

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

Paparella, F., Lucas, C., Cenedese, C., & Salazar, M. (2025). Analysis of Nash and Stackelberg Equilibria of Autonomous Mobility-on-Demand Systems in Mixed Traffic. In *Proceedings of the 23rd European Control Conference (ECC 2025)* (pp. 2964-2969). IEEE. <https://doi.org/10.23919/ECC65951.2025.11187086>

Important note

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

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

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

Takedown policy

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

Analysis of Nash and Stackelberg Equilibria of Autonomous Mobility-on-Demand Systems in Mixed Traffic

Fabio Paparella^{1*}, Clim Lucas^{1*}, Carlo Cenedese^{2,3}, Mauro Salazar¹

Abstract— This paper analyzes the differences between Nash and Stackelberg equilibria of Autonomous Mobility-on-Demand (AMoD) systems in mixed traffic conditions, whereby self-driving robotaxis provide on-demand mobility, possibly pooling users together, while sharing the road with selfish private cars. In particular, we first introduce the optimal fleet routing problem in mixed traffic conditions, considering a car-road network where also private, selfish vehicles are present. Second, we model the interactions between the centrally controlled AMoD fleet and the private cars with two equilibrium formulations: the first corresponds to a Nash equilibrium, where each agent (the fleet and the private users) reacts to the other agent's action until convergence is reached. For the second approach, corresponding to a Stackelberg equilibrium, the leader (the fleet) can predict the best response of the follower (private users) and plan the strategy accordingly. The results of a case study of Sioux Falls, USA, indicate that the two equilibria are very similar in terms of the fleet's objective function, suggesting that even if the operator can predict the best response of the private users, no benefit arises.

I. INTRODUCTION

Recently, companies such as Waymo [1] have begun deploying AMoD systems in cities like San Francisco and Los Angeles. These systems consist of centrally managed fleets of vehicles that provide on-demand transportation to users. By 2024, Waymo alone was operating over 100,000 trips per week, highlighting the influence that these innovative mobility solutions can have on urban areas. To improve efficiency, AMoD systems often encourage ride-pooling, where multiple users who share portions of their routes travel together in the same vehicle. These autonomous fleets share the roads with private drivers, who act selfishly to minimize their travel times. This interaction has been observed to lead to an equilibrium between the AMoD and the private users. While AMoD systems have the potential to improve the efficiency of trips via ride-pooling, they could also produce unintended negative consequences, such as increment of congestion due to deadhead trips and rebalancing. To assess the impact of AMoD systems, it is crucial to develop tools that analyze the effect that AMoD systems have on the driving habits of private vehicle users, and vice versa.

* These authors contributed equally to this paper.

This publication is part of the project NEON, project number 17628, financed by the Dutch Research Council. This work was supported as a part of NCCR Automation, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 51NF40.225155).

¹ Control Systems Technology section, Eindhoven University of Technology, Eindhoven, The Netherlands {f.paparella,m.r.u.salazar}@tue.nl, c.lucas@student.tue.nl

² Department of Mechanical and Process Engineering at ETH Zürich, Zürich, Switzerland. Email: ccenedese@ethz.ch

³ the Delft Center for Systems and Controls, Delft University of Technology, 2628 CD Delft, The Netherlands, ccenedese@tudelft.nl

Related works: In literature, several works such as [2]–[4] investigated Stackelberg equilibria within a mobility network. These works, in combination with centrally operated AMoD systems [5], pave the way to mixed-traffic systems, whereby a centrally controlled operator coexists with selfish users. By tactically routing the fleet, the overall congestion, and thereby the total travel time, was significantly decreased, as well as the number of cars needed to serve a group of commuters. From here, different extensions have been analyzed, from the interaction with the power grid [6] to ride-pooling [7], [8] and mobility justice [9].

Bi-level (or Stackelberg) games capture games in which the Nash Equilibrium (NE) arising from the decision-making process of the agents (or followers) involved is influenced directly by an external operator (the leader) [10]. Finding an exact Stackelberg Equilibrium (SE) requires solving a complex optimization problem that is usually ill-posed thus few solutions are available, and they are usually restricted to very particular scenarios. Thus, in the literature on routing problems an auxiliary NE seeking problem is iteratively solved instead, where the leader is modeled as a special agent taking part in the decision-making process of the followers. This allows us to employ scalable algorithms, where the AMoD operator and the private users iteratively find their best routing action, which allows them to efficiently compute the NE [11]. In this context, some works analyzed the presence of private vehicles and their interaction with (intermodal) AMoD systems computing a NE in an iterative manners [12]. This results in a purely reactive strategy. However, it is reasonable to assume that the AMoD operator, e.g., Waymo, may plan the fleet routes while predicting the response of the private users, essentially creating a Stackelberg game where the fleet operator is the leader and the private users are the following agent [13]. The Best Intervention in Games using Hypergradients (BIG Hype) algorithm has been proposed in [14] to find local SEs that maintain the hierarchical structure of the problem and satisfy the constraints of both the operator and the users. BIG Hype addresses the convergence challenges through iterative and implicit differentiation, using fixed-point iteration to approximate the equilibrium in the lower-level game. Linear equations derived from these approximations guide the hypergradient descent, while termination criteria mitigate inexactness, ensuring convergence via non-smooth analysis tools [15]. This unique approach is particularly promising for bi-level games where traditional methods fall short, such as designing incentive schemes for optimal routing policies [16]. Another relevant work has been recently proposed [17], where the authors study the leader's strategy in Stackelberg pricing games. Yet these methods were not applied to the operational planning

