

Optimization-based incentivization and control scheme for autonomous traffic

Kalabic, Uros; Grover, Piyush; Aeron, Shuchin

TR2020-079 June 24, 2020

Abstract

We consider the problem of incentivization and optimal control of autonomous vehicles for improving traffic congestion. In our scenario, autonomous vehicles must be incentivized in order to participate in traffic improvement. Using the theory and methods of optimal transport, we propose a constrained optimization framework over dynamics governed by partial differential equations, so that we can optimally select a portion of vehicles to be incentivized and controlled. The goal of the optimization is to obtain a uniform distribution of vehicles over the spatial domain. To achieve this, we consider two types of penalties on vehicle density, one is the L2 cost and the other is a multiscale-norm cost, commonly used in fluid-mixing problems. To solve this nonconvex optimization problem, we introduce a novel algorithm, which iterates between solving a convex optimization problem and propagating the flow of uncontrolled vehicles according to the Lighthill-Whitham-Richards model. We perform numerical simulations, which suggest that the optimization of the L2 cost is ineffective while optimization of the multiscale norm is effective. The results also suggest the use of a dedicated lane for this type of control in practice.

IEEE Intelligent Vehicle Symposium

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Optimization-based incentivization and control scheme for autonomous traffic

Uroš Kalabić

Piyush Grover

Shuchin Aeron

Abstract—We consider the problem of incentivization and optimal control of autonomous vehicles for improving traffic congestion. In our scenario, autonomous vehicles must be incentivized in order to participate in traffic improvement. Using the theory and methods of optimal transport, we propose a constrained optimization framework over dynamics governed by partial differential equations, so that we can optimally select a portion of vehicles to be incentivized and controlled.

The goal of the optimization is to obtain a uniform distribution of vehicles over the spatial domain. To achieve this, we consider two types of penalties on vehicle density, one is the L^2 cost and the other is a multiscale-norm cost, commonly used in fluid-mixing problems. To solve this non-convex optimization problem, we introduce a novel algorithm, which iterates between solving a convex optimization problem and propagating the flow of uncontrolled vehicles according to the Lighthill-Whitham-Richards model. We perform numerical simulations, which suggest that the optimization of the L^2 cost is ineffective while optimization of the multiscale norm is effective. The results also suggest the use of a dedicated lane for this type of control in practice.

I. INTRODUCTION

Recent results show that the price of anarchy in vehicle traffic is essentially unbounded, *i.e.*, traffic congestion can be arbitrarily worse as compared to the case when all vehicles on the road are routed by a central controller that aims to minimize congestion or average travel time [1], [2]. The development of autonomous vehicles (AVs) is creating new possibilities for traffic management and control [3] and can be utilized for reducing the cost of anarchy. Keeping in mind that, in the autonomous future, vehicles will retain their primary function of transportation of individuals, we must consider that the desire of AV passengers may sometimes be antagonistic to the goals of traffic improvement. For example, given a choice of whether or not to let their vehicle participate in traffic improvement, passengers would likely consider the cost of their own time or comfort and, for this reason, would likely require some sort of incentive in order to participate.

In this work, we present an incentive and control scheme for the use of AVs to control traffic on a single road segment.

This work was primarily supported by Mitsubishi Electric Research Laboratories. S. Aeron performed some of this work while on an academic sabbatical with MERL in Spring 2019 and Summer 2019. S. Aeron would like to acknowledge the support of NSF CAREER award CCF:1553075.

U. Kalabić is with Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA kalabic@merl.com

P. Grover is with the Department of Mechanical & Materials Engineering, University of Nebraska, Lincoln, NE 68588, USA piyush.grover@unl.edu

S. Aeron is with the Department of Electrical and Computer Engineering, Tufts University, Medford, MA 02155, USA shuchin@ece.tufts.edu

We contrast this to most previous works that have considered incentivization for traffic management by improved routing on multiple road segments, *e.g.*, [4], [5], [6]; we instead focus on dynamic control of vehicles. Specifically, we model traffic as a flow and assume that a certain part of this flow can be incentivized to switch and cooperate in order to control traffic. This turns the problem of controlling traffic into both an incentivization and control problem. AV passengers must first be incentivized to cooperate, and then their vehicles must be controlled. We note that a related mixed traffic model using mean-field ideas was recently introduced in [7], however the focus in that work is on stability and not on incentivization.

