

# Centralised Vehicle Routing for Optimising Urban Traffic: A Scalability Perspective

Lukáš Chrpa

*Czech Institute of Informatics, Robotics and Cybernetics*  
*Czech Technical University in Prague*  
Prague, Czech Republic  
chrpaluk@cvut.cz

Mauro Vallati

*School of Computing and Engineering*  
*University of Huddersfield*  
Huddersfield, United Kingdom  
m.vallati@hud.ac.uk

**Abstract**—In the light of revolutionary technologies such as *connected autonomous vehicles*, centralised vehicle routing is attracting a growing interest as an effective method to tackle traffic congestion in urban areas, which causes enormous economic losses. Whereas potential benefits of centralised vehicle routing techniques are huge, they are not yet mature enough to be deployed in (large) urban areas. The major issue preventing their deployment being the lack of scalability.

This paper aims at providing an all encompassing discussion around how the scalability issue for centralised vehicle routing approaches can be addressed. In particular, we elaborate on how the model of the environment (the road network and the traffic) can be reasonably abstracted to allow simplified yet meaningful reasoning. Then, we provide an overview of relevant classes of decision-making techniques and elaborate how they can be applied to tackle the problem. At the end, we present our perspective on how different types of decision-making techniques can be effectively combined such that they can deal with the scalability issue while maintaining reasonable quality of assigned routes.

**Index Terms**—urban traffic management, intelligent vehicle routing, scalability

## I. INTRODUCTION

Congestion in urban areas has become one of the major economical problems due to losses directly caused by accidents and travel delays. The problem is far more apparent in rush hours in which traffic jams and road congestion have become a norm. According to recent estimates, the cost of traffic congestion in London has exceeded £5 billion in 2020 in lost time and fuel consumption<sup>1</sup>, and has also become a major health threat [1]. With continuing growth of global urbanisation, traffic congestion is expected to exacerbate. There is therefore a critical need for intelligent traffic management techniques, that can prevent or mitigate congestion and resulting negative effects on economy and society. The introduction of disruptive technologies such as *Connected Autonomous Vehicles* (CAV) has the potential to revolutionise the field by providing innovative means to design intelligent traffic control techniques [2]. A pivotal aspect of connected autonomous vehicles is their ability

to communicate with other vehicles (V2V) and with the infrastructure (V2I) via dedicated protocols and networks (VANET) [3]–[5]. This ability to communicate provides the ideal ground for Artificial Intelligence (AI) techniques to support the duties of urban traffic controllers. V2I communication allows the traffic controller to collect real-time traffic information from CAVs in the area, hence obtaining an up-to-date understanding of the overall traffic conditions, and to broadcast commands back to CAVs (e.g. their routes through the region).

The ability to communicate allows to explore new avenues of *intelligent vehicle routing*, that can play a crucial role in effective traffic management aiming at minimising travel time, fuel consumption or traffic density, to mention some. In particular, the problem deals with finding and assigning routes to vehicles such that they reach their destination while minimising (or maximising) a given objective function. Distributed approaches (e.g. [6]) tackle the problem from the perspective of individual vehicles, hence focus on dynamic user optimal (DUO) [7]. Instead, centralised approaches take the global perspective of the controlled network and aim to optimise traffic for the whole region, and are therefore well positioned to consider global objective functions (e.g., average travel time, average fuel cost, etc.) Notably, centralised approaches allow to consider both the dynamic system optimal (DSO) principle [8] and the mentioned dynamic user optimal (DUO) principle.

Centralised approaches for intelligent vehicle routing pose significant challenges, due to the need to consider a vast amount of information to provide useful routing instructions. For this reason, AI-based techniques are on the forefront of research in the field, due to their ability to perform systematic search. In particular, Automated Planning [9] is an example of a class of approaches that holds the promise to support the generation of centralised routing plans. Planning techniques have achieved promising results in traffic signal control [10], [11] and in intelligent vehicle routing [12], [13] have shown that the quality of solutions was higher than for traditional techniques. However, planning approaches, or in general approaches based on systematic search, are computationally demanding. In consequence, such techniques usually do not scale well and can be currently applied only in very small road networks [12].

