



UNIVERSITY OF LEEDS

This is a repository copy of *Decentralized Optimization of Vehicle Route Planning - A Cross-City Comparative Study*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/172731/>

Version: Accepted Version

Article:

Davis, B, Jennings, G, Pothast, T et al. (3 more authors) (2021) Decentralized Optimization of Vehicle Route Planning - A Cross-City Comparative Study. IEEE Internet Computing. ISSN 1089-7801

<https://doi.org/10.1109/mic.2021.3058928>

This item is protected by copyright. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Department: Head
Editor: Name, xxxx@email

Decentralized Optimization of Vehicle Route Planning – A Cross-City Comparative Study

Brionna Davis

Department of Civil and Environmental Engineering, Vanderbilt University, Nashville, USA

Grace Jennings

Department of Civil and Environmental Engineering, Vanderbilt University, Nashville, USA

Taylor Pothast

Department of Human and Organizational Development, Vanderbilt University, Nashville, USA

Ilias Gerostathopoulos

Department of Computer Science, Vrije Universiteit Amsterdam, Amsterdam, Netherlands

Evangelos Pournaras

School of Computing, University of Leeds, UK

Raphael E. Stern

Department of Civil, Environmental, and Geo- Engineering, University of Minnesota, Minneapolis, USA

Abstract—The introduction of connected and autonomous vehicles enables new possibilities in vehicle routing: Knowing the origin and destination of each vehicle in the network can allow for coordinated real-time routing of the vehicles to optimize network performance. However, this relies on individual vehicles being “altruistic” i.e., willing to accept alternative less-preferred routes. We conduct a study to compare different levels of agent altruism in decentralized vehicles coordination and the effect on the network-level traffic performance. This work introduces novel load-balancing scenarios of traffic flow in real-world cities for varied levels of agent altruism. We show evidence that the new decentralized optimization router is more effective with networks of high load.

■ **NEW MOBILITY** concepts are at the forefront of research and innovation in smart cities. They are enabled by advances in intelligent infrastructures [1]. The shift toward an *autonomous vehicle*

(AV) fleet means that we will soon have the possibility to control the routes that individual vehicles take. Even before AVs are prevalent on our roadways, vehicle connectivity via smart-

phone apps (e.g., Waze, Google Maps, Apple Maps, Nokia HERE, etc.) make it possible to suggest individualized routes for each driver in the network, and optimize these routes based on some desired network state: the *system optimal* (SO) route assignment [2]. However, compliance to these routes cannot be enforced. Furthermore, coordination solutions between AVs at the system level are limited but required to guarantee that traffic systems remain in a viable equilibrium.

If routes are assigned to drivers to achieve a network state under some SO criteria, these may not be routes that are best for each individual driver. Instead, depending on how selfish they are, an individual driver may select the *user equilibrium* (UE) route, i.e. the route that is optimal for each individual driver based on a greedy assessment of the route options [2]. While selfish drivers may select the route that is best for them, more altruistic drivers may be willing to accept the SO route assignment, while some drivers who have both selfish and altruistic traits may select a route that makes a trade-off.

Dynamic traffic assignment (DTA) dates back to the late 1970s [3] and studies how vehicles are routed in networks. An excellent summary of early DTA efforts is provided by Peeta and Ziliakopoulos [4]. Much of the existing early literature on DTA focuses on mathematical programming-based solutions to routing vehicles in fixed time steps [3], while other approaches include the formulation as an optimal control problem [5]. Some consider DTA with AVs (e.g., [6]). With respect to prior work, this study focuses on the use of a decentralized approach to compare alternative routing plans and cooperatively select one that optimizes the network-level traffic flow with respect to roadway utilization.

Our main thesis is that the efficiency of the entire network can be improved by pursuing a global objective such as reducing CO₂ emissions or balancing the traffic flow assigned to each link in the network. In particular, we assume based on earlier evidence [7] that even if *some* of the agents pursue a mix of UE and SO routing, this can benefit all agents in the system, including the selfish ones, and still improve the efficiency of the network. To support this thesis, we explore the trade-off between optimizing complex non-

linear global and local objectives¹ that the agents consider when selecting a route. In this context, a global objective is a system-level objective such as balancing the traffic flow in the network, while a local objective is one that is specific to an individual vehicle, such as travel time.

