Traffic21 Policy Brief **Mixed-Autonomy Era of Transportation**

Resilience & Autonomous Fleet Management

By I-Cheng Lin, Osman Yağan and Carlee Joe-Wong





Carnegie Mellon University

Summer 2022

Acknowledgments

This work was supported by the National University Transportation Center for Improving the Mobility of People and Goods (Mobility21) at Carnegie Mellon University. We would like to express our deep and sincere gratitude to Sean Qian and the community stakeholders, especially Chuck Imbrogno from the Southwestern Pennsylvania Commission, for their feedback on our work. We would also like to thank Shaival Chetan Parikh and Chen Zhao for helping to build the simulator for the dynamic fleet management.

Introduction

T

Future Environment: Mixed-Autonomy





Recent advances in technology, including systems like the Internet of Things (IoT) connecting a significant number of devices, increasingly autonomous vehicles, and widespread deployments of machine learning, promise to form the core of newly intelligent future transportation systems. The advent of autonomous vehicles (AVs) in particular, which have recently started to be deployed (with safety drivers) to the public in California [1], may dramatically change the dynamics of the vehicles in the transportation network and the interactions between them. Prior analyses of AVs' effects on traffic flow, however, often do not account for realistic mixed autonomy settings, in which autonomous vehicles interact with humandriven vehicles (HVs). Studies that do focus on the interactions between AVs and HVs, however, mostly focus on microscopic behaviors, so the effect on the whole transportation network is not clear. Our study demonstrates the potential for autonomous vehicles to improve traffic flow for both autonomous and nonautonomous vehicles in mixed-autonomy settings within a transportation network system.

[1] https://www.reuters.com/technology/california-issues-permitscruise-waymo-autonomous-vehicle-service-2022-02-28/

Autonomous Vehicles: Benefits

The adaptation of autonomous vehicles (AVs) can benefit the operation of the whole transportation network in terms of better utilization, lower travel time, and energy consumption. Many of these potential benefits stem from three core AV capabilities. First, the AVs can be programmed to make collaborative decisions that will improve the overall driving experience for all users, and look for ways to reduce traffic congestion even if it means the AVs might experience a longer travel time or a less efficient trip. Second, AVs are equipped with multiple sensors and use computers to automatically detect and respond to changes in the environment. Thus, the adaptation of AVs could enhance the safety of the transportation network and accommodate a denser traffic flow by allowing autonomous vehicles to follow each other more closely on roads without compromising safety. Third, since the computation of AVs takes less time than humans to respond, they generally have a lower reaction time than the HVs. Moreover, since the AVs can share the information between them, they can respond to incidents a few blocks away that are not visible to human sight, further reducing the reaction time, which we emphasize next.



Autonomous Vehicles: Collaboration and Information Exchange



The third benefit arises from the fact that the AVs can exchange information with other AVs in real-time via vehicle-to-vehicle (V2V) or vehicle-to-infrastructure (V2I) communication and can also obtain information about the transportation network in general by communicating with other monitoring services. Thus, AVs will have a more complete picture of road network conditions (e.g., congestion along their routes) and can react and adjust their chosen routes, travel speeds, etc. accordingly. This capability is similar to that of Google Maps and similar services, but it is available all the time and can respond automatically without needing to open a navigation app and have a human choose between the offered routes. Since AVs can communicate with each other, they could negotiate and make collaborative decisions, preventing them from competing for the same road resource. In fact, HVs could also offer some of these benefits via navigation apps, which could guide different vehicles to use different roads and inform them about real-time traffic conditions. However, the reaction times for HVs will still be longer than they are for AVs, and in some cases, HVs may not follow the recommendations given by navigation apps. Thus, we expect that coming deployments of AVs, or at least semiautonomous vehicles that exhibit these three capabilities, will most likely be required to realize these benefits. Indeed, AVs are likely to be initially deployed in fleets (e.g., of taxis or delivery vehicles), increasing their ability to cooperate with one another.



Interaction Between AVs and HVs: How Decisions of AVs Influence HVs

In the big picture, HVs (without navigation apps) generally try to behave selfishly, to minimize their own travel times, maximize their fuel efficiency, etc., while the AVs could make collaborative decisions that seek to benefit all vehicles. Interactions between AVs and HVs occur on several different timescales, due to the different timescales of decisions taken by AVs and HVs. At a macroscopic scale, AVs and HVs must choose routes to take to their destinations. For example, AVs could try to utilize the roads that are relatively empty to induce lower traffic levels on the often-used road. At the microscopic scale, AVs and HVs choose their velocities along these routes.

