Paper ID 18056

A Centralised Framework for Maximising the Utilisation of Urban Road Networks by Leveraging on Connected Vehicles

Mauro Vallati[1*], **Zeyn Saigol**[2]

1. School of Computing and Engineering, University of Huddersfield, United Kingdom

2. Connected Places Catapult, United Kingdom

*Email: m.vallati@hud.ac.uk

Abstract

One of the pivotal challenges presented to urban road traffic controllers is the effective utilisation of transport infrastructure, as a result of growing urbanisation, the finite network capacity, and of the increasing number of road vehicles. In this context, the arrival of connected autonomous vehicles (CAVs) represents a unique opportunity for a fundamental change in urban mobility, and urban traffic control should take an active role in integrating CAVs into the mobility ecosystem in order to maximise benefits.

To support this integration, we propose to leverage on a centralised architecture that can exploit, for instance, Artificial Intelligence techniques to distribute vehicles in a controlled urban region, with the aims of reducing congestion and fostering a balanced use of the available road network. In this paper, we describe the overall framework that is under development in the AI4ME project, funded by the UK Engineering and Physical Sciences Research Council. Further, we demonstrate the impact of the proposed approach using real-world historical data of a large UK town.

Keywords:
Artificial Intelligence, Connected Vehicles.

Introduction

Over half of the world's population now lives in cities and global urbanisation continues at a steady pace. In the UK alone, the cost of congestion has reached nearly £8 billion in 2018 in lost time and fuel consumption, and has become a major health threat that goes beyond the cardiac and respiratory systems [1].

In this context, the arrival of Connected Autonomous Vehicles (CAVs) present a unique opportunity for a fundamental change in urban mobility and urban traffic control. CAVs can communicate with other vehicles and with the infrastructure, via dedicated protocols, to take better informed decisions. While the general importance and improved capabilities that can be obtained via different types of communication have been well argued [2], and a number of protocols and technologies to implement such communications has been presented [3], here we focus on how urban traffic control can exploit communication with vehicles to improve the use of the controlled network. Nowadays, it is often the case that during rush hours most main roads are congested, at least in one direction, while many other roads are underused. Such under-use of the network is caused by the fact that traffic often navigates via the same route between given way- points, and is often results from similar behaviour and habits of vehicle drivers. Considering the available network, it is often the case that the exploitation of alternative routes can lead to a better distribution of vehicles and a better use of the capacity of the network.

This paper introduces a framework for maximising the utilisation of the urban road network with the main aim of minimising traffic congestion and reduce average journey times. The framework exploits a centralised approach: connected vehicles in the controlled area communicate their destination and route to a controller, that, given its understanding of the current and predicted status of the network, can decide whether to propose alternative routes to a vehicle. The centralised perspective allows to perform network-wise reasoning, that would be impossible in a decentralised architecture.

The proposed approach can foster the use of advanced Artificial Intelligence (AI) techniques; a number of AI-based techniques have been proposed to deal with, for instance for controlling traffic lights [4,5], for finding static routes for vehicles traversing a network [10], or for predicting traffic conditions in the immediate future [8,9]. The architecture we propose in this work is agnostic with regards to the technique used, and provides a clear definition of expected input and output. To demonstrate its effectiveness, we describe a technique based on road occupancy, and we test it, in simulation, on real-world traffic data of the Milton Keynes centre area.

Overall Architecture

The architecture of the proposed framework is depicted in Figure 1. The main aim of the framework is to support the exploitation of techniques for maximising the use of the controlled urban network by generating personalised routes, in the sense that each vehicle will receive re-routing instructions based on its current destination and assigned route. In Figure 1, blue is used to indicate data sources, white indicates input and output of the centralised controller, and yellow is used to highlight the controller module. It should be noted that the proposed architecture can be modified according to the available data sources, and that real-time data can be integrated with historical data.

