally compute Jacobians $\mathbf{D}_{\pi} x_i^*$ by applying it to the KKT mapping $l_i(z_i, \pi \mid x_{-i}^*)$. For every $i \in \mathcal{I}$, let the set-valued map $\Xi_i^* : \mathcal{P} \rightrightarrows \mathcal{Z}_i$, with $\mathcal{Z}_i := \mathbb{R}^{m_f} \times \mathbb{R}_{\geq 0}^{m_{\text{ineq}, i}} \times \mathbb{R}^{m_{\text{eq}, i}}$ be

$$\Xi_i^*(\pi) := \{z_i \in \mathcal{Z}_i \mid l_i(z_i, \pi \mid x_{-i}^*) = 0\}. \quad (9)$$

Moreover, let $\Theta = [1, m_{\text{ineq}, i}] \cap \mathbb{N}$, and the set of non-strongly active inequality constraints $I_i^{\pi}(x_i^*, \lambda_i^*)$ be

$$I_i^{\pi}(x_i^*, \lambda_i^*) := \left\{ j \in \Theta \mid \lambda_i^{j*} = 0 \wedge \mathbf{g}_i^{\text{ineq}}(x_i^*, \pi)_j = 0 \right\},$$

where $\mathbf{g}_i^{\text{ineq}}(x_i, \pi)_j$ is the j th inequality constraint and λ_i^{j*} represents the corresponding dual variable of the best-response optimization problem. With a slight abuse of notation, the Implicit Function Theorem from [Dontchev and Rockafellar \(2009\)](#) adapted to our problem reads as the following theorem.

Theorem 1 (Theorem 1.B1 of *Dontchev and Rockafellar (2009)*). Let *Standing Assumptions 1–5* hold and $x^* \in \mathcal{X}(\pi)$ be the unique NE of the game $G^0(\mathcal{I}; \pi)$ for some $\pi \in \mathcal{P}$. Furthermore, let the best-response optimization problem of each agent $i \in \mathcal{I}$ be defined via (1), its KKT mapping $l_i(z_i, \pi \mid x_{-i}^*)$ via (2), and $\mathcal{E}_i^*(\pi)$ be defined via (9). If $l_i(\hat{z}_i, \pi \mid x_{-i}^*) = 0$, $\Gamma_i^\pi(\hat{x}_i, \hat{\lambda}_i)$ is empty and $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ is non-singular for some \hat{z}_i , then the solution mapping $\mathcal{E}_i^*(\pi)$ has a single-valued localization z_i^* around $\hat{z}_i = (\hat{x}_i, \hat{\lambda}_i, \hat{v}_i)$, that is continuously differentiable in a neighborhood Ω_π of π , with the Jacobian satisfying for every $\pi \in \Omega_\pi$

$$\mathbf{D}_\pi z_i^*(\pi) = -\mathbf{D}_{z_i}^{-1} l_i(\hat{z}_i, \pi \mid x_{-i}^*) \mathbf{D}_\pi l_i(\hat{z}_i, \pi \mid x_{-i}^*),$$

where $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ and $\mathbf{D}_\pi l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ satisfy

$$\mathbf{D}_{z_i} l_i := \begin{bmatrix} \mathbf{D}_{x_i} \nabla_{x_i} \mathcal{L}_i, & \mathbf{D}_{x_i}^T \mathbf{g}_i^{\text{inq}}, & \mathbf{D}_{x_i}^T \mathbf{g}_i^{\text{eq}} \\ \text{Dg}(\hat{\lambda}_i) \mathbf{D}_{x_i} \mathbf{g}_i^{\text{inq}}, & \text{Dg}(\mathbf{g}_i^{\text{inq}}), & \mathbf{0} \\ \mathbf{D}_{x_i} \mathbf{g}_i^{\text{eq}}, & \mathbf{0}, & \mathbf{0} \end{bmatrix}$$

$$\mathbf{D}_\pi l_i := \begin{bmatrix} \mathbf{D}_\pi \nabla_{x_i} \mathcal{L}_i \\ \text{Dg}(\hat{\lambda}_i) \mathbf{D}_\pi \mathbf{g}_i^{\text{inq}} \\ \mathbf{D}_\pi \mathbf{g}_i^{\text{eq}} \end{bmatrix}.$$

The triplet $\hat{z}_i = (x_i^*, \lambda_i^*, v_i^*)$, with x_i^* , λ_i^* and v_i^* being the solution of $G_i^0(\pi, x_{-i}^*) = \min_{x_i \in \mathcal{X}_i(\pi)} J_i(x_i, x_{-i}^*, \pi)$, satisfies $l_i(\hat{z}_i, \pi \mid x_{-i}^*) = 0$. However, based on the Implicit Function Theorem, extracting the derivative $\mathbf{D}_\pi x_i^*$ from $\mathbf{D}_\pi z_i^*$ requires that the matrix $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ be invertible. The necessary condition for this to hold is that the set $\Gamma_i^\pi(x_i^*, \lambda_i^*)$ be empty. Namely, observe that $\Gamma_i^\pi(x_i^*, \lambda_i^*) \neq \emptyset$ implies that there would exist a zero row in $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$, hence making it singular. On the other hand, the sufficient condition for the Implicit Function Theorem to hold directly depends on the structure of the game resulting from the nature of the application and the NE computed prior to the leader's strategy update step. To ensure this, in Section 3.3, we reorganize the constraints of the original best-response optimization problem to form an equivalent one whose KKT map $l_i(z_i, \pi \mid x_{-i}^*)$ yields invertible $\mathbf{D}_{z_i} l_i(z_i, \pi \mid x_{-i}^*)$.

3.3. Equivalent best-response optimization problem

For any $\pi \in \mathcal{P}$, let the unique v-NE of the lower-level game $G^0(\mathcal{I}; \pi)$ be $x^* \in \mathcal{X}(\pi)$. For every follower $i \in \mathcal{I}$, its NE strategy $x_i^* \in \mathcal{X}_i(\pi)$ explicitly provides information on what inequality constraints are active for a particular leader's strategy. Let $\mathcal{A}_i(x_i^*)$ represent the set of all active inequality constraints at x_i^* , i.e.,

$$\mathcal{A}_i(x_i^*) := \{j \in [1, m_{\text{inq},i}] \cap \mathbb{N} \mid \mathbf{g}_i^{\text{inq}}(x_i^*, \pi)_j = 0\}. \quad (10)$$

Consequently, let the complement of $\mathcal{A}_i(x_i^*)$ be

$$\mathcal{A}_i^\dagger(x_i^*) = ([1, m_{\text{inq},i}] \cap \mathbb{N}) \setminus \mathcal{A}_i(x_i^*). \quad (11)$$

If $|\mathcal{A}_i(x_i^*)| = m_{\text{act},i} > 0$, then we can define

$$\bar{\mathbf{g}}_i^{\text{inq}}(x_i, \pi) = \text{col}(\mathbf{g}_i^{\text{inq}}(x_i, \pi)_{j \in \mathcal{A}_i(x_i^*)}), \quad (12)$$

$$\underline{\mathbf{g}}_i^{\text{inq}}(x_i, \pi) = \text{col}(\mathbf{g}_i^{\text{inq}}(x_i, \pi)_{j \in \mathcal{A}_i^\dagger(x_i^*)}), \quad (13)$$

that effectively split the inequality constraints into a set of active and inactive ones. This allows us to formulate an auxiliary best-response optimization problem equivalent to (1) whose corresponding set $\Gamma_i^\pi(x_i^*, \lambda_i^*)$ is empty.

Lemma 2. Let *Standing Assumptions 1–5* hold and $x^* \in \mathcal{X}(\pi)$ be the unique NE of the game $G^0(\mathcal{I}; \pi)$ for some $\pi \in \mathcal{P}$. Moreover, let $\mathcal{A}_i(x_i^*)$ and $\mathcal{A}_i^\dagger(x_i^*)$ be defined as (10) and (11) and $|\mathcal{A}_i(x_i^*)| = m_{\text{act},i} \neq 0$. If $\bar{\mathbf{g}}_i^{\text{inq}}(x_i, \pi)$ and $\underline{\mathbf{g}}_i^{\text{inq}}(x_i, \pi)$ are given by (12) and (13),

then $x_i^* \in \mathcal{X}_i(\pi)$ solves the best-response problem $G_i^0(\pi, x_{-i}^*)$ given by (1) if and only if it solves the surrogate problem

$$\bar{G}_i^0(\pi, x_{-i}^*) := \left\{ \begin{array}{l} \min_{x_i \in \mathbb{R}^{\text{mf}}} J_i(x_i, x_{-i}^*, \pi) \\ \text{s. t. } \bar{\mathbf{g}}_i^{\text{inq}}(x_i, \pi) \leq \mathbf{0}_{m_{\text{inq},i} - m_{\text{act},i}} \\ \underline{\mathbf{g}}_i^{\text{eq}}(x_i, \pi) = \mathbf{0}_{m_{\text{eq},i} + m_{\text{act},i}} \end{array} \right\}, \quad (14)$$

where $\bar{\mathbf{g}}_i^{\text{eq}}(x_i, \pi) = [\mathbf{g}_i^{\text{eq}T}(x_i, \pi), \bar{\mathbf{g}}_i^{\text{inq}T}(x_i, \pi)]^T$.

