

Efficiently Solving Mixed-Hierarchy Games with Quasi-Policy Approximations

Hamzah I. Khan^{1*}[0000-0001-5481-6388], Dong Ho Lee¹[0000-0002-9045-2574],
Jingqi Li¹[0000-0002-3731-3807], Tianyu Qiu¹[0009-0008-5189-3722], Christian
Ellis^{1,2}[0000-0002-0544-8627], Jesse Milzman²[0000-0003-4937-8912], Wesley
Suttle²[0000-0003-1234-7151], and David Fridovich-Keil¹[0000-0002-5866-6441]

¹ University of Texas at Austin, Austin, TX 78712, USA
Corresponding Author*: hamzah@utexas.edu

² U.S. Army Research Laboratory, Adelphi, MD 20783, USA

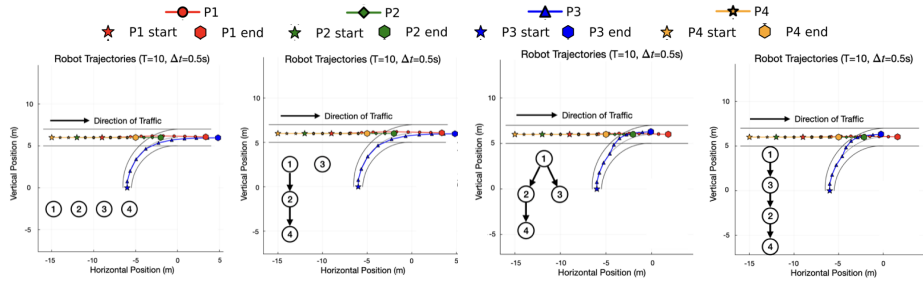
Abstract. Multi-robot coordination often exhibits hierarchical structure, with some robots’ decisions depending on the planned behaviors of others. While game theory provides a principled framework for such interactions, existing solvers struggle to handle mixed information structures that combine simultaneous (Nash) and hierarchical (Stackelberg) decision-making. We study N -robot forest-structured mixed-hierarchy games, in which each robot acts as a Stackelberg leader over its subtree while robots in different branches interact via Nash equilibria. We derive the Karush–Kuhn–Tucker (KKT) first-order optimality conditions for this class of games and show that they involve increasingly high-order derivatives of robots’ best-response policies as the hierarchy depth grows, rendering a direct solution intractable. To overcome this challenge, we introduce a quasi-policy approximation that removes higher-order policy derivatives and develop an inexact Newton method for efficiently solving the resulting approximated KKT systems. We prove local exponential convergence of the proposed algorithm for games with nonquadratic objectives and nonlinear constraints. The approach is implemented in a highly optimized Julia library (`MixedHierarchyGames.jl`) and evaluated in hardware and simulated multi-agent experiments, demonstrating real-time convergence for complex mixed-hierarchy information structures.

Keywords: game theory, information structure, planning, optimization

1 Introduction

Strategic multi-robot decision-making often involves hierarchical information structures that define leader and follower roles among different robots, each

Code available: <https://github.com/CLeARoboticsLab/MixedHierarchyGames.jl>.
This work was sponsored by the Army Research Laboratory under Cooperative Agreement Number W911NF-25-2-0021, and by the National Science Foundation under Grants 2211548 and 2336840.



(a) Nash hierarchy (b) Mixed hierarchy (c) Mixed hierarchy (d) Stackelberg structure. structure A. structure B. chain structure.

Fig. 1: We experiment on four hierarchy structures on a merging scenario and report positions over time. In Fig. 1a, vehicles play a pure Nash game. In Fig. 1b, vehicle 1 leads vehicles 2 and 4, and vehicle 2 leads vehicle 4. In Fig. 1c, vehicle 1 leads every other agent and vehicle 2 leads vehicle 4. Vehicles 3 and (2,4) are negotiating the merge simultaneously. In Fig. 1d, vehicles play a Stackelberg chain with hierarchy $1 \rightarrow 3 \rightarrow 2 \rightarrow 4$. In the absence of leadership from vehicle 1 to vehicle 3, vehicle 3 merges first (Figs. 1a and 1b). When that leadership exists, vehicle 3 merges into the convoy after vehicle 1 (Figs. 1c and 1d).

of which makes a constrained decision to optimize a distinct objective with access to different information. Mathematical games, particularly the Nash and Stackelberg equilibrium concepts, provide a powerful framework to model such strategic interactions. For example, consider a convoy of three vehicles driving in a single lane. In such a scenario, we may model each vehicle as a leader of each subsequent vehicle in the lane, resulting in a leadership “chain” that prior works model within the formalism of Stackelberg games [4,37,27].

However, many scenarios require modeling a mix of Nash and Stackelberg relationships between groups of robots [28,32]. Beyond convoy merging, mixed-hierarchy structures arise in multi-robot patrolling, social navigation, and human-robot teams, where designated leaders influence subgroups whose members interact as peers. Consider a fourth car seeking to merge into the convoy’s lane between the first and second convoy vehicles, as in Fig. 1. Merging into this position in the convoy requires the fourth car to coordinate with both other vehicles. First, the merging vehicle must respond to signaled behavior from the first car, such as slowing down for other vehicles or speeding up to make space. The same type of signaling does not necessarily exist between the merging robot and the second car in the convoy, making the Nash equilibrium model of simultaneous play a natural model for their interaction.

While some approaches model this type of mixed Nash-Stackelberg information structure [25,28,32], they typically solve problems with linear constraints and quadratic costs, limiting their applicability to real-world robotic systems that feature nonlinear constraints (often arising from nonlinear motion dynamics) and nonquadratic objectives (i.e., not restricted to quadratic polynomials in the decision variables) that encode complex tasks. To address this gap, we develop a more general solver that efficiently handles nonquadratic objectives

and nonlinear equality constraints for a subclass of mixed-hierarchy games with leadership structures that are forests. This subclass of games includes Nash and Stackelberg games as special cases.

A key difficulty arises, however, in expressing the first-order optimality conditions for these mixed-hierarchy games: these conditions depend upon high-order derivatives of robots’ best-response maps, or policies. To maintain tractability, we develop a “quasi-policy” approximation which extends recent work in the feedback Nash [26] and Stackelberg [27] settings, wherein each robot’s first-order optimality conditions are derived by ignoring higher-order derivatives of downstream followers’ policies. This approximation enables us to solve mixed-hierarchy games efficiently, with real-time computation demonstrating the value of the quasi-policy approximation in enabling real-world motion planning and robot coordination applications. Ultimately, we claim three contributions.

1. We propose a principled approach for constructing approximate optimality conditions for a class of nonlinear, equality-constrained mixed-hierarchy games with forest-structured information patterns, including a novel method for embedding followers’ best responses within leaders’ optimality conditions as *locally first-order* policy approximations (which we refer to as a quasi-policy approximation, following [26,27]).
2. We introduce an efficient, flexible, and scalable algorithm for settings in which robots have arbitrary nonlinear, differentiable costs and (equality) constraints, that uses these optimality conditions to approximate local equilibria for arbitrary forest-structured mixed-hierarchy games. We provide a theoretical proof and an empirical study for the local exponential convergence of our algorithm to an approximate mixed-hierarchy equilibrium.
3. We provide an open-source, optimized Julia library for solving these problems, which supports simulations of the merging scenario in Fig. 1 as well as in hardware with robots playing a target-guarding game in Fig. 4. We demonstrate real-time motion planning under complex hierarchical leadership structures.

2 Related Work

Nash and Stackelberg Games. Classical Nash and Stackelberg equilibria capture mutual best responses [33] and leader-follower interactions via sequential choices [41], respectively. Although closed-form expressions for the solutions to these problems are available in highly-structured cases—e.g., when agents’ objectives are quadratic and only linear equality constraints are present [3]—more general settings typically have no clear analytic solution and must be approached computationally, from the perspective of mathematical programming and complementarity-based formulations [17,31]. These formulations enable *local* equilibrium computation for Nash and Stackelberg games with complex objectives and constraints [18,26,27]. However, they inherently assume simple information structures such as simultaneous interaction (Nash) or a direct leader-follower hierarchy (Stackelberg). Extensions exist for multi-leader

or multi-follower settings [39,24,32], but do not address nonlinear games under mixed Nash-Stackelberg leadership structures. Our work targets this gap.