The specific mechanisms to incentivize travel behavior or nudge drivers to become “altruistic” and take SO routes are the focus of other works (e.g. [8]), and are beyond the scope of this paper. Instead, assuming an incentivization mechanism exists, we try to answer the question: *What degree of altruism is required by the agents to observe system-level benefits and to what extent is the required altruism dependent on the traffic level in the network?*

In the context of employing multi-agent learning to optimize route planning, we make the following contributions:

- We present a study that compares different altruism levels of agents (autonomous vehicles) and their effect on the overall traffic performance optimization.
- The understanding of how different traffic levels in the network influence the effectiveness of alternative optimized routes.
- The first application of an open-source software framework² [9] to different large-scale urban transport networks.

Technical Background

In order to investigate the behavioral influence of selfish vs altruistic agents on local and global objectives under dynamic traffic assignment, we performed an agent-based simulation study with real-world urban traffic networks. Note that we do not study the effect of a particular agents’ behavior observed in reality, but rather we profile a spectrum of agents’ behavior to understand how it influences the distributed optimization performance. In this section, we outline a novel applicability scenario of the traffic simulator–SUMO, integrated with the agent-based planning framework–EPOS, for the purpose of our study.

¹Such objectives make agents’ route choices dependent on each other and as a result agents need to coordinate these choices.

²TRAPP, available at <https://github.com/iliasger/TRAPP>

Traffic Simulation with SUMO

SUMO (Simulation of Urban MObility) is a well-known open-source microscopic traffic simulator³. It can be used for simulating up to hundreds of thousands of vehicles in complex, realistic city networks. High realism is achieved by simulating acceleration and deceleration of cars in traffic lights and intersections, vehicle manoeuvres, and different driving styles, etc. In our experiments, we relied on a module of SUMO that can be used for controlling a SUMO simulation via Python.

In order to experiment with dynamic route assignment, we implemented three different routers in Python. All routers rely on the internal representation of a city network as a graph and perform a Dijkstra algorithm to find the shortest path between an initial position A and a destination B. The difference of the routers is what they consider as cost of an edge (street). In the first router such a cost is the length of the street, resulting in routes of minimal overall length. In the second router, the cost is the inverse of the maximum speed allowed on a street, resulting in routes of maximum overall speed. In the third router, the cost is formed by the length of a street divided by the maximum speed allowed on the street, resulting in routes of minimal length *and* maximum overall speed.

In our experiments, each router produced a single route for each trip, hence each car could select among three different routes to navigate from A to B. Which route to choose was a decision that involved agent-based planning via EPOS, described next.

Traffic Optimization with EPOS

EPOS (Economic Planning and Optimized Selections) is a decentralized multi-agent optimization framework written in Java [10]. EPOS can be used for efficiently solving complex multi-objective combinatorial problems via participatory collective learning. EPOS is designed for problems in which a number of agents needs to coordinate their decisions in order to effectively use a shared medium such a power grid or a set of streets. Each agent’s decision may influence the decision of other agents, i.e. agents make choices

based on non-linear cost functions. The problem that EPOS solves is to allow each agent to take decisions that considers both local and global objectives with the minimum number of interactions with the other agents. This is achieved by having agents in EPOS self-organize in tree topologies over which they can perform efficient aggregation and decision-making in an iterative fashion: consecutive child-parent interactions in the bottom-up phase, followed by parent-child interactions in the top-down phase. In the following, we will describe EPOS only to extent necessary for this study, we refer the interested reader to [10] for more details.

In this study, an EPOS agent is an AV, i.e. a vehicle with a decision-support systems in routing selection. Decision-making in EPOS involves selecting a *plan* from a finite set of plans for each agent. In our setting, a plan corresponds to a *route* from a position A to a destination B. As explained in the previous subsection, we equipped each AV with the ability to select among three possible routes, each corresponding to one of the three available routers (“minimum length”, “maximum speed”, and “combined length and speed” router). Nevertheless, our setting can be easily extended to accommodate more routes and routers, even routers that only serve specific cars (essentially creating agents that have more options). EPOS is then used by the AVs so that each car selects one route to follow out of the three options they have (Figure 1).

A plan is represented as a vector of real values in EPOS, each representing the “contribution” of the agent to the shared medium. In our study, the shared medium is the set of all streets in a city. A plan hence becomes a vector of real values containing the expected utilization of each street by the car for a specific planning horizon (e.g. 30 minutes). For instance, assuming a city consisting of only four streets, A, B, C, D, a route that only uses A is encoded as $X, 0, 0, 0$, where X depends on the expected occupancy of A (calculated based on the street length, vehicle length and expected time spent on the street). Plan 2 of agent n in Figure 1 is a concrete such example with $X = 0.3$.