The velocity of the AVs could then influence the velocity of HVs. Imagine there are many AVs on the road; if all of them operate at a similar speed, the HVs will likely follow the same speed. The more AVs on the road, the more the HVs will follow the decisions of the AV fleet. For example, if AVs have extra



information due to their network connectivity, they could react to the incidents happening a few blocks away, preventing sudden braking, which may also help the HVs to change their velocity smoothly.

Our Goals

There are three main goals of this research. We first consider the optimal strategy of the autonomous vehicle fleets and the equilibrium, or average steady-state, traffic conditions. Since the traffic flow in practice may vary over time, we then consider the resilience and robustness of the whole transportation system.

- When the traffic flow is stable, vehicles should eventually reach an equilibrium state (e.g., level of congestion) on a road network. This equilibrium, however, may be suboptimal if the vehicles do not have full information of the city or only aim to optimize their own delays. We model the coexistence of HVs and AVs in routes chosen, finding that the system reaches a **close-to-optimal equilibrium** when only 20-50% of vehicles are programmed to make collaborative decisions that aim to optimize the overall delay of the whole transportation network.
 - The adaptation of AVs and their co-existence with HVs form a multi-agent mixed-autonomy environment. This work aims to find the **optimal operating policies** for the fleet of AVs, considering HVs behave selfishly based on their own information.

We find out how AVs could prevent cascading network congestion, which may result in large-scale congestion and failure initiated by a disaster or a particular event – reaching the **optimal robustness** of the whole transportation system.

Equilibrium Analysis

IL

F

UNI

:/

I

÷.

B

. .

ı

CP

EB

2

;

1

1011

((((()))))))

< <

Equilibrium with Selfish Behaviors

The selfish behavior of the vehicle means that it only aims to minimize its own travel delay regardless of the available resources and the capacity of the roads. When all vehicles behave selfishly, they often reach equilibrium traffic conditions that are not collectively optimal for all vehicles on the road. Take the famous Braess's Paradox as an example. The figure on the right illustrates the original network, which reaches a steady state travel delay on each road. If we add a road with low delay connecting the middle two nodes A and B, the traffic will likely go through the path O->A->B->D utilizing this new road, eventually causing traffic jams on this path. This equilibrium is then suboptimal: many vehicles on this path will experience long delays throughout their journey, while the original equilibrium with traffic using paths O->A->D or O->B->D induces a low delay on roads OA and BD with a high delay on roads AD and OB. Thus, the equilibrium state reached by selfish behaviors with this new road is not the optimal one since on average, vehicles experience high delays. Sometimes, this equilibrium is far from optimum.





Illustration of Braess's Paradox. O is the origin of traffic and D is the destination of the traffic. Roads OA and BD are shorter and have lower capacity. On the other hand, roads OB and AD are longer while having higher capacity.

Benefits from Collaborative Decisions

ILLUSTRATION OF BENEFITS FROM ALTRUISM







Sometimes a collaborative decision could play an essential role in optimizing the traffic and avoiding these suboptimal equilibria. Let us take this toy example. If all the vehicles are selfish, they will try to use the shorter route with lower travel time for them. Once the delay of this route is the same as that of the longer ones, the selfish agent will start to use the longer one. However, this means that during the equilibrium state, the delay of the shorter route will be the same as that of the longer route. If some vehicles have information on the road usage and could make collaborative decisions that are not selfish, they could utilize the longer route with higher capacity while leaving the shorter route to the selfish agent. We could compensate the users using the longer route by waiving road tolls or other means to guarantee fairness. By doing so, the equilibrium reached could have a much lower overall travel time than the equilibrium state induced by fully selfish behavior.

How Do Ratios of AVs Influence the Travel Time?

If too few vehicles are not selfish, their decisions will not effectively influence the equilibrium state of the traffic. After reaching a certain percentage of vehicles that do collaborate in aiming to optimize the overall delay of the whole transportation system, the equilibrium state starts to change, and the overall performance increases. When the percentage of collaborative AVs exceeds a certain level, the optimal equilibrium is reached, and increasing the ratio of AVs will no longer decrease total travel time. However, the AVs could maintain a shorter following distance from other vehicles. If we consider this, the overall travel time will keep decreasing slightly with the increasing ratios of the AVs.