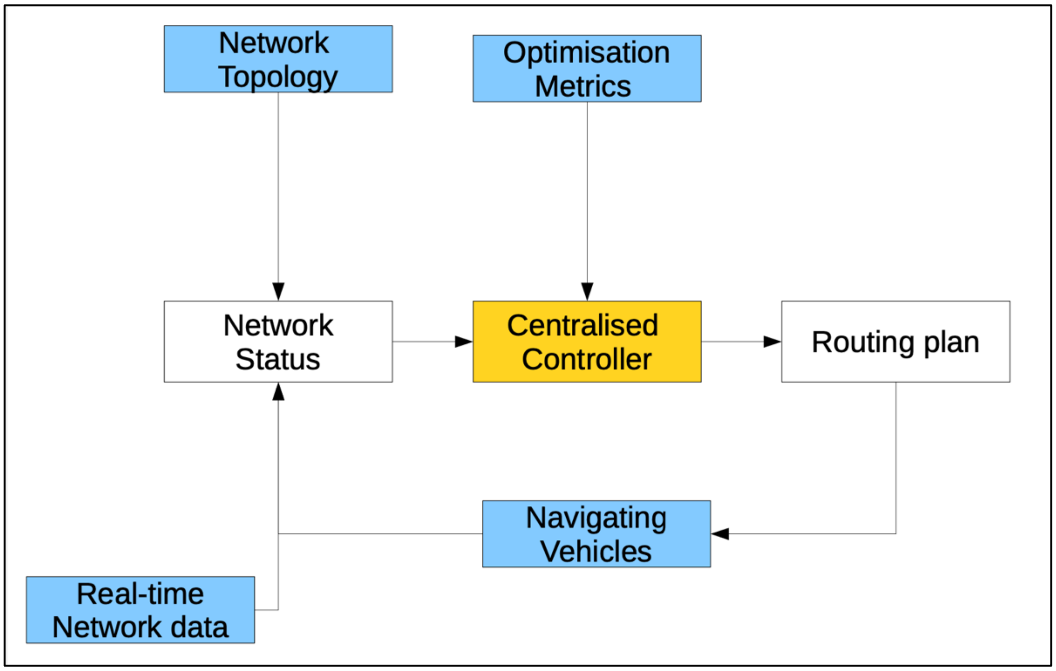

Figure 1. The proposed architecture, composed of data and knowledge (blue), the AI-based controller (yellow), and input/output (white) of the controller – which come under the form of aggregated data.

The main element of the framework is the controller; by taking into account the vast amount of available information and data, the controller is in charge of deciding what are the best routes to be followed by the connected vehicles that are navigating the network, according to the considered optimisation metrics. In particular, the network status is defined by:

- *Network topology*. This includes a description of the topology of the network in terms of links, junctions, etc.
- *Real-time data*. Available sensors distributed in the controlled region can provide additional information about the status, for instance in terms of air quality, congestion, accidents, etc. In fact, it is worth noting that there may be vehicles that are not connected, and therefore their presence can only be identified via traditional sensors (pressures, cameras, etc.).
- *Navigating connected vehicles*. Data acquired from vehicles that are navigating or approaching the controlled region include: current location, final destination, and the currently assigned route.

The defined status is then provided as input to the centralised controller that, by taking into account some predefined optimisation metrics, is in the best position for assessing the network conditions and deciding if there is the need for re-routing vehicles in order to maintain a required level of service, or to counteract unexpected network conditions. Here we focus on the controller's capability of re-routing connected vehicles that are navigating in the controlled urban region. We define as a routing plan the set of routes that are assigned to vehicles that are navigating the network. If the current route of a vehicle is modified in the routing plan, the affected vehicles have to be notified accordingly. Here we assume that an adequate communication infrastructure is in place for supporting the necessary two-ways communication between controller and the vehicles.

The controller

The centralised controller plays a pivotal role in the depicted architecture. In a nutshell, its main duties are:

a. Analyse the current status of the network, that is provided under the form of aggregated data.

b. Predict the likely evolution of the network, on the basis of the available information.

c. Given a. and b. above, decide if any of the vehicles entering the controlled region need to be re-routed. In case of re-routing, an appropriate alternative has to be communicated in a timely fashion to the affected vehicles.

Notably, one of the major strengths of a centralised controller is the ability to take into account the condition of the network and predict likely future evolutions of traffic conditions, to support better informed decisions. This cannot be performed in a decentralised fashion, where a complete overview of the network’s condition is not available, but instead each vehicle or intelligent agent has its own limited knowledge of a subpart of the network.