Each plan comes with a cost denoting the agent dislike for this plan. Agents can express preferences for their plans by lowering the cost

³<https://www.eclipse.org/sumo/>

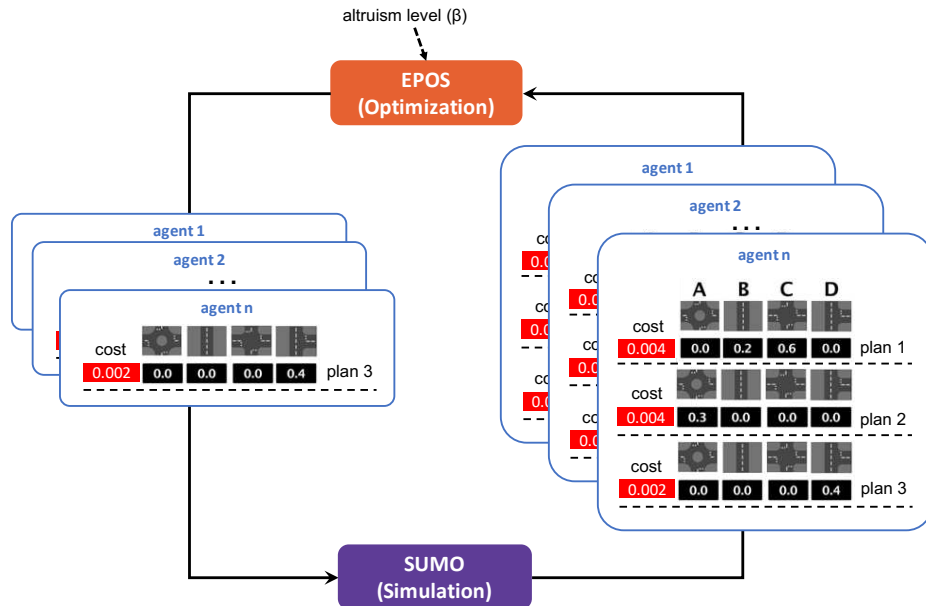


Figure 1: Plan specification and selection via inter-operation of SUMO and EPOS in our study.

of preferred plans. For instance, plan 3 for agent n in Figure 1 is preferred over plans 1 and 2. In our study, we assign costs using historical (simulated) runs of the AVs (see Study Design section).

In EPOS, an agent makes a plan selection based on two criteria: (i) global cost (GC) and (ii) local cost (LC). We model global cost as the variance of street utilization by all cars (by summing up element-wise all the selected plans and then calculating the variance of the resulting vector). Our model of global cost is a natural one, since global cost represents what needs to be optimized at system-level. This is a commonly used metric for quality of load balancing in computer and communication networks, but can also apply for load distribution on infrastructure or roadway traffic networks [11]. In our case, this is the variance of street utilization, since we would like to balance the AVs in the available streets to avoid congestion on any particular part of the network. The variance is a quadratic cost function and, as a result, agents need to coordinate their route selections [10]. EPOS is performing this coordination in a fully decentralized and privacy-preserving way. Local cost represents the preference for each plan, as provided by each agent (see Figure 1). The final cost of a plan is a weighted

sum of the two criteria:

$$(1 - \beta)GC + \beta LC \quad (1)$$

where β is real values in $[0,1]$.

The objective of each agent is to select the plan with the lowest final cost. High values of β represent more selfish agents, since they care more about their local cost than the global one. Conversely, lower values of β represent more altruistic agents.

Integration of SUMO and EPOS

In this study, we couple each car present in SUMO to an agent present in EPOS. Using the TRAPP framework [9] developed in our earlier work, we are able to run SUMO simulations which involve invoking EPOS at predefined time points (e.g. at the beginning of the simulation and periodically). At each invocation, the Python-controlled simulation pauses and waits for the EPOS run to complete. After completion, the selected routes are applied to the active simulation and the simulation resumes.

From user's perspective, the simulation can be configured with different maps, different number of cars, and different values of the β parameter that controls the altruism level of agents in EPOS. In order to evaluate the effect of balancing the cars in a city network—the main performance

indicator for the passengers of self-driving cars—we log the duration of all completed trips.