The problem of computing appropriate incentives falls under the umbrella of mechanism design [8]. We are interested in developing an incentivization framework in the context of mixed-traffic control, where both cooperative and non-cooperative vehicles are present. As a first step towards this goal, in this work we consider the problem of selecting a subset of vehicles, periodically updated, that can be potentially incentivized, and hence optimally controlled to drive the overall vehicle density to be close to uniform. As part of the solution, the corresponding optimal controls for this subset are also computed. This problem is generally motivated by recent results [9], [10] showing that intelligent control of a small fraction of vehicles can significantly reduce congestion and attenuate traffic waves.

We propose a periodically updated, receding-horizon control scheme. During each period, we determine the distribution of incentive, *i.e.*, which subset of the vehicles to control, as well as the control inputs to these vehicles over the fixed time-period. The receding horizon approach is used to address undesirable edge effects that typically appear at the end of the control time-horizon. The control history is determined according to the assumption that controlled AV flow can be directly controlled and that uncontrolled flow follows the conventional Lighthill-Whitham-Richards (LWR) model. The resulting problem is a variation of the mean-field optimal control problem [11] with congestion, where only a subset of the agents are controlled.

We consider two types of cost on vehicle density. The first is an L^2 cost, which penalizes the distance to the uniform density, and the second is an \dot{H}^{-1} cost, which is a type of multiscale norm and used in fluid mixing to ensure uniformity. To show how one can implement our scheme, we perform a numerical simulation on a single, circular road scenario and use the optimization software `cvx` to solve the optimization problem. Our results show that the penalization

of L^2 cost does not effectively achieve convergence to a uniform distribution while penalization of \dot{H}^{-1} cost does. For this reason, we suggest that approaches from fluid mixing can be helpful in solving traffic optimization. Notably, we find that the scheme requires flow rates that are too high to achieve in a single lane; we therefore suggest that the practical implementation of our approach may require use of a dedicated lane.

The rest of the paper is organized as follows. In Section II, we present a detailed problem formulation and introduce the optimization framework. In Section III, we present a computationally feasible algorithm for solving the proposed optimization problem. In Section IV, we present detailed simulation results and provide a discussion on how one may implement the resulting scheme in practice. Section V is the conclusion.

II. PROBLEM FORMULATION

We consider a two-population traffic scenario, consisting of heteronomous (non-cooperative) vehicles (HVs) and autonomous (cooperative) vehicles (AVs). The densities of HVs and AVs are given by $\rho_1(x, t), \rho_2(x, t) \in \mathbb{R}$, respectively, where $(x, t) \in \Omega \times [0, T]$ is the position x of a vehicle at time t , and Ω and $[0, T]$ are the spatial and time domains, respectively. The vehicular dynamics of HVs are governed by the Lighthill-Whitham-Richards (LWR) model,¹ with the dynamics given by,

$$\partial_t \rho_1 + \nabla \cdot (\rho_1 v_1(\rho_1, \rho_2)) = 0, \quad (1)$$

where,

$$v_1(\rho_1, \rho_2) = u_0 \left(1 - \frac{\rho_1 + \rho_2}{\rho^*} \right), \quad (2)$$

is the HV velocity field and the parameters u_0 and ρ^* are the free-flow speed and the maximum density, respectively. The density evolution of AVs is governed by the continuity equation,

$$\partial_t \rho_2 + \nabla \cdot m_2(\rho_2, v_2) = 0, \quad (3)$$

where $m_2(x, t) \in \mathbb{R}$ is the controlled flux, given by,

$$m_2(\rho_2, v_2) = \rho_2 v_2,$$

where $v_2(x, t) \in \mathbb{R}$ is the controlled, AV velocity field. Note that the HV dynamics are affected by the AV dynamics but not *vice versa*.