The centralised controller has to operate in near real-time, in order to perform the above listed steps a. and b., and provide timely information to vehicles. Therefore, the analysis and predictions have to consider the trade-off between accuracy and computational complexity (in terms of time needed).

In this paper, we do not advocate for a specific approach to be used by the controller in order to guarantee the best results. Instead, we highlight that the proposed architecture can accommodate a wide range of approaches.

In order to demonstrate the usefulness of the described framework, we designed a controller that is based on the notion of occupancy, as the number of passenger car units (PCUs) that are on a specific road link at the same time. Occupancy is calculated regardless of the status of the vehicles, i.e., it does not include information about the fact that vehicles are moving or waiting, but includes information about the direction of vehicles. In other words, for each link there can be two occupancy values, one per direction (only one is needed in the case of one-way roads). Given this notion of occupancy, the current condition of the network (point a. above) is summarised by the occupancy value of all the road links in the controlled region. The future evolution of the network (point b. above) is then assessed in terms of “cumulative occupancy”. Given the current condition of the network, and the routes assigned to vehicles that are navigating the region, the predicted future condition of the network is given by assuming that each vehicle is occupying –at the same time– all the road links included in its route. In a sense, this approach is taking a very conservative vision of the traffic, by assuming that each vehicle is "booking" all the links included in the assigned route. Intuitively, this heuristic works well in busy conditions, where the number of vehicles that are entering the network is similar to vehicles exiting, and the network is overall busy.

The idea behind the proposed approach is to force the controller to take pre-emptive re-routing decisions, that reduce the risk of congestion by maximising the distribution of vehicles in the network. Given this prediction, approaching vehicles communicate their destination and the current route to the controller. The controller can then exploit the prediction made to identify the most promising route to be taken by the vehicle in order to reach the destination, by considering the route with the minimum cumulative occupancy. Alternative routes can either be pre-calculated for an origin-destination pair, or can be calculated during the re-routing process. The latter approach can further increase the computational burden on the controller.

For the sake of this analysis, and in order to minimise the burden on vehicles and limit the computational complexity on the controller side, we assume that vehicles can be rerouted only when they are on the boundaries of the controlled region. In other words, the controller can only consider the option of re-routing vehicles that are on links directly connected to an entry point.

Empirical Evaluation

To assess the usefulness of the proposed centralised architecture, here we consider a SUMO [6] microsimulation model of Milton Keynes centre. The network is shown in Figure 2. Milton Keynes is a town of the United Kingdom, located about 80 kilometres north-west of London. Milton Keynes has a population of circa 230,000. The model covers an area of approximately 2.9 square kilometres, and includes more than 25 junctions and more than 50 links.

The model simulates the morning rush hour, and has been built by considering historical traffic data collected between 8am and 9am on non-holiday weekdays. Data has been provided by the Milton Keynes Council, and gathered by sensors distributed in the region between December 2015 and December 2016. Traffic signal control information has been provided again by the Council. The model has been calibrated and validated.

During the modelled period, 1,900 vehicles enter the controlled region either to navigate through it, or for reaching one of the parking slots in the area. No vehicle is initially in the area, but they are all injected over the simulation time. All vehicles have a pre-defined destination and route. The largest flow of traffic comes from the west entry points, via the large North Grafton roundabout (the largest roundabout of the map). Many residential areas are connected with the centre of Milton Keynes through those entry points.