Proof. Observe that both problems are convex, so it suffices to look at their KKT optimality conditions. If $x_i^* \in \mathcal{X}_i(\pi)$ solves (1) for some π , then $\nabla_{x_i} [J_i(x_i, x_{-i}^*, \pi) + \lambda_i^T \mathbf{g}_i^{\text{inq}}(x_i, \pi) + v_i^T \mathbf{g}_i^{\text{eq}}(x_i, \pi)] = 0$ is satisfied for x_i^* and some feasible λ_i^* and v_i^* . We can partition λ^* into $\bar{\lambda}_i$ and $\underline{\lambda}_i$ and rewrite $\nabla_{x_i} [J_i(x_i, x_{-i}^*, \pi) + \bar{\lambda}_i^T \mathbf{g}_i^{\text{inq}}(x_i, \pi) + v_i^T \mathbf{g}_i^{\text{eq}}(x_i, \pi) + \underline{\lambda}_i^T \bar{\mathbf{g}}_i^{\text{inq}}(x_i, \pi)] = 0$. However, this is exactly the KKT stationarity condition of the surrogate best-response problem (14) for $\bar{v}_i^T = [v_i^T, \bar{\lambda}_i^T]$. Since the primal and dual feasibility conditions are equivalent, the proof is completed. \square

We can now postulate the following results regarding the differentiability of the KKT mapping $l_i(z_i, \pi \mid x_{-i}^*)$.

Theorem 2. Let *Standing Assumptions 1–5* hold and $x^* \in \mathcal{X}(\pi)$ be the unique NE of the game $G^0(\mathcal{I}; \pi)$ for some $\pi \in \mathcal{P}$. Let the auxiliary best-response optimization problem $\bar{G}_i^0(\pi, x_{-i}^*)$ be defined as in Lemma 2, $\hat{z}_i = (x_i^*, \lambda_i^*, v_i^*)$ be its solution and $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ be defined as in Theorem 1. If $\mathbf{D}_{x_i} \nabla_{x_i} \mathcal{L}_i(\hat{z}_i, \pi) > 0$ and $\mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{eq}}(x_i, \pi)$ has full row rank for \hat{x}_i , then the matrix $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ associated with $\bar{G}_i^0(\pi, x_{-i}^*)$ is invertible and the Jacobian $\mathbf{D}_\pi x_i^*$ is given by

$$\mathbf{D}_\pi x_i^* = -\Sigma_1^{-1} [\Sigma_3 - \Sigma_2^T (\Sigma_2 \Sigma_1^{-1} \Sigma_2^T)^{-1} (\Sigma_2 \Sigma_1^{-1} \Sigma_3 - \Sigma_4)],$$

where $\Sigma_1 = \mathbf{D}_{x_i} \nabla_{x_i} \mathcal{L}_i$, $\Sigma_2 = \mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{eq}}$, $\Sigma_3 = \mathbf{D}_\pi \nabla_{x_i} \mathcal{L}_i$ and $\Sigma_4 = \mathbf{D}_\pi \underline{\mathbf{g}}_i^{\text{eq}}$ are all evaluated at \hat{z}_i, π, x_{-i}^* .

Proof. We start by noting that $\bar{\mathbf{g}}_i^{\text{inq}}(x_i^*, \pi) < 0$ guarantees that $\lambda_i^* = \mathbf{0}$ due to complementary slackness, and hence $\bar{\Gamma}_i^\pi = \emptyset$. In order to prove invertibility of $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$, we invoke Lemma 3 listed in the Appendix. Namely, we can partition $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ into blocks M_1, M_2, M_3 and M_4 , evaluated at \hat{z}_i, π and x_{-i}^* , such that $M_1 = \mathbf{D}_{x_i} \nabla_{x_i} \mathcal{L}_i(z_i, \pi)$,

$$M_2^T = \begin{bmatrix} \mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{inq}}(x_i, \pi) \\ \mathbf{D}_{x_i} \underline{\mathbf{g}}_i^{\text{eq}}(x_i, \pi) \end{bmatrix}, \quad M_3 = \begin{bmatrix} \mathbf{0} \\ \mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{eq}}(x_i, \pi) \end{bmatrix},$$

$$M_4 = \begin{bmatrix} \text{Dg}(\underline{\mathbf{g}}_i^{\text{inq}}(x_i, \pi)) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

For \hat{z}_i, π, x_{-i}^* , the Schur complement of M_1 is given by

$$\text{Sh}(M_1) := \begin{bmatrix} \text{Dg}(\bar{\mathbf{g}}_i^{\text{inq}}) & \mathbf{0} \\ \star & -\mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{eq}} M_1^{-1} \mathbf{D}_{x_i}^T \bar{\mathbf{g}}_i^{\text{eq}} \end{bmatrix}.$$

Since $\bar{\mathbf{g}}_i^{\text{inq}}(x_i^*, \pi)$ encompasses inactive inequality constraints, we have that $\text{Dg}(\bar{\mathbf{g}}_i^{\text{inq}}(x_i^*, \pi)) < 0$. Similarly, because $\mathbf{D}_{x_i} \nabla_{x_i} \mathcal{L}_i(\hat{z}_i, \pi) > 0$ and $\mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{eq}}(x_i, \pi)$ has full row rank, we have that $\mathbf{D}_{x_i} \bar{\mathbf{g}}_i^{\text{eq}} M_1^{-1} \mathbf{D}_{x_i}^T \bar{\mathbf{g}}_i^{\text{eq}} > 0$, making $\text{Sh}(M_1)$, and hence $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$, nonsingular. Moreover, based on Theorem 1, we have $\mathbf{D}_\pi x_i^* = -\bar{M}_1 \Sigma_3 - \bar{M}_2 [\mathbf{0}^T, \Sigma_4^T]^T$, where Lemma 3 gives $\bar{M}_1 = \Sigma_1^{-1} (\mathbb{I} - \Sigma_2^T (\Sigma_2 \Sigma_1^{-1} \Sigma_2^T)^{-1} \Sigma_2 \Sigma_1^{-1})$ and $\bar{M}_2 = -\Sigma_1^{-1} [\star, -\Sigma_2^T (\Sigma_2 \Sigma_1^{-1} \Sigma_2^T)^{-1}]$. Direct computation of the right-hand side completes the proof. \square

Theorem 2 offers two general conditions that can be used to assess the invertibility of $\mathbf{D}_{z_i} l_i(\hat{z}_i, \pi \mid x_{-i}^*)$ and, consequently, establish the well-posedness of the Jacobian $\mathbf{D}_\pi x_i^*$. This essentially involves confirming the typical structural characteristics of

Algorithm 1 Finding leader's optimal strategy

```

1: Input:  $\gamma, \beta, \bar{s}, \delta, \varepsilon, T$ 
2: Output:  $\pi$ 
3:  $\pi_0 = \text{Initialize}()$ ;
4: for  $t \leftarrow 0$  to  $T$  do
5:    $x_{\pi_t}^* = \text{ComputeVariationalNE}(\pi_t)$ ;
6:   for  $i \in \mathcal{I}$  do                                      $\triangleright$  In parallel
7:     Define  $\underline{g}_i^{\text{inq}}(x_{\pi_t, i}, \pi_t), \bar{g}_i^{\text{eq}}(x_{\pi_t, i}, \pi_t)$ ;
8:     Obtain  $\mathbf{D}_\pi x_{\pi_t, i}^*$  using Theorem 1 on  $\bar{G}_i^0$ ;
9:   end for
10:  Leader:
11:    $\frac{dJ_L(\cdot)}{d\pi} = \frac{\partial J_L(\cdot)}{\partial \pi} + \sum_{i \in \mathcal{I}} \mathbf{D}_\pi^T x_i^* \frac{\partial J_L^L(\cdot)}{\partial x_i^*}$ ;
12:    $s_t = \text{ArmijoStep}\left(\beta, \bar{s}, \delta, \pi_t, \frac{dJ_L(\cdot)}{d\pi}\right)$ ;
13:    $\pi_{t+1} = \pi^+(\pi_t, s_t)$ ;
14: end for

```

the followers' cost functions and constraint sets for commonly encountered instances of Stackelberg games. On the other hand, the closed form of the Jacobian is a direct consequence of Lemma 3 and shows that the Jacobian retains constant functional form during the segments of the leader's update procedure with the same sets of active inequality constraints. In the following section, we will further discuss the applicability of Theorem 2 in particular cases. However, before we proceed, we will first present the formal convergence guarantees for the more general case.

Theorem 3. *Let the Stackelberg game be defined as (3) under Standing Assumptions 1–5. At every update step $t \in \mathbb{N}$ of the leader, let $x_t^* \in \mathcal{X}(\pi_t)$ be the unique v-NE of the lower level game and the surrogate best-response optimization problem of the i th follower be defined as in Lemma 2. If the sequence $\{\pi_t\}$ generated by the projected gradient descent method defined by (6) and (7) fulfills the conditions of Theorem 2, then it holds that $\lim_{t \rightarrow +\infty} J_L(x_{\pi_{t+1}}^*, \pi_{t+1}) - J_L(x_{\pi_t}^*, \pi_t) = 0$ and every limit point of $\{\pi_t\}$ is stationary.*

Proof. First, note that based on the Armijo rule, the sequence $\{J_L(x_{\pi_t}^*, \pi_t)\}_{t=1}^\infty$ is monotonically nonincreasing. Because $J_L(x_\pi^*, \pi)$ is continuous in $z^T = [(x_\pi^*)^T, \pi^T]$ and $\cup_{\pi \in \mathcal{P}} \mathcal{X}(\pi) \times \mathcal{P}$ is compact, there exists $J_L^{\min} \in \mathbb{R}$ such that $J_L(x_\pi^*, \pi) \geq J_L^{\min}$ for all $z \in \cup_{\pi \in \mathcal{P}} \mathcal{X}(\pi) \times \mathcal{P}$. Since the sequence $\{J_L(x_{\pi_t}^*, \pi_t)\}_{t=1}^\infty$ is monotonically nonincreasing and bounded, it converges to a finite value implying $\lim_{t \rightarrow +\infty} [J_L(x_{\pi_{t+1}}^*, \pi_{t+1}) - J_L(x_{\pi_t}^*, \pi_t)] = 0$. Since Theorem 2 guarantees that $J_L(x_\pi^*, \pi)$ is continuously differentiable at every $\pi \in \mathcal{P}$, every limit point of $\{\pi_t\}$ is stationary based on Bertsekas (1999, P2.3.3). \square

The complete iterative procedure for finding a l-SE is outlined in Algorithm 1. The Jacobian computations required for the upper-level update can be fully parallelized, making the complexity of this step independent of the number of players. As such, the distributed nature of the process is completely driven by the subproblem of computing the v-NE for a given leader's strategy. It is known that for lower-level games with a L_μ -strongly monotone and L_l -Lipschitz pseudo-gradient, the v-NE can be found as a fixed-point of the projected pseudo-gradient mapping (Facchinei & Pang; Maljkovic et al., 2023a; Paccagnan et al., 2019; Paccagnan, Kamgarpour, & Lygeros, 2016), with a linear convergence rate depending on L_μ and L_l (Bauschke & Combettes, 2017). It is important to note that in the current setting, Algorithm 1 does not offer convergence guarantees when an equilibrium-seeking mechanism is used, meaning that the lower-level NE must be

fully attained before each update of the leader's decision. In Section 5, however, we present a numerical analysis suggesting that Algorithm 1 can still identify the optimal leader's strategy in certain scenarios, even when the inner procedure for computing the lower-level Nash equilibrium is truncated after a fixed number of iterations.