Computing Equilibria under Hierarchical Information Structures. Optimization problems which encode a follower’s decision problem as a constraint upon the leader’s decision are naturally bilevel in nature [7,40]. Moreover, as longer leader-follower chains arise, these hierarchical, multilevel problems’ [4,9] computational complexity grows rapidly [20,6]. Recent work explores monotone operator formulations [37], gradient-based methods [36], and parametric techniques [14] for these problems, but these approaches largely focus on low-dimensional settings. More expressive frameworks include quadratic constrained games under bilevel tree-structured hierarchies [10] and the mathematical programming network (MPN) framework [25], which encodes arbitrary mixtures of Nash and Stackelberg relations via symbolic graphs. While the MPN framework fundamentally describes the same type of mixed Nash-Stackelberg hierarchies we consider in this work, existing MPN solvers are restricted to quadratic objectives and linear inequality constraints, and face scalability challenges due to polytope-based feasible-set representations and iterative active-set updates that must be repeatedly refined during a recursive graph traversal. In contrast, our work solves the nonquadratic MPN problem in the nonlinear equality-constrained setting and supports tree-structured hierarchies of arbitrary depth.

Stochastic Game Perspectives on Information Structure. A large body of work models information asymmetry and partial observability in sequential decision-making using stochastic and Markov games [38,29,16]. Within this framework, mixed Nash-Stackelberg and incomplete-information settings have been studied via regret-based and value-based learning [42,12,15,22,1], as well as belief-state dynamic programming for partially observable stochastic games [13,19,30,5]. Another line of work studies correlated equilibria under various information structures [2,8,11]. These approaches primarily focus on stochastic dynamics and discrete state-action spaces. In contrast, we consider deterministic dynamic games with continuous actions under mixed Nash and Stackelberg information structures, and efficiently compute equilibria using differentiable mathematical programming rather than belief-space planning or model-free learning.

3 Problem Formulation

Set Notation. For a positive integer d , let $[d] = \{1, 2, \dots, d\}$. Let \mathcal{S} be an arbitrary finite set with cardinality denoted $|\mathcal{S}|$. For a set of indices $E \subseteq [|\mathcal{S}|]$, we define set-based element lookup $\mathcal{S}^E = \{e^i \mid i \in E, e^i \in \mathcal{S}\}$ and exclusion by $\mathcal{S}^{-E} = \mathcal{S} \setminus \mathcal{S}^E$.

Decision Variables. We consider N interacting robots. Each robot $i \in [N]$ has trajectory $\mathbf{z}^i \in \mathbb{R}^{n_z^i}$. We denote the joint trajectories by $\mathbf{z} = (\mathbf{z}^1, \dots, \mathbf{z}^N)^\top \in \mathbb{R}^{n_z}$, where $n_z = \sum_{i=1}^N n_z^i$.

3.1 Mixed Hierarchy Games

In an N -robot *mixed hierarchy game*, each robot $i \in [N]$ minimizes an objective function of the strategies \mathbf{z} of all robots,

$$J^i : \mathbb{R}^{n_z} \rightarrow \mathbb{R}, \quad (1)$$

subject to a set of equality constraints which restrict its strategy, e.g. by enforcing dynamic feasibility:

$$0 = g^i(\mathbf{z}^i) \in \mathbb{R}^{n_g^i}. \quad (2)$$

Assumption 1. We require (1) and (2) to satisfy higher-order differentiability, i.e. $J^i, g^i \in C^\infty$.

Finally, we model the information structure (i.e., the set of hierarchical leader-follower relationships between robots) as a directed acyclic graph $\mathcal{G} = ([N], \mathcal{E})$, i.e., *without any directed cycles*. In \mathcal{G} , each robot $i \in [N]$ corresponds to a node and $(i, j) \in \mathcal{E}$ defines a direct *leader-follower relationship* between leader i and follower j ; if (i, j) are connected by a longer path, we say that they have an *indirect* leader-follower relationship. Thus, a mixed-hierarchy game consists of graph \mathcal{G} and the set of all robots' objectives and equality constraints.

3.2 Notation for Defining Mixed-Hierarchy Equilibria

We will restrict our attention to graphs in which each node has at most one parent, i.e., each robot has at most one direct leader. To this end, we introduce additional notation and an assumption on the topology of graph \mathcal{G} . First, we define the sets of *direct leaders* of a robot j and *direct followers* of a robot i as

$$\mathbf{L}(j) = \{i \mid (i, j) \in \mathcal{G}\} \quad \text{and} \quad \mathbf{F}(i) = \{j \mid (i, j) \in \mathcal{G}\}. \quad (3)$$

Next, we define the direct leaders of a set of robot indices \mathcal{S} as the union of the direct leaders of each robot, $\mathbf{L}(\mathcal{S}) = \bigcup_{e \in \mathcal{S}} \mathbf{L}(e)$. We define the direct followers of \mathcal{S} similarly, $\mathbf{F}(\mathcal{S}) = \bigcup_{e \in \mathcal{S}} \mathbf{F}(e)$. Finally, we recursively define their transitive closures, i.e., the sets of all leaders and followers (whether direct or indirect) of robot i to be

$$\mathbf{L}(i) = \mathbf{L}(i) \cup \mathbf{L}(\mathbf{L}(i)) \cup \dots \quad \text{and} \quad \mathbf{F}(i) = \mathbf{F}(i) \cup \mathbf{F}(\mathbf{F}(i)) \cup \dots. \quad (4)$$

Since \mathcal{G} is finite and acyclic, then these sets are finite too. Finally, we define a set corresponding to the indices of robots with no other relationship to robot i ,

$$\mathbf{N}(i) = [N] \setminus \mathbf{L}(i) \setminus \mathbf{F}(i) \setminus \{i\}. \quad (5)$$

These correspond to robots that have Nash relationships with robot i , as we discuss in further detail below.

To make the problem tractable and ensure a clear leadership hierarchy, we restrict the topology of the information structure to one made of a union of

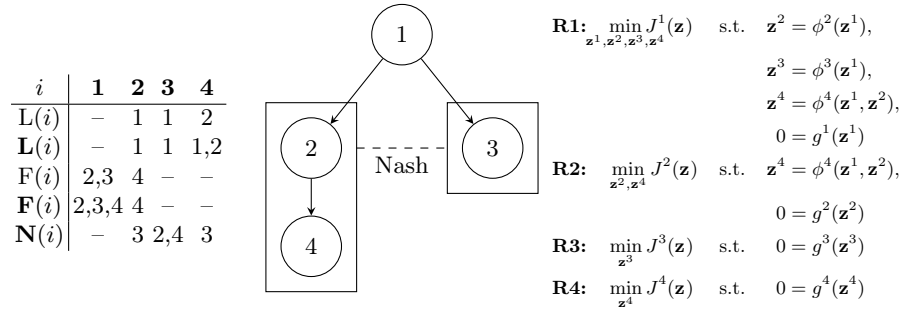




Fig. 2: We introduce three representations of a mixed-hierarchy game which models the information structure of the example from the introduction and satisfies Assumption 2: We include a set-based representation which lists elements of each robot’s leader/follower/Nash sets (left), a graphical representation of \mathcal{G} (middle), and a coupled optimization problem (right). In this game, robot 1 leads all other robots (though robot 4 is led indirectly), robot 2 leads robot 4, and robot 2’s subtree (including robot 4) is in a Nash relationship with robot 3.

a finite number of disconnected trees, i.e., a forest. Such a topology remains versatile and is able to model a wide variety of hierarchical structures of interest, including pure Nash (for which $\mathcal{E} = \{\}$) and Stackelberg chains (for which $\mathcal{E} = \{(i, i + 1) \mid i \in [N - 1]\}$).