of transport systems such as AMoD.

To the best of the authors' knowledge, comparing Nash and Stackelberg equilibria of AMoD systems in mixed car traffic is an open question due to the complexity and dimension of the problem. This study can prove the reliability of the classical bi-level solutions in capturing the hierarchical nature of the ride-pooling network flow problem when AMoD is involved.

Main contributions: The contributions of this paper are threefold: i) We cast the AMoD problem with ride-pooling in mixed traffic conditions as a SE finding problem, emphasizing the hierarchical structure of the problem; ii) We specialize the BIG Hype algorithm to compute the SE of the AMoD system; iii) We compare the numerical results obtained with a bi-level optimization method, showing that the two formulations give comparable results.

Organization: The remainder of the paper is organized as follows: In Section II we briefly introduce the optimization problem to characterize the AMoD system in a mixed-traffic environment [18]. Then, in Section III we leverage the BIG Hype algorithm to solve this problem and compute the SE. Finally, in Section IV we present a case study of Sioux Falls, USA. Lastly, we conclude in Section V with some remarks and future work.

Notation: We denote a column vector of appropriate dimension with $\mathbf{1}$, while $\mathbf{0}$ denotes a vector of zeros. The i th component of a vector v is denoted by v_i and the entry (i, j) of a matrix A is denoted by A_{ij} . Given two matrices A, B of compatible dimensions, $[A, B]$ and $[A; B]$ respectively denote their horizontal and vertical concatenations. $\text{vec}(A)$ gives the column concatenation of all $\text{col}_i(A)$, where $\text{col}_i(A)$ gives the i th column of a matrix A . $|\mathcal{S}|$ is used to denote the cardinality of set \mathcal{S} . We use $\mathbf{J}_1 F(x, y)$ and $\mathbf{J}_2 F(x, y)$ to denote the partial Jacobians of a function F with respect to the first and second argument respectively. Similarly, the partial gradients are denoted as $\nabla_1 F$ and $\nabla_2 F$.

II. PROBLEM FORMULATION

In this section, we briefly revise the network flow model of a mixed-traffic AMoD system first introduced in [18].

A. Preliminaries

The mobility network is modeled as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where the set of nodes \mathcal{V} , and of edges \mathcal{A} of the graph represent the road intersections and links, respectively. The incidence matrix of \mathcal{G} is denoted by the matrix $B \in \{-1, 0, 1\}^{|\mathcal{V}| \times |\mathcal{A}|}$, see [19, Ch. 8], where entry B_{ia} is -1 or 1 for outgoing and incoming arcs $a \in \mathcal{A}$ to node $i \in \mathcal{V}$, respectively. The function $t : \mathbf{R}_+^{|\mathcal{A}|} \rightarrow \mathbf{R}_+^{|\mathcal{A}|}$ computes the time necessary to traverse the arcs in the network given for each arc $a \in \mathcal{A}$ the associated flow of vehicles s_a , and capacity m_a . Thus, it reads as

$$t_a(s_a) = t_a^0 \left(1 + \eta \left(\frac{s_a}{m_a} \right)^\zeta \right), \quad (1)$$

where $t_a^0, \eta, \zeta > 0$. Next, we define the concept of travel request.

Definition 2.1 (Travel Requests): A travel request is defined as the tuple $r = (o, d, \lambda) \in \mathcal{V} \times \mathcal{V} \times \mathbb{R}_{>0}$, in which λ is the number of users traveling from the origin o to the

destination $d \neq o$ per unit time. Define the set of requests as $\mathcal{R} := \{r_m\}_{m \in \mathcal{M}}$, where $\mathcal{M} = \{1, \dots, M\}$ is the set of indices of those travel requests.