In terms of the convergence rate of Algorithm 1, within the context of classical optimization, projected gradient descent algorithms using the Armijo step-size rule typically exhibit a convergence rate that is no better than linear (Bertsekas, 1999, Ch. 2.3). Therefore, in more general bi-level setups with equilibrium constraints as described by (3), it becomes increasingly difficult to guarantee a specific convergence rate due to the lack of a closed-form relationship between the lower-level v-NE and the leader's decision. Interestingly, for a special case of games introduced in Definition 2, involving only affine equality constraints, we can establish such a connection under certain common structural assumptions, and hence formally prove linear convergence.

In essence, Section 4 is entirely devoted to Quadratic Aggregative games with general polytopic constraints that are in line with Definition 2. Apart from demonstrating linear convergence in a restricted scenario, for the general case with \mathcal{P} that is also polytopic, we show how the communication hub, i.e., a central aggregator entity discussed in Maljkovic et al. (2023a), Paccagnan et al. (2019), Paccagnan, Kamgarpour, and Lygeros (2016), also supports designing a decentralized warm-start procedure that may help mitigate issues arising from an unfavorable initial value of the leader's action $\pi_0 \in \mathcal{P}$.

4. Quadratic aggregative Stackelberg games

We focus on a particular instance of games in Definition 2 where $\pi_i = \pi_j = \pi$ for all $i, j \in \mathcal{I}$, and the lower-level agents minimize a quadratic cost of the form

$$J_i = \frac{1}{2} x_i^T P_i x_i + x_i^T Q_i \sigma(x_{-i}) + r_i^T x_i + x_i^T S_i \pi, \quad (15)$$

under π -parametrized constraints given by $\underline{g}_i^{\text{inq}}(x_i, \pi) = G_i(\pi)x_i - b_i(\pi)$ and $\bar{g}_i^{\text{eq}}(x_i, \pi) = A_i(\pi)x_i - b_i(\pi)$. In light of Theorem 2, if $x^* \in \mathcal{X}(\pi)$ denotes the v-NE of the lower level game and $|\mathcal{A}_i(x_i^*)| \neq 0$, we proceed to formulate the surrogate best-response optimization problem by letting $\bar{G}_i(\pi) \in \mathbb{R}^{\text{mact}, i \times \text{mf}}$ be a matrix whose rows are the rows of $G_i(\pi)$ listed in $\mathcal{A}_i(x_i^*)$. Moreover, we let $\underline{G}_i(\pi)$ encompass all the remaining rows of $G_i(\pi)$ and decompose the vector $h_i(\pi)$ into $\bar{h}_i(\pi)$ and $\underline{h}_i(\pi)$ such that $\bar{G}_i(\pi)x_i^* = \bar{h}_i(\pi)$ and $\underline{G}_i(\pi)x_i^* < \underline{h}_i(\pi)$. For the cost (15), the conditions of Theorem 2 ensuring the invertibility of

$$\mathbf{D}_{z_i} l_i = \begin{bmatrix} -P_i & \mathbf{0} & \bar{A}_i^T(\pi) \\ \mathbf{0} & \text{Dg}(\underline{G}_i(\pi)x_i^* - \underline{h}_i(\pi)) & \mathbf{0} \\ \bar{A}_i(\pi) & \mathbf{0} & \mathbf{0} \end{bmatrix},$$

related to $\bar{G}_i^0(\pi, x_{-i}^*)$ reduce to $\mathbf{D}_{x_i} \nabla_{x_i} L(\hat{z}_i, \pi) = P_i \succ 0$ and ensuring that $\bar{A}_i(\pi) = [A_i^T(\pi), \bar{G}_i^T(\pi)]^T$ is full row rank. Note that the latter can be easily accounted for during the construction step of the surrogate best-response problem. When designing $\bar{A}_i(\pi)$, it suffices to exclude the active inequality constraints for which there already exists a linearly dependent equality constraint.

As mentioned earlier, quadratic aggregative games allow us to tackle the initialization problem to a certain extent. Our previous work (Maljkovic et al., 2023c) empirically demonstrated that by adjusting the $\pi_0 \in \mathcal{P}$ value used to initiate the iterative update procedure of the leader, we can give rise to local Stackelberg equilibria of varying quality with respect to the value of the leader's objective. While it is possible to sample the leader's action space and repeat the complete iterative process multiple

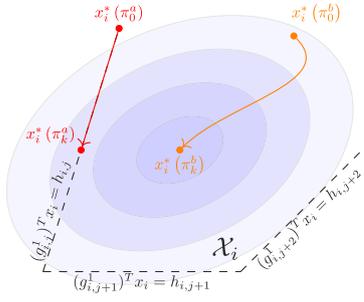


Fig. 2. Illustrative example of the lower-level NE evolution when the projected gradient descent algorithm is initialized by π_0^a , i.e., such that the NE is on the boundary of the feasible set \mathcal{X}_i (red), and by π_0^b , i.e., such that the NE is in the interior (orange). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

times to find a better solution, opting for initial values π_0 that immediately lead to a v-NE that causes certain inequality constraints to become active can unnecessarily hinder the subsequent steps of the procedure. As illustrated in Fig. 2, since the Jacobian is calculated with respect to the surrogate best-response problem, poorly choosing the initial value π_0 means that we effectively start with a lower-level game where each follower has more equality constraints than originally postulated. Although later on we do not have precise control over the trajectory of the leader's iterative procedure, the structure of the analyzed games allows to increase the flexibility of the algorithm by at least providing an adequate point π_0 in the first iteration of the algorithm. More specifically, if \mathcal{P} and $\mathcal{X}_i(\pi)$ are polytopes, with $\mathcal{X}_i(\pi)$ being implicitly governed by a scaled linear combination of π and x_i , then we can design a decentralized method for computing a feasible point π_0 , should such a point exist, that results in a v-NE of the lower-level game that renders no inequality constraint active.

Finally, when the sets $\mathcal{X}_i(\pi)$ encompass only equality constraints, then we can show that Algorithm 1 achieves linear convergence for strongly convex leader objectives.

4.1. Decentralized initialization procedure for agents with polytopic actions spaces

As mentioned earlier, a decentralized initialization procedure can be devised for specific configurations of the agents' constraint sets. Before we go deeper into details about the procedure, we summarize all structural requirements on \mathcal{P} and $\mathcal{X}_i(\pi)$ in the following assumption.

Assumption 1. Sets $\mathcal{P} \subset \mathbb{R}^{m_L}$ and $\mathcal{X}_i(\pi) \subset \mathbb{R}^{m_F}$ are given by $\mathcal{P} := \{\pi \mid \mathbf{A}_\pi \pi = \mathbf{b}_\pi \wedge \mathbf{G}_\pi \pi \leq \mathbf{h}_\pi\}$ and $\mathcal{X}_i(\pi) := \{x_i \mid A_i x_i + A_i^\pi \pi = b_i \wedge G_i x_i + G_i^\pi \pi \leq h_i\}$ that are assumed to be bounded.

To find a $\pi_0 \in \mathcal{P}$ that yields an interior point v-NE $x^* \in \mathcal{X}(\pi_0)$, we are essentially interested in ensuring that such a π_0 entails existence of a positive slack vector δ_i such that $G_i x_i^* + G_i^\pi \pi_0 + \delta_i \leq h_i$. However, as we hope to avoid imposing information exchange regarding the personal constraint sets of the followers, we anticipate a consensus-based mechanism to compute π_0 .

For every $i \in \mathcal{I}$, we look at the KKT conditions of the best-response optimization problem and define $\psi_i \in \mathbb{R}^{m_{\psi_i}}$, with $m_{\psi_i} = Nm_F + m_L + m_{\text{eq},i} + m_{\text{in},i}$ and $\psi_i := [\tilde{\mathbf{x}}_i^T, \tilde{\mathbf{p}}_i^T, v_i^T, \delta_i^T]^T$, where $\tilde{\mathbf{x}}_i = \text{col}(\tilde{x}_i^j)_{j \in \mathcal{I}} \in \mathcal{X}(\pi_0)$ denotes the i th follower's local copy of the complete v-NE x^* , $\tilde{\mathbf{p}}_i \in \mathcal{P}$ being the local copy of the vector π_0 , v_i being the Lagrangian multiplier associated with equality constraints of the best response optimization problem and δ_i is the slack vector that we are looking for. Furthermore, let $\Lambda_{\tilde{x}_i}$,

$\Lambda_{\tilde{p}_i}$, Λ_{δ_i} and Λ_i be selection matrices such that $\Lambda_{\tilde{x}_i} \psi_i = \tilde{x}_i^i$, $\Lambda_{\tilde{p}_i} \psi_i = \tilde{\mathbf{p}}_i$, $\Lambda_{\delta_i} \psi_i = \delta_i$ and $\Lambda_i \psi_i = [\tilde{\mathbf{x}}_i^T, \tilde{\mathbf{p}}_i^T]^T$. We can now postulate a necessary and sufficient feasibility test based on linear programming.