Study Design

This paper focuses on the following research questions:

- 1) What altruism level is required to coordinate agents' route choices such that the trip times are optimized?
- 2) How does a varied traffic level in the network influence the agents' coordination?

To investigate the above questions, we used the TRAPP framework that integrates SUMO and TraCI, a well-known traffic simulator, with EPOS, a decentralized optimization framework, as described in the previous section. We performed several simulation runs using different⁴ values of the β parameter of EPOS on different traffic settings—city maps and traffic levels—as described below. We also provide the open-source software TRAPP, together with concrete replication instructions online⁵.

For each *traffic setting* (explained next), we performed a systematic parameter sweep of β starting from 1 (corresponding to completely selfish agents) to 0 (corresponding to completely altruistic agents) with a step of 0.1. We kept all other TRAPP parameters constant to observe the effect of changing β alone. For each β value, we performed 5 simulation runs with different random seeds, which affect the initial positions and destinations of cars, to obtain statistical validity. Hence, for each traffic setting we performed a total of $5 \times 11 = 55$ simulation runs.

We consider a *simulation run* as the simulation of a specific number of cars on the specific city network for a time duration—the *simulation horizon*. The initial position of cars (each vehicle's origin) are selected to be proportional to the population distribution of city districts. This simulates morning commute where each vehicle begins its trip at a home location and

drives to a place of work that may be located anywhere within the city. Since no city-specific origin/destination data is available for the cities studied in this numerical example, a uniform distribution of trip destinations is assumed for the purpose of this numerical example. Note, however, that the underlying method of conducting traffic assignment using EPOS would also be applicable if additional data were available such as in the examples presented by [12], [13]. EPOS is invoked only at the beginning of each run, with a planning horizon equal to the simulation horizon, and selects one route for each car. Cars follow their routes without further planning. Once a car completes its trip, it picks another random destination and retrieves a new route via its preferred low-cost router. This ensures that the number of cars remains constant for the whole duration of a run. After a run completes, the duration of the first completed trip of each car are analyzed to determine the effect of EPOS optimization with the particular β value on traffic. In particular, we compute the average of all logged trip overheads, where a trip overhead is an actual trip duration divided by the theoretical trip duration if the car would drive alone in maximum speed.

A *traffic setting* in our study is defined by the city map and the number of cars in the simulation. The study considers four cities, Annapolis, Boulder, Duluth, and the borough of Manhattan in New York City. All of the cities used for the comparative study are located in the United States for the consistency of ZIP codes, census data, and commute data. These cities were also chosen for their diversity of urban infrastructure, including their area, population size, population density, and street organization.

The number of cars used for each simulation run was determined by comparing cities based on morning commute time. Hence, we first set the simulation horizon of a simulation run to 30 minutes, which roughly corresponds to the average commute time in the US of 25.5 minutes. Then, we calculate the total number of commuting trips using the percentage of drivers in each city that drive alone and the total population of the city from the 2010 census. This total number of commuting trips is divided by six, as typical morning peak traffic is from 6:30 am to 9:30 am, covering six half-hour time periods. Thus,

⁴We assume that all vehicles adopt the same *beta* value for each different *beta* we assess, i.e. homogeneous population. While such homogeneity is hard to achieve in practice and may require calibrated incentive mechanisms, we use it here as an approximation to avoid the explosion of the parameter space. Note also that earlier work indicates that such homogeneous cases can approximate well heterogeneous ones [7].

⁵https://github.com/iliager/TRAPP/blob/experiments/Beta_Alpha_Testing_README.md

the assumption that commuters are uniformly distributed in time over the morning commute is implicitly made. The number of cars and the simulation horizon thus represent a period of peak traffic of a morning commute.

To make the study more realistic, it is important to consider additional factors that contribute to traffic flow, like traffic route patterns. We used a simplified model for traffic patterns based on the assumption that during peak morning traffic people are driving from their places of residence to work. In our model, cities are divided into districts (squares in the map) with certain number of residents—population. We computed a district’s population by adding up the population of all ZIP codes whose centroids lie within the district. (Population per ZIP code is also obtained by the US 2010 Census data.) Then, we calculated the distribution of a city’s population in districts. Using this distribution, cars’ initial positions are assigned to districts; their specific position (a street within a district) is selected uniformly at random. Each trip destination is randomly selected within the city using a uniform distribution in space (without considering districts).