Assumption 2 (The Hierarchical Information Structure Graph \mathcal{G} is a Forest). *We assume that every robot $i \in [N]$ has either zero or one direct leader, i.e. $|\mathbf{L}(i)| \in \{0, 1\}$ (though it can have multiple indirect leaders). This condition ensures that every connected component of \mathcal{G} is a tree (i.e. \mathcal{G} is a forest) and that every node $i \in [N]$ is the root of a subtree, and therefore a leader to every robot in that subtree $\mathcal{T}_i = (\{i\} \cup \mathbf{F}(i), \{(j, k) \in \mathcal{E} \mid j, k \in \{i\} \cup \mathbf{F}(i)\})$. We provide examples of (in)valid hierarchy structures in Figs. 2 and 3.*



(a) Hierarchy structure graph can not contain a directed cycle.



(b) Hierarchy structure is sound but violates Assumption 2.

Fig. 3: Examples of invalid hierarchical information structure graphs.

3.3 Equilibrium Conditions for Mixed-Hierarchy Games

Next, we define the equilibrium conditions associated with the mixed-hierarchy information structure under Assumption 2.

In the two-robot setting, if a robot is a Stackelberg leader, then the leader communicates its strategy to the follower, who uses that knowledge when making its own decision. The leader anticipates the followers' best response and makes a decision to steer the follower's decision to a desired outcome. Mathematically, this type of relationship is enforced by defining an optimal follower response map that can be embedded within the leader's problem, cf. [3, Ch. 3, 7].

In an N -robot mixed hierarchy game, when robot i is an equality-constrained Stackelberg follower of (direct and indirect) leaders $\mathbf{L}(i)$, the set of all optimal follower responses $\mathcal{R}_F^{i*}(\mathbf{z}^{\mathbf{L}(i)})$ given its leaders' communicated strategies $\mathbf{z}^{\mathbf{L}(i)}$ is

$$\mathcal{R}_F^{i*}(\mathbf{z}^{\mathbf{L}(i)}) \equiv \underset{\mathbf{z}^i, \mathbf{z}^{\mathbf{F}(i)}}{\operatorname{argmin}} J^i(\mathbf{z}^{\mathbf{L}(i)}, \mathbf{z}^i, \mathbf{z}^{\mathbf{F}(i)}, \mathbf{z}^{\mathbf{N}(i)}) \quad (6a)$$

$$\text{s.t. } g^i(\mathbf{z}^i) = 0, \quad (6b)$$

$$(\mathbf{z}^j, \mathbf{z}^{\mathbf{F}(j)}) \in \mathcal{R}_F^{j*}(\underbrace{\mathbf{z}^i, \mathbf{z}^{\mathbf{L}(i)}}_{\mathbf{z}^{\mathbf{L}(j)}}), \quad \forall j \in \mathbf{F}(i). \quad (6c)$$

Eq. (6a) describes the minimization of robot i 's objective function with respect to its own strategy and its followers' strategies (which it can influence). Since each robot i is the leader of the robots in its own subtree \mathcal{T}_i , the optimal follower response includes robot i 's strategy \mathbf{z}^{i*} and those of all of its followers, $\mathbf{z}^{\mathbf{F}(i)*}$. This best response optimization is subject to two constraints: first, the equality constraint (6b) corresponds to (2). Second, (6c) accounts for followers' reactions to robot i 's choice of strategy, because the best responses affect robot i 's objective. Thus, (6c) recursively collects the optimal responses of robots within \mathcal{T}_i that make up $\mathbf{z}^{\mathbf{F}(i)*}$. We highlight that (6a) relies on Nash-related robots' trajectories $\mathbf{z}^{\mathbf{N}(i)}$ which robot i cannot directly influence. Thus, each robot's equilibrium conditions must incorporate both Stackelberg leader conditions that embed (6) and Nash conditions on the relationship between \mathbf{z}^{i*} , $\mathbf{z}^{\mathbf{F}(i)*}$, and $\mathbf{z}^{\mathbf{N}(i)*}$.

Generally, a best response set $\mathcal{R}_F^{i*}(\mathbf{z}^{\mathbf{L}(i)})$ can contain multiple optimal combinations of values for \mathbf{z}^{i*} and $\mathbf{z}^{\mathbf{F}(i)*}$, even in the simple two-robot quadratic case [3, Ch. 7]. Nevertheless, for ease of analysis, it is common to assume $|\mathcal{R}_F^{i*}(\mathbf{z}^{\mathbf{L}(i)*})| = 1$ [3,21], i.e., that *equilibrium* leader strategies result in a unique optimal response for all followers of robot i . We call a function that generates an element $(\mathbf{z}^{i*}, \mathbf{z}^{\mathbf{F}(i)*}) \in \mathcal{R}_F^{i*}(\mathbf{z}^{\mathbf{L}(i)})$ a *policy* $\Phi^i(\mathbf{z}^{\mathbf{L}(i)})$, where

$$(\mathbf{z}^{i*}, \mathbf{z}^{\mathbf{F}(i)*}) \equiv \Phi^i(\mathbf{z}^{\mathbf{L}(i)}) \quad \text{and, at } \mathbf{z}^{\mathbf{L}(i)*}, \quad \mathcal{R}_F^{i*}(\mathbf{z}^{\mathbf{L}(i)*}) = \left\{ \Phi^i(\mathbf{z}^{\mathbf{L}(i)*}) \right\}. \quad (7)$$

We equivalently define Φ^i with component sub-policies $\mathbf{z}^{j*} = \phi^j(\mathbf{z}^{\mathbf{L}(j)})$ for each follower's individual best-responses,

$$\Phi^i(\mathbf{z}^{\mathbf{L}(i)}) := \left\{ \mathbf{z}^{j*} = \phi^j(\mathbf{z}^{\mathbf{L}(i)}, \mathbf{z}^{\mathbf{L}(j) \setminus \mathbf{L}(i)*}) \mid j \in \{i\} \cup \mathbf{F}(i) \right\}. \quad (8)$$

Only leader strategies $\mathbf{z}^{\mathbf{L}(i)}$ are provided to Φ^i , so the remaining leader strategies, $\mathbf{z}^{\mathbf{L}(j) \setminus \mathbf{L}(i)*}$, are computed from other subpolicies in (8).

Definition 1 (Mixed-Hierarchy Equilibrium). *At a mixed-hierarchy equilibrium, each robot i must satisfy both Stackelberg leadership conditions to (direct and indirect) followers $\mathbf{F}(i)$ and Nash conditions with $\mathbf{N}(i)$. Given equilibrium leader strategies $\mathbf{z}^{\mathbf{L}(i)*}$ and equilibrium Nash-related robot strategies $\mathbf{z}^{\mathbf{N}(i)*}$, we introduce the following mixed-hierarchy equilibrium condition for each robot i :*

$$\left(\mathbf{z}^{i*}, \mathbf{z}^{\mathbf{F}(i)*}\right) \in \underset{\mathbf{z}^i, \mathbf{z}^{\mathbf{F}(i)}}{\operatorname{argmin}} J^i\left(\mathbf{z}^{\mathbf{L}(i)*}, \mathbf{z}^i, \mathbf{z}^{\mathbf{F}(i)}, \mathbf{z}^{\mathbf{N}(i)*}\right) \quad (9a)$$

$$\text{s.t. } g^i(\mathbf{z}^i) = 0, \quad (9b)$$

$$\mathbf{z}^j = \phi^j(\mathbf{z}^{\mathbf{L}(i)*}, \mathbf{z}^{\mathbf{L}(j) \setminus \mathbf{L}(i)}), \quad \forall j \in \mathbf{F}(i). \quad (9c)$$

Condition (9) incorporates both Stackelberg leadership and Nash conditions. It requires each robot i to choose its optimal leader strategy to guide its followers $\mathbf{F}(i)$ to benefit the leading robot, as enforced by (9a) and (9c). Simultaneously, the optimal strategies of robot i and its followers, $\mathbf{z}^{\{\{i\} \cup \mathbf{F}(i)\}^*}$, must be unilaterally optimal with respect to the equilibrium strategies of Nash-related robots, $\mathbf{z}^{\mathbf{N}(i)*}$, as enforced by (9a). When condition (9) holds for all robots $i \in [N]$ within some neighborhood of a point in the joint strategy space for all robots, \mathbf{z}^* , then we say that \mathbf{z}^* is a *local (equality-constrained) mixed-hierarchy equilibrium*.