In a ride-pooling setting, whereby users can either ride-pool or ride alone, we denote the ride-pooling penetration rate by $\psi \in [0, 1]$. Hence, $\psi = 0$ means no user is willing to ride-pool, and $\psi = 1$ that every user is willing to ride-pool. Following [7], it is possible to create an equivalent set of travel requests, which is a function of the maximum waiting time $\bar{t} > 0$ that users are willing to wait for another user, and the maximum delay $\bar{\delta} > 0$ experienced by users with respect to the shortest path [7]. The new set $\mathcal{R}^{\text{mo}} = f(\mathcal{R}, \psi, \bar{t}, \bar{\delta})$ thus represents the expected number of users that can be ride-pooled together, including the ones that cannot be ride-pooled together because of waiting and delay constraints. The definition of f , as well as the full procedure for computing the new ride-pooling, set \mathcal{R}^{mo} is in [7]. We also define the users that do not want to ride-pool by $\mathcal{R}^{\text{so}} = (1 - \psi)\mathcal{R}$, i.e., the travel requests are the same, but with a coefficient equal to $(1 - \psi)\lambda_m$.

Then, we define two optimization variables, the flow of vehicles with users on board, and the flow of vehicles without users on board, called the rebalancing flow, which are denoted as $X = [X^{\text{so}}, X^{\text{mo}}]$ and x^r respectively. The flow of single-occupancy rides is $X^{\text{so}} \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{V}|}$, which is further detailed as $X^{\text{so}} := [x^{\text{so},1}, x^{\text{so},2}, \dots, x^{\text{so},|\mathcal{V}|}]$, where $x^{\text{so},i} \in \mathbb{R}^{|\mathcal{A}|}$ is the flow of the users that share the same origin $i \in \mathcal{V}$. The flow of and multi-occupancy rides is defined analogously, $X^{\text{mo}} \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{V}|}$. The entry a of a vector flow $x_a^{\text{so},i}$ represents the number of single-occupancy vehicles per unit time traversing arc a that originated from node i . Then, the rebalancing flow is denoted by $x^r \in \mathbb{R}^{|\mathcal{A}|}$, where the entries represent the rebalancing flow on each arc.

In a similar fashion to the vehicle flows, we define an additional optimization variable, the private vehicle flows, $X^p \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{V}|}$ with $X^p := [x^{p,1}, x^{p,2}, \dots, x^{p,|\mathcal{V}|}]$.

The total flow of single-occupancy and multi-occupancy vehicles is denoted as $x^u := \sum_i x^{\text{so},i} + x^{\text{mo},i}$. Finally, the total flow on the arcs is defined as $\bar{x} \in \mathbb{R}^{|\mathcal{A}|}$:

$$\bar{x} := x^u + x^r + \sum_{i \in \mathcal{V}} x^{p,i}.$$

B. AMoD Formulation in Mixed-traffic Conditions

Following utilitarian distributive principles, we assume that the main goal of the AMoD operator is to find the car flows X, x^r that minimize the total time traveled by the fleet or the users. To achieve this, we define the network flow problem as follows:

Problem 1 (Ride-pooling Network Flow Problem):

Given a road graph \mathcal{G} and an origin-destination matrix D , the optimal vehicle flow X and rebalancing flow x^r result from

$$\min_{X, x^r} J(X, x^r) = t(X, x^r, X^p)^\top (X \mathbf{1} + \rho x^r)$$

$$\text{s.t. } BX = D$$

$$B(X \mathbf{1} + x^r) = \mathbf{0}$$

$$X, x^r \geq \mathbf{0}$$

$$X^p \in \text{TAP}(\mathcal{R}^p, X, x^r).$$

where $\mathbf{1}$ and $\mathbf{0}$ are column vectors of appropriate dimensions of ones and zeros respectively. Note that compared to (1), the total flow on each arc a , s_a , is decomposed in AMoD flow X , rebalancing flow x^r and private flow X^p , e.g., $s = X\mathbf{1} + x^r + X^p\mathbf{1}$. The factor $\rho \in [0, 1]$ is a parameter that weights the operational costs of the vehicles with respect to the cost of the users' time. The OD matrix is defined as $D := [D^{\text{so}}, D^{\text{mo}}]$, where the two matrices $D^{\text{so}}, D^{\text{mo}}$ are fully characterized given the two sets, $\mathcal{R}^{\text{so}}, \mathcal{R}^{\text{mo}}$. Each matrix can be computed by using the following definition [7]:

$$D_{ij} = \begin{cases} \alpha_m, & \exists m \in \mathcal{M} : o_m = j \wedge d_m = i \\ -\sum_{k \neq j} D_{kj}, & i = j \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where m is an index that spans through the whole set of travel requests. D^p is the OD matrix of the private users, obtained in the same way starting from a set of travel requests of private users \mathcal{R}^p . The objective represents the vehicle hours traveled by the fleet, which has also been proven to be equivalent to the minimum number of vehicles required to implement the flows [5]. The first constraint imposes the conservation of flow throughout the system, the second constraint takes care of the rebalancing of vehicles, and the third constraint makes sure that the flows are non-negative. Lastly, the operator's strategy is implicitly influenced by the self-centered strategy of private vehicles via the final constraint. The private vehicles are assumed to follow a User-Centric (UC) routing policy, similar to the Traffic Assignment Problem (TAP) [20]. Such formulation finds the flows that reach a Wardrop equilibrium [21], i.e., every user selfishly minimizes its own travel time. Following [12], the TAP is thus expressed as

$$\min_{X^p} J(X^p) = \int_{x^u + x^r}^{\bar{x}} t(s)^\top ds \quad (3a)$$

$$\text{s.t. } BX^p = D^p \quad (3b)$$

$$X^p \geq \mathbf{0}. \quad (3c)$$

The TAP is subjected to the same conservation of flow and feasibility constraints as Problem 1. However, rebalancing is not necessary as private vehicles do not rebalance.