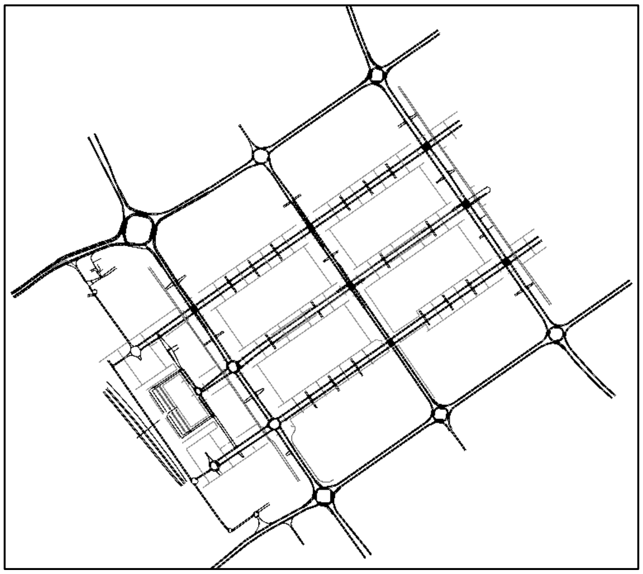

Figure 2. The modelled central Milton Keynes urban area. Avebury, Midsummer, and Silbury Boulevard are the three parallel road links traversing the area SW to NE.

The framework presented in the previous section has been implemented in Python, and uses the TRaCI interface to interact with the SUMO simulation environment, in order to get the current network status, communicate with approaching vehicles, and inform vehicles of re-routing. By focusing on the considered notion of occupancy, the controller is aiming at reducing congestion and maximising the use of the network. For every origin-destination, between 2 and 3 alternative routes are considered for distributing traffic. Such routes have been provided by a human expert. While alternatives can, in principle, be automatically calculated (see, for instance, [7]), relying on human expertise can allow to exploit some insights that are based on knowledge that is not captured by the symbolic model of the network. The simulation is run for 1 hour and then stopped. For each set of experiments, the simulation is run five times and results are averaged, to account for non-determinism.

Results

In a first set of experiments, we consider the ideal scenario: all the vehicles that are navigating through the network are CAVs, and follow the re-routing instructions provided by the centralised controller. Figure 3 presents a comparison of the average travelling times of vehicles navigating the network with (green) and without (black) the use of the proposed framework. When on “default”, vehicles are following their pre-computed route regardless of the traffic conditions. At the start of the simulation the average travel time is unknown, as no vehicle has yet reached its destination.

As it is apparent, the centralised controller is capable of effectively exploiting the available road links of the network for reducing travel times. The analysis of the results indicate also that the use of the framework allows vehicles to reach their destination earlier, and can significantly reduce waiting times, i.e. time wasted in queuing. In terms of affected origin-destination routes, we observed that reductions in travel time tend to be evenly distributed; all vehicles are therefore benefitting from the improvement.

Remarkably, despite the simplicity of the metric used for assessing the state of the network and predicting the future evolution, the network-wise improvement can be significant. Figure 3 indicates that up to 400 seconds (more than 5 minutes) can be saved on average per journey. Results are expected to generalise well in networks that offer alternatives for most of the usually congested routes. If alternative routes are not available, it is expected that the impact of the proposed framework is reduced by some degree.

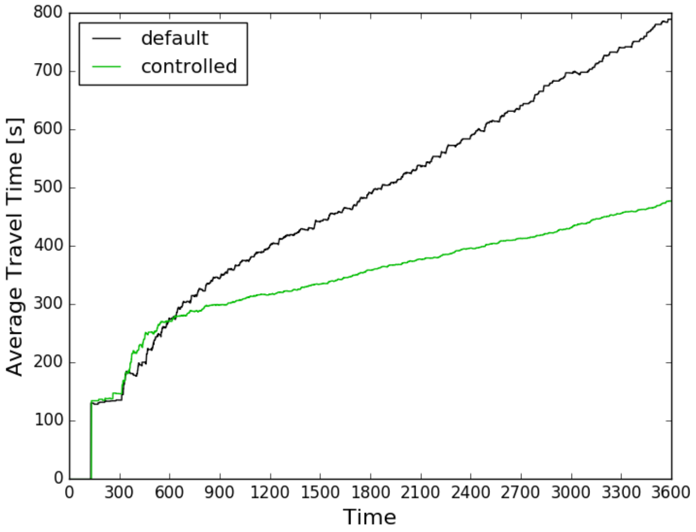

Figure 3. The average travel time of vehicles navigating through the Milton Keynes centre area with (green) and without (black) the use of the proposed framework.

To assess the importance of the penetration rate, we also considered experiments where an increasing percentage of vehicles are not communicating with the centralised controller, and their routes cannot therefore be modified. We considered four penetration rate values: 10%, 30%, 50%, and 100%. Results, in terms of average travel runtime, are presented in Figure 4.