Lukáš Chrpa was funded by the Czech Science Foundation (project no. 23-05575S) and the OP VVV project no. EF15\_003/0000470 “Robotics 4 Industry 4.0”. Mauro Vallati is supported by the UKRI Future Leaders Fellowship [grant number MR/T041196/1].

<sup>1</sup><https://www.london.gov.uk/press-releases/mayoral/cost-of-congestion-in-capital-revealed>

This paper aims at providing a discussion concerning the possibilities to mitigate scalability issues associated with centralised approaches for intelligent vehicle routing in road networks. Specifically, we consider two complementary perspectives, (i) simplifying the model of the environment while maintaining its reasonable accuracy and (ii) leveraging and combining different types of decision-making techniques to do the reasoning with the emphasis on their applicability in larger (and real) urban areas as well as their capability of delivering reasonable quality solutions. Regarding the model simplification, we discuss the possibility of using meso simulation (to route group of vehicles rather than individual vehicles) and the possibility of precomputing promising routes up front. Then, we discuss relevant types of decision-making techniques, ranging from reactive to deliberative ones, and its possible use for the centralised routing problem. At the end, inspired by the study of Cerny et al. [14] regarding the effective use of reactive and deliberative techniques, we present our perspective on how different types of decision-making techniques can be combined for effective centralised vehicle routing.

## II. RELATED WORK

The underlying idea of vehicle routing is to mitigate traffic congestion by assigning routes to vehicles moving through a controlled area. The assignment involves both a spatial and a temporal perspective, as it should consider not only a path reach the destination that is less congested at the moment, but also according to the changeable travel time and temporal-spatial distribution of the traffic flows [15].

Recent advances in CAV technology provides a once-in-a-lifetime opportunity for routing approaches to be increasingly practicable and to revolutionise the field [16]. Vehicle to Infrastructure (V2I) communication can support a one or two way exchange of information, where the vehicle provides the roadside agents with real-time traffic information, and the infrastructure can provide traffic or route suggestions to the vehicle. There is a wide range of information that can be communicated by the vehicle, for instance vehicle's origin and destination (OD), location, and currently planned route. The potential to exchange complex information offers the possibility to take into account, if possible, individual heterogeneity, leading to the increase of travel utilities for individuals [17]. Further, vehicle to vehicle (V2V) communication can encourage cooperation between vehicles, supporting the self organisation and self forecasting of traffic and traffic patterns.

Turning our attention to intelligent vehicle routing, in urban areas there is the possibility to consider either decentralised or centralised approaches [18]. The former case refers to cases where the system generating the route does not have a complete overview of the future evolution of the traffic in the region, that can reduce the overall quality of the routing but can lead to significant benefits in terms of computational resources and communication loads [6], [19], [20]. With V2V and V2I communication capabilities, the infrastructure and the vehicles can detect joint states and actions of the other

CAVs; the information is aggregated to estimate the travel time; thereby new routes are re-planned to avoid other CAVs' routes. It is worth noting that the routing decision-making of decentralised approaches is made at the level of individual vehicle based on traffic predictions. Due to the lack of efficient coordination between CAVs, the existing studies have applied the Monte Carlo tree search algorithm in short planning horizons [21], [22]. Other studies investigate the Dijkstra algorithm when a CAV reach a junction: to prevent the simultaneous route updating of CAVs the first-in-first-out (FIFO) principle [23] is also exploited. Others further extend the Dijkstra algorithm into a hierarchical and regional approach, reducing the search space into a limited range [6]. Another algorithm that also depends on ongoing communications among CAV clusters [24] is based on the A\* algorithm [25].