III. LEVERAGING BIG HYPE

Previous work has often tackled Problem 1 as a bi-level optimization problem [12], [18], [22]. In other words, each agent reacts to the response of the other agent until convergence, that is, NE. This approach does not have any convergence guarantees, but in practice it always converges to a consistent solution in a small number of iterations. This method results in purely reactive strategies. It is interesting to analyze how the result changes if the operator can predict what the best response of the private users is, and plan accordingly. This translates into a SE, where the AMoD operator is the leader and the private users are the followers. To tackle the problem as a Stackelberg game it is possible to leverage the BIG Hype methodology [14], which we briefly describe.

First, we propose a generalized way of writing a Stackelberg game. Then, we apply it to Problem 1 and we

comment on the applicability of the algorithm by evaluating the necessary assumptions.

A. General Stackelberg Formulation

The problem definition of a Stackelberg optimization problem can be generalized as

$$\min_{\mathbf{x}, \mathbf{y}} \varphi(\mathbf{x}, \mathbf{y}) \quad (4a)$$

$$\text{s.t. } \mathbf{x} \in \mathcal{X} \quad (4b)$$

$$\mathbf{y} \in \text{SOL}(F(\mathbf{x}, \cdot), \mathcal{Y}(\mathbf{x})), \quad (4c)$$

where \mathbf{x} and \mathbf{y} are the decision variables of the leader and followers respectively, \mathcal{X} is the leader's strategy set, $\mathcal{Y}(\mathbf{x})$ is the followers' strategy set, SOL is the solution of the variational inequality problem, and F is the *pseudo-gradient* (PG) mapping defined as

$$F(\mathbf{x}, \mathbf{y}) := (\nabla_{\mathbf{y}_i} f_i(\mathbf{x}, \mathbf{y}))_{i \in \mathcal{N}}, \quad (5)$$

where f is the objective function of the lower level. The solution to the variational inequality finds the unique NE of the lower level, given the strategy of the leader and the resulting feasible set. It acts as a function that maps strategies \mathbf{x} into a General Nash Equilibrium (GNE) \mathbf{y}^* . Contrary to the bi-level Nash optimization approach, the leader calculates the hypergradient and predicts the behavior of the follower when leveraging BIG Hype [14]. The hypergradient of (4) is derived using the chain rule and can be defined as

$$\hat{\nabla} \varphi(\mathbf{x}, \mathbf{y}^*(\mathbf{x})) = \nabla_1 \varphi(\mathbf{x}, \mathbf{y}^*(\mathbf{x})) + \mathbf{J}\mathbf{y}^*(\mathbf{x})^\top \nabla_2 \varphi(\mathbf{x}, \mathbf{y}^*(\mathbf{x})), \quad (6)$$

where $\mathbf{J}\mathbf{y}^*(\mathbf{x}) = \mathbf{s}$ is the sensitivity. Because \mathbf{s} can become computationally expensive for large-scale games [23], such as this mobility network, the sensitivity is approximated which results in an approximated hypergradient as well. To mitigate the errors arising from inexact estimates, accuracy bounds are imposed. To give insight into how the Stackelberg equilibrium is obtained, an overview of the critical components of the algorithm is highlighted.

B. BIG Hype Algorithm

The BIG Hype algorithm works by nesting two iterative loops. The inner loop uses the leader's strategy to update the local estimates of the lower-level equilibrium and sensitivity of the problem until the termination criteria are met. The equilibrium is found by using the fixed-point iteration [24, Alg. 12.1.1], indexed by $l \in \mathbb{L}$, for each agent¹

$$(\forall l \in \mathbb{L}) \quad \tilde{\mathbf{y}}_i^{l+1} = \mathbb{P}_{\mathcal{Y}_i(\mathbf{x})} \left[\tilde{\mathbf{y}}_i^l - \gamma F_i(\mathbf{x}, \tilde{\mathbf{y}}^l) \right], \quad (7)$$

where γ is a step size and

$$h(\mathbf{x}, \cdot) := \mathbb{P}_{\mathcal{Y}_i(\mathbf{x})} [\cdot - \gamma F(\mathbf{x}, \cdot)] \quad (8)$$

indicates the projection pseudo-gradient (PPG) mapping. As a validation, the solution to the variational inequality $\mathbf{y}^*(\mathbf{x})$ is also a fixed point of (7) [25]. The sensitivity is then approximated by differentiating (7) at $\mathbf{y}^*(\mathbf{x})$, such that

$$(\forall l \in \mathbb{L}) \quad \tilde{\mathbf{s}}_i^{l+1} = S_{2,i} \tilde{\mathbf{s}}_i^l + S_{1,i}, \quad (9)$$