Even in the case of very limited penetration rates, i.e. when only few vehicles are following re-routing instructions, there can be significant improvement. With a 10% penetration rate, the average travel time is reduced by approximately 150 seconds (more than 2 minutes). Better results are of course achieved when a larger percentage of vehicles are communicating with the controller and following instructions. In particular, there is a remarkable performance difference between 30% and 50% penetration rates. This may be due to specific conditions of the traffic, and to the structure of the considered network. In different scenarios gaps may be different.

More interestingly, the difference between 50% and 100% is not extremely significant. In fact, the difference lies in few tens of seconds. The gap between the two penetration rates is significantly reduced by the end of the simulation time, because most of the vehicles have reached their destination and the vast majority of the network is free from congestion.

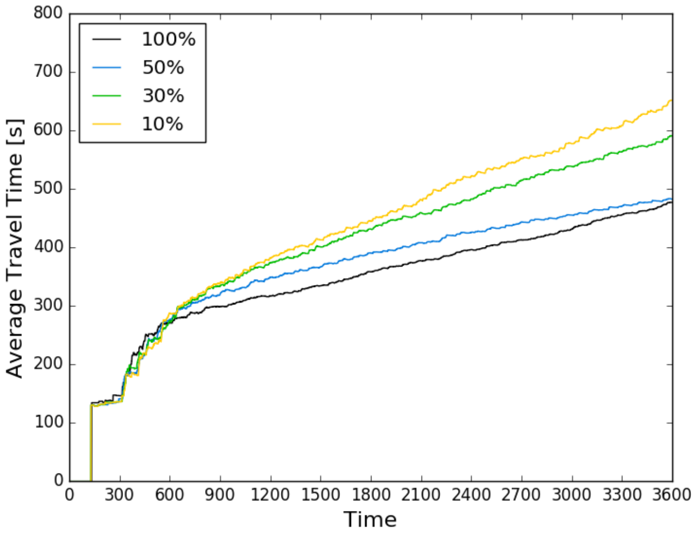

Figure 4. The impact of penetration rate on the average travel time of vehicles navigating through the Milton Keynes centre area, when the proposed framework is in use.

The presented results suggest that, at least for the considered network and traffic flows, significant improvements can be achieved in the near future, as soon as communication between centralised urban traffic controllers and vehicles can be established.

Discussion

In its current form, the proposed framework minimises the communication between vehicles and the centralised controller. However, it is envisaged that vehicles can also share preferences or constraints with the controller. For instance, a vehicle may need to avoid highly polluted areas due to the presence on board of a person with lung conditions. The centralised controller can then take preferences into account when considering re-routing, in order to provide a better, and more personalised, re-routing service.

With regards to re-routing, a more comprehensive set of features can be taken into account by the controller. Information about pollution, for instance, can be taken into account, to allow the controller to reason in terms of pollution and to reduce traffic in areas where air quality is already low. Network conditions can also be updated by considering car accidents or unexpected events. More sophisticated approaches can also take into account traffic light settings, in order to accurately estimate journey times.

Finally, the proposed approach can coordinate with emergency response services. For instance, it can be used to minimise traffic on some road links, to ensure that emergency vehicles can quickly reach their destinations.

Conclusions and future work

In this paper, we presented a framework for performing centralised urban traffic control. The aim of the framework is to maximise the use of network resources by re-routing navigating connected vehicles. An experimental analysis, performed on real-world data collected for the Milton Keynes area, demonstrated the benefit of the framework. A distinctive feature of this approach is the ability to support network-wise reasoning, and to accommodate a wide range of techniques for assessing the current and future status of the network, as well as the techniques for re-routing vehicles.

As well as exploration of the most effective algorithms to use for these aspects, future work should consider the path to deployment of this framework. Roll-out will require coordination of both public authorities with responsibility for the road network and vehicle manufacturers. As such, collaborative, international R&D and standardisation projects are likely to be very beneficial to the process. While standardisation efforts are well underway for the basic vehicle-to-infrastructure communications layer, see for instance LTE-V2X [12] and ETSI ITS-G5 [13], standards will be needed for the specific routing messages that need to be passed to vehicles. The full benefits of our framework will only be realised if all vehicles communicate using the same protocol and with the same traffic management system, and common standards will make this proposition more likely.

