

Exploiting Automated Planning for Efficient Centralized Vehicle Routing and Mitigating Congestion in Urban Road Networks

Lukáš Chrpa
Charles University &
Czech Technical University in Prague
Prague, Czech Republic
chrpaluk@fel.cvut.cz

Mauro Vallati
Simon Parkinson
University of Huddersfield
Huddersfield, United Kingdom
{m.vallati,s.parkinson}@hud.ac.uk

ABSTRACT

One of the pivotal challenges presented to urban traffic controllers is the effective utilization of the transport infrastructure, as a result of the growing urbanization and of the increasing number of vehicles.

In response to this challenge, here we investigate the use of Automated Planning for reducing congestion and maximizing the exploitation of urban networks, by generating routes for the vehicles navigating through the controlled region. Empirical analysis is performed on four case studies (including maps of actual UK urban areas) with different complexities and volume of vehicles. The results demonstrate the potential of the presented approach, leading towards an intelligent navigation system, which shows to be able to provide effective personalized directions to vehicles.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; • **Applied computing** → **Transportation**;

KEYWORDS

Automated Planning, Urban Traffic Control, Connected Vehicles

ACM Reference Format:

Lukáš Chrpa, Mauro Vallati, and Simon Parkinson. 2019. Exploiting Automated Planning for Efficient Centralized Vehicle Routing and Mitigating Congestion in Urban Road Networks. In *The 34th ACM/SIGAPP Symposium on Applied Computing (SAC '19), April 8–12, 2019, Limassol, Cyprus*. ACM, New York, NY, USA, Article 4, 4 pages. <https://doi.org/10.1145/3297280.3297539>

1 INTRODUCTION

Urban Traffic Control (UTC) aims to mitigate issues in urban traffic by controlling mechanisms available in urban areas. Intelligent control systems in UTC nowadays focus on scheduling traffic lights in order to minimize delays in (busy) intersections. Such techniques are, for example, SCOOT [1], MOVA [12] and SCATS [7]. These traffic light control systems are able to adapt their own green phase lengths depending on current traffic conditions. Triggers which can rapidly (in a couple of minutes) adjust the behaviour of a small set of traffic lights are embedded in such systems. Such adjustment

to *fixed plans* can be made in rush hours in order to cope with increased traffic on main roads. Two recent lines of work focus on controlling traffic lights using AI planning techniques. The first one [8, 10] exploited continuous planning techniques and developed a mixed discrete-continuous model for controlling traffic lights in urban areas that aims to decrease queue lengths on traffic lights. It has been demonstrated that this approach is efficient especially in cases when unanticipated events occur. The second work [5] involves combining an AI planning engine and a traffic simulator. The simulator is used for monitoring the traffic conditions, and to decide whether some actions are needed. Traffic concentrations on road links are represented by relative density descriptors, and possible traffic light changes are enumerated to cover all the ways that a given configuration can affect the road links of the region.

Looking towards the future, UTC's function is to make better use of the physical infrastructure of a city and utilise Vehicle-to-Infrastructure communication. Nowadays, it is often the case that during rush hours, or in case of large events, most main roads are congested, at least in one direction, while many other roads are underused. This is caused by the fact that traffic is navigated via the same route between given way-points, and is often resulting 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. AI planning has been exploited as a centralized tool for navigating vehicles in the road network, although with different aims – Jimoh et al. [6] focused on exceptional circumstances (e.g. traffic accidents) while Chrpa et al. [2] focused on enforcing air quality constraints.

In this paper, we address the problem of navigating vehicles in urban areas with the clear aim of maximizing the utilization of the network and minimizing overall congestion. Addressing such a problem requires a global view of the situation and being able to find a solution in an order of seconds. Domain-independent Automated Planning provides a required machinery for addressing such complex conditions: PDDL, which is a language for describing planning problems [4], and off-the-shelf planning engines accommodating advanced planning techniques (e.g. Mercury [3]). Advantages of the exploitation of Automated Planning in a centralized architecture include having a global view of the situation, i.e., positions and intentions (goals) of vehicles in the network, and can thus take better (globally motivated) informed decisions than individual vehicles (or their drivers). In contrast to related works [2, 6] that use temporal PDDL models to route vehicles (while optimizing for different objectives), we specify a planning domain model in the classical subset of PDDL (typed STRIPS). Exploiting temporal planning has shown

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SAC '19, April 8–12, 2019, Limassol, Cyprus

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5933-7/19/04.