Compared to decentralised approaches, the centralised approach relies on the presence of a traffic management centre (TMC) that is in charge of optimising traffic for a controlled region. This comes with the clear benefit of an holistic vision of the behaviour and level of service of the network, hence the ability to generate optimal routes and to take into account both the dynamic system and the dynamic user optimal. Of course, the holistic perspective and the centralisation of computation comes with a significant cost, as the TMC needs to perform all the tasks involved in the vehicle routing problem, such as data collection, storage, and distribution. Specifically, this approach plays a vital role in the mixed traffic flow (CAVs and human-driven vehicles) network, where the central controller guides the CAVs to achieve dynamic system optimal in a specific zone [26]. Generating the dynamic best route for each individual CAV is the most computationally expensive part of the routing problem, and this is exacerbated in centralised approaches. The TMC needs to detect and predict traffic intensity and check whether the arrival time at each road intersection is delayed from the initially estimated time. If so, the rerouting process will be activated [27]. Recent works exploit domain-independent planning as means for centralised vehicle routing [12], [13]. These approaches achieved promising results in terms of routing quality but lack scalability.

## III. PROBLEM SPECIFICATION

Centralised intelligent vehicle routing deals with the problem of assigning routes to vehicles approaching or navigating the controlled region, such that each vehicle is provided with a path from its origin to its destination location that optimises a given global objective function (e.g. average travel time, average fuel consumption, CO<sub>2</sub> emissions).

A road network can be represented by a directed graph where vertices represent junctions and edges represent road segments. Then, we have a set of vehicles, where for each vehicle a pair of vertices representing Origin and Destination locations (hereinafter, OD, for short) is specified. The goal is to find a path for each vehicle in that graph representing a route that goes from the location of origin to the destination location while optimising a given objective function.

In an ideal case, the objective function would reflect how the traffic would evolve if vehicles have taken the assigned routes. That said, the objective function has to consider dynamic aspects such as movement of all vehicles in the network as they influence each other (e.g., in more intense traffic, the vehicles move more slowly, or vehicles on the side road have to give way to vehicles on the main road etc.).

#### IV. MODEL GRANULARITY AND ACCURACY

The model of the road network and traffic has to capture the topology and the dynamics associated with traffic flows in the considered region. In literature, the traffic simulation models are traditionally divided into three main categories.

- **Micro simulation** — The model considers traffic at the level of individual vehicles, that is, it models how an individual vehicle behaves in the road network (e.g. which route it navigates).
- **Macro simulation** — The model considers traffic at the level of traffic flows, that is, it models how the traffic evolve in the road network (e.g., where traffic is heavier or lighter).
- **Meso simulation** — The model combines the aspects of both micro and macro simulation models. In other words, it is possible to count and measure vehicles, but not to identify individual vehicles.

Although traffic light control approaches traditionally operate over macro-simulation models (see e.g. [10], [11]), for intelligent vehicle routing approaches the use of micro-simulation models seems to be more appropriate. On the other hand, micro-simulation-based approaches tend to scale poorly when the number of vehicles they have to consider is high (see e.g. [12]). The number of vehicles considered for routing can be reduced by taking into account only some ODs and/or by considering a small time span for which data of vehicles that need routes are collected [13].

Such a strategy might tackle the scalability issue only to some extent as it might be useful to remove some ODs (due to insignificant amount of traffic or due to existence of only one feasible route) as well as having too large time span to consider would be impractical (since the vehicles would have to wait for a long time to get their routes after broadcasting their intentions into the system). On the other hand, considering very small time span would likely lead to inaccurate results as only a few vehicles will be routed at time which would have similar effect than using distributed routing approaches. Also, in larger metropolitan areas, the number of trips that should be considered would still be very high (hundreds or thousands per minute).