Assumption 3 (Linear-Independence Constraint Qualification). *To ensure that necessary KKT conditions hold, we require the linear independence constraint qualification (LICQ is a standard regularity condition for these problems), cf. [34, Def. 12.4], to hold in (9), for each robot i .*

We note that conditions (9), when aggregated across all robots, define an equality-constrained variant of the mathematical programming network problem introduced by [25], under Assumption 2.

3.4 Running Example: N -Robot Convoy Merging Game

To make the problem formulation more concrete, consider the following example of a game in which N robots each select trajectories in a convoy merging scenario, taking place over finite horizon $T \in \mathbb{N} \setminus \{0\}$, and with time discretization $\Delta t > 0$. At each time step $t \in [T]$, each robot $i \in [N]$ has state vector $\mathbf{x}_t^i \in \mathbb{R}^{n_x^i}$ and control vector $\mathbf{u}_t^i \in \mathbb{R}^{n_u^i}$. Concretely, each vehicle i 's state contains planar horizontal and vertical position $\mathbf{p}_t^i = (p_{x,t}^i, p_{y,t}^i)^\top \in \mathbb{R}^2$, heading $\psi_t^i \in \mathbb{R}$, and scalar velocity $v_t^i \in \mathbb{R}$, and vehicle i 's controls contain longitudinal acceleration $a_t^i \in \mathbb{R}$ and rotational velocity $\omega_t^i \in \mathbb{R}$, i.e. $\mathbf{x}_t^i = (p_{x,t}^i, p_{y,t}^i, \psi_t^i, v_t^i)^\top$, $\mathbf{u}_t^i = (a_t^i, \omega_t^i)^\top$. Each robot i 's strategy takes the form

$$\mathbf{z}^i = (\mathbf{x}_1^{i\top}, \mathbf{u}_1^{i\top}, \dots, \mathbf{x}_T^{i\top}, \mathbf{u}_T^{i\top})^\top. \quad (10)$$

Robot i 's trajectory \mathbf{z}^i is constrained to start at an initial state $\mathbf{x}_{\text{init}}^i$ and for its state to evolve in time according to robot dynamics f^i , defined such that

$$\mathbf{x}_{t+1}^i = f^i(\mathbf{x}_t^i, \mathbf{u}_t^i), \quad t \in [T-1]. \quad (11)$$

Concretely, for this example, we take these to encode unicycle dynamics, i.e.

$$\mathbf{x}_{t+1}^i = f^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = \begin{bmatrix} p_{x,t+1}^i \\ p_{y,t+1}^i \\ \psi_{t+1}^i \\ v_{t+1}^i \end{bmatrix} = \begin{bmatrix} p_{x,t}^i + \Delta t v_t^i \cos \psi_t^i \\ p_{y,t}^i + \Delta t v_t^i \sin \psi_t^i \\ \psi_t^i + \Delta t \omega_t^i \\ v_t^i + \Delta t a_t^i \end{bmatrix}. \quad (12)$$

Ultimately, we can then transcribe each robot’s initial state and dynamics constraints as

$$g^i(\mathbf{z}^i) = \begin{bmatrix} \mathbf{x}_{\text{init}}^i - \mathbf{x}_1^i \\ \mathbf{x}_2^i - f^i(\mathbf{x}_1^i, \mathbf{u}_1^i) \\ \vdots \\ \mathbf{x}_T^i - f^i(\mathbf{x}_{T-1}^i, \mathbf{u}_{T-1}^i) \end{bmatrix} = 0. \quad (13)$$

Additional equality constraints—e.g., encoding a desired terminal state—can be appended to (13) as needed.

Lastly, we construct convoy merging scenarios in which N_c vehicles have formed a convoy on the road, and $N_m = N - N_c$ vehicles wish to merge. All vehicles wish to track the center of the merged lane (with their vertical positions and headings), maintain a consistent speed v_c^i , minimize their control effort, and avoid collision. The objective, given below, is parameterized by positive weights $\mathbf{w}^i = (w^{i,1}, w^{i,2}, w^{i,3}, w^{i,4}, w^{i,5})$, collision safe distance $d_s > 0$ beyond which the collision cost is 0, and collision sensitivity $c_s > 0$:

$$J^i(\mathbf{z}) = \sum_{t=1}^T \left(w^{i,1} \|\mathbf{u}_t^i\|_2^2 + w^{i,2} (p_{y,t}^i)^2 + w^{i,3} (\psi_t^i)^2 + w^{i,4} (v_t^i - v_c^i)^2 \right. \quad (14) \\ \left. + w^{i,5} \sum_{j=1}^{N-1} \sum_{k=j+1}^N \left[\frac{1}{c_s} \log \left(1 + \exp \left(c_s (d_s^2 - \|\mathbf{p}_t^j - \mathbf{p}_t^k\|_2^2) \right) \right) \right]^2 \right)$$

The merging vehicles start off the merging lane and are incentivized to track the center of the round entry ramp and merge safely onto the highway, as shown in Fig. 1. We enforce curve-following by altering (14) for vehicle 3 to include a component term, $(\|\mathbf{p}_t^3\|_2^2 - R^2)^2$, and we alter vehicle 1’s cost to include a component term $(p_{x,t}^1 - p_{x,t}^3 - d_{13})^2$, which incentivizes vehicle 1 to stay d_{13} meters ahead of vehicle 3. In Section 5.2, we experiment with different hierarchical structures (e.g., Nash, two mixed hierarchies, and Stackelberg chain) and report their impact on vehicles’ decisions.

4 Solving Mixed-Hierarchy Games

In order to build an algorithm to solve mixed-hierarchy games (i.e., find a set of robot strategies that satisfy Definition 1), we must first develop a procedure for constructing appropriate Karush-Kuhn-Tucker (KKT) conditions for first-order optimality. As we will show, these conditions involve high-order derivatives of follower robots’ policies; to circumvent the associated computational challenges,

we describe a *quasi-policy* approximation, which embeds local first-order approximations of followers' policies within each leader's optimization problem in order to reduce the recursive complexity of evaluating high-order derivatives. Finally, we propose an inexact Newton's method to iteratively solve the approximate KKT system and prove local exponential convergence.

4.1 Backward Induction for Deriving KKT Conditions

We begin by constructing KKT conditions for robots without followers (leaves in the graph \mathcal{G}) and then proceed in reverse topological order.

Constructing KKT Conditions for Leaves. Equilibrium conditions for robots without followers do not include the follower constraints (9c). Thus, a leaf robot i 's optimization problem becomes

$$J^i(\mathbf{z}^*) = \min_{\mathbf{z}^i} J^i(\mathbf{z}^{\mathbf{L}(i)*}, \mathbf{z}^i, \mathbf{z}^{\mathbf{N}(i)*}) \quad \text{s.t. } g^i(\mathbf{z}^i) = 0. \quad (15)$$

We associate the equality constraint $g^i(\mathbf{z}^i) = 0$ with a Lagrange multiplier $\lambda^i \in \mathbb{R}^{n_g^i}$, and introduce the corresponding Lagrangian:

$$\mathcal{L}^i(\mathbf{z}, \lambda^i) = J^i(\mathbf{z}) - \lambda^{i\top} g^i(\mathbf{z}^i). \quad (16)$$

We then compactly formulate robot i 's KKT conditions π^i by concatenating its Lagrangian stationarity condition with the original equality constraint,

$$\pi^i(\mathbf{z}, \lambda^i) = \begin{bmatrix} \nabla_{\mathbf{z}^i} \mathcal{L}^i(\mathbf{z}, \lambda^i) \\ g^i(\mathbf{z}^i) \end{bmatrix} = 0. \quad (17)$$

From Leaf Robots' KKT Conditions to Their Policies. Consider a robot j with (direct or indirect) leader robot $i \in \mathbf{L}(j)$, but no followers. Robot j 's KKT conditions π^j implicitly specify its *policy* ϕ^j , which maps leaders' strategies $\mathbf{z}^{\mathbf{L}(j)}$ to robot j 's best response \mathbf{z}^j , i.e. so that $\mathbf{z}^j = \phi^j(\mathbf{z}^{\mathbf{L}(j)})$. For notational convenience, we denote the primal and dual variables making up the leaf robot's decision variables by $\mathbf{w}^j = (\mathbf{z}^{j\top}, \lambda^{j\top})^\top$, and compactly represent all variables relevant to leaf robot j 's optimization problem by $\mathbf{s}^j = (\mathbf{z}^{\mathbf{L}(j)\top}, \mathbf{w}^{j\top})^\top$.