Theorem 4 (Internal v-NE Feasibility Check). Let the Stackelberg game be defined as (3) under Standing Assumptions 1–5, the structural Assumption 1 and objective functions given by (15). There exists a $\pi_0 \in \mathcal{P}$ such that the corresponding v-NE, $x^* \in \mathcal{X}(\pi_0)$, of the lower-level game $G^0(\mathcal{I}; \pi_0)$ renders no inequality constraint active if and only if there exists $\varepsilon > 0$ such that the following linear optimization problem has a solution

$$\begin{aligned} \text{minimize} \quad & - \sum_{i \in \mathcal{I}} \mathbf{1}^T \Lambda_{\delta_i} \psi_i \\ \text{subject to} \quad & \Lambda_i \psi_i - \beta = \mathbf{0} \end{aligned} \quad (16a)$$

$$[W_i \ S_i^T \ A_i^T \ \mathbf{0}] \psi_i + r_i = \mathbf{0} \quad (16b)$$

$$(A_i \Lambda_{\tilde{x}_i} + A_i^\pi \Lambda_{\tilde{p}_i}) \psi_i = b_i \quad (16c)$$

$$(G_i \Lambda_{\tilde{x}_i} + G_i^\pi \Lambda_{\tilde{p}_i} + \Lambda_{\delta_i}) \psi_i \leq h_i, \quad (16d)$$

$$- \Lambda_{\delta_i} \psi_i \leq -\varepsilon \mathbf{1} \quad (16e)$$

$$\mathbf{A}_\pi \Lambda_{\tilde{p}_i} \psi_i = \mathbf{b}_\pi,$$

$$\mathbf{G}_\pi \Lambda_{\tilde{p}_i} \psi_i \leq \mathbf{h}_\pi$$

where $W_i = [P_i, \ \mathbf{1}_{N-1}^T \otimes \mathbb{I}]$.

Proof. Firstly, observe that the dummy variable $\beta \in \mathbb{R}^{m_\beta}$, where $m_\beta = Nm_F + m_L$, ensures through (16a) that all the followers have equal local copies of the pricing vector and the v-NE. Moreover, since $\varepsilon > 0$, (16d) and (16e) ensure that any optimal solution of this problem renders no inequality constraint active. Due to complementarity slackness, we search for solutions where the dual variables satisfy $\lambda_i = \mathbf{0}$. Hence, (16b) represents the stationarity condition of the KKT system for the convex best-response optimization problem. By adding (16c) and (16d), we form a complete set of KKT optimality conditions, so any optimal β corresponds to a π_0 for which the v-NE is an interior point. \square

Firstly, we note that (16) could have been posed as a feasibility problem for simplicity. However, we opt for the proposed functional form as it places the attained v-NE further away from the boundaries. Clearly, the optimal solution β^* encodes a viable π_0 and its corresponding v-NE, i.e., $\beta^* = [(x_{\pi_0}^*)^T, \ \pi_0^T]^T$. Thanks to the separable objective function and constraints of the optimization problem (16), we can preserve the privacy of the followers and solve (16) in a decentralized fashion through the consensus alternating direction method of multipliers (ADMM) (Boyd, Parikh, Chu, Peleato, & Eckstein, 2011). Here, β acts as a global variable to be shared among all the followers and is the only one that needs to be updated in a centralized manner, e.g., in case studies presented in Maljkovic et al. (2023a, 2023c), Paccagnan et al. (2019), Paccagnan, Kamgarpour, and Lygeros (2016), this could be served by the same central entity required for computing the v-NE. If we let the polytope $\Omega_{\psi_i} := \{\psi_i \in \mathbb{R}^{m_{\psi_i}} \mid M_i^1 \psi_i = v_i^1 \wedge M_i^2 \psi_i \leq v_i^2\}$ encode all the constraints of (16) except (16a), then the augmented Lagrangian of (16) is

$$\begin{aligned} \mathcal{L}_\rho(\{\psi_i\}, \beta, \{y_i\}) = & \sum_{i \in \mathcal{I}} -\mathbf{1}^T \Lambda_{\delta_i} \psi_i + I_{\Omega_{\psi_i}}(\psi_i) \\ & + y_i^T (\Lambda_i \psi_i - \beta) + \frac{\rho}{2} \|\Lambda_i \psi_i - \beta\|_2^2, \end{aligned}$$

where $I_{\Omega_{\psi_i}}(\psi_i)$ denotes the indicator function and $\rho > 0$ is an a priori chosen parameter. The consensus ADMM consists of

repeating the following three steps

$$\begin{aligned}\psi_i^{k+1} &= \arg \min_{\psi_i \in \Omega_{\psi_i}} \mathcal{L}_\rho(\psi_i, \{\psi_j^k\}_{j \in \mathcal{I} \setminus \{i\}}, \beta^k, \{y_i^k\}), \\ \beta^{k+1} &= \arg \min_{\beta \in \mathbb{R}^{m_\beta}} \mathcal{L}_\rho(\{\psi_i^{k+1}\}, \beta, \{y_i^k\}), \\ y_i^{k+1} &= y_i^k + \rho(A_i \psi_i^{k+1} - \beta^{k+1}).\end{aligned}\quad (17)$$

Due to the separability of the augmented Lagrangian, solving the N convex quadratic optimization problems for updating individual ψ_i can be done in parallel. The same holds for updating the dual variables y_i . On the other hand, the unconstrained quadratic minimization problem to be solved to obtain β^{k+1} yields

$$\beta^{k+1} = \frac{1}{N} \left[\frac{1}{\rho} \sum_{i \in \mathcal{I}} y_i^k + \sum_{i \in \mathcal{I}} A_i \psi_i^{k+1} \right],$$

and requires that the followers communicate their updated local estimate of π_0 and the corresponding v-NE, both encoded in ψ_i^{k+1} , to the central aggregator who will then update the consensus variable β . Formal convergence guarantees are given in the following theorem.

Theorem 5. Let $\Delta_i^k = A_i \psi_i^k - \beta^k$ denote the residual at each iteration of the consensus ADMM given by (17). If A_i is full row rank then $\Delta_i^k \rightarrow 0$ when $k \rightarrow \infty$.

Proof. We aim to directly invoke [Chen, Sun, and Toh \(2017, Th 4.1\)](#). Namely, for [Chen et al. \(2017, Th 4.1\)](#) to hold, we first observe that the extended, real-valued function $\bar{f} = \sum_{i \in \mathcal{I}} -\mathbf{1}^T A_i \psi_i + I_{\Omega_{\psi_i}}(\psi_i)$ is closed, proper and convex. Secondly, we need to make sure that the solution set of (16) is bounded. For this, it suffices to prove that Ω_{ψ_i} is bounded as then $\beta = A_i \psi_i$ is bounded as well. Clearly, both $\tilde{\mathbf{x}}_i \in \mathcal{X}(\pi)$ and $\tilde{\mathbf{p}}_i \in \mathcal{P}$ are bounded due to imposed assumptions. From (16b), we have $A_i^T v_i = \gamma_i$ for $\gamma_i := -r_i - W_i \tilde{\mathbf{x}}_i - S_i^T \tilde{\mathbf{p}}_i$. If γ_i^j denotes the j th element of the vector, then for every $j \in [1, m_F] \cap \mathbb{N}$ there exist $\underline{\gamma}_i^j, \bar{\gamma}_i^j \in \mathbb{R}$ such that $\underline{\gamma}_i^j \leq \gamma_i^j \leq \bar{\gamma}_i^j$ since both $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{p}}_i$ are bounded. If we set $\gamma_i^{\min} = \min_j \underline{\gamma}_i^j$ and $\gamma_i^{\max} = \max_j \bar{\gamma}_i^j$, then $\gamma_i^{\min} \mathbf{1} \leq A_i^T v_i \leq \gamma_i^{\max} \mathbf{1}$. If A_i is full row rank, then the polytope $\gamma_i^{\min} \mathbf{1} \leq A_i^T v_i \leq \gamma_i^{\max} \mathbf{1}$ is bounded. Similarly, we can establish that δ_i is bounded based on (16d) and (16e), which completes the proof. \square

Before delving into the case studies that demonstrate the performance of the proposed method and its warm-start extension, we examine one final, more restricted, version of the games described in [Definition 2](#). Namely, we aim to show that when the sets $\mathcal{X}_i(\pi)$ in Quadratic Aggregative games include only equality constraints, sufficient conditions on $J_L(x, \pi)$ can be established such that the overall procedure achieves linear convergence rate.

4.2. On convergence rate for quadratic aggregative games with equality constraints

We start by observing that for sets $\mathcal{X}_i(\pi)$ of the form $\mathcal{X}_i(\pi) := \{x_i \mid A_i x_i = b_i - A_i^T \pi\}$, for every $\pi \in \mathcal{P}$, a closed-form description of the lower-level game's v-NE exists, provided that certain structural properties hold.

Proposition 2. Let the Stackelberg game be defined as (3) under [Standing Assumptions 1–3](#), and in line with [Definition 2](#). Moreover, let the followers' constraint sets be given by $\mathcal{X}_i(\pi) := \{x_i \mid A_i x_i = b_i - A_i^T \pi\}$, with A_i that is full row rank for every $i \in \mathcal{I}$. Then, the unique NE of $G^0(\mathcal{I}; \pi)$ is given by $x^*(\pi) = \mathbf{T}\pi + \tau$ for some specific $\mathbf{T} \in \mathbb{R}^{N m_F \times m_L}$ and $\tau \in \mathbb{R}^{N m_F}$.