Apart from being able to control the AV velocity field, we are able to assign the initial densities of AVs $\rho_2(x, 0)$ over Ω . We do this by assuming that a certain distribution of vehicles can be either an HV or AV, but must be incentivized to turn on AV capability. We assume that the total amount available for incentivization is fixed and an offer can be made only once over a time period $[0, T]$. As a consequence of optimality, there is no reason to not make an offer after the initial time $t = 0$ since, if $v_2(x, t) = v_1(x, t)$ is optimal over

$t \in [0, a]$ but not over $t \in [a, T]$, it will be equivalent to make an offer at any time $t \leq a$. In this work, we assume that each vehicle operator's cost preference $\beta(x)$ is equal, i.e., $\beta(x)$ is constant over x .

We assume the initial density distribution $\rho_0 = \rho_1(\cdot, 0) + \rho_2(\cdot, 0)$ is given and its average is $\bar{\rho}$, i.e., $\int_{\Omega} \rho_0(x) dx = \bar{\rho}$. We also assume that the initial density distribution of potential AVs $\hat{\rho}_0 \leq \rho_0$, with average $\hat{\rho}$, is given and therefore the initial distribution of AVs $\rho_{2,0} = \rho_2(\cdot, 0)$ must satisfy the incentivization condition,

$$\int_{\Omega} \rho_2(x, 0) dx \leq \int_{\Omega} \beta(x) \hat{\rho}_0(x) dx = \beta \hat{\rho}. \quad (4)$$

A. Optimization

Given an initial density distribution $\rho_0(x)$, we wish to determine the initial AV distribution $\rho_{2,0} : \Omega \rightarrow \mathbb{R}$ satisfying (4) and a controlled velocity profile $v_2 : \Omega \times [0, T] \rightarrow \mathbb{R}$ to transport the density from ρ_0 to a density that is close to uniform over Ω .

We consider two approaches in order to achieve this end. Both minimize a cost function of the form,

$$\iint \rho_2(x, t) v_2(x, t)^2 dx dt + cV(\rho_1, \rho_2), \quad (5)$$

so that the first term in the weighted sum is a cost on control and the second term is a cost on the state, and $c > 0$ is a scalar weight.

In the first case, the cost function is chosen with the aim of minimizing the L^2 cost of the density over $\Omega \times [0, T]$,

$$V(\rho_1, \rho_2) = \iint (\rho_1(x, t) + \rho_2(x, t))^2 dx dt.$$

which is equivalent to minimizing the distance to the uniform density at all times. By performing the Benamou-Brenier change of variables [12] $(\rho_2, v_2) \mapsto (\rho_2, m_2)$, the optimization problem can be written as,

$$\min_{\rho_{2,0}, m_2} \iint \frac{m_2^2}{\rho_2} dx dt + c \iint (\rho_1 + \rho_2)^2 dx dt, \quad (6)$$

subject to the constraints,

$$\rho_1(x, t), \rho_2(x, t), m_2(x, t) \geq 0, \quad (7a)$$

$$\rho_1(x, t) + \rho_2(x, t) \leq \rho^*, \quad (7b)$$

$$\int_{\Omega} \rho_{2,0}(x) dx \leq \beta \hat{\rho}, \quad (7c)$$

$$\rho_1(x, 0) + \rho_{2,0}(x) = \rho_0(x), \quad (7d)$$

for all $(x, t) \in \Omega \times [0, T]$, and dynamics (1)-(3).