Finally, for each traffic setting, we conducted a number of baseline simulation runs to mine the cost that each agent associates with each router. The difference from the normal simulation runs that were used to derive the results is that in the baseline runs EPOS was not invoked. Instead, each car was selecting a router at random to perform a trip and logging the trip’s overhead. For each traffic setting, we performed 100 baseline runs with different random seeds. We then calculated the average trip overhead per car per router, which became the *cost* of that particular router for that car. These costs characterize the local objective of each agent in EPOS. They represent the effectiveness of a certain router in a certain traffic network.

Study Results

The *local cost* of a run represents the average cost of the selected routes. This gives an estimate of how much agents are dissatisfied — lower values of local cost are better. Figure 2a depicts the (normalized) local costs. For high β values, the local cost remains unaffected, however, for β close to 0, local cost increases sharply. Due

to different numeric values, trade-offs between global vs. local cost show sensitivity only in low β values⁶.

The *global cost* of a run represents the expected variance of street utilization. This gives an estimate of an important system-level objective of EPOS, i.e. to balance the presence of cars in the streets. Figure 2b depicts the (normalized) global costs. With varying β , we see the inverse trend than local costs: global cost decreases going from $\beta = 1$ to $\beta = 0$. As higher priority is given to the reduction of global cost with the reduction of β , the variance decreases, indicating more load-balanced street utilization.

The *trip overhead* is the main⁷ metric used to evaluate the effectiveness of our overall approach. For each setting we calculated the mean of the trip overheads corresponding to the first trip of each car. This provides an estimate of the overall utility of the system – with lower mean trip overhead corresponding to faster trips and hence higher system utility. Figure 2c depicts the (normalized) mean trip overheads of all runs performed for each traffic setting and for each β value.


The first observation is that trip overhead clearly depends on the setting, with Annapolis having the highest values and Duluth the lowest. With varying β values, trip overhead shows no discernible trend, except for two cases: Manhattan with $\beta = 0$ and Duluth with $\beta = 0$ both show a statistically significant decrease in the trip overhead. This trend is however not observed for Boulder and Annapolis. We show next that the traffic level of the network can influence the travel times equilibrium.

Influence of traffic level


As a next step, we hypothesized that the effect of optimization is stronger when the network load, i.e. traffic level, is at a critical state. The intuition is that balancing of traffic flows decreases trip overhead only if there is a certain amount of traffic congestion in the network. Alternative

⁶Techniques for Pareto front optimality and solutions to such numerical problems in weighted multi-objective optimization are studied in literature [14], [15], however, applying such schemes in a decentralized multi-agent context remains an open question.


⁷It measures the optimization effect more explicitly on the traffic network, in contrast to the global and local cost that are performance indicators of the (traffic-agnostic) EPOS optimization heuristic.




(a) Boxplots of local cost for all four cities as a function of selfishness β . Blue dots represent averages (over 5 runs). The trend shows that in all four cities, the local cost is constant except for the completely altruistic case ($\beta = 0$), where local cost is significantly increased.



(b) Boxplots of global cost for all four cities as a function of selfishness β , showing an increasing global cost with increased selfishness, meaning that when agents are more selfish, the overall cost for all agents is increased. Blue dots represent averages (over 5 runs).



(c) Boxplots of trip overheads for all four cities showing an increasing trend with selfishness β for Manhattan and Duluth, and no discernible trend with respect to selfishness β for Boulder and Annapolis. Blue dots represent averages (over 5 runs).



(d) Distribution of selected routers that route cars based on minimum length of streets (*Min length*), maximum speed allowed on streets (*Max speed*), or based on a combination of the two (*Balanced*) for different β values and different traffic settings.

Figure 2: Local costs, global costs, trip overheads and selected routers in our study.

routes in a network free of traffic congestion are likely to increase travel times, while alternative

routes in a congested networks are likely to reduce them.




Figure 3: Median results (over 5 random runs) for different number of cars in Eichstätt.

Given the large spectrum of the different performed experiments on large-scale networks, simulations are very computationally costly. For this reason, the aforementioned hypothesis is assessed on a smaller-scale network. We use the city of Eichstätt and varying number of cars ranging from 100 to 1500 with a step of 100 (see Figure 3). Cars have random starting points and random destinations. Similar to the other experiments, when a car reaches a destination it picks another one to drive to; this ensures that the total number of cars in a run remains constant. Each run has a duration of 800 SUMO ticks, which is long enough to ensure that almost all the initial trips of each car are completed. In these experiments, we investigate the influence of total altruism ($\beta=0$) versus total selfishness ($\beta=1$, the baseline) under different number of cars.