Proof. If a pair $(x^*(\pi), v^*)$ constitutes the NE of $G^0(\mathcal{I}; \pi)$, then it has to solve the joint KKT system

$$\begin{bmatrix} \bar{K} & \bar{A}^T \\ \bar{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} x^* \\ v^* \end{bmatrix} + \begin{bmatrix} \bar{S} \\ \bar{A} \end{bmatrix} \pi = \begin{bmatrix} -\bar{r} \\ \bar{b} \end{bmatrix}, \quad (18)$$

where $\bar{A} = \text{blkdiag}(A_i)_{i \in \mathcal{I}}$, $\bar{S} = \text{vstack}(S_i)_{i \in \mathcal{I}}$, $\bar{A}^T = \text{vstack}(A_i^T)_{i \in \mathcal{I}}$, $\bar{r} = \text{vstack}(r_i)_{i \in \mathcal{I}}$, $\bar{b} = \text{vstack}(b_i)_{i \in \mathcal{I}}$ and $\bar{K} = \text{blkdiag}(P_i - Q_i)_{i \in \mathcal{I}} + \text{vstack}(\mathbf{1}_N^T \otimes Q_i)_{i \in \mathcal{I}}$. Observe now that because each A_i has full row rank so does \bar{A} . Let $Z \in \mathbb{R}^{N m_F \times (N m_F - \bar{m}_{\text{eq}})}$, with $\bar{m}_{\text{eq}} = \sum_{i \in \mathcal{I}} m_{\text{eq}, i}$, be a matrix whose columns form a basis for the null-space of \bar{A} , i.e., $\bar{A}Z = \mathbf{0}$. Moreover, let $Y \in \mathbb{R}^{N m_F \times \bar{m}_{\text{eq}}}$ be any matrix such that $[Y \mid Z]$ is invertible. Then, every x^* that solves (18) can be represented as $x^* = Yx_r + Zx_n$, where x_r is any vector such that $\bar{A}x^* = \bar{b} - \bar{A}^T \pi$, which is guaranteed to exist due to [Standing Assumption 1](#). Similarly, by multiplying the first equation of (18) by Z^T we get $Z^T \bar{K} Z x_n = -Z^T \bar{K} Y x_r - Z^T \bar{S} \pi - Z^T \bar{r}$. Observe that $\bar{A}Y$ is a square invertible matrix so x_r can be expressed as $x_r = \mathbf{T}_r \pi + \tau_r$, where $\mathbf{T}_r = -(\bar{A}Y)^{-1} \bar{A}^T \pi$ and $\tau_r = (\bar{A}Y)^{-1} \bar{b}$. On the other hand, due to [Standing Assumption 3](#), matrix \bar{K} has to be positive definite as the game's pseudo-gradient given by $F(x, \pi) = \bar{K}x + \bar{S}\pi + \bar{r}$ is strongly monotone. Therefore, $Z^T \bar{K} Z > 0$ also holds since Z forms a basis for the null-space of \bar{A} , and we can write $x_n = \mathbf{T}_n \pi + \tau_n$, where $\mathbf{T}_n = -(Z^T \bar{K} Z)^{-1} Z^T (\bar{S} + \bar{K} Y \mathbf{T}_r)$ and $\tau_n = -(Z^T \bar{K} Z)^{-1} Z^T (\bar{r} + \bar{K} Y \tau_r)$. Finally, it is clear that $x^*(\pi) = \mathbf{T}\pi + \tau$ for $\mathbf{T} = Y \mathbf{T}_r + Z \mathbf{T}_n$ and $\tau = Y \tau_r + Z \tau_n$. \square

With [Proposition 2](#) in place, we can proceed to identify a particular setup of bi-level game G_1 that achieves linear convergence when [Algorithm 1](#) is used.

Proposition 3. Let the Stackelberg game be defined such that [Proposition 2](#) holds. Moreover, let the leader's objective $J_L(x, \pi)$ be strongly convex in $p^T = [x^T, \pi^T]$. Then, [Algorithm 1](#) exhibits linear convergence rate.

Proof. If [Proposition 2](#) holds, then $J_L(x^*(\pi), \pi)$ is a composition of a linear map and a strongly convex function, i.e., $J_L(x^*(\pi), \pi) := J_L(\bar{\mathbf{T}}\pi + \bar{\tau})$, where $\bar{\mathbf{T}}^T = [\mathbf{T}^T, \mathbb{I}_{m_L}^T]$ and $\bar{\tau}^T = [\mathbf{0}_{N m_F}^T, \tau^T]$. The chain rule now directly yields $\text{Hess}_{\pi} J_L = \bar{\mathbf{T}}^T \cdot \text{Hess}_{p} J_L(\bar{\mathbf{T}}\pi + \bar{\tau}) \cdot \bar{\mathbf{T}} > 0$, as J_L is strongly convex and $\bar{\mathbf{T}}$ has a full column rank. Hence, under [Proposition 2](#), [Algorithm 1](#) reduces to projected gradient descent with strongly convex objective for which we can directly invoke ([Bauschke & Combettes, 2017, Prop. 26.16](#)) to establish linear convergence rate. \square

[Proposition 3](#) presents a rare instance of (3) where the problem reduces to a standard optimization setting, enabling the use of traditional constrained optimization tools. This provides valuable insight into the algorithm's ability to recover similar properties as those achieved by algorithms that rely on polytopic constraints of the followers, such as in [Grontas et al. \(2024\)](#). For more general setups of G_1 , the inherent non-convexity arising from the equilibrium constraints makes it challenging to establish convergence rate guarantees and is a topic for future research.

In the following section, we will introduce in detail the two numerical case studies showcasing the performance of the main decentralized algorithm and its corresponding warm-start procedure. Additionally, as discussed in [Section 3](#), we will demonstrate the algorithm's behavior when the lower-level procedure for computing the v-NE is halted after a fixed number of iterations, rather than being allowed to fully converge.

5. Numerical examples

We consider two scenarios of a case study in the smart mobility domain previously introduced in [Maljkovic et al. \(2023a\)](#),

2023c). In particular, we analyze a market model depicted in Fig. 3, where ride-hailing companies $\mathcal{I} = \{I_1, I_2, I_3\}$ compete to meet demand requests that are distributed heterogeneously across the city of Shenzhen (Bellocchi, Latora, & Geroliminis, 2021). The city is inherently partitioned into four Voronoi-based regions by the available charging infrastructure that consists of stations $\mathcal{M} = \{M_1, M_2, M_3, M_4\}$ and is controlled by the central authority L through adjustable electricity prices $\pi \in [p_{\min}, p_{\max}]^4$. At a particular point in time, we assume that each company $i \in \mathcal{I}$ wants to recharge its N_i vehicles by distributing them among charging stations \mathcal{M} . Namely, we let the vector $x_i \in \mathcal{X}_i \subseteq \mathbb{R}^4$ denote the strategic decision of company i that describes the fleet split among charging stations, i.e., $\|x_i\|_1 = N_i$ and $x_i^j \geq 0$ represents what fraction of the fleet is to be directed to a particular station $j \in \mathcal{M}$.

The central authority, which can for example be the power-providing company or the government, may have an interest in balancing the demand on the power grid or it might aim to design pricing incentives to enhance coverage and encourage idle taxi drivers to avoid flocking to the more demand-attractive areas. We assume that the nominal prices of charging are encoded in $\pi_{\text{base}} \in \mathbb{R}^4$ and that the central authority is then interested in determining the optimal discount $\Delta\pi \in \mathbb{R}^4$, all while adhering to the total monetary discount budgets $B_i \in \mathbb{R}$ assigned to companies \mathcal{I} based on external subsidies that they receive. Upon the announcement of the pricing vector $\pi := \pi_{\text{base}} - \Delta\pi$, every company operator is interested in minimizing its operational cost under the feasibility constraints imposed by the battery status of its vehicles. Similar to the objectives analyzed in Maljkovic, Nilsson, and Geroliminis (2022), Maljkovic et al. (2023c), Tushar, Saad, Poor, and Smith (2012), the operator's cost is assumed in the form of a sum of three terms, i.e., $J_i(x_i, \sigma(x_{-i}), \pi) = J_i^1(x_i, \sigma(x_{-i})) + J_i^2(x_i) + J_i^3(x_i, \pi)$, where $J_i^1(x_i, \sigma(x_{-i}))$ denotes the expected queuing cost at different charging stations due to their limited capacities, $J_i^2(x_i)$ denotes the negative expected revenue in the regions around charging stations, $J_i^3(x_i, \pi)$ denotes the charging cost and $\sigma(\cdot)$ is defined as in Definition 2. The resulting form is quadratic and given by $J_i(x_i, \sigma(x_{-i}), \pi) = \frac{1}{2}x_i^T P_i x_i + x_i^T Q_i \sigma(x_{-i}) + r_i^T x_i + \pi^T S_i x_i$. On the other hand, we assume that the central authority chooses a desired vehicle distribution vector $\mathcal{Z} \in [0, 1]^4$ satisfying $\|\mathcal{Z}\|_1 = 1$ and plays the game with the ride-hailing companies in an attempt to minimize the cost $J_G(\sigma(x)) = \frac{1}{2}\|\sigma(x) - \mathbf{1}^T n \mathcal{Z}\|_2^2$, with $n = \text{col}(N_i)_{i \in \mathcal{I}}$ being the vector containing the number of vehicles per company that need to be recharged.