In the second case, the cost function is chosen with the aim of minimizing the square of the \dot{H}^{-1} seminorm [13] of the density over $\Omega \times [0, T]$. This cost function is used in fluid mixing problems [13] to optimize mixing. Setting a penalty on the deviation from the mean, the cost is given by,

$$V(\rho_1, \rho_2) = \frac{1}{\bar{\rho} T} \iint \left| (-\Delta)^{-\frac{1}{2}} (\rho_1(x, t) + \rho_2(x, t)) \right|^2 dx dt - 1.$$

¹While in this paper we only report results for the LWR model, we note that it is straightforward to extend our algorithm to other models of uncontrolled traffic flow such as the Payne-Whitham and Aw-Rascle-Zhang (ARZ) models.

The cost can be reformulated in terms of Fourier coefficients to obtain the full optimization problem,

$$\min_{\rho_{2,0}, m_2} \iint \frac{m_2^2}{\rho_2} dx dt + \frac{c}{\bar{\rho}} \int \sum_{k=2}^{\infty} \frac{(\hat{\rho}_{1,k} + \hat{\rho}_{2,k})^2}{k^2} dt, \quad (8)$$

subject to the constraints (7), where $\hat{\rho}_{1,k}$ and $\hat{\rho}_{2,k}$ are the k -th Fourier coefficients of $\rho_1(\cdot, t)$ and $\rho_2(\cdot, t)$, respectively. Hence, this cost weighs the spatial low-frequency components of fluctuations in the density higher than the high-frequency components. It has been shown to have good numerical properties in practical optimization problems [14].

The optimization problem can be considered as a variation of a mean-field optimal control problem [11] with congestion, where only a subset of the agents are directly controlled. Since the problem is non-convex due to the LWR dynamics in (2), the conventional approaches to solving convex mean-field control problems do not apply, and we must design a new algorithm.

III. ALGORITHM

In both cases above, the nonlinear dynamics (1)-(3) result in non-convex problems. This is because the flux expression obtained by multiplying (2) by the density ρ_1 cannot be converted to a linear PDE. We can see this when we discretize the dynamics using a conservative scheme, like that of Godunov, where the flux term becomes dependent on logical conditions.

The problem can be somewhat alleviated by using a discretization of the PDEs that is dissipative, such as that of Lax-Friedrichs. In this way, the map from the state $(\rho_1(\cdot, t_k), \rho_2(\cdot, t_k))$ at one time instant t_k to the next time instant t_{k+1} becomes linear. The drawback, however, is that the use of a dissipative discretization does not lead to a physically correct solution [15].

We propose to use a two step iterated approach in solving the optimization problem.

- Step 1: Fix the density ρ_1 of the HVs (in space and time), and use the Lax-Friedrichs discretization to obtain an optimized control v_2 and AV density ρ_2 .
- Step 2: Propagate the HV dynamics using a first-order Godunov scheme to determine a new value for HV density, using the AV density obtained in step 1.

In this way, by keeping ρ_1 static in step 1, the corresponding optimization problem becomes convex and the dissipative components of Lax-Friedrichs can be compensated through vehicle actuation.

At the beginning of the algorithm, we set initial conditions $\rho_1^{(0)}$ and $\rho_2^{(0)}$ that satisfy the constraints (7). Then, during each iteration $i \geq 1$, we perform the two steps.

In the first step, we solve the optimization problem,

$$\min_{\rho_{2,0}, m_2} \iint \frac{m_2(x, t)^2}{\rho_2(x, t)} dx dt + cV \left(\rho_1^{(i-1)}, \rho_2 \right), \quad (9)$$

TABLE I
SIMULATION PARAMETERS

Parameter	Value	Units
$ \Omega $	2	mi
T	8	min
u_0	1	mi/min
ρ^*	10	50 veh/mi
c	0.1	

subject to the constraints,

$$\rho_2(x, t), m_2(x, t) \geq 0, \quad (10a)$$

$$\rho_2(x, t) \leq \bar{\rho} - \rho_1^{(i-1)}(x, t), \quad (10b)$$

$$\int_{\Omega} \rho_{2,0}(x) dx \leq \beta \hat{\rho}, \quad (10c)$$

$$\rho_{2,0}(x) = \rho_0(x) - \rho_1^{(i-1)}(x, 0), \quad (10d)$$

for all $(x, t) \in \Omega \times [0, T]$, and dynamics (3). Upon discretization, the dynamics (3) are linear, of the form,