Before constructing the leader i 's KKT conditions π^i , we must first understand how follower j 's best response strategy \mathbf{z}^j locally changes in response to small changes in the leader's decision. To this end, we express the first-order Taylor approximation of robot j 's KKT conditions as follows:

$$0 = \pi^j(\mathbf{s}^j) + \nabla_{\mathbf{w}^j} \pi^j(\mathbf{s}^j) \cdot \Delta \mathbf{w}^j + \nabla_{\mathbf{z}^{\mathbf{L}(j)}} \pi^j(\mathbf{s}^j) \cdot \Delta \mathbf{z}^{\mathbf{L}(j)}, \quad (18)$$

where $\Delta \mathbf{z}^{\mathbf{L}(j)}$ is a small perturbation in the leaders' strategies and $\Delta \mathbf{w}^j$ is the corresponding change in j 's best response (including both primal and dual variables). We note that $\nabla_{\mathbf{w}^j} \pi^j(\mathbf{s}^j)$ is a square matrix; if it has full rank, we can apply the implicit function theorem [23] and obtain that, locally:

$$\Delta \mathbf{w}^j = \Phi^j(\Delta \mathbf{z}^{\mathbf{L}(j)}) = -(\nabla_{\mathbf{w}^j} \pi^j(\mathbf{s}^j))^{-1} (\pi^j(\mathbf{s}^j) + \nabla_{\mathbf{z}^{\mathbf{L}(j)}} \pi^j(\mathbf{s}^j) \cdot \Delta \mathbf{z}^{\mathbf{L}(j)}). \quad (19)$$

From (19), we extract the policy $\mathbf{z}^j = \phi^j(\mathbf{z}^{\mathbf{L}(j)})$ for robot j 's strategy by selecting the corresponding rows of $\Phi^j(\Delta\mathbf{z}^{\mathbf{L}(j)})$. This explicit construction can then replace the implicit constraint (9c) in the optimization problem of each leader $i \in \mathbf{L}(j)$, which will shortly allow us to build an explicit expression for the leader's approximate KKT conditions and ultimately an efficient solver.

Constructing KKT Conditions for Non-Leaf Robots. Proceeding in reverse topological order on graph \mathcal{G} from a leaf robot, we consider a robot i whose followers $\mathbf{F}(i)$ are all leaves. Let robot $j \in \mathbf{F}(i)$ and denote its local policy by ϕ^j , per the construction in (17). As discussed above, we replace the implicit constraints (9c) on followers' best responses in robot i 's problem (9) with an explicit approximation constructed from (19).

To this end, we introduce dual variable $\lambda^i \in \mathbb{R}^{n_g^i}$ associated with the equality constraint (9b) and $\mu^{i,j} \in \mathbb{R}^{n_\mu^j}$ corresponding to the best-response constraint (9c), for each $j \in \mathbf{F}(i)$. We compactly represent all policy duals associated with robot i as $\boldsymbol{\mu}^i = (\mu^{i,j^\top} \mid j \in \mathbf{F}(i))^\top$, and collect all decision variables for robot i as $\mathbf{w}^i = (\mathbf{z}^{i^\top}, \mathbf{z}^{\mathbf{F}(i)\top}, \lambda^{i^\top}, \boldsymbol{\mu}^{i^\top})^\top$. We also collect all variables relevant to robot i 's optimization problem as $\mathbf{s}^i = (\mathbf{z}^{\mathbf{L}(i)\top}, \mathbf{w}^{i^\top})^\top$. We then write robot i 's Lagrangian

$$\mathcal{L}^i(\mathbf{s}^i) = J^i(\mathbf{z}) - \lambda^{i^\top} g^i(\mathbf{z}^i) - \sum_{j \in \mathbf{F}(i)} \mu^{i,j^\top} (\mathbf{z}^j - \phi^j(\mathbf{z}^{\mathbf{L}(j)})). \quad (20)$$

Finally, we define robot i 's KKT conditions by concatenating Lagrangian stationarity conditions for robot i and its followers, its followers' best-response constraints, and its own equality constraints,

$$\pi^i(\mathbf{s}^i) = \begin{bmatrix} \nabla_{\mathbf{z}^i} \mathcal{L}^i(\mathbf{s}^i) \\ \nabla_{\mathbf{z}^j} \mathcal{L}^i(\mathbf{s}^i) \\ \mathbf{z}^j - \phi^j(\mathbf{z}^{\mathbf{L}(j)}) \\ g^i(\mathbf{z}^i) \end{bmatrix}_{j \in \mathbf{F}(i)} = 0. \quad (21)$$

Evaluating a direct follower j 's stationarity condition in (21) requires computing $\nabla_{\mathbf{z}^j} \mathcal{L}^i(\mathbf{s}^i)$, which depends on the policy gradient $\nabla_{\mathbf{z}^j} \phi^j$:

$$\nabla_{\mathbf{z}^j} \mathcal{L}^i(\mathbf{s}^i) = \nabla_{\mathbf{z}^j} J^i(\mathbf{z}) - \sum_{j \in \mathbf{F}(i)} (I - \nabla_{\mathbf{z}^j} \phi^j(\mathbf{z}^{\mathbf{L}(j)}))^\top \mu^{i,j}. \quad (22)$$

By the same mechanism, a depth-3 chain (robot i leading j leading k) introduces the second-order policy derivative $\nabla_{\mathbf{z}^k}^2 \phi^k$ into robot i 's KKT system via differentiation of ϕ^j (which itself depends on $\nabla_{\mathbf{z}^j} \phi^j$). The order of policy gradients increases with hierarchy depth, introducing tractability challenges. To make such computation tractable, we take inspiration from [26,27] and propose to replace leaders' exact optimality conditions with an approximate set of KKT conditions where we locally keep the first-order policy gradient (for leader-follower relationships) and zero out higher-order policy gradients wherever they appear in robot i 's KKT conditions. This *quasi-policy* approximation makes it straightforward to

recursively continue the aforementioned procedure in reverse topological order. Once each of robot i 's followers' quasi-policy approximations is computed, it can be embedded in robot i 's problem, even if the followers are not leaves.

Encoding Nash Relationships Every robot's optimal choice of strategy depends on its leaders', followers', and Nash relatives' strategies. While the KKT construction process described above encodes leader-follower relationships, Nash relationships are captured by solving KKT conditions jointly across sets of multiple robots, i.e. by concatenating KKT conditions of robots with the same leaders.

Constructing Joint KKT Conditions. Ultimately, the full set of KKT conditions for an equality-constrained mixed-hierarchy equilibrium has size n_s (the number of variables match the number of nonlinear equations) and is given by

$$\pi(\mathbf{s}) = \left[\pi^1(\mathbf{s}^1)^\top, \pi^2(\mathbf{s}^2)^\top, \dots, \pi^N(\mathbf{s}^N)^\top \right]^\top = 0. \quad (23)$$

4.2 Inexact Newton's Method for Solving Mixed Hierarchy Games

We propose to solve the joint KKT system (23) via an inexact Newton's method (Alg. 1). At iteration k , we first construct the KKT system $\pi(\mathbf{s}_k)$ under the quasi-policy approximation. Then, for computational tractability, we compute an inexact Newton step for $\pi(\mathbf{s}_k)$ using an approximate Jacobian (denoted $\nabla_{\mathbf{s}} \pi(\mathbf{s}_k)$), which we construct under the quasi-policy approximation by setting the high-order policy gradient terms to zero as discussed earlier. Accordingly, we form an approximate first-order Taylor expansion of the KKT conditions about the current iterate \mathbf{s}_k (Alg. 1), resulting in the inexact Newton system

$$\nabla_{\mathbf{s}} \pi(\mathbf{s}_k) \Delta \mathbf{s}_k = -\pi(\mathbf{s}_k). \quad (24)$$

Algorithm 1: General Solver for Mixed Hierarchy Games

Input: Hierarchy structure \mathcal{G} , initial guess \mathbf{s}_0 , tolerance $\tau > 0$, maximum iterations $K_{\max} > 0$.