Concerning the constraint sets of ride-hailing companies, they encompass information about the number of vehicles that can reach a certain station under a linear battery discharge model and given the current battery level after the rush-hour period simulation. It has been shown in Maljkovic et al. (2023a) that a specifically designed polytopic constraint allows for the consistent matching of each ride-hailing vehicle with precisely one charging station in an attempt to respect the allocation given by the split x . For every $i \in \mathcal{I}$, the matching constraints in accordance with Maljkovic et al. (2023a) are given by $\mathcal{X}_i^m := \{x_i \in \mathbb{R}^4 \mid A_i x_i = b_i \wedge G_i x_i \leq h_i\}$, for some properly chosen A_i, b_i, G_i, h_i . Apart from them, we also account for the limited discount budget B_i through the constraint $\mathcal{X}_i^b(\pi) := \{x_i \in \mathbb{R}^4 \mid (\pi_{\text{base}} - \pi)^T S_i x_i \leq B_i\}$. Hence, for any pricing strategy $\pi \in \mathcal{P}$ and for every $i \in \mathcal{I}$, the resulting constraint set is given by $\mathcal{X}_i(\pi) := \mathcal{X}_i^m \cap \mathcal{X}_i^b(\pi)$. Generally speaking, $\mathcal{X}_i(\pi)$ is a polytopic constraint in x_i but does not comply with the structure proposed in Assumption 1 of Section 4. Therefore, we test two scenarios:

- (1) To illustrate the performance of the algorithm in a more general scenario, e.g., where the sets $\mathcal{X}_i(\pi)$ are not governed by a scaled linear combination of π and x_i , we shift

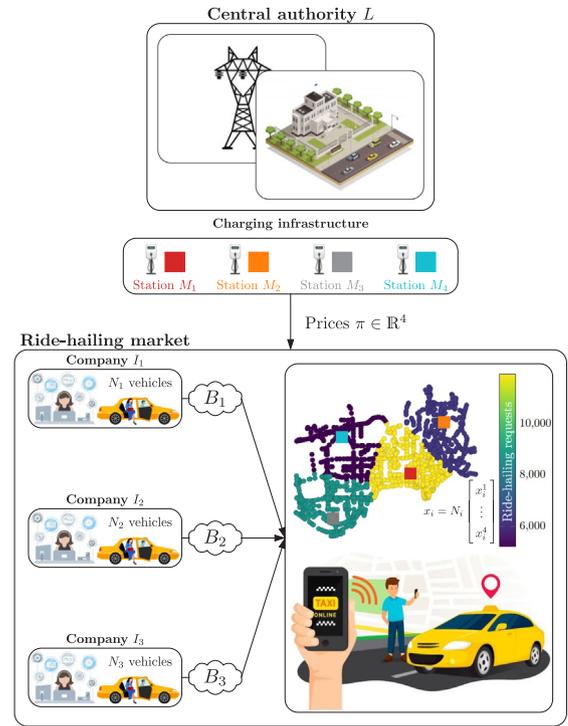


Fig. 3. Illustration of the problem setup with ride-hailing companies $\mathcal{I} = \{I_1, I_2, I_3\}$ operating in a region with charging stations $\mathcal{M} = \{M_1, M_2, M_3, M_4\}$. The central authority L chooses the electricity price $\pi \in \mathcal{P} \subseteq \mathbb{R}^4$ so as to respect the discount budget B_i of each company $i \in \mathcal{I}$.

away from the original setup in Maljkovic et al. (2023c) and assume that the discount budgets are finite, i.e., $B_i < \infty$ for all $i \in \mathcal{I}$;

- (2) To demonstrate the effects of the warm-start procedure and early stopping of the lower-level v-NE computation, we set $B_i = \infty$ for all $i \in \mathcal{I}$, which yields $\mathcal{X}_i^b(\pi) = \mathbb{R}^4$ and gives rise to an identical problem setup as the one analyzed in Maljkovic et al. (2023c).

The number of vehicles per company that want to recharge is given by $n = [194, 181, 157]$ and \mathcal{Z} is chosen to correspond to the total number of requests in each cell. For the analyzed case study, \mathcal{Z} is such that $\mathbf{1}^T n \mathcal{Z} = [198, 103, 144, 87]$ and we set $p_{\min} = 0.0$ and $p_{\max} = 5.0$. For the extensive list of all remaining parameters in the simulation, we refer the reader to Maljkovic et al. (2023a).

5.1. Finite discount budgets

In this case study, the finite discount budgets are given by vector $B = [14000, 13000, 12000]$ and the base price is given by $\pi_{\text{base}} = [5.0, 3.0, 5.0, 3.0]$. Before each update step of the pricing policy, we perform $k_{v\text{-NE}} = 5000$ of the Picard–Banach fixed point iteration procedure to compute the v-NE of the lower-level game (Maljkovic et al., 2023c) for the current value of the pricing vector. For the outer loop, we set the number of iterations to $k_{l\text{-SE}} = 350$ and observe the average duration of one update step of approximately $\tau_{\text{avg}} \approx 0.5$ sec. For the given number of iterations and the pricing vector $\pi_{\text{init}}^1 = [4.0, 2.0, 3.0, 1.0]$ used to initialize the outer loop of the procedure, the system manages to achieve perfect matching with respect to the desired vehicle distribution and attains $J_L(x^*(\pi^*), \pi^*) = 0.022$. This is further supported by plots in Fig. 4. The three upper plots show the evolution of the attained vehicle accumulations at different charging stations, the evolution of the pricing vector π , and the

Table 1
Vehicle distribution and charging prices.

Stations \mathcal{M}	Vehicle distribution		Charging prices	
	$\mathbf{1}^T n \mathcal{Z}$	$\sigma(x^*)$	π_j^*	$\Delta \pi_j^*$
M_1	198	197.81	3.39	1.61
M_2	103	103.04	2.20	0.80
M_3	144	144.09	2.83	2.17
M_4	87	87.06	1.58	1.42

Table 2
Performance comparison for different initializations.

Initial π	Resulting prices at M_j				Attained J_L
	π_1	π_2	π_3	π_4	
π_{init}^1	3.39	2.20	2.83	1.58	0.0222
π_{init}^2	3.57	2.38	3.01	3.00	2000.0
π_{init}^w	4.58	3.36	3.99	2.75	0.0039

corresponding value of the central authority's objective function. The lower three plots demonstrate that no discount budget constraint has been violated during the iterative procedure, i.e., the used discount budget for company I_1 is $B_1^{used} \approx 13523$, for company I_2 is $B_2^{used} \approx 12860$, and for company I_3 is $B_3^{used} \approx 11371$. A full overview of the relevant numerical values is presented in Table 1. It is important to note that the initial value π_{init} has been obtained via sparse grid search as the setup does not comply with the structure of inequality constraints in Assumption 1 of Section 4. Since the complexity of the grid-search procedure grows exponentially in the size of π and polynomially in the granularity of the grid, it is evident that this kind of heuristic is in general not suitable for larger problem sizes. However, for a broad class of bi-level games where the agent's constraints are given as in Assumption 1, we can deploy our iterative warm-up procedure. Therefore, in the following subsection, we shift back our focus to the original setup of Maljkovic et al. (2023c).

5.2. Infinite discount budgets

When discount budgets are infinite for every ride-hailing company, starting the outer loop with the initial value $\pi_{init}^2 = [3.0, 3.0, 3.0, 3.0]$ yields that the value of the central authority's objective converges to $J_L(x^*(\pi^*), \pi^*) = 2001$ (Maljkovic et al., 2023c). As previously discussed, π_{init}^2 already renders certain inequality constraints active which immediately creates a distinction between the original and the surrogate best-response optimization problems. Instead, we let the warm-start procedure with $\rho = 1.0$ run for $k_w = 500$ iterations to obtain the initial pricing vector $\pi_{init}^w = [4.8, 3.3, 3.9, 2.7]$. Starting the outer loop with π_{init}^w induces an interior v-NE in the first iteration and the complete algorithm is later capable of recovering the perfect matching attained when starting from π_{init}^1 . Since the theoretical optimal value for the central authority's objective is zero, the generation of two distinct pricing vectors from the initial states π_{init}^1 and π_{init}^w indicates the general non-uniqueness of the solution in these bi-level games. In Fig. 5, we depict the evolution of the complete algorithm for different initial pricing vectors while Table 2 lists all the relevant numerical values. It is interesting to note that the warm-start procedure provides a significantly smaller starting value of the central authority's objective compared to π_{init}^1 and π_{init}^2 . However, from the perspective of ride-hailing company operators, starting from π_{init}^1 results in more favorable charging prices in terms of the pricing vector's magnitude and hence, the total charging costs.

Finally, we test the robustness of the proposed method by exploring the impact of varying the number of Picard-Banach iterations, $k_{v-NE} \in [10, 5000]$, performed before each update step

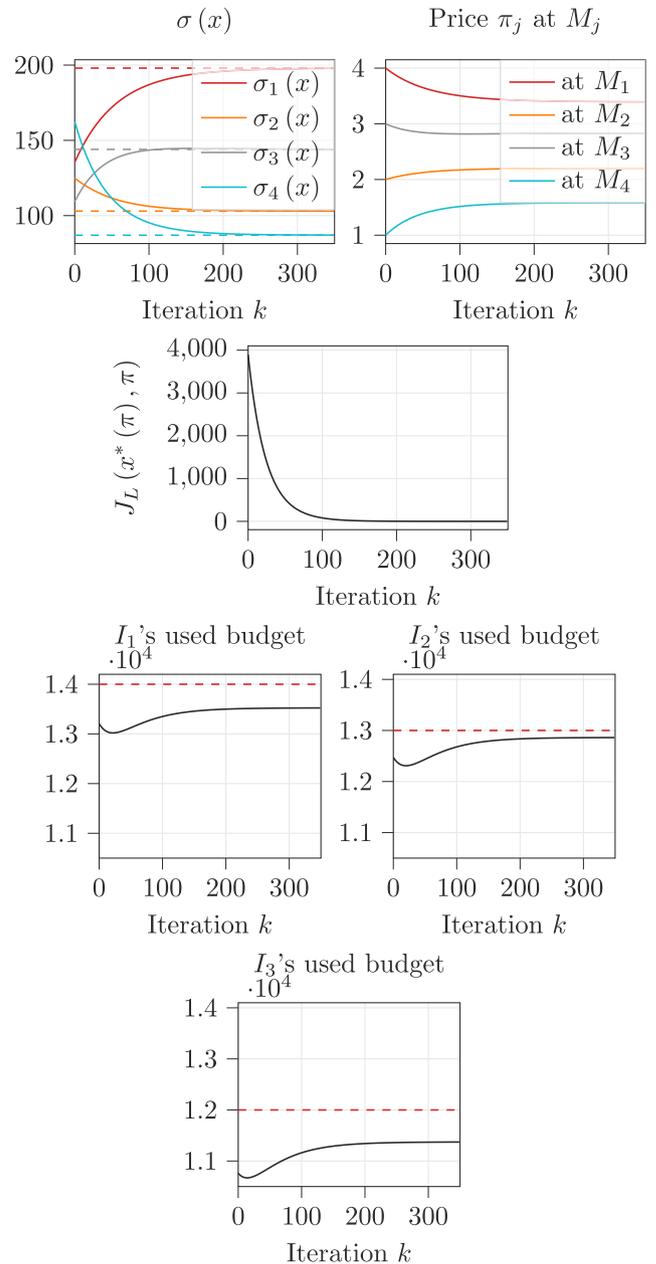


Fig. 4. The plots show the evolution of the total vehicle accumulation at the charging stations $\sigma(x)$, the price of charging π_j at the station M_j , the leader's objective $J_L(x^*(\pi), \pi)$, and the portion of the budgets used at each iteration.