$$\rho_2(x_k, t_{k+1}) = D_{LF} \begin{bmatrix} \rho_2(\cdot, t_k) \\ m_2(\cdot, t_k) \end{bmatrix}, \quad (11)$$

where D_{LF} is the fixed, spatial discretization matrix obtained using the Lax-Friedrichs method and t_k is the k -th time step.

In the second step, we use a first-order Godunov method to solve the dynamic equation,

$$\partial_t \rho_1 + u_0 \nabla \cdot \rho_1 \left(1 - \frac{\rho_1 + \rho_2^{(i)}}{\rho^*} \right) = 0, \quad (12)$$

with initial condition $\rho_1(x, 0) = \rho_0(x) - \rho_{2,0}^{(i)}(x)$, alternating between steps until either a stopping criterion or the maximum number of iterations have been reached.

1) *Convergence*: The problem (9)-(10) is convex and the solution to the dynamics (12) is unique due to continuity [16]. These are desirable properties, but do not guarantee convergence. The constraints prevent the algorithm from blowing up but do not prevent limit cycles from occurring. For this reason, we terminate the algorithm after a number of prescribed iterations.

A. Receding-Horizon Approach

To ensure that the control scheme operates as seamlessly as possible, and because optimizing over a finite time-horizon can often result in undesirable effects as the time reaches the end of the optimization, we consider the use of a receding-horizon approach.

Specifically, we solve the problem (9)-(12) over the time-interval $[0, T]$ but only implement the control over $[0, T/N]$ where N is a design parameter. We then solve the same problem over the time-interval $[T/N, T + T/N]$ with initial conditions set to the values obtained at time T/N and implement the control over $[T/N, 2T/N]$. We repeat this procedure for times $2T/N, 3T/N, \dots$ until we obtain a solution over the entire time-interval $[0, T]$.




Fig. 1. Initial distribution of vehicles ρ_0 (solid) and average (dotted)




Fig. 2. Total vehicle density at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of L^2 cost with $N = 1$

IV. NUMERICAL SIMULATIONS

We perform numerical simulations to test and exhibit the properties of the algorithm just presented. In all simulations, the simulation parameters are set to the values given in Table I and the optimization problems are solved using the software package `cvx`. We begin with the initial distribution ρ_0 over $\Omega = S^1$ given in Fig. 1 as,

$$\rho_0(x) = 3.5 + 2 \sin(\pi x).$$

For comparison, the figure also shows the evolution of vehicle density without control, *i.e.*, $v_2 = v_1(\rho_1, \rho_2)$. We begin by considering the optimization of the L^2 cost on density and then the \dot{H}^{-1} cost.

A. Optimization of L^2 cost

In the first simulation, we determine the control input according to the solution of the optimization problem (6)-(7), which penalizes density according to the L^2 cost. We set the incentivization parameter $\beta = 0.2$ so that 0.2 of vehicles




Fig. 3. Total vehicle density at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of L^2 cost with $N = 2$




Fig. 4. L^2 cost density component corresponding to uncontrolled (solid, light), $N = 1$ (solid), $N = 2$ (dot-dashed), $N = 4$ (dashed), and $N = 8$ (dotted)

can be AVs and we set $N = 1$ so that we only solve the optimization once before implementing the control.

The time history is shown in Fig. 2 and a qualitative evaluation suggests that we have not achieved a substantial improvement over the baseline. In particular, the density distribution at the final time is very close to the density distribution in the uncontrolled case. Shortening the control horizon by setting $N > 1$ achieves an even worse result; this is shown in Fig. 3, for which we set $N = 2$. Investigating this further, we ran simulations for $N = 1, 2, 4, 8, 16$ for which we plot in Fig. 4 the normalized time history of the density component of the cost (6),

$$\frac{1}{3.5^2 |\Omega|} \int_{\Omega} (\rho_1(x, t) + \rho_2(x, t))^2 dx.$$

The plot shows that we are able to initially, greatly reduce the density component of the cost in all cases but that this does not imply that we evenly distribute traffic in the steady state. In fact, for $N > 1$, we have worsened performance.