EXAMPLES OF PARALLEL ROUTES



AVERAGE TRAVEL TIME DIAGRAM



TRAVEL TIME OF THE SHORTEST PATH





Mechanism to Incentivize a Better Equilibrium

We can find the optimal equilibrium state by solving a two-stage convex optimization problem. In the first stage, if we view each route as a bucket of water that fills up as more vehicles take that route, then we can treat selfish HVs as acting like water that flows into all buckets uniformly, eventually equalizing the delays on each route. This problem is analogous to a power allocation problem for the Gaussian noise channel in communication engineering. By applying specific methods, we could solve the optimal equilibrium state of parallel routes and figure out the portion of AVs that can reach this optimum by adopting water filling algorithms for the inner optimization problem for the HVs and then solving the AVs' optimization problem.



POLICY SUGGESTED

Identify the optimal ratios of collaborative vehicles and ensure that this proportion of vehicles follows the collaborative decisions. In our simulations, usually, collaborative vehicles start to be effective after the ratio is between 20%~50%. We could reach such ratios with a combination of AVs and HVs following the collaborative decisions by compensating the expense of the vehicles or making a certain degree of obedience to the collaborative decisions a prerequisite for AV registration with local or state transportation agencies. This could be realized in practice by exemptions from tolls, or a dynamic tolling system with discounts for the collaborative vehicles. Agencies like PennDOT, for example, may impose vehicle-miles-traveled tolls on state highways that could be waived for collaborative vehicles. Other example incentives more suited to urban areas might include discounts at parking garages.

ILLUSTRATION OF WATER-FILLING ALGORITHMS



Dynamic Fleet Management

19 I

Balance Between Centralized and Distributed Decision Making

The AVs and their co-existence with HVs form a multiagent, mixed-autonomy environment. To monitor the dynamics of such a transportation system, we need to consider the different effects of decisions and actions of each agent made in different times. This is different from the equilibrium analysis which does not involve the concept of "time." This work aims to find the optimal operating policies for the fleet of AVs, considering that HVs behave selfishly based on their own information. Optimizing a single AV will fail to reach the optimal operation of the whole system. On the other hand, centralized control of thousands of AVs in the city will face computation issues and severe communication delays. We strike a balance between centralized and decentralized policies using state-ofthe-art reinforcement learning: Centralized training and decentralized actions using counterfactual multi-agent policy gradient methods (COMA) [1].



[1] Counterfactual Multi-Agent Policy Gradient https://dl.acm.org/doi/abs/10.5555/3504035.3504398

Logic Behind the Model: Small Group Management

Our work uses a mesoscopic view, considering both network-level routing decisions and road level velocity decisions. We group the AVs in the same area as a small fleet, using multi-agent reinforcement learning (MARL) and the cell-transmission model (CTM) [1] of the traffic to find the optimal policies for the AVs that maximize the social welfare — the average travel time of all vehicles. In each cell or region, vehicles can communicate locally with each other via V2V communication to the cellular base station in that area. Our trained model can be stored at the base station, gathering all the information in the block and determining the best actions.



[1] The Cell Transmission Model: Network Traffic https:// escholarship.org/content/ qt9pz309w7/qt9pz309w7.pdf

Velocity Control: Reason and Interactions Between AVs and HVs

When a particular event ends (e.g., concert, football game, etc.), many people leave a specific area; or when an accident happens, many vehicles may detour to adjacent areas. These kinds of incidents may trigger sudden traffic in the transportation network. Without velocity control, all this traffic might feed into the transportation network to the same area at the same time. Thus, a traffic jam will occur. Once the traffic jam starts, the travel speed of the whole network will start to decrease drastically as the traffic spreads; thus, the average travel time of the whole system might increase! If we control the vehicles' velocity in the network, we can try to prevent the vast traffic flows from meeting in a specific area. The control of the speed of the vehicles is a key factor in the scheduling of traffic control.