Output: Approximate solution \mathbf{s}^* satisfying first-order optimality conditions.

```

1 for  $k \leftarrow 1$  to  $K_{\max}$  do
2   // Construct approx. KKT conditions in reverse topological order.
3   for each robot  $i$  in reverse topological order of  $\mathcal{G}$  do
4     └ Construct robot  $i$ 's KKT conditions  $\pi^i$  per Eqs. (16) to (21).
5     Assemble and evaluate  $\pi(\mathbf{s}_k)$  (23), and solve for inexact Newton step  $\Delta \mathbf{s}_k$ .
6     Select step size  $\alpha_k$  along direction  $\Delta \mathbf{s}_k$  and generate  $\mathbf{s}_{k+1}$  per (25)
7     if  $\|\pi(\mathbf{s}_{k+1})\|_2^2 \leq \tau$  then
8       └ return  $\mathbf{s}_{k+1}$ 
9 return  $\mathbf{s}_{K_{\max}}$ 

```

Solving (24) produces an inexact Newton direction $\Delta \mathbf{s}_k$, which can be used to update the current iterate \mathbf{s}_k (Alg. 1) as follows:

$$\mathbf{s}_{k+1} = \mathbf{s}_k - \alpha_k \Delta \mathbf{s}_k, \quad (25)$$

with step size $\alpha_k \in (0, 1]$ selected (Alg. 1) using a standard backtracking line search, geometric decrease, etc. (cf. [34, Ch. 3]). Iterations terminate when $\|\pi(\mathbf{s}_k)\|_2^2 \leq \tau$ (a user-specified tolerance for KKT error; Alg. 1), yielding an approximate mixed-hierarchy equilibrium; it is not exact due to the use of the quasi-policy approximation in constructing the KKT conditions.

In practice, we can accelerate mixed-hierarchy KKT computation (Alg. 1) under the quasi-policy approximation by caching reusable symbolic components (generated via a symbolic algebra toolchain [35]), exploiting sparsity induced by the hierarchy graph, etc. Next, we comment on the quality of the approximation (Remark 1) and present a local convergence analysis of our method (Theorem 1).

Remark 1 (Quality of Quasi-Policy Approximations). Provided quadratic objectives and linear constraints, the optimal policy is linear [25], and the quasi-policy approximation is exact. In this case, our method converges in a single iteration for step size $\alpha = 1$. For games with nonquadratic objectives and nonlinear constraints, the quasi-policy provides a good approximation when higher-order policy gradients of the optimal policy are small relative to first-order gradients, as suggested by analyses of both feedback Nash [26] and Stackelberg games [27]. On an equality-constrained Stackelberg-chain example problem from [27], our solver reproduces the equilibrium of that paper's specialized solver.

Theorem 1. *Let \mathbf{s}_0 be the initialization of Alg. 1 and define the sub-level set $\mathcal{F} := \{\mathbf{s} : \|\pi(\mathbf{s})\|_2 \leq \|\pi(\mathbf{s}_0)\|_2\}$. Denote by $\nabla\pi(\mathbf{s})$ the approximate Jacobian under the quasi-policy assumption, i.e., setting the high-order policy gradient terms to zeros. Suppose that \mathcal{F} is closed, that $\nabla\pi(\mathbf{s})$ is invertible for all $\mathbf{s} \in \mathcal{F}$, and that there exist constants $D, C > 0$ such that $\|(\nabla\pi(\mathbf{s}))^{-1}\|_2 \leq D$, $\|\nabla^*\pi(\mathbf{s}) - \nabla^*\pi(\tilde{\mathbf{s}})\|_2 \leq C\|\mathbf{s} - \tilde{\mathbf{s}}\|_2$, $\forall \mathbf{s}, \tilde{\mathbf{s}} \in \mathcal{F}$. Denote by $\nabla^*\pi(\mathbf{s})$ the exact Jacobian of π , and assume that there exists $\delta > 0$ such that $\|\nabla^*\pi(\mathbf{s}) - \nabla\pi(\mathbf{s})\|_2 \leq \delta$, $\forall \mathbf{s} \in \mathcal{F}$, and $D \cdot \delta < 1$. Let $\Delta \mathbf{s} = -(\nabla\pi(\mathbf{s}))^{-1}\pi(\mathbf{s})$ be the inexact Newton step in Alg. 1. Then, for the k -th iteration solution $\mathbf{s}_k \in \mathcal{F}$, there exists $\alpha \in [0, 1]$ such that*

1. *if $\|\pi(\mathbf{s}_k)\|_2 > \frac{1-D\delta}{D^2C}$, then $\|\pi(\mathbf{s}_k + \alpha\Delta\mathbf{s}_k)\|_2 \leq \|\pi(\mathbf{s}_k)\|_2 - \frac{1-D\delta}{2D^2C}$;*
2. *if $\|\pi(\mathbf{s}_k)\|_2 \leq \frac{1-D\delta}{D^2C}$, then $\|\pi(\mathbf{s}_k + \alpha\Delta\mathbf{s}_k)\|_2 \leq \frac{1}{2}(1 + D\delta)\|\pi(\mathbf{s}_k)\|_2$.*

Proof. By the fundamental theorem of calculus, we have

$$\begin{aligned} \|\pi(\mathbf{s} + \alpha\Delta\mathbf{s})\|_2 &= \left\| \pi(\mathbf{s}) + \int_0^1 \nabla^*\pi(\mathbf{s} + \tau\alpha\Delta\mathbf{s})\alpha\Delta\mathbf{s}d\tau \right\|_2 \\ &\leq \|\pi(\mathbf{s}) + \alpha\nabla^*\pi(\mathbf{s})\Delta\mathbf{s}\|_2 + \left\| \int_0^1 (\nabla^*\pi(\mathbf{s} + \tau\alpha\Delta\mathbf{s}) - \nabla^*\pi(\mathbf{s}))\alpha\Delta\mathbf{s}d\tau \right\|_2 \\ &\leq \|\pi(\mathbf{s}) - \alpha\nabla^*\pi(\mathbf{s})(\nabla\pi(\mathbf{s}))^{-1}\pi(\mathbf{s})\|_2 + \|\alpha\Delta\mathbf{s}\|_2 \cdot \int_0^1 C\|\alpha\tau\Delta\mathbf{s}\|_2d\tau \quad (26) \\ &\leq (1 - \alpha)\|\pi(\mathbf{s})\|_2 + \alpha\delta D\|\pi(\mathbf{s})\|_2 + \|\alpha\Delta\mathbf{s}\|_2 \cdot \int_0^1 C\|\alpha\tau\Delta\mathbf{s}\|_2d\tau \\ &\leq (1 - \alpha(1 - \delta D))\|\pi(\mathbf{s})\|_2 + \frac{1}{2}\alpha^2CD^2\|\pi(\mathbf{s})\|_2^2. \end{aligned}$$

The right-hand side is a quadratic function of α and is minimized by $\alpha^* = \frac{1-\delta D}{D^2 C \|\pi(\mathbf{s})\|_2}$. This quadratic term also ensures that for all $\alpha \in [0, \alpha^*]$, we have $\mathbf{s} + \alpha \Delta \mathbf{s} \in \mathcal{F}$, provided $\mathbf{s} \in \mathcal{F}$. Suppose that $\|\pi(\mathbf{s})\|_2 > \frac{1-\delta D}{D^2 C}$, then we have $\alpha^* < 1$ is a feasible stepsize and therefore $\|\pi(\mathbf{s} + \alpha^* \Delta \mathbf{s})\|_2 \leq \|\pi(\mathbf{s})\|_2 - \frac{(1-\delta D)^2}{2D^2 C}$. Moreover, when $\|\pi(\mathbf{s})\|_2 \leq \frac{1-\delta D}{D^2 C}$, the right handside of (26) is minimized by the maximum allowable step size $\alpha = 1$, so we can strengthen the bound by substituting $\|\pi(\mathbf{s})\|_2 \leq \frac{1-\delta D}{D^2 C}$ to find $\|\pi(\mathbf{s} + \alpha \Delta \mathbf{s})\|_2 \leq \frac{1}{2}(1 + \delta D)\|\pi(\mathbf{s})\|_2$. \square

Remark 2 (Implications of Theorem 1). Theorem 1 clarifies structural conditions under which Alg. 1 is expected to converge locally. In particular, when robot objectives include strong quadratic control regularization and the dynamics are locally close to linear, higher-order policy sensitivities are small, making the assumptions of Theorem 1 well justified. This observation provides practical guidance for modeling objectives and constraints to promote convergence of our solver in mixed-hierarchy games. The constants in Theorem 1 are regularity conditions on the problem and need not be computed by the user.