A possible remedy to reduce the number of vehicles in vehicle routing episodes is to cluster vehicles with similar ODs and timings, and generate routes for these clusters such that each individual vehicle will get the same assigned route. Hence, the idea is to consider meso simulation rather than micro simulation, where aspects of flows and individual vehicles are mixed. However, it should be noted that how the clusters are made affects the scalability as well as the accuracy of the

routing method. In particular, having too small clusters might not help to improve scalability while having too large clusters might lead to too suboptimal routes. The latter case might prohibit distributing traffic among different routes if vehicles have the same ODs (benefits of distributing traffic among different routes has been elaborated in e.g. [12], [28]). Hence, it is important to maintain a tradeoff between scalability of the method and quality of solutions by considering appropriate sizes of vehicle clusters (sizes of clusters might differ for different ODs). We believe that an anytime algorithm that starts with large vehicle clusters and iteratively reduce their sizes could be a promising direction for achieving that tradeoff. Alternatively, a machine learning method might be considered to train an estimator of vehicle cluster sizes based on the road network topology and traffic intensity (here, however, we might be lacking useful training data).

The size of the considered road network, i.e., the number of junctions and the number of road segments, as well as the geographical distribution, affects scalability of the centralised vehicle routing methods as well. In larger metropolitan areas, these numbers can easily reach tens or hundreds thousands. Although it is usually sufficient to consider more significant roads, i.e., discarding those that are very local (e.g., narrow roads to houses or apartment buildings), the number of remaining road segments (and junctions) can still be very high. A possible remedy is to precompute several feasible (and promising) routes for each OD. This can be done offline and can considerably reduce the number of road segments and junctions that the centralised routing techniques have to consider. Precomputed routes should not be very suboptimal and should be diverse enough. The reason behind reasonably small suboptimality (in terms of expected travel time or distance, for example) is straightforward as very suboptimal routes would not be beneficial. The reason behind route diversity is to increase chances of avoiding the same bottlenecks, i.e., road segments and junctions that are likely to get congested if the traffic intensity increases.

Another aspect that affects scalability of centralised vehicle routing techniques is to what detail the dynamics of the road network is considered. Technically, partial solutions can be evaluated by using detailed mathematical models of traffic or by using realistic traffic simulators such as SUMO [29]. However, such an approach might be too expensive as during the generation of the solutions (especially for larger problems) there might be simply way too many such evaluations. Automated Planning based approaches simplify the dynamics model by estimating the level of traffic intensity for each road segment and optimise vehicle routes such that the number of road segments with more intense traffic is minimised [12], [30]. Such a simplification of the dynamics model does help to considerably reduce time needed for determining cost of partial solutions as well as to calculate heuristics. However, dynamics models of [12], [30] derive traffic intensity levels from the physical capacities of given road segments and relax other aspects of the network topology (e.g., junctions, traffic on nearby road segments). Consequently, the solutions might not

be of a good quality due to lack of accuracy of such simplified models of dynamics. A possible remedy to enhance accuracy of such simplified models is to leverage machine learning techniques that can dynamically readjust traffic intensity levels according to observation of the current traffic situation (real or simulated). A similar idea has been used in planning of traffic signal control [31].

## V. DELIBERATIVE AND REACTIVE DECISION-MAKING TECHNIQUES

Decision-making techniques, in a nutshell, deal with generating (sequences of) elementary decisions that aim to achieve given goals or optimise a given objective function. We distinguish between two main categories of decision-making techniques.

- **Reactive techniques** — Such techniques aim at short-term goals or rewards and the solutions often contain only a few elementary decisions. These techniques are fast and can usually operate in (almost) real time.
- **Deliberative techniques** — Such techniques aim at long-term goals and are capable of generating more complex solutions involving hundreds elementary decisions. These techniques are, however, time consuming and are often used offline.

Deliberative techniques, in the context of intelligent vehicle routing, are capable of effectively distribute traffic among the road network while optimising a given “global” objective function and hence operate from the centralised perspective. Reactive techniques, on the other hand, can operate from a rather decentralised perspective. Whereas reactive techniques can provide (decentralised) individual vehicle paths in almost real-time as techniques such as Contraction Hierarchies [32] can scale up to enormously large networks [33], deliberative techniques do not scale very well and can be used only on small road networks (so far) [12], [30].