¹The two agents are the fleet operator and the private users.

where $S_{1,i}$ and $S_{2,i}$ contain linear combinations of the partial Jacobians of the PG as a result of differentiating through the projection. The computational methods for this task are included in Appendix I and are abbreviated to

$$[S_{1,i}, S_{2,i}] \in \mathbf{J}^c h_i(\mathbf{x}, \tilde{\mathbf{y}}_i^{l+1}). \quad (10)$$

Using the estimated equilibrium and its corresponding sensitivity, the outer loop then approximates the hypergradient as shown in (6) and performs a projected hypergradient step. The BIG Hype algorithm is summarized in Alg. 1.

C. Optimal Stackelberg Routing

For the network flow problem of Problem 1, the decision variables of the leader, i.e. the AMoD operator, are defined as $\mathbf{x} := [\text{vec}(X); x^r]$, and the ones of the followers, $\mathbf{y} := \text{vec}(X^p)$. To improve the scalability of the problem, the private users are divided into multiple agents by grouping them by their origin. Thus, the variables for the individual agents become $\mathbf{y}_i = x^{p,i}$. We note \mathbf{y}_{-i} as the variables of the other agents. The upper-level objective function φ of (4a) is equal to the objective function of Problem 1. The constraints of (4b) coincide with the conservation of flow, the rebalancing of vehicles, and the feasibility of the AMoD decision variables, whereas (4c) captures the TAP constraint. Next, we comment on the properties of the upper- and lower-level problems.

Lemma 3.1: For any fixed AMoD decision vector $\mathbf{x} \in \mathcal{X}$, (3) admits a unique NE.

Algorithm 1 Summary of BIG Hype algorithm

Parameters : step sizes $\{\alpha^k, \beta^k\}_{k \in \mathbb{N}}$ and γ , tolerances $\{\sigma^k\}_{k \in \mathbb{N}}$, outer loop iterations \mathbb{K} .

Input: $\mathbf{x}, \mathbf{y}, s$.

Initialization: $\{k, l\} \leftarrow 0, \mathbf{x}^k \in \mathcal{X}, \mathbf{y}^k \in \mathbb{R}^{n_y}, s^k \in \mathbb{R}^{n_y \times n_x}$, termination = false.

Outer Loop :

for $k \in \mathbb{K}$ **do**

Leader's projected hypergradient step

$$\begin{aligned} \hat{\nabla} \varphi^k &= \nabla_1 \varphi(\mathbf{x}^k, \mathbf{y}^k) + (s^k)^\top \nabla_2 \varphi(\mathbf{x}^k, \mathbf{y}^k) \\ \mathbf{x}^{k+1} &= \mathbf{x}^k + \beta^k (\mathbb{P}_{\mathcal{X}}[\mathbf{x}^k - \alpha^k \hat{\nabla} \varphi^k] - \mathbf{x}^k) \end{aligned}$$

Inner Loop :

Approximate lower-level equilibrium and sensitivity

while termination=false **do**

for Agents $i \in \mathcal{N}$: **do**

Equilibrium seeking step: Solve (7)

PPG Jacobian update: Solve (10)

Sensitivity learning step: Solve (9)

end

termination = $\max\{\|\tilde{\mathbf{y}}^{l+1} - \tilde{\mathbf{y}}^l\|, \|\tilde{\mathbf{s}}^{l+1} - \tilde{\mathbf{s}}^l\|\} \leq \sigma^k$

$l \leftarrow l + 1$

end

$k \leftarrow k + 1$

end

Output: $\mathbf{x} = \mathbf{x}^k, \mathbf{y} = \tilde{\mathbf{y}}^{k,l}, s = \tilde{\mathbf{s}}^{k,l}$.

The proof is based on the fact that the lower-level objective function is continuously differentiable and convex on the




Fig. 1: VHT by the fleet (left), as well as normalized per number of AMoD users (right), both in minutes, as a function of the penetration rates ϕ, ψ .

feasible domain since typically $\eta = 0.15$ and $\zeta = 4$. On top of that, the feasible domain $\mathcal{Y}_i(\mathbf{x})$ is convex by definition, due to the combination of convex inequality constraints with the affine equality constraints. As a result, when the strategy profile \mathbf{y}^* is a solution of the variational inequality, it is also a NE [24, Prop. 1.4.2]. Moreover, the PG is μ -strongly monotone and L_F -Lipschitz continuous, which ensures the existence and uniqueness of the NE [24, Th. 2.3.3(b)].