Routing Control: Choices of AVs and HVs

Routing is always a critical problem that also influences the traffic a lot. Every single vehicle's selfish goal is to find the routes that minimize their costs (i.e., travel time, energy consumption, toll, etc.). However, if all vehicles attempt to choose the route that minimizes their costs, all of them might try to use the shortest path and make specific paths congested or even blocked. With a collaborative routing control of AVs that considers the selfish behavior of individual HVs, we could try to let the fleet of AVs optimize the whole traffic network by incentivizing the optimal behavior of the HVs.





Our methods show that under different ratios of AVs, the model could lower the average travel time of the whole transportation system. With the control of the velocity and routing of AVs, our model could learn the traffic pattern and react to the sudden traffic that may potentially fail the network and incentivize a better behavior of HVs. In short, reasonable control of the AVs could drastically change the performance of the whole transportation system even when the portion of AVs is much smaller than HVs.

0 -200000 -400000 -400000 -800000 -1000000 -12000000 0 50 100 150 200 k epochs

Learning curve of the MARL framework

LEARNING CURVE

POLICY SUGGESTED

An AV fleet operator can apply our proposed method to train the model to manage AV fleets locally. The operator should (i) execute such a model in each area to determine the actions of micro fleets and (ii) gather the observations and actions of each micro fleet in each time interval and return them to the centralized controller to judge the operations. The update frequency can be determined by monitoring the traffic situation. If some new traffic pattern occurs, the performance of the whole network may decrease drastically; we can update the model more frequently to deal with such a situation. On the other hand, if the current trained models work well enough, we could lower the update frequency to ease the communication between local actors and the centralized critic. Such models can be adopted by the Incident response teams who need to direct traffic away from accidents.

The Resilience of the Transportation Network



Accidents, Natural Disasters in the Transportation Network



We finally consider the resilience of transportation systems to sudden disruptions. If a severe accident occurs on a particular road, the traffic should be detoured. The extra traffic on other roads might cause a traffic jam in the adjacent area where the accident initially happened. The recent bridge collapse in Frick Park, Pittsburgh, is an example of such a sudden disruption. Moreover, if a disaster like an earthquake or flood occurs, more than one area will be impacted. With more than one spot in the transportation network failed, we could expect the whole transportation network to suffer from significant additional traffic. These additional traffics might potentially trigger a process called cascading failure which eventually might cause the whole system to break down as congestion spreads to the entire network.

Traffic Redistribution: How AVs Could Help

We conjecture that AVs can help to prevent cascading failures by taking advantage of the benefits we mentioned earlier in this report. By pulling up-todate traffic information from the internet, AVs could respond to failures quicker than HVs, preventing the vehicle from realizing the situation at the last moment. With this information, the AVs could plan their routing based on the information of the failure. In other words, the AVs could avoid the low traffic capacity area, which the extra traffic could potentially jam. Finally, selfish agents might want to minimize travel time and be reluctant to use longer routes. The altruism of the AVs might disperse the traffic to the roads and the areas that are used less often, thus ensuring that the traffic burden is more equitably spread throughout the city.



The Phenomenon of Cascading Failures

Much research [1-3] on the robustness of the network systems focuses on the phenomenon called cascading failures, where failures in a small part of the network (e.g., congestion on a few critical roads) may trigger a process of sequential failures due to traffic in the failed part of the network being transferred elsewhere, eventually causing a disastrous impact on the whole system. Such a cascading failure process can also make other systems like railway networks, communication networks, etc., break down. The blackout in Italy on September 28, 2003, is an example where the initial failure of a few power lines eventually caused a nationwide power outage.



ILLUSTRATION OF CASCADING FAILURE



[1] Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization https://aip.scitation.org/doi/10.1063/1.2737822

[2] Model for cascading failures in complex networks https://arxiv.org/abs/cond-mat/0309141

[3] Cascading failures in interdependent systems under a flow redistribution model https://journals.aps.org/pre/abstract/10.1103/PhysRevE.97.022307

A Better Redistribution Decision by AVs

Simulations were done to verify whether AVs could help enhance the transportation network's resilience. We evaluate the AVs' benefits with different AV-to-HV ratios. The AVs could collaborate to determine the areas to detour with the complete information of the system. The results in the figure below show that the collaboration of the AVs could prevent the system from failing. With a higher portion of AVs, eventually, there will be fewer congested areas on the map.