With respect to trip overheads (Figure 3), there is a critical state between 400 and 500 cars after which altruistic agents following the alternative routes consistently reduce the trip overhead. The differences in median of trip overheads observed between altruistic and selfish agents is on average 20 to 40 seconds across all settings with a different number of cars⁸. In a similar vein, we observe a sharp increase in the difference of global cost between altruistic and selfish agents for a number of cars starting from 700 on. On the contrary, local costs show a consistent difference between altruistic and selfish agents across all number of cars.

⁸For instance, for 200 vehicles, these calculations consider the median simulation duration (300 ticks in this case), the observed overheads, and the fact that one simulation tick in SUMO corresponds to one second’s duration in real-world.

Interpretation of results

The effect of changing the level of drivers’ altruism (β value) is both clear and consistent across city settings for both the local and global cost. Local cost is practically unaffected for β values other than 0 and increases sharply when complete altruism is in place ($\beta=0$). In complete altruism, optimization in EPOS takes into account only the global objective (“reduce the variance of street utilization”) without taking the agents’ preferences into account. Even slight consideration of agents’ preferences (e.g. $\beta = 0.1$ or 0.2) drastically reduces the cost that the agents pay for the optimization to take place. The same “all or nothing” pattern is present in the evolution of global costs: even slight introduction of selfish behavior is enough to increase the global cost considerably. Still, in contrast to local costs, the global costs show a more gradual value change by increasing the altruism level.

Looking at the results on trip overheads, we conclude that it is possible to use EPOS with altruistic agents, distribute vehicles more evenly in the streets and, as a result, reduce the overall trip overheads, especially when the network is at a critical high traffic flow state.

We clearly see such a positive traffic effect on average overhead values when setting $\beta=0$ for Manhattan and Duluth. However, such an effect (i) is only present for the case of complete altruism ($\beta=0$), and (ii) is not present in Boulder and Annapolis. This provides another indication that the traffic load on the network (and not the structural properties⁹ or the particular plans)

⁹We measured correlations of density without confirming influence.

explains the insignificant effect on trip overheads in the cases of large cities. Future work will confirm this using a large-scale computational infrastructure to support networks larger than Eichstätt.

In all settings, the distribution of router utilization changes as moving from altruistic ($\beta=0$) to selfish ($\beta=1$) agents, signaling the selection of the more preferred balanced router over the min-length one.

Conclusion and Outlook

In this paper, we focused on new mobility concepts in smart cities and investigated the use of multi-agent learning in optimizing route planning. Considering each (potentially autonomous) car as an agent that has several plans, i.e. routes to a destination, we investigated whether increasing the altruism of the agents can have a positive effect on the overall performance of traffic under varying traffic levels. We performed rigorous measurements to answer the above question using a simulation framework that integrates SUMO, a well-known traffic simulator, and EPOS, a decentralized agent-based framework. Our study focused on rush hour traffic in four US cities and found that (i) load-balancing is indeed achieved by increasing the agents' altruism, (ii) whether a positive effect on network performance can be observed depends primarily on the traffic load.

As future work, we would like to compare further cities under varying traffic levels. In our future experiments, we would like to consider cities for which real traces for traffic demand is available, as in [12] to further increase the validity of our results. Finally, a very interesting direction of research concerns the addition of fairness objectives in the decision making of agents.