of the outer procedure. Using $\pi_{init}^1 = [4.0, 2.0, 3.0, 1.0]$ as the initial value of the leader's decision and keeping the number of outer loop iterations constant, Fig. 6 illustrates the evolution of the leader's objective. Even without full convergence of the inner loop, the outer loop consistently improves the leader's decision-making, albeit at a slower pace. Notably, $k_{v-NE} = 5000$ achieves the lowest value of the leader's objective significantly faster than the other tested values. However, after 350 iterations, the difference between the objective values attained for $k_{v-NE} = 5000$ and $k_{v-NE} = 2500$ becomes negligible, suggesting also the possibility to reduce the number of inner-loop iterations and still end up in the global minimum. Hence, at least for certain cases, it seems like the method could potentially be extended to incorporate an equilibrium-seeking mechanism on the lower level, which is a promising direction for future research.

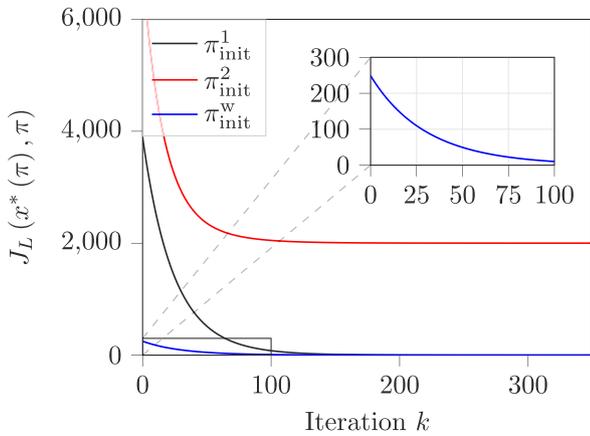


Fig. 5. Evolution of the central authority's objective for different initializations of the upper-level iterative loop.

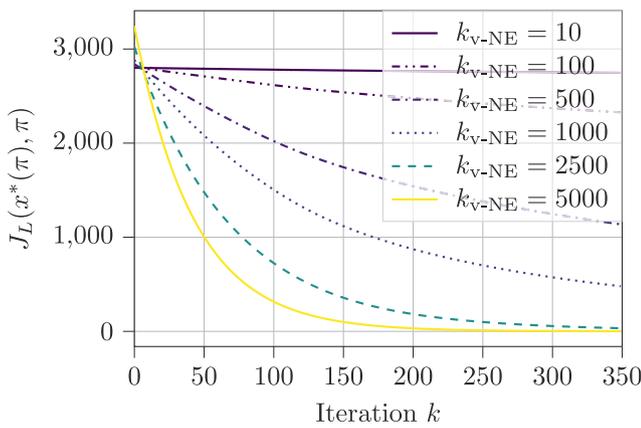


Fig. 6. Evolution of the central authority's objective per iteration for different numbers of inner loop iterations, k_{v-NE} .

6. Conclusion

In this work, we formalized an iterative framework based on the Implicit Function Theorem for tackling the problem of computing the local Stackelberg equilibrium in bi-level games. Apart from generalizing the idea introduced in Maljkovic et al. (2023c) to suit a broader class of games, we address the reported shortcomings of random initialization for a class of quadratic, aggregative games with polytopic constraints. In light of the overall decentralized nature of the approach, we formulate an internal v-NE feasibility test in the form of a linear program that can be efficiently solved by a distributed alternating direction method of multipliers. In addition to theoretical guarantees, we provide an experimental demonstration of the performance improvement in the previously analyzed case study in the smart mobility domain.

In the future, it would be interesting to explore the development of other distributed initialization methods capable of handling convex constraints beyond polyhedral forms. Moreover, the proposed approach holds practical significance, as it can be applied to various real-world problems in domains such as energy management and transportation. In light of the conducted robustness analysis in the numerical section, it would be valuable to investigate whether convergence guarantees could also be established for systems with rapidly changing dynamics, particularly if lower-level agents employ an equilibrium-seeking mechanism.

Appendix

Lemma 3 (Chapter 2.17 of Bernstein (2009)). Let $M_1 \in \mathbb{R}^{p \times p}$, $M_2, M_3^T \in \mathbb{R}^{p \times q}$ and $M_4 \in \mathbb{R}^{q \times q}$. If M_1 is nonsingular, then the inverse

$$M^{-1} := \begin{bmatrix} \bar{M}_1 & \bar{M}_2 \\ \bar{M}_3 & \bar{M}_4 \end{bmatrix} \text{ of } M := \begin{bmatrix} M_1 & M_2 \\ M_3 & M_4 \end{bmatrix}$$

exists if and only if Schur complement of M_1 in M , i.e., $\text{Sh}(M_1) = M_4 - M_3 M_1^{-1} M_2$, is nonsingular. The blocks are given by $\bar{M}_1 = M_1^{-1} + M_1^{-1} M_2 \text{Sh}(M_1)^{-1} M_3 M_1^{-1}$, $\bar{M}_2 = -M_1^{-1} M_2 \text{Sh}(M_1)^{-1}$, $\bar{M}_3 = -\text{Sh}(M_1)^{-1} M_3 M_1^{-1}$ and $\bar{M}_4 = \text{Sh}(M_1)^{-1}$.

References

- Aussel, Didier, Brotcorne, Luce, Lepaul, Sébastien, & von Niederhäusern, Léonard (2020). A trilevel model for best response in energy demand-side management. *European Journal of Operational Research*, 281(2), 299–315.
- Bauschke, Heinz H., & Combettes, Patrick L. (2017). *Convex analysis and monotone operator theory in Hilbert spaces*. Springer New York, NY.
- Bellocchi, Leonardo, Latora, Vito, & Geroliminis, Nikolas (2021). Dynamical efficiency for multimodal time-varying transportation networks. *Scientific Reports*, 11.
- Bengio, Yoshua (2000). Gradient-based optimization of hyperparameters. *Neural Computation*, 12(8), 1889–1900.
- Bernstein, Dennis S. (2009). *Matrix mathematics: Theory, facts, and formulas* (2nd ed.) Princeton: Princeton University Press.
- Bertsekas, D. P. (1999). *Nonlinear programming*. Athena Scientific.
- Bianchi, Mattia, Belgioioso, Giuseppe, & Grammatico, Sergio (2022). Fast generalized Nash equilibrium seeking under partial-decision information. *Automatica*, 136.
- Boyd, Stephen, Parikh, Neal, Chu, Eric, Peleato, Borja, & Eckstein, Jonathan (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers.
- Chen, Liang, Sun, Defeng, & Toh, Kim-Chuan (2017). A note on the convergence of ADMM for linearly constrained convex optimization problems. *Computational Optimization and Applications*, 66.
- Dontchev, Asen L., & Rockafellar, R. Tyrrell (2009). *Implicit functions and solution mappings*. Springer New York, NY.
- Fabiani, Filippo, Tajeddini, Mohammad Amin, Kebriaei, Hamed, & Grammatico, Sergio (2022). Local Stackelberg equilibrium seeking in generalized aggregative games. *IEEE Transactions on Automatic Control*, 67(2), 965–970.
- Facchinei, Francisco, & Pang, Jong-Shi *Finite-dimensional variational inequalities and complementarity problems*. Springer New York, NY.
- Fochesato, Marta, Cenedese, Carlo, & Lygeros, John (2022). A Stackelberg game for incentive-based demand response in energy markets. In *2022 IEEE 61st conference on decision and control* (pp. 2487–2492).
- Fortuny-Amat, José, & McCarl, Bruce (1981). A representation and economic interpretation of a two-level programming problem. *Journal of the Operational Research Society*, 32(9), 783–792.
- Golrezaei, Negin, Jaillet, Patrick, Cheuk Nam Liang, Jason, & Mirrokni, Vahab (2023). Pricing against a budget and ROI constrained buyer. In Francisco Ruiz, Jennifer Dy, & Jan-Willem van de Meent (Eds.), *Proceedings of machine learning research: vol. 206, Proceedings of the 26th international conference on artificial intelligence and statistics* (pp. 9282–9307). PMLR.
- Grazzi, Riccardo, Franceschi, Luca, Pontil, Massimiliano, & Salzo, Saverio (2020). On the iteration complexity of hypergradient computation. In Hal Daumé III, & Aarti Singh (Eds.), *Proceedings of machine learning research: vol. 119, Proceedings of the 37th international conference on machine learning* (pp. 3748–3758). PMLR.
- Grontas, Panagiotis D., Belgioioso, Giuseppe, Cenedese, Carlo, Fochesato, Marta, Lygeros, John, & Dörfler, Florian (2024). BIG hype: Best intervention in games via distributed hypergradient descent. *IEEE Transactions on Automatic Control*, 1–16.
- Groot, Noortje, Zaccour, Georges, & de Schutter, Bart (2017). Hierarchical game theory for system-optimal control: Applications of reverse Stackelberg games in regulating marketing channels and traffic routing. *IEEE Control Systems*, 37, 129–152.
- Hirose, Kosuke, & Matsumura, Toshihiro (2019). Comparing welfare and profit in quantity and price competition within Stackelberg mixed duopolies. *Journal of Economics*.
- Ho, Yu-Chi, Luh, P., & Muralidharan, R. (1981). Information structure, Stackelberg games, and incentive controllability. *IEEE Transactions on Automatic Control*, 26(2), 454–460.
- Ho, Yu-Chi, Luh, Peter B., & Olsder, Geert Jan (1982). A control-theoretic view on incentives. *Automatica*, 18(2), 167–179.