Fig. 5. Total vehicle density at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of \dot{H}^{-1} cost with $N = 2$




Fig. 7. AV density at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of \dot{H}^{-1} cost with $N = 2$




Fig. 6. \dot{H}^{-1} cost density component corresponding to uncontrolled (solid, light), $N = 1$ (solid), and $N = 2$ (dot-dashed)




Fig. 8. AV flow at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of \dot{H}^{-1} cost with $N = 2$

For this reason, we consider the minimization of \dot{H}^{-1} .

B. Optimization of \dot{H}^{-1} cost

The reason we believe that the L^2 cost optimization was not able to achieve good steady-state convergence is that being near to the optimal in terms of L^2 does not mean that the distribution is near uniform; due to the equivalence of norms, we expected that another choice of L^p cost would have similar performance and this has been confirmed by simulations that are not reported here. An alternative is to consider the family of multiscale norms, typically used in fluid mixing.

In the first simulation, we determine the control input according to the solution of the optimization problem (8), (7), which penalizes density according to the \dot{H}^{-1} cost. We keep the incentivization parameter at $\beta = 0.2$ and set $N = 1$.

In the results, not shown here, the density reaches close-to-uniform convergence at time $T/2$ but, by the end of the simulation, it drifts away from uniform. We suspect that the

drift away from uniform is due to the fact that we optimize over a finite time-horizon, and so do not sufficiently penalize steady state. This is supported by the results of Figs. 5-6, for which we have set $N = 2$ and obtained sustained convergence to uniform. Fig. 5 shows sustained convergence to the uniform and Fig. 6 shows that the drift away from uniform in the case of $N = 1$ occurs toward the end of the simulation.

C. Discussion

Having obtained a satisfactory numerical results, we now discuss the practicality and implementability of this approach. In Fig. 7, we show the AV density, and in Fig. 8, we show the AV flow. The AV density shows that the optimal approach is to choose 0.2 of the slower vehicles and switch these into autonomous mode, recomputing the distribution again halfway through the time-period.

From the plot of AV flow, we see that the required flow to obtain the desired result is higher than the maximum




Fig. 9. Total vehicle density at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of \dot{H}^{-1} cost with $N = 2$ and an upper limit on AV flow




Fig. 10. AV flow at $t = 0$ (solid, light), $t = T/4$ (solid), $t = T/2$ (dot-dashed), and $t = T$ (dotted) corresponding to optimization of \dot{H}^{-1} cost with $N = 2$ and an upper limit at 2.5 (dotted, light)

achievable in a single lane; to achieve it, it would require AVs to drive at the maximum rate in four separate lanes. This suggests the use of at least one dedicated lane for AVs in order to implement the scheme, along with a limit on the flow m_2 . We perform an additional simulation, setting the limit at the maximum rate according to the LWR model (2),

$$\frac{u_0 \rho^*}{4} = 2.5, \quad (13)$$

which corresponds to the use of at most one additional dedicated AV lane. The results are plotted in Figs. 9-10. Fig. 10 shows a plot of the AV flow at different times, exhibiting adherence to the prescribed limit. In Fig. 9, the plot of total vehicle flow is not much qualitatively different than that of the unconstrained case, plotted in Fig. 5, validating our results. Note that we did not consider lane-change interactions between vehicles and assumed a bulk fluid model.

Before concluding, we note that the general scheme pre-

sented in this work is broadly applicable to most traffic scenarios and not limited to the single, circular road example on which we performed the numerical simulation. Nevertheless, this requires more future work as different scenarios will exhibit different topologies Ω and likely require a more specialized approach towards optimization.