<https://doi.org/10.1145/3297280.3297539>

to be computationally expensive (e.g. about 20 vehicles can be handled within the required time bounds), while classical planning, as we show in this paper, can provide routes for a large number of vehicles in very short time windows. The proposed approach can be easily integrated into existing UTC frameworks, since it can receive information from autonomous vehicles, or vehicles' SATNAVs, and then can affect their routes according to the current condition of the road network. The experimental analysis, which considers four different road networks – out of which two represent the road network of actual UK urban area centres – with increasing traffic loads, demonstrates the usefulness of the approach. In the presence of significant amount of vehicles entering the network, the centralized system demonstrates to be able to distribute vehicles in order to exploit the road network and reduce congestion.

2 DOMAIN MODEL SPECIFICATION

In a nutshell, the problem we are addressing in this paper is to effectively and efficiently distribute traffic in the road network in a controlled (urban) region. Initially, vehicles communicate their entry and exit points to and from the controlled region to the system. Then, the vehicles obtain the route they should take to pass through the controlled region. Specifically, the emphasis is given to mitigate congestion on the roads within the controlled region. That, in consequence, might result in situations where two individual vehicles with the same entry and exit points will receive different recommended routes (which is contrary to distributed approaches).

A controlled region of the *road network* can be represented by a directed graph, where edges stand for *roads* (two-way roads are represented as two one-way roads) and vertices stand for *junctions*, some of which are *entry* or *exit* points. Intuitively, vehicles enter the network in entry points, and exit the network in exit points. Each road has its *length* and *bandwidth*. Road length determines how long (under normal circumstances) it takes for a vehicle to pass the road through. Road bandwidth represents how many vehicles can pass through road in a given time unit. It can be determined from sensory data considering current road traffic, weather conditions, current limitations (e.g. a lane closure) to mention some. Road bandwidth is further divided into three categories – *light*, *medium* and *heavy* – referring to traffic intensity on the given road. Since the destination (exit point) of each vehicle is known a priori – exit point has been determined by the satellite navigation system –, such an information can be exploited to put restrictions on roads that vehicles going to a certain destination can take. This helps to filter out suboptimal alternatives before the actual planning and thus considerably reduce the search space for the planner.

The routing of vehicles in the road network is handled by a *drive-to* type of action that navigates a vehicle v going to an exit point x from a junction j_1 to a junction j_2 via a road r whose use increases from u_1 to u_2 . It is the case that j_1 and j_2 are connected by r , r can be used by vehicles going to x and $u_2 = u_1 + 1$. Since we want to minimize the level of congestion of the roads, we “split” the *drive-to* action into four variants: *drive-to-light*, *drive-to-medium*, *drive-to-heavy* and *drive-to-congested*. Each variant is applicable depending on the level of congestion of the road.

The cost of the *drive-to* (for the road r) is $length(r)$, $10 * length(r)$, or $100 * length(r)$ for *drive-to-light*, *drive-to-medium* and *drive-to-heavy*

```
(drive-to-light car4 js ringsw1 jw jw use0 use1)
(drive-to-light car2 jw ringw2 jn jn use0 use1)
(drive-to-light car1 js ringse1 je je use0 use1)
(drive-to-light car0 jw centw1 jc je use0 use1)
(drive-to-light car0 jc cente2 je je use0 use1)
(drive-to-light car3 jw centw1 jc je use1 use2)
(drive-to-light car3 jc cente2 je je use1 use2)
```

Figure 1: A plan showing the routing of 5 vehicles in map 1 (Fig 3 top). To exemplify, the first action is used to route vehicle car4 from junction js to junction jw, via ringsw1 road: this increases the use of the road from 0 to 1. The final destination of car4 is the exit point junction jw.

```
(drive-to-L car0 jw centw1 jc je use0 use1)
(drive-to-L car0 jc cente2 je je use0 use1)
```

car0: jw → centw1 → jc → cente2 → je

Figure 2: Acquiring the route for the vehicle named as car0, from the plan depicted in Figure 1

respectively. The *drive-to-congested* action costs 100,000 as it must be much more expensive than the other *drive-to* actions.