- Hobbs, B. E. (2001). Linear complementarity models of Nash-Cournot competition in bilateral and POOLCO power markets. *IEEE Transactions on Power Systems*, 16(2), 194–202.
- Hoheisel, Tim, Kanzow, Christian, & Schwartz, Alexandra (2011). Theoretical and numerical comparison of relaxation methods for mathematical programs with complementarity constraints. *Mathematical Programming*, 137, 1–32.
- Hu, Xinmin, & Ralph, Daniel (2007). Using EPECs to model bilevel games in restructured electricity markets with locational prices. *Operations Research*, 55(5), 809–827.
- Jara-Moroni, Francisco, Pang, Jong-Shi, & Wächter, Andreas (2018). A study of the difference-of-convex approach for solving linear programs with complementarity constraints. *Mathematical Programming*, 169.
- Kebriaei, Hamed, & Iannelli, Luigi (2018). Discrete-time robust hierarchical linear-quadratic dynamic games. *IEEE Transactions on Automatic Control*, 63(3), 902–909.
- Kleinert, Thomas, Labbé, Martine, Ljubić, Ivana, & Schmidt, Martin (2021). A survey on mixed-integer programming techniques in bilevel optimization. *EURO Journal on Computational Optimization*, 9, Article 100007.
- Kulkarni, Ankur A., & Shanbhag, Uday V. (2015). An existence result for hierarchical Stackelberg v/s Stackelberg games. *IEEE Transactions on Automatic Control*, 60(12), 3379–3384.
- Luo, Zhi-Quan, Pang, Jong-Shi, & Ralph, Daniel (1996). *Mathematical programs with equilibrium constraints*. Cambridge University Press.
- Maljkovic, Marko, Nilsson, Gustav, & Geroliminis, Nikolas (2022). A pricing mechanism for balancing the charging of ride-hailing electric vehicle fleets. In *2022 European control conference* (pp. 1976–1981).
- Maljkovic, Marko, Nilsson, Gustav, & Geroliminis, Nikolas (2023a). Hierarchical pricing game for balancing the charging of ride-hailing electric fleets. *IEEE Transactions on Control Systems Technology*, 31(6), 2728–2743.
- Maljkovic, Marko, Nilsson, Gustav, & Geroliminis, Nikolas (2023b). Learning how to price charging in electric ride-hailing markets. In *62nd IEEE conference on decision and control* (pp. 3136–3141). IEEE.
- Maljkovic, Marko, Nilsson, Gustav, & Geroliminis, Nikolas (2023c). On finding the leader's strategy in quadratic aggregative Stackelberg pricing games. In *2023 European control conference* (pp. 1–6).
- Mayer, Stefan, & Steinhardt, Claudius (2016). Optimal product line pricing in the presence of budget-constrained consumers. *European Journal of Operational Research*, 248(1), 219–233.
- Monderer, Dov, & Shapley, Lloyd (1996). Potential games. *Games and Economic Behavior*, 14, 124–143.
- Motalleb, Mahdi, Siano, Pierluigi, & Ghorbani, Reza (2019). Networked Stackelberg competition in a demand response market. *Applied Energy*, 239(C), 680–691.
- Nabetani, Koichi, Tseng, Paul, & Fukushima, Masao (2011). Parametrized variational inequality approaches to generalized Nash equilibrium problems with shared constraints. *Computational Optimization and Applications*, 48, 423–452.
- Paccagnan, Dario, Gentile, Basilio, Parise, Francesca, Kamgarpour, Maryam, & Lygeros, John (2016). Distributed computation of generalized Nash equilibria in quadratic aggregative games with affine coupling constraints. In *2016 IEEE 55th conference on decision and control* (pp. 6123–6128).
- Paccagnan, Dario, Gentile, Basilio, Parise, Francesca, Kamgarpour, Maryam, & Lygeros, John (2019). Nash and Wardrop equilibria in aggregative games with coupling constraints. *IEEE Transactions on Automatic Control*, 64(4), 1373–1388.
- Paccagnan, Dario, Kamgarpour, Maryam, & Lygeros, John (2016). On aggregative and mean field games with applications to electricity markets. In *2016 European control conference* (pp. 196–201).
- Pedregosa, Fabian (2016a). Hyperparameter optimization with approximate gradient. In Maria Florina Balcan, & Kilian Q. Weinberger (Eds.), *Proceedings of machine learning research: vol. 48, Proceedings of the 33rd international conference on machine learning* (pp. 737–746). New York, New York, USA: PMLR.
- Pedregosa, Fabian (2016b). Hyperparameter optimization with approximate gradient. In Maria Florina Balcan, & Kilian Q. Weinberger (Eds.), *Proceedings of machine learning research: vol. 48, Proceedings of the 33rd international conference on machine learning* (pp. 737–746). New York, New York, USA: PMLR.
- Rajeswaran, Aravind, Finn, Chelsea, Kakade, Sham M., & Levine, Sergey (2019). Meta-learning with implicit gradients. In *Proceedings of the 33rd international conference on neural information processing systems*. Red Hook, NY, USA: Curran Associates Inc.
- Scheel, Holger, & Scholtes, Stefan (2000). Mathematical programs with complementarity constraints: Stationarity, optimality, and sensitivity. *Mathematics of Operations Research*, 25(1), 1–22.
- Staňková, Kateřina, Olsder, G. J., & Bliemer, M (2011). *Bi-level optimal toll design problem solved by the inverse Stackelberg games approach: vol. 89*. (pp. 871–880).
- Subotić, Irina, Hauswirth, Adrian, & Dörfler, Florian (2022). Quantitative sensitivity bounds for nonlinear programming and time-varying optimization. *IEEE Transactions on Automatic Control*, 67(6), 2829–2842.
- Toubeau, Jean-François, Bottieau, Jérémie, De Grève, Zacharie, Vallée, François, & Bruninx, Kenneth (2021). Data-driven scheduling of energy storage in day-ahead energy and reserve markets with probabilistic guarantees on real-time delivery. *IEEE Transactions on Power Systems*, 36(4), 2815–2828.
- Tushar, Wayes, Saad, Walid, Poor, H. Vincent, & Smith, David B. (2012). Economics of electric vehicle charging: A game theoretic approach. *IEEE Transactions on Smart Grid*, 3(4), 1767–1778.
- Von Stackelberg, Heinrich (1952). *The theory of the market economy*. William Hodge.
- Wang, Mingshen, Li, Xue, Dong, Chaoyu, Mu, Yunfei, Jia, Hongjie, & Li, Fangxing (2023). Day-ahead optimal bidding for a retailer with flexible participation of electric vehicles. *IEEE Transactions on Smart Grid*, 14(2), 1482–1494.
- Wang, Kai, Xu, Lily, Perrault, Andrew, Reiter, Michael K., & Tambe, Milind (2022). Coordinating followers to reach better equilibria: End-to-end gradient descent for Stackelberg games. In *Proceedings of the AAAI conference on artificial intelligence*, vol. 36, no. 5 (pp. 5219–5227).
- Zardini, Gioele, Lanzetti, Nicolas, Guerrini, Laura, Frazzoli, Emilio, & Dörfler, Florian (2021). Game theory to study interactions between mobility stakeholders. In *2021 IEEE international intelligent transportation systems conference* (pp. 2054–2061).



Marko Maljkovic — received his B.Sc. in Electrical and Computer Engineering from the School of Electrical Engineering, University of Belgrade, and his M.Sc. in Robotics, Systems, and Control from ETH Zurich in 2018 and 2021, respectively. He is currently a doctoral student at the Urban Transport Systems Laboratory (LUTS) at EPFL. His research interests lie in game-theoretic and learning-based mechanism design and control for multi-agent systems, with applications in future mobility systems.



Gustav Nilsson — received his M.Sc. in Engineering Physics and Ph.D. in Automatic Control from Lund University in 2013 and 2019, respectively. He is currently a Postdoctoral researcher at University of Trento. Prior to his current position, he has been a Postdoctoral Researcher at EPFL and a Postdoctoral Associate at GeorgiaTech, GA, USA. His primary research interest lies in modeling and control of dynamical flow networks with applications in transportation networks.



Nikolas Geroliminis — is a Full Professor at EPFL and the head of the Urban Transport Systems Laboratory (LUTS). Before joining EPFL he was an Assistant Professor on the faculty of the Department of Civil Engineering at the University of Minnesota. He has a diploma in Civil Engineering from the National Technical University of Athens (NTUA) and a M.Sc. and Ph.D. in civil engineering from University of California, Berkeley. His research interests focus primarily on urban transportation systems, traffic flow theory and control, public transportation and on-demand transport, car sharing, Optimization and Large Scale Networks. Among his recent initiatives is the creation of an open-science large-scale dataset of naturalistic urban trajectories of half a million vehicles that have been collected by one-of-a-kind experiment by a swarm of drones (<https://open-traffic.epfl.ch>). Among other editorial responsibilities, he is currently the Editor-In-Chief of Transportation Research part C: Emerging Technologies.