Computational Complexity. The total runtime complexity is $\mathcal{O}(K_{\max} \cdot (N^2 + H(n_s)))$, where N^2 accounts for worst-case (Stackelberg chain) KKT construction and $H(n_s)$ for the runtime of the linear system solver for the KKT system.

5 Experiments

We demonstrate our method’s efficacy in hardware and simulated experiments.

5.1 Target Guarding: A Three-Robot Hardware Experiment

We describe a three-robot ($N = 3$) dynamic mixed-hierarchy game between a pursuer (R1), a guard (R2), and a target (R3). The pursuer (R1) wants to achieve a similar position to the target (R3), but wishes to stay away from the guard (R2). The guard (R2) wants to keep the pursuer (R1) away from the target (R3), and the target wants to stay close to the guard (R2) while reaching a goal state $\mathbf{x}_g \in \mathbb{R}^{n_s}$ by the end of the planning horizon. All robots have a control effort cost, and component terms are weighted by $\mathbf{w}^i \in \mathbb{R}^3$.

$$J^1(\mathbf{z}) = \sum_{t=1}^T \mathbf{w}^1 \cdot [\|\mathbf{x}_t^3 - \mathbf{x}_t^1\|_2^2, -\|\mathbf{x}_t^2 - \mathbf{x}_t^1\|_2^2, \|\mathbf{u}_t^1\|_2^2]^\top, \quad \mathbf{w}^1 = [2, 1, 1.25]^\top,$$

$$J^2(\mathbf{z}) = \sum_{t=1}^T \mathbf{w}^2 \cdot [\|\mathbf{x}_t^3 - \mathbf{x}_t^2\|_2^2, -\|\mathbf{x}_t^3 - \mathbf{x}_t^1\|_2^2, \|\mathbf{u}_t^2\|_2^2]^\top, \quad \mathbf{w}^2 = [0.5, 1, 0.25]^\top,$$

$$J^3(\mathbf{z}) = \sum_{t=1}^T \mathbf{w}^3 \cdot \left[\frac{1}{T} \|\mathbf{x}_T^3 - \mathbf{x}_g\|_2^2, \|\mathbf{x}_t^3 - \mathbf{x}_t^2\|_2^2, \|\mathbf{u}_t^3\|_2^2 \right]^\top, \quad \mathbf{w}^3 = [10, 1.25, 0.1]^\top.$$

We define hierarchy structure $\mathcal{G} = ([3], \{(2, 1)\})$ to model the guard leading the target, while both have a Nash relationship with the pursuer.

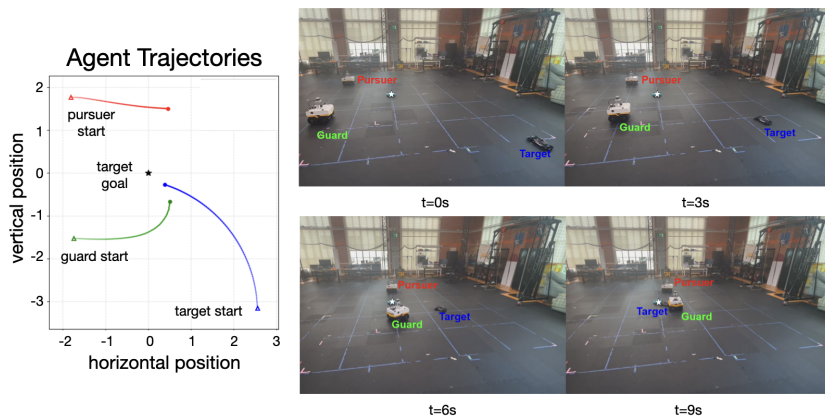


Fig. 4: We play a target-guarding game in hardware with three robots. Robot 2 (Clearpath Jackal, green) acts as the guard, leading Robot 1 (Clearpath Jackal, pursuer, red) and Robot 3 (NVIDIA JetRacer, target, blue), which seeks a goal. All robots execute plans generated by our solver in a receding-horizon fashion.

Experimental Setup We evaluate this example in hardware with three ground robots: two agents (Clearpath Jackals) acting as the pursuer (R1) and guard (R2), and an NVIDIA JetRacer as the target (R3). The interaction is modeled with horizon $T = 10$ and sampling period $\Delta t = 0.1$ s, and solved³ in a receding-horizon manner using simulator-provided state measurements. We convert the resulting double-integrator controls to differential drive controls for each robot.

We highlight a number of noteworthy behaviors in Fig. 4. R3 makes a wide turn to move towards the goal and get closer to R2. R1 pursues a course to meeting R3, but diverts from it due to the presence of R2. R2, using knowledge of the target’s actions, chooses to intersect with R1 and R3’s paths towards the origin, leading R1 to continue its course and R3 diverting its course to avoid R2. The solver is called 93 ± 2.8 times during execution and each call takes 13.6 ± 2.8 ms (mean \pm standard deviation). This demonstration shows that our solver runs in real-time, making it practical for real-world deployment.

5.2 Merging into a Convoy: A Four-Vehicle Simulation

We run the experiments described in the running example from Section 3.4 on a game with horizon $T = 10$, and discretization period $\Delta t = 0.5$ s. We choose the following hierarchy structures: Nash, mixed hierarchy A ($1 \rightarrow 2, 2 \rightarrow 4$), mixed hierarchy B ($1 \rightarrow 3, 1 \rightarrow 2, 2 \rightarrow 4$), and Stackelberg chain ($1 \rightarrow 3 \rightarrow 2 \rightarrow 4$) as depicted in Fig. 2. The first hierarchy structure defines simultaneous play among all vehicles (Nash). The second structure describes a situation in which the entire convoy is a Stackelberg chain led by vehicle 1, and they play a Nash game with vehicle 3. In the third structure, vehicle 1 plays first and then its direct followers 2 and 3 play simultaneously in response. This could simulate

³ CPU version: 13th Gen Intel Core i9-13900HX (2.20 GHz). Memory: 32 GB.




Fig. 5: We report the mean and standard deviation of the merit function value $\|\pi(\mathbf{s}_k)\|_2$ over 10 runs with random initial states under a fixed hierarchy structure, corroborating Theorem 1 via monotone decay and local exponential convergence of $\|\pi(\mathbf{s}_k)\|$.

a real-world scenario in which the merging vehicle responds to leader vehicle 1’s strategy but has not yet resolved whether to yield to vehicle 2. The fourth describes turn-based play among vehicles (Stackelberg chain).

Our experimental results are shown in Fig. 1. These scenarios illustrate how hierarchy structure influences merging behavior. When vehicle 3 is a follower of vehicle 1, vehicle 1 exploits its leadership position by accelerating to remain ahead after the merge. In contrast, under a Nash relationship between vehicles 1 and 3, vehicle 1 behaves more conservatively and yields, allowing vehicle 3 to merge ahead. Despite these differing outcomes, all equilibria in our experiments (e.g., Nash, Stackelberg, and mixed hierarchies) result in safe interactions among all vehicles, including vehicles 2 and 4. Finally, Fig. 5 provides empirical corroboration of the convergence results in Theorem 1, showing constant decrease further from the equilibrium and exponential convergence near it.

Overall, this convoy example demonstrates that mixed-hierarchy games naturally encode asymmetric information access within an optimization-based framework for decision-making. As the hierarchy structure changes, our solver converges to distinct, interpretable behavior for this simulated example with hundreds of variables, nonquadratic objective functions, and nonlinear constraints.