Proposition 3.1: Let $\{\alpha^k\}_{k \in \mathbb{N}}$ be non-negative, non-summable and square-summable², let $\{\sigma^k\}_{k \in \mathbb{N}}$ be non-negative and satisfy $\sum_{k=0}^{\infty} \alpha^k \sigma^k < \infty$, and let γ be sufficiently small. Then any limit point of the sequence $\{\mathbf{x}^k\}_{k \in \mathbb{N}}$, generated by the outer loop of Alg. 1, is a composite critical point of (4).

To prove this claim, we refer back to the convex properties of the lower-level objective, as well as the monotone and Lipschitz continuous properties of the PG. Finally, because \mathcal{X} is convex for the same reason as \mathcal{Y} , and φ is continuously differentiable in (\mathbf{x}, \mathbf{y}) with its hypergradient being Lipschitz continuous, the BIG Hype is guaranteed to converge to a minimum [14]. However, φ is non-convex due to the product of $t(X, x^r, x^p)$ and $(X\mathbb{1} + \rho x^r)$. Because of this nonlinearity, the BIG Hype algorithm only ensures local optimality convergence of the leader's strategy set for this problem definition.

IV. NUMERICAL RESULTS

This section showcases a case study of Sioux Falls, USA, where we compare the numerical results of the two equilibria obtained, i.e. NE and SE. The data is from the Transportation Networks for Research repository [26]. To compute $\mathcal{R}, \mathcal{R}^p$, we scale the OD matrix by the AMoD penetration rate $\phi \in [0, 1]$ and $1 - \phi$ respectively. The total user demand on the network is independent from ϕ .

A. Case Study

This section first showcases the results obtained by solving Problem 1 as a SE with BIG Hype, and then analyzes the differences with respect to computing a NE, as already done in previous works by iteratively solving Problem 1 fixing the flow of private users X^p and the TAP [12], [18].

First, we analyze the Vehicle Hours Traveled (VHT) by the fleet, obtained by computing the SE of Problem 1 with

²The sequence $\{\alpha^k\}_{k \in \mathbb{N}}$ is non-summable if $\sum_{k=0}^{\infty} \alpha^k = \infty$ and square-summable if $\sum_{k=0}^{\infty} (\alpha^k)^2 < \infty$.




Fig. 2: Average travel time in minutes experienced by users as a function of penetration rates ϕ, ψ for AMoD users that do not ride-pool, private users, and AMoD users that ride-pool. The last subplot depicts the difference between AMoD users that ride-pool and the ones that do not.

BIG Hype, as shown in Fig. 1. For higher penetration rate ϕ , the VHT of the fleet increases because more users need to be served. However, at the same time, the VHT decreases for an increasing ψ . This is the result of more travel requests being matched together and therefore needing fewer resources. We observe that for a high ψ the relative improvement from increasing ϕ is greater, which is attributed to the fact that more people use the system, making it easier to efficiently match users [7]. Fig. 2 shows the average travel time of the different types of users separated, i.e., AMoD users who are not willing to ride-pool (*individual ride-sharing* users) and those who are willing, but depending on the ride-pooling allocation might still also travel alone (*ride-pooling* users), and *private* users. We highlight that the selfish nature of private users will always make them the quickest of all network users. Lastly, from the right plot of Fig. 1, as well as the upper plots of Fig. 2, we recognize that an increase in ϕ does not necessarily coincide with an improvement in VHT or average travel time, even if they are centrally operated. This characteristic is associated with the need to rebalance the AMoD fleet, contrary to the private vehicles.

When comparing the found SE with the NE in Fig. 3, the results have strong similarities. Table I further emphasizes the agreement between the two equilibria. Here, the VHT per individual group, as well as the weighted total VHT on the network, for given values of ϕ, ψ are compared using the NE as benchmark. The numerical results indicate that the two methodologies are substantially equivalent in terms of the vehicle flows achieved. These results suggest that due to the problem's non-convexity, the SE found is either a local optimum or that understanding the users' best response may offer only limited benefits, as deviating from the fastest route is rarely advantageous. In the latter case, the hierarchical nature of the Stackelberg problem is adequately captured in traditional bi-level formulations.




Fig. 3: Overall objective of the AMoD operator (top), and average travel time of all AMoD users (bottom), as a function of the AMoD penetration rate ϕ , ride-pooling rate ψ , computed as a Stackelberg game with Big Hype (left) and as a Nash with bi-level optimization (right).

TABLE I: Relative difference between the VHT of both groups of users in equilibrium for Stackelberg and Nash at different ϕ and ψ values.