V. CONCLUSION

In this work, we presented an optimization-based control scheme for stabilizing traffic to a uniform state by incentivizing and controlling a group of autonomous vehicles. We considered two approaches, the first optimizing an L^2 cost and the second optimizing an \dot{H}^{-1} cost.

We performed numerical simulations to determine the efficacy of the approach. The results showed that, by optimizing the \dot{H}^{-1} cost, we are able to stabilize traffic to the uniform in finite time while, by optimizing the L^2 cost, we are not. We discussed the practicality of this approach and recommend the addition of a dedicated lane for implementability.

ACKNOWLEDGMENT

The authors acknowledge Dr. Saleh Nabi of Mitsubishi Electric Research Laboratories for technical discussions.

REFERENCES

- [1] J. Zhang, S. Pourazarm, C. G. Cassandras, and I. Ch. Paschalidis, "The price of anarchy in transportation networks by estimating user cost functions from actual traffic data," in *Proc. IEEE Conf. Decision and Control*, Las Vegas, 2016, pp. 789–794.
- [2] —, "The price of anarchy in transportation networks: Data-driven evaluation and reduction strategies," *Proc. IEEE*, vol. 106, no. 4, pp. 538–553, 2018.
- [3] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and decision-making for autonomous vehicles," *Annu. Rev. Control Robot. Auton. Syst.*, pp. 187–210, 2018.
- [4] N. Mehr and R. Horowitz, "How will the presence of autonomous vehicles affect the equilibrium state of traffic networks?" *IEEE T. Contr. Netw. Syst.*, to be published.
- [5] —, "Pricing traffic networks with mixed vehicle autonomy," in *Proc. American Control Conf.*, Philadelphia, 2019, pp. 2676–2682.
- [6] T. Tanaka, E. Nekouei, A. R. Pedram, and K. H. Johansson, "Linearly solvable mean-field traffic routing games," in *Proc. Allerton Conf.*, Monticello, IL, 2018, pp. 283–289.
- [7] K. Huang, X. Di, Q. Du, and X. Chen, "Stabilizing traffic via autonomous vehicles: A continuum mean field game approach," in *Proc. IEEE Intell. Transp. Syst. Conf.*, Auckland, New Zealand, 2019, pp. 3269–3274.
- [8] T. Börgers, *An Introduction to the Theory of Mechanism Design*. New York: Oxford University Press, 2015.
- [9] C. Wu, A. M. Bayen, and A. Mehta, "Stabilizing traffic with autonomous vehicles," in *Proc. IEEE Int. Conf. Robotics and Automation*, Brisbane, Australia, 2018, pp. 6012–6018.
- [10] C. Wu, A. Kreidieh, E. Vinitzky, and A. M. Bayen, "Emergent behaviors in mixed-autonomy traffic," in *Proc. Conf. Robot Learning*, Mountain View, CA, 2017, pp. 398–407.
- [11] A. Bensoussan, F. Jens, and P. Yam, *Mean Field Games and Mean Field Type Control Theory*. New York: Springer, 2013.
- [12] J.-D. Benamou and Y. Brenier, "A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem," *Numer. Math.*, vol. 84, no. 3, pp. 375–393, 2000.
- [13] J.-L. Thiffeault, "Using multiscale norms to quantify mixing and transport," *Nonlinearity*, vol. 25, no. 2, pp. 538–553, 2012.
- [14] D. Fournes, C. Caulfield, and P. J. Schmid, "Optimal mixing in two-dimensional plane poiseuille flow at finite Péclet number," *J. Fluid Mech.*, vol. 748, pp. 241–277, 2014.
- [15] R. J. LeVeque, *Finite Volume Methods for Hyperbolic Problems*. New York: Cambridge University Press, 2002.
- [16] M. Garavello, K. Han, and B. Piccoli, *Models for Vehicular Traffic on Networks*. Springfield, MO: AIMS, 2016.