Decision-making techniques that can be leveraged for intelligent vehicle routing can be divided into the following main categories.

- **Rule-based Control** — Predefined or learnt rules are used to make immediate elementary decisions according to the current situation. Such techniques are very fast and can be used in real-time.
- **Local Search** — Local search techniques, in a nutshell, explore neighbourhood of the current state as long as the value of a given objective can be improved. They provide shorter sequences of elementary decisions. These techniques are also fast and do not consume much memory (only to keep the current state and its immediate neighbourhood).
- **Reinforcement Learning** — The aim of such techniques, in a nutshell, is to learn which elementary decision is the most promising in each state. If the training data are representative enough, i.e., cover a large portion of the state space, then they can work reasonably well for longer-term goals or rewards.

- **Monte-Carlo Tree Search** — Such techniques provide a tradeoff between exploration and exploitation with the aim to work in a limited amount of time. These techniques expand only a part of the search tree and then from each leaf node they perform a number of simulations (usually to a terminal state) to estimate an expected reward (or penalty) which is backpropagated to the particular leaf node. Based on such a reward estimate, the most promising elementary decision (in the current state) is performed.
- **Systematic Search** — Such techniques perform an exhaustive search in order to provide a sequence of decisions that achieves a specified goal. If the state space is finite, then these techniques (specifically, their optimal variants) are able to find sequences of elementary decisions that minimises (or maximises) the value of a given objective function. However, systematic search techniques are very demanding on runtime (as well as on memory).

Systematic search techniques such as automated planning, which are representative of deliberative reasoning, have the potential to (nearly) optimally tackle the problem of centralised vehicle routing. To do so, the model of the environment would have to be (nearly) accurate and elementary decisions would have to refer to choosing to which adjacent road segment a single vehicle would move. That said, each route for a single vehicle would consist of a sequence of elementary decisions reflecting the path in the graph representing a given road network. As indicated in literature such an approach does not scale well even with (much) simpler model of the environment [12]. Although there might be some potential in improving scalability by developing an effective domain-dependent heuristic, it might not be sufficient for larger road networks (and a larger number of vehicles). For the other categories of decision-making techniques, keeping elementary decisions on the road segment level does not seem to be very practical as the state space can be huge and accurately evaluating states with partial paths might be expensive (and not that straightforward).

As elaborated in the previous section, a promising simplification of the model is to precompute promising (and diverse) routes for each OD. Although it might affect the accuracy of the model (some better routes might be omitted), it can reduce the size of the state space as in this context, the elementary decisions would consider assigning (or reassigning) a whole route for a single vehicle. Such a simplification of the model can improve scalability of the systematic search methods as the state space would be (much) smaller and solutions, sequences of elementary decisions, would be shorter. Also, the simplified mode would be more suitable for the other (more) reactive categories of decision-making techniques. For example, local search techniques might quickly reassign routes for vehicles if their original routes became block by a traffic accident.

With regards to (deep) reinforcement learning techniques that in recent years achieved impressive results in domains such as the game of Go [34] or Starcraft [35] their applicability

to centralised vehicle routing depends on the availability of training data and “complexity” of good quality solutions. The latter aspect is derived from an observation that if more complex plans or policies are required to achieve longer-term goals, such as in the Montezuma’s Revenge game, deep reinforcement learning is not very efficient [36]. Regarding the availability of training data, we can assume that regular patterns of the traffic can be acquired from sensory data, for instance. However, unusual and unexpected circumstances such as flooded roads, blockage of some area by a social event are unlikely to be covered by training data and might limit the (general) use of reinforcement learning techniques.