ϕ	ψ	AMoD	Private	Total
0.1	0.1	0.0003	-0.0003	-0.0003
0.1	0.9	-0.0000	-0.0000	-0.0000
0.9	0.1	0.0017	0.0012	0.0016
0.9	0.9	0.0003	0.0001	0.0002

V. CONCLUSION

In this paper, we analyzed the different types of equilibria of an AMoD system in mixed-traffic conditions with selfish private commuters, a Nash and a Stackelberg equilibrium. We analyzed the Stackelberg nature of these problems, where the fleet operator anticipates the best response of the private users, leveraging BIG Hype. We compared the obtained results with the previously extensively studied iterative Nash approximations, where the fleet and private users only react to each other's actions. The results of the case study of Sioux Falls, USA, show that there is a strong agreement between the two kinds of equilibria. We conclude that no substantial benefit arises from the perspective of the AMoD fleet by knowing the best response of the private users.

In the future, we would like to extend this analysis to intermodal settings [22] and explore more transport-justice-oriented objective functions like accessibility fairness [9].

APPENDIX I SENSITIVITY LEARNING STEP

To estimate \mathbf{Jy}^* , we derive the linear combinations of the partial Jacobians of the PG needed to evaluate the fixed-point iteration of (9). Intuitively, differentiating (7) at \mathbf{y}^* allows us to denote $S_{1,i} := \mathbf{J}_1 h(\mathbf{x}, \mathbf{y}^*)$ and $S_{2,i} := \mathbf{J}_2 h(\mathbf{x}, \mathbf{y}^*)$. Then,

the partial Jacobians of $h(\mathbf{x}, \mathbf{y})$ are given by

$$\begin{aligned} \mathbf{J}_1 h(\mathbf{x}, \mathbf{y}) &= \mathbf{J}_1 g(\mathbf{x}, \mathbf{y} - \gamma F(\mathbf{x}, \mathbf{y})) \\ &\quad - \gamma \mathbf{J}_2 g(\mathbf{x}, \mathbf{y} - \gamma F(\mathbf{x}, \mathbf{y})) \mathbf{J}_1 F(\mathbf{x}, \mathbf{y}), \\ \mathbf{J}_2 h(\mathbf{x}, \mathbf{y}) &= \mathbf{J}_2 g(\mathbf{x}, \mathbf{y} - \gamma F(\mathbf{x}, \mathbf{y})) (I - \gamma \mathbf{J}_2 F(\mathbf{x}, \mathbf{y})), \end{aligned} \quad (11)$$

where $g(\mathbf{x}, \mathbf{y}) := \mathbb{P}_{\mathcal{Y}_i(\mathbf{x})}[\mathbf{y}]$. In practice, computing these matrices requires differentiating through the projection operator $\mathbb{P}_{\mathcal{Y}_i(\mathbf{x})}[\cdot]$. However, as explained in [27], from a computational perspective, this operation amounts to solving systems of linear equations of the form

$$\begin{bmatrix} I & A_i^\top & C_i^\top \\ \text{diag}(\omega^*) A_i & \text{diag}(b_i + G_i x - A_i z^*) & 0 \\ C_i & 0 & 0 \end{bmatrix} \begin{bmatrix} dz \\ d\omega \\ d\nu \end{bmatrix} = \begin{bmatrix} dy \\ \text{diag}(\omega^*) G_i dx \\ H_i dx \end{bmatrix}, \quad (12)$$

where the matrices A_i, C_i, G_i, H_i, b_i originate from the polyhedral description of the feasible domain for the lower-level variables

$$\mathcal{Y}_i(\mathbf{x}) := \{\mathbf{y}_i \in \mathbb{R}^{n_{y_i}} \mid A_i \mathbf{y}_i \leq b_i + G_i \mathbf{x}, C_i \mathbf{y}_i = d_i + H_i \mathbf{x}\},$$

which, after vectorization of (3), are equal to

$$\begin{aligned} A_i &:= B, & b_i &:= \text{col}_i(D^p), \\ C_i &:= -I_{n_{y_i}}, & G_i = H_i &:= \mathbf{0}_{n_x}, \end{aligned} \quad (13)$$

where $n_x = |\mathcal{A}|(2|\mathcal{V}| + 1)$ and $n_{y_i} = |\mathcal{A}|$ are the number of decision variables for the leader and individual follower agents respectively.