REFERENCES

1. R. Mehmood, R. Meriton, G. Graham, P. Hennelly, and M. Kumar, "Exploring the influence of big data on city transport operations: a Markovian approach," *International Journal of Operations & Production Management*, vol. 37, no. 1, pp. 75–104, Jan. 2017. [Online]. Available: <https://www.emerald.com/insight/content/doi/10.1108/IJOPM-03-2015-0179/full/html>
2. M. G. H. Bell and Y. Iida, *Transportation network analysis*. Wiley, 1997.
3. D. K. Merchant and G. L. Nemhauser, "A model and an algorithm for the dynamic traffic assignment problems," *Transportation Science*, vol. 12, no. 3, pp. 183–199, 1978.
4. S. Peeta and A. K. Ziliaskopoulos, "Foundations of dynamic traffic assignment: The past, the present and the future," *Networks and Spatial Economics*, vol. 1, no. 3-4, pp. 233–265, 2001.
5. S. Samaranayake, W. Krichene, J. Reilly, M. L. D. Monache, P. Goatin, and A. Bayen, "Discrete-time system optimal dynamic traffic assignment (so-dta) with partial control for physical queuing networks," *Transportation Science*, vol. 52, no. 4, pp. 982–1001, 2018.
6. J. Wang, S. Peeta, and X. He, "Multiclass traffic assignment model for mixed traffic flow of human-driven vehicles and connected and autonomous vehicles," *Transportation Research Part B: Methodological*, vol. 126, pp. 139–168, 2019.
7. E. Pournaras, S. Jung, S. Yadhunathan, H. Zhang, and X. Fang, "Socio-technical smart grid optimization via decentralized charge control of electric vehicles," *Applied Soft Computing*, vol. 82, p. 105573, 2019.
8. I. Klein and E. Ben-Elia, "Emergence of cooperative route-choice: A model and experiment of compliance with system-optimal atis," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 59, pp. 348–364, 2018.
9. I. Gerostathopoulos and E. Pournaras, "TRAPPED in traffic?: a self-adaptive framework for decentralized traffic optimization," in *Proceedings of the 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS@ICSE 2019, Montreal, QC, Canada, May 25-31, 2019*, 2019, pp. 32–38. [Online]. Available: <https://dl.acm.org/citation.cfm?id=3341532>
10. E. Pournaras, P. Pilgerstorfer, and T. Asikis, "Decentralized collective learning for self-managed sharing economies," *ACM Trans. Auton. Adapt. Syst.*, vol. 13, no. 2, pp. 10:1–10:33, Nov. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3277668>
11. A. M. de Souza, R. S. Yokoyama, G. Maia, A. Loureiro, and L. Villas, "Real-time path planning to prevent traffic jam through an intelligent transportation system," in *2016 IEEE Symposium on Computers and Communication*. IEEE, 2016, pp. 726–731.
12. M. Gramaglia, O. Trullols-Cruces, D. Naboulsi, M. Fiore, and M. Calderon, "Mobility and connectivity in highway vehicular networks: A case study in madrid," *Computer Communications*, vol. 78, pp. 28–44, 2016.
13. C. Chen and Y. Zhu, "An evaluation of vehicular net-

Department Head

- works with real traces,” in *2012 IEEE 18th International Conference on Parallel and Distributed Systems*. IEEE, 2012, pp. 716–717.
14. K. Van Moffaert, T. Brys, A. Chandra, L. Esterle, P. R. Lewis, and A. Nowé, “A novel adaptive weight selection algorithm for multi-objective multi-agent reinforcement learning,” in *2014 International joint conference on neural networks (IJCNN)*. IEEE, 2014, pp. 2306–2314.
 15. J. Nikolic and E. Pournaras, “Structural self-adaptation for decentralized pervasive intelligence,” in *2019 22nd Euromicro Conference on Digital System Design (DSD)*. IEEE, 2019, pp. 562–571.

Ms. Brionna Davis is an undergraduate student in the School of Engineering at Vanderbilt University. Contact her at brionna.e.davis@vanderbilt.edu.

Ms. Grace Jennings is an undergraduate student in the Department of Electrical Engineering and Computer Science at Vanderbilt University. Contact her at grace.e.jennings@vanderbilt.edu.

Ms. Taylor Pothast is a student at Vanderbilt University studying to obtain both a Bachelors of Computers Science and a Bachelors of Human and Organizational Development by 2021. Contact her at taylor.m.pothast@vanderbilt.edu.

Dr. Ilias Gerostathopoulos is an Assistant Professor of Computer Science at Vrije Universiteit Amsterdam. His research interests lie in adaptive systems and software architectures. Contact him at i.g.gerostathopoulos@vu.nl.

Dr. Evangelos Pournaras is an Associate Professor at Distributed Systems and Services group, School of Computing, University of Leeds, UK. His research interest focus on distributed and intelligent social computing systems with expertise in the inter-disciplinary application domains of Smart Cities and Smart Grids. Contact him at E.Pournaras@leeds.ac.uk.

Dr. Raphael Stern is an Assistant Professor in the Department of Civil, Environmental, and Geo- Engineering at the University of Minnesota. His research interests are in the area of traffic control and estimation with autonomous vehicles in the flow. Contact him at rstern@umn.edu.