Monte-Carlo Tree Search techniques might be an interesting alternative for centralised vehicle routing with models considering (re)assigning whole routes to the vehicles (recent results concerning highway driving have shown some promise [21]). The advantage of such techniques is a limited runtime, so the routes can be always assigned in time, and no need for training data. A possible drawback might be lower quality of solutions. We believe that Monte-Carlo Tree Search techniques might work as a backup for systematic search techniques (e.g. planning) such that if the latter technique fails to generate a solution in time, the former technique is invoked (the techniques can also run in parallel to leave more time for both).

## VI. DISCUSSION AND CONCLUSION

Centralised intelligent vehicle routing techniques are very important for optimising urban traffic and thus mitigating economic losses caused by traffic congestion. In the light of new revolutionary technologies such as CAV and VANET, the need for scalable techniques for intelligent vehicle routing is even more apparent. Centralised routing techniques are capable of taking a global perspective on the situation and have the potential to route traffic more effectively and thus minimising the risk of traffic congestion.

Jimoh et al. [28] pioneered the idea of using automated planning for vehicle routing as a component for an autonomic architecture that self-manage the traffic in a controlled urban region. More recent works [12], [13], [30] have shown that planning techniques have a potential to tackle the centralised vehicle routing problem albeit they do not scale well. Taking from a general perspective, planning techniques belong to the umbrella of deliberative decision-making techniques based on systematic search. Even though they are theoretically capable of finding an optimal solution if the state space is finite, they are computationally very demanding (as they are intractable) and might be applicable only on (very) small road networks.

As elaborated in Section IV, the model of the environment (road network and traffic) can be simplified such that promising routes between particular ODs are precomputed, meso-simulation is used instead of micro-simulation and simplifying the model of traffic dynamics. Although such a model simplification might affect its accuracy, it can considerably reduce the size of the state space and, possibly together with an informative domain-dependent heuristic, it can make

systematic search techniques more scalable. The question how large road networks can be handled by such improved systematic search techniques has yet to be examined.

It, however, might be the case that a systematic search technique fails to generate a solution in a given time limit (usually a few seconds). It should be noted that there is no guarantee that such techniques will finish in a given time limit (on a given problem) and runtimes might vary even on similar problems (in size). As we pointed out in Section V it might be possible to leverage Monte-Carlo Tree Search methods as a backup since, we believe, they might be capable to deliver somewhat reasonable solutions. In situations, when some event (e.g. a traffic accident) suddenly changes some aspects in the road network, local search-based techniques might be leveraged to quickly tackle the situation (e.g., reassign routes to vehicles that were routed through an accident site). Resorting to more reactive decision-making techniques in cases in which deliberative techniques might not deliver results quickly enough seems to be an effective strategy to deal with more uncertain situations (e.g., traffic accidents, rapid change of traffic flows) as studied by Cerny et al. [14]. To take a more general perspective, lack of scalability of deliberative (and systematic) methods can be backed up by (more) reactive methods albeit at cost of (possibly) more suboptimal solutions.

We would like to note that we argue more strongly towards traditional AI techniques to tackle the problem of centralised vehicle routing, mainly because we feel that obtaining training data that are representative enough might not be possible. In our opinion, (deep) learning methods can be effectively used in more common traffic situations that might be well covered by acquired historical traffic data. The training process might be enhanced by computing (nearly) optimal solutions for some situations by deliberative techniques, which can be done offline. Less common and unusual traffic situations might still be better handled by traditional AI techniques as these might not be well covered by training data. However, the question whether we can train (deep) neural networks that generalize well even for less common traffic situations has yet to be answered.

This paper presents the authors’ perspective on how the centralised vehicle routing problem can be effectively addressed. We believe that effective combination of deliberative and reactive techniques alongside with a proper simplification of the traffic models provides a tradeoff between (lack of) scalability on one side and quality of routing on the other side. In future, we plan to follow the ideas we presented in this paper in order to integrate the aforementioned decision-making techniques into a framework that would route vehicles in larger urban networks.