REFERENCES

- [1] Waymo. (2018) Early ride program. Available at <https://waymo.com/apply/>. Retrieved on March 2, 2018.
- [2] Y. A. Korilis, A. A. Lazar, and A. Orda, "Achieving network optima using stackelberg routing strategies," *IEEE/ACM Transactions on Networking*, vol. 5, no. 1, pp. 161–173, 1997.
- [3] W. Krichene, J. Reilly, S. Amin, and A. Bayen, "Stackelberg routing on parallel networks with horizontal queues," *IEEE Transactions on Automatic Control*, vol. 59, no. 3, pp. 714–727, 2014.
- [4] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, p. 236–259, Mar. 2002.
- [5] M. Pavone, S. L. Smith, E. Frazzoli, and D. Rus, "Robotic load balancing for Mobility-on-Demand systems," *Proc. of the Inst. of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 31, no. 7, pp. 839–854, 2012.
- [6] F. Rossi, R. Iglesias, M. Alizadeh, and M. Pavone, "On the interaction between Autonomous Mobility-on-Demand systems and the power network: Models and coordination algorithms," in *Robotics: Science and Systems*, 2018, Extended version available at <https://arxiv.org/abs/1709.04906>.
- [7] F. Paparella, L. Pedroso, T. Hofman, and M. Salazar, "A time-invariant network flow model for ride-pooling in mobility-on-demand systems," *IEEE Transactions on Control of Network Systems*, 2024.
- [8] M. Tsao, R. Iglesias, and M. Pavone, "Stochastic model predictive control for autonomous mobility on demand," in *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2018, in Press. Extended Version, Available at <https://arxiv.org/pdf/1804.11074>.
- [9] M. Salazar, S. Betancur Giraldo, F. Paparella, and L. Pedroso, "On accessibility fairness in intermodal mobility-on-demand systems," in *Proc. of the IFAC Symposium on Control in Transportation Systems*, 2024.
- [10] S. Dempe, *Foundations of Bilevel Programming*. Kluwer Academic Publishers, 2002, vol. 61.
- [11] J. Nash, "Non-Cooperative Games," *The Annals of Mathematics*, vol. 54, no. 2, p. 286, 1951.
- [12] S. Wollenstein-Betech, M. Salazar, A. Houshmand, M. Pavone, C. G. Cassandras, and I. C. Paschalidis, "Routing and rebalancing intermodal autonomous mobility-on-demand systems in mixed traffic," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 12263–12275, 2021.
- [13] Z. Yin, D. Korzhlyk, C. Kiekintveld, V. Conitzer, and M. Tambe, "Stackelberg vs. Nash in Security Games: Interchangeability, Equivalence, and Uniqueness."
- [14] P. D. Grontas, G. Belgioioso, C. Cenedese, M. Fochesato, J. Lygeros, and F. Dörfler, "Big hype: Best intervention in games via distributed hypergradient descent," *IEEE Transactions on Automatic Control*, pp. 1–16, 2024.
- [15] J. Bolte, T. Le, E. Pauwels, and T. Silveti-Falls, "Nonsmooth Implicit Differentiation for Machine-Learning and Optimization," in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021, pp. 13537–13549.
- [16] P. D. Grontas, C. Cenedese, M. Fochesato, G. Belgioioso, J. Lygeros, and F. Dörfler, "Designing optimal personalized incentive for traffic routing using big hype," in *Proc. IEEE Conf. on Decision and Control*, 2023.
- [17] M. Maljkovic, G. Nilsson, and N. Geroliminis, "On finding the leader's strategy in quadratic aggregative stackelberg pricing games," in *European Control Conference*, 2023, pp. 1–6.
- [18] F. Paparella, L. Pedroso, T. Hofman, and M. Salazar, "Congestion-aware ride-pooling in mixed traffic for autonomous mobility-on-demand systems," in *European Control Conference*, 2024.
- [19] F. Bullo, *Lectures on Network Systems*, 1st ed. Kindle Direct Publishing, 2020, available online at <http://motion.me.ucsb.edu/book-Ins>, with contributions by J. Cortes, F. Dorfler, and S. Martinez.
- [20] M. Patriksson, *The Traffic Assignment Problem: Models and Methods*. Courier Dover Publications, 2015, google-Books-ID: PDhkBgAAQBAJ.
- [21] J. G. WARDROP, "Road Paper. Some Theoretical Aspects of Road Traffic Research." *Proceedings of the Institution of Civil Engineers*, vol. 1, no. 3, pp. 325–362, 1952.
- [22] M. Salazar, N. Lanzetti, F. Rossi, M. Schiffer, and M. Pavone, "Intermodal autonomous mobility-on-demand," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 9, pp. 3946–3960, 2020.
- [23] F. Parise and A. Ozdaglar, "Analysis and Interventions in Large Network Games," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, no. 1, pp. 455–486, 2021.
- [24] "Methods for Monotone Problems," in *Finite-Dimensional Variational Inequalities and Complementarity Problems*, F. Facchinei and J.-S. Pang, Eds. New York, NY: Springer New York, 2003, pp. 1107–1234.
- [25] G. Scutari, D. Palomar, F. Facchinei, and J.-s. Pang, "Convex Optimization, Game Theory, and Variational Inequality Theory," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 35–49, 2010.
- [26] Transportation Networks for Research. (2023) Available online at <https://github.com/bstabler/TransportationNetworks>.
- [27] B. Amos and J. Z. Kolter, "OptNet: Differentiable optimization as a layer in neural networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 136–145.