INITIAL FAILURE



50% AVs





NO AVs



100% AVs

Extension to Multiple Networks: Transportation in Metropolitan Areas

Finally, we consider the presence of alternate transportation modalities, which can similarly help to relieve congestion on road networks. The urban transportation system typically includes road networks, bus, subway, railway, and bikesharing systems. These networks are interdependent in the sense that passengers may transfer between them. The robustness of such interdependent systems was studied, with results showing that interdependent networks could be more vulnerable to cascading failures than an isolated network. We study the coupling between the networks, which indicates the portion of the passengers to use another network system. We show that dynamically adjusting the coupling [1] (i.e., passenger transfers) between networks can optimize the usage of the whole system and let each network share the load of the passengers, maximizing the utilization of the whole system and preventing it from overloading and failing!

[1] Dynamic Coupling Strategy for Interdependent Network Systems Against Cascading Failures https://arxiv.org/abs/2203.01295



ILLUSTRATION OF COUPLING BETWEEN INTERDEPENDENT NETWORKS



A Dynamic Strategy to Prevent Failure of Transportation Networks

We introduce the dynamic coupling strategy with a step-wise optimization (SWO) framework between networks, which could determine the optimal coupling coefficients at each stage of the cascading failure process. We show that (under certain conditions) the SWO strategy reaches the global optimum regarding the final surviving number of areas. Thus, our step-wise dynamic optimization can lead to the optimal operation of the whole system. The figure below shows the results of our SWO strategy compared to other fixed coupling strategies. Under the same initial failure size (the portion of the network that fails), the dynamic strategy could yield the largest surviving portion of the whole system.



SURVIVING CURVES FOR UNIFORM DISTRIBUTED FREE SPACE

Performance of the SWO strategy compared to other fixed coupling strategies. The x-axis indicates the initial failure portion, and the y-axis indicates the final surviving portion. Under the same initial attack, better strategies lead to higher surviving portions.

POLICY SUGGESTED

For the road transportation system, AVs could identify the roads that are close to the capacity limits; or foresee the roads that may potentially receive lots of extra traffic due to the failures in adjacent areas. Based on such information, planning agencies could compute traffic flow redistribution policies for AVs to avoid traffic jams and cascading failures.

Applying the proposed framework to multiple interdependent transportation systems can optimize

the dynamic flow adjustment between different networks. Recommendations with transportation apps, e.g., the Port Authority mobile app, could achieve such traffic flow transferring between different network systems. Furthermore, ensuring the availability of public transportations and providing discounted parking areas, discounted tickets, etc. in these locations could also be a good policy to encourage passenger transfers.

Summary

AVs in the mixed-autonomy era could act as coordinators in the transportation system. With collaboration and information sharing between the AVs and the centralized server, we find that AV fleets can take the actions that incentivize the best response of the HVs. Even a partial adaptation of AVs not only increases the utilization of the transportation system, minimizing vehicles' delays, but also increases the resilience of the transportation network against failures and sudden traffic. Thus, drivers' on-road experience could be elevated to the next level in the new era of mixed-autonomy.

Summary of Policies Suggested

- 1. Identify the optimal ratios (usually between 20%~50%) of collaborative vehicles and ensure that this proportion of vehicles follows the collaborative decisions.
- 2. Execute the dynamic micromanagement model in each area to determine the actions of micro AV fleets. Gather the observations and actions of each micro fleet every specific time interval and return to the centralized controller to judge the operations and update the model.
- 3. Compute traffic flow redistribution policies for AVs to avoid traffic jams and cascading failures on road networks. In multiple interdependent transportation systems, agencies may apply dynamic coupling strategies to optimize the robustness of the system against cascading failures of the transportation system.

Future Works

- 1. The equilibrium analysis of the transportation network could be extended to different road topologies. A more sophisticated operation policy can be obtained for different network structures.
- 2. Various models for the MARL settings of AV fleet management could be applied in addition to the COMA used in this report. Based on the centralized/distributed level of the environment, we can find the model that meets the communication resource constraints.
- 3. We can search for an optimal strategy for setting the coupling coefficients under different network topologies. For example, many real networks follow the power law and could be generated by preferential attachment.



Heinz College Carnegie Mellon University 5000 Forbes Ave. Pittsburgh, PA 15213 412-268-9505 traffic21.heinz.cmu.edu

Carnegie Mellon University

Carnegie Mellon University 5000 Forbes Ave. Pittsburgh, PA 15213

412-268-2000 cmu.edu