6 Conclusion

We define mixed-hierarchy games, which contain agents with both Nash and Stackelberg relationships to one another. We present a formal definition of the relevant equilibrium concept for this type of game, and develop a tractable procedure for constructing the KKT conditions of mixed-hierarchy games using quasi-policy approximations to embed followers’ best response policies into their leaders’ optimization problems. We propose an inexact Newton’s method for efficiently solving those approximate KKT systems. We provide a theoretical proof and an empirical study for the local exponential convergence of our algorithm to an approximate mixed-hierarchy equilibrium. We evaluate our method on two complex multi-agent scenarios, and showcase its broad functionality and real-time performance using our open-source solver implementation. Future work will explore the impact of different information structures on robot behavior, incorporate inequality constraints via barrier methods, and relax the forest structure assumption (Assumption 2) to allow multi-leader mixed-hierarchies.

References

1. Altabaa, A., Yang, Z.: On the role of information structure in reinforcement learning for partially-observable sequential teams and games. *Advances in Neural Information Processing Systems* **37**, 6648–6710 (2024)
2. Aumann, R.J.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica: Journal of the Econometric Society* pp. 1–18 (1987)
3. Başar, T., Olsder, G.J.: *Dynamic noncooperative game theory*. SIAM (1998)
4. Başar, T.: Equilibrium strategies in dynamic games with multi-levels of hierarchy. *Automatica* **17**(5), 749–754 (1981)
5. Becker, T., Sunberg, Z.: Bridging the gap between partially observable stochastic games and sparse POMDP methods. arXiv preprint arXiv:2405.18703 (2024)
6. Blair, C.: The computational complexity of multi-level linear programs. *Annals of Operations Research* **34** (1992)
7. Bracken, J., McGill, J.T.: Mathematical programs with optimization problems in the constraints. *Operations Research* **21**(1), 37–44 (1973)
8. Brandenburger, A., Dekel, E., Geanakoplos, J.: Correlated equilibrium with generalized information structures. *Games and Economic Behavior* **4**(2), 182–201 (1992)
9. Candler, W., Norton, R.D.: Multi-level programming: theory, algorithms, and applications. *European Journal of Operational Research* **1**(1), 1–11 (1977)
10. Carvalho, M., Dragotto, G., Feijoo, F., Lodi, A., Sankaranarayanan, S.: When Nash meets Stackelberg. *Management Science* **70**(10), 7308–7324 (2024)
11. Castiglioni, M., Marchesi, A., Gatti, N., et al.: Be a leader or become a follower: The strategy to commit to with multiple leaders. In: *IJCAI*. pp. 123–129 (2019)
12. Chang, Y., Erera, A.L., White III, C.C.: Value of information for a leader–follower partially observed Markov game. *Annals of Operations Research* **235**(1), 129–153 (2015)
13. Emery-Montemerlo, R., Gordon, G., Schneider, J., Thrun, S.: Approximate solutions for partially observable stochastic games with common payoffs. In: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004*. pp. 136–143. IEEE (2004)
14. Faísca, N.P., Saraiva, P.M., Rustem, B., Pistikopoulos, E.N.: A multi-parametric programming approach for multilevel hierarchical and decentralised optimisation problems. *Computational Management Science* **6**(4), 377–397 (2009)
15. Farina, G., Sandholm, T.: Model-free online learning in unknown sequential decision making problems and games. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 35, pp. 5381–5390 (2021)
16. Filar, J., Vrieze, K.: *Competitive Markov decision processes*. Springer Science & Business Media (2012)
17. Fortuny-Amat, J., McCarl, B.A.: A representation of linear complementarity problems as linear programs. *Journal of the Operational Research Society* **32**(9), 783–792 (1981)
18. Fridovich-Keil, D., Ratner, E., Peters, L., Dragan, A.D., Tomlin, C.J.: Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In: *IEEE International Conference on Robotics and Automation (ICRA)* (2020)
19. Hansen, E.A., Bernstein, D.S., Zilberstein, S.: Dynamic programming for partially observable stochastic games. In: *AAAI*. vol. 4, pp. 709–715 (2004)
20. Jeroslow, R.G.: The polynomial hierarchy and a simple model for competitive analysis. *Mathematical Programming* **32**(2), 146–164 (1985)

21. Khan, H., Fridovich-Keil, D.: Leadership inference for multi-agent interactions. *IEEE Robotics and Automation Letters (RA-L)* (2024)
22. Kozuno, T., Ménard, P., Munos, R., Valko, M.: Learning in two-player zero-sum partially observable Markov games with perfect recall. *Advances in Neural Information Processing Systems* **34**, 11987–11998 (2021)
23. Krantz, S.G., Parks, H.R.: The implicit function theorem: history, theory, and applications. Springer Science & Business Media (2002)
24. Kulkarni, A.A., Shanbhag, U.V.: A shared-constraint approach to multi-leader multi-follower games. *Set-Valued and Variational Analysis* **22**(4), 691–720 (2014). <https://doi.org/10.1007/s11228-013-0271-4>
25. Laine, F.: Mathematical program networks. arXiv preprint arXiv:2404.03767 (2024)
26. Laine, F., Fridovich-Keil, D., Chiu, C.Y., Tomlin, C.: The computation of approximate generalized feedback Nash equilibria. *SIAM Journal on Optimization* (2023)
27. Li, J., Sojoudi, S., Tomlin, C., Fridovich-Keil, D.: The Computation of Approximate Feedback Stackelberg Equilibria in Multi-Player Nonlinear Constrained Dynamic Games. *SIAM Journal on Optimization* (2024)
28. Li, X., Wu, N., Lei, L.: Nash-Stackelberg-Nash three-layer mixed game optimal control strategy for multi-integrated energy systems considering multiple uncertainties. *Energy* **320**, 135418 (2025)
29. Littman, M.L.: Markov games as a framework for multi-agent reinforcement learning. In: *Machine Learning Proceedings*, pp. 157–163. Elsevier (1994)
30. Liu, Q., Szepesvári, C., Jin, C.: Sample-efficient reinforcement learning of partially observable Markov games. *Advances in Neural Information Processing Systems* **35**, 18296–18308 (2022)
31. Luo, Z.Q., Pang, J.S., Ralph, D.: *Mathematical Programs with Equilibrium Constraints*. Cambridge University Press (1996)
32. Mohapatra, P.S., Reddy, P.V., Zaccour, G.: Feedback Stackelberg-Nash equilibria in difference games with quasi-hierarchical interactions and inequality constraints. arXiv preprint arXiv:2504.15019 (2025)
33. Nash, J.F.: Non-cooperative games. In: *The Foundations of Price Theory Vol 4*, pp. 329–340. Routledge (2024)
34. Nocedal, J., Wright, S.J.: *Numerical optimization*. Springer (1999)
35. Peters, L.: *SymbolicTracingUtils.jl*. <https://github.com/JuliaGameTheoreticPlanning/SymbolicTracingUtils.jl> (2025), GitHub repository, last accessed 2025-12-02
36. Sato, R., Tanaka, M., Takeda, A.: A gradient method for multilevel optimization. *Advances in Neural Information Processing Systems* **34**, 7522–7533 (2021)
37. Shafiei, A., Kungurtsev, V., Marecek, J.: Trilevel and multilevel optimization using monotone operator theory. *Mathematical Methods of Operations Research* **99**(1), 77–114 (2024)
38. Shapley, L.S.: Stochastic games. *Proceedings of the National Academy of Sciences* **39**(10), 1095–1100 (1953)
39. Sherali, H.D., Soyster, A.L., Murphy, F.H.: A Stackelberg-Nash-Cournot equilibrium model for oligopolistic markets. *Operations Research* **31**(2), 254–272 (1983)
40. Shi, W., Shimizu, K.: A general formulation of the bilevel programming problem. *Journal of Information and Optimization Sciences* **2**(1), 51–68 (1981)
41. Stackelberg, H.v.: *The Theory of the Market Economy*. William Hodge (1952)
42. Zinkevich, M., Johanson, M., Bowling, M., Piccione, C.: Regret minimization in games with incomplete information. *Advances in Neural Information Processing Systems* **20** (2007)