## REFERENCES

- [1] K.-H. Chang, P.-Y. Hsu, C.-J. Lin, C.-L. Lin, S.-H. H. Juo, and C.-L. Liang, “Traffic-related air pollutants increase the risk for age-related macular degeneration,” *Journal of Investigative Medicine*, vol. 67, no. 7, pp. 1076–1081, 2019.

- [2] M. Vallati and L. Chrpá, "A principled analysis of the interrelation between vehicular communication and reasoning capabilities of autonomous vehicles," in *Proceedings of ITSC*, 2018, pp. 3761–3766.
- [3] P. S. Narayanan and C. S. Joice, "Vehicle-to-vehicle (v2v) communication using routing protocols: a review," in *2019 International Conference on Smart Structures and Systems (ICSSS)*. IEEE, 2019, pp. 1–10.
- [4] A. Rasheed, S. Gillani, S. Ajmal, and A. Qayyum, "Vehicular ad hoc network (vanet): A survey, challenges, and applications," in *Vehicular Ad-Hoc Networks for Smart Cities*. Springer, 2017, pp. 39–51.
- [5] F. Arena and G. Pau, "An overview of vehicular communications," *Future Internet*, vol. 11, no. 2, p. 27, 2019.
- [6] J. Wang and H. Niu, "A distributed dynamic route guidance approach based on short-term forecasts in cooperative infrastructure-vehicle systems," *Transportation Research Part D: Transport and Environment*, vol. 66, pp. 23–34, 2019.
- [7] T. L. Friesz, J. Luque, R. L. Tobin, and B.-W. Wie, "Dynamic network traffic assignment considered as a continuous time optimal control problem," *Operations Research*, vol. 37, no. 6, pp. 893–901, 1989.
- [8] 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.
- [9] M. Ghallab, D. Nau, and P. Traverso, *Automated Planning: theory and practice*. Elsevier, 2004.
- [10] M. Vallati, D. Magazzeni, B. D. Schutter, L. Chrpá, and T. L. McCluskey, "Efficient macroscopic urban traffic models for reducing congestion: A PDDL+ planning approach," in *Proceedings of AAAI*, 2016, pp. 3188–3194.
- [11] T. L. McCluskey and M. Vallati, "Embedding automated planning within urban traffic management operations," in *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS 2017, Pittsburgh, Pennsylvania, USA, June 18-23, 2017*, 2017, pp. 391–399.
- [12] L. Chrpá, M. Vallati, and S. Parkinson, "Exploiting automated planning for efficient centralized vehicle routing and mitigating congestion in urban road networks," in *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC 2019*, 2019.
- [13] M. Vallati, E. Scala, and L. Chrpá, "A hybrid automated planning approach for urban real-time routing of connected vehicles," in *24th IEEE International Intelligent Transportation Systems Conference, ITSC*. IEEE, 2021, pp. 3821–3826.
- [14] M. Cerný, R. Barták, C. Brom, and J. Gemrot, "To plan or to simply react? an experimental study of action planning in a game environment," *Comput. Intell.*, vol. 32, no. 4, pp. 668–710, 2016.
- [15] C. Tang, W. Hu, S. Hu, and M. E. Stettler, "Urban traffic route guidance method with high adaptive learning ability under diverse traffic scenarios," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 5, pp. 2956–2968, 2020.
- [16] X. Lin, M. Li, Z.-J. M. Shen, Y. Yin, and F. He, "Rhythmic control of automated traffic—part ii: Grid network rhythm and online routing," *Transportation Science*, vol. 55, no. 5, pp. 988–1009, 2021.
- [17] C. Wang, S. Peeta, and J. Wang, "Incentive-based decentralized routing for connected and autonomous vehicles using information propagation," *Transportation Research Part B: Methodological*, vol. 149, pp. 138–161, 2021.
- [18] G. S. Shahi, R. S. Bath, and S. Egerton, "Mrgm: an adaptive mechanism for congestion control in smart vehicular network," *International Journal of Communication Networks and Information Security*, vol. 12, no. 2, pp. 273–280, 2020.