In the near term, some of the benefits can likely be achieved before highly automated vehicles are fully deployed – both due to the effectiveness of the framework with low penetration rates of participating vehicles, and because routing information can be provided as advice to human drivers. Many drivers today rely on a connected navigation system to supply routes, providing a natural pathway to the introduction of a centralised system, albeit with the challenges of integrating with existing services and/or persuading the public to change the services they use. Further issues may arise from human drivers failing to follow the provided route, either deliberately or accidentally, which is likely to mean the framework will be more effective when used with highly automated vehicles.

The current high levels of interest and investment in CAVs has led to a recognition of the potential value of traffic management systems, for example “intelligent network management" is one of the top-level categories in the Zenzic roadmap [11]. The framework we have presented has significant promise to fulfil this need, given that even with relatively simple algorithms for predicting congestion and re-routing vehicles, it has produced improvements of up to 40% in average journey times.

Acknowledgment

The Artificial Intelligence for Maximising the Exploitation of Urban Networks via Personalised AVs Route Generation (AI4ME) project is supported by the UK Engineering and Physical Sciences Research Council under grant EP/R51343X/1.

References

1. Chang, K., Hsu, P.Y., Lin, C., Lin, C., Juo, S.H., Liang, C. (2019). Traffic-related air pollutants increase the risk for age-related macular degeneration. *Journal of Investigative Medicine*. Vol. 67, pp.1076-1081.
2. Vallati, M., Chrpa, L. (2018). A Principled Analysis of the Interrelation between Vehicular Communication and Reasoning Capabilities of Autonomous Vehicles, In Proceedings *21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, IEEE.
3. Arena, F., Pau, G. (2019). An Overview of Vehicular Communications. *Future Internet*. Vol 11.
4. Antoniou, G., Batsakis, S., Davies, J., Duke, A., McCluskey, T., Peytchev, E., Tachmazidis, I., Vallati, M. (2019). Enabling the Use of a Planning Agent for Urban Traffic Management via Enriched and Integrated Urban Data. *Transportation Research Part C: Emerging Technologies*. Vol. 98, pp.284-297.
5. Hu, H. C., Smith, S. F. (2019). *Using Bi-Directional Information Exchange to Improve Decentralized Schedule-Driven Traffic Control*, In Proceedings 29th International Conference on Automated Planning and Scheduling. AAAI.
6. Krajzewicz, D., Erdmann, J., Behrisch, M., Bieker, L. (2012). Recent Development and Applications of SUMO - Simulation of Urban MObility. *International Journal On Advances in Systems and Measurements*. Vol. 5, pp.128-138.
7. Kabraham, I., Delling, D., Goldberg, A.V., Werneck, R.F. (2010) Alternative Routes in Road Networks. In: Festa P. (eds) Experimental Algorithms. SEA 2010. Lecture Notes in Computer Science, vol 6049. Springer.
8. Chen, M., Yu, X., Liu, Y. (2018). PCNN: Deep convolutional networks for short-term traffic congestion prediction. IEEE Transactions on Intelligent Transportation Systems. Vol. 19, pp.3550-3559.
9. Fouladgar, M., Parchami, M., Elmasri, R., Ghaderi, A. (2017). Scalable deep traffic flow neural networks for urban traffic congestion prediction. In Proceedings 2017 International Joint Conference on Neural Networks (IJCNN) (pp. 2251-2258). IEEE.
10. Chrpa, L., Vallati, M., Parkinson, S. (2019). *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 (pp. 191-194). ACM.
11. Zenzic (2019). *UK Connected and Automated Mobility Roadmap to 2030*. https://zenzic.io/roadmap/

12. 3GPP TR 36.885 V0.4.0 (2016). *Study on LTE-based V2X Services*. http://www.3gpp.org/dynareport/36885.htm
13. Intelligent Transport Systems (2013). *Access Layer Specification for Intelligent Transport Systems Operating in the 5 Ghz Frequency Band*, ETSI ITS 102 687 V1.1.1.