- [19] L. Du, L. Han, and X.-Y. Li, "Distributed coordinated in-vehicle online routing using mixed-strategy congestion game," *Transportation Research Part B: Methodological*, vol. 67, pp. 1–17, 2014.
- [20] R. Claes, T. Holvoet, and D. Weyns, "A decentralized approach for anticipatory vehicle routing using delegate multiagent systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 2, pp. 364–373, 2011.
- [21] D. Lenz, T. Kessler, and A. Knoll, "Tactical cooperative planning for autonomous highway driving using monte-carlo tree search," in *2016 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2016, pp. 447–453.
- [22] K. Kurzer, C. Zhou, and J. Marius Zöllner, "Decentralized cooperative planning for automated vehicles with hierarchical monte carlo tree search," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 2018, pp. 529–536.
- [23] A. Mostafazi, C. Koll, and H. Wang, "A decentralized and coordinated routing algorithm for connected and autonomous vehicles," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [24] A. Mostafazi, S. Dong, and H. Wang, "Percolation phenomenon in connected vehicle network through a multi-agent approach: Mobility benefits and market penetration," *Transportation Research Part C: Emerging Technologies*, vol. 85, pp. 312–333, 2017.
- [25] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Trans. Syst. Sci. Cybern.*, vol. 4, no. 2, pp. 100–107, 1968.
- [26] M. H. Mansourianfar, Z. Gu, S. T. Waller, and M. Saberi, "Joint routing and pricing control in congested mixed autonomy networks," *Transportation Research Part C: Emerging Technologies*, vol. 131, pp. 103–338, 2021.
- [27] G. P. Rocha Filho, R. I. Meneguette, J. R. T. Neto, A. Valejo, L. Weigang, J. Ueyama, G. Pessin, and L. A. Villas, "Enhancing intelligence in traffic management systems to aid in vehicle traffic congestion problems in smart cities," *Ad Hoc Networks*, vol. 107, pp. 102–265, 2020.
- [28] F. Jimoh, L. Chrpá, T. L. McCluskey, and S. Shah, "Towards application of automated planning in urban traffic control," in *16th International IEEE Conference on Intelligent Transportation Systems, ITSC 2013, The Hague, The Netherlands, October 6-9, 2013*. IEEE, 2013, pp. 985–990.
- [29] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *Proceedings of ITSC*, 2018.
- [30] M. Vallati and L. Chrpá, "Effective real-time urban traffic routing: An automated planning approach," in *7th International Conference on Models and Technologies for Intelligent Transportation Systems, MT-ITS*, 2021, pp. 1–6.
- [31] A. Pozanco, S. Fernández, and D. Borrajo, "On-line modelling and planning for urban traffic control," *Expert Systems*, 2021.
- [32] R. Geisberger, P. Sanders, D. Schultes, and D. Delling, "Contraction hierarchies: Faster and simpler hierarchical routing in road networks," in *Experimental Algorithms, 7th International Workshop, WEA 2008, Provincetown, MA, USA, May 30-June 1, 2008, Proceedings*, ser. Lecture Notes in Computer Science, C. C. McGeoch, Ed., vol. 5038. Springer, 2008, pp. 319–333.
- [33] J. Blum and S. Storandt, "Scalability of route planning techniques," in *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling, ICAPS 2018, Delft, The Netherlands, June 24-29, 2018*, M. de Weerd, S. Koenig, G. Röger, and M. T. J. Spaan, Eds. AAAI Press, 2018, pp. 20–28.
- [34] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. P. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, pp. 484–489, 2016.
- [35] O. Vinyals, I. Babuschkin, W. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. Agapiou, M. Jaderberg, and D. Silver, "Grandmaster level in starcraft ii using multi-agent reinforcement learning," *Nature*, vol. 575, 11 2019.
- [36] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.