A problem instance is specified by an initial state, i.e., the road network topology (as described above) and initial positions of vehicles, and by the goal, i.e., final positions of the vehicles. Solution plans are optimized for minimizing (total action cost), so we minimize the congestion level in the first case and then we minimize the average travel length or duration.

It should be noted that the road use can only be incremented in the presented model. This is because we deal with flows of vehicles: although we model individual vehicles in a planning problem, they represent vehicles entering the controlled area in a given –very short– time span (e.g. 30 seconds). Hence, vehicles cannot wait anywhere in the area for some roads to be de-congested, and the increasing road use values represent the fact that incoming vehicles are queued with previous traffic (and even with intra-regional traffic). The *standard* domain model, without the *drive-to-congested* action, does not allow vehicles to be routed to congested roads, so the bandwidth of any road in the area cannot be exceeded. The *congested* domain model relaxes this constraint, so vehicles can be routed to congested roads – this might be necessary for finding at least some solution for situations in which congestion of some roads cannot be avoided. Given the optimization criterion for the models, we at first minimize the number of *drive-to-congested* actions, then *drive-to-heavy* actions and then *drive-to-medium* actions. Such a criterion enforces the planning engine to balance traffic in the controlled area with respect to road bandwidth and, hence, minimizes congestion of the roads in the network.

The domain model and a problem instance, encoded in PDDL, can then be used as an input to a planning engine, in order to generate a plan for distributing the traffic across the network. In fact, the plan determines the behaviour of individual vehicles, i.e., which roads they take to navigate throughout the road network from their given entry points to their desired exit points. From the plans, we can find out which modelled vehicle is taking which route (for illustration see Figure 1), so they can be straightforwardly distributed to the particular vehicles as shown in Figure 2.

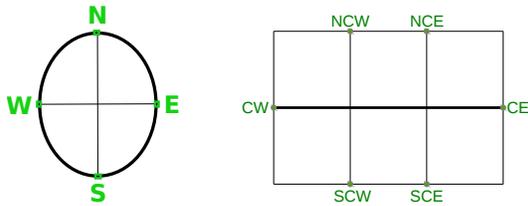


Figure 3: The generic European (left) and American (right) city maps. Entry/exit points are depicted as green dots. Thick lines represents roads with higher bandwidth.

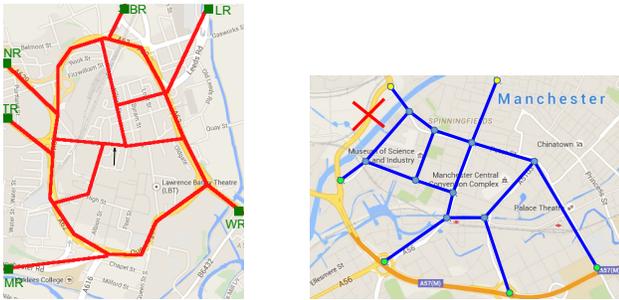


Figure 4: The Huddersfield town centre (top), and the considered roads of Manchester city centre (bottom), taken as examples of real urban areas.

3 EXPERIMENTAL EVALUATION

The aim of our experimental analysis is to demonstrate that by using the proposed centralized AI Planning technique, the traffic is effectively distributed across the whole road network and thus the chances for road congestion, especially in rush hours, are reduced.

We synthesized two generic maps (see Figure 3) representing abstracted versions of the typical European city centre (denoted as Map1) and the typical American city centre (Map2). In order to gain some insights into performance on more realistic scenarios, we also consider two real-world maps, Huddersfield (UK) town centre (denoted as Map3) and the Manchester (UK) city centre scenario introduced by Vallati et al. [10], where parts of the main ring road were closed because of floods (Map4). Maps are shown in Figure 4. Note that incoming/outgoing roads are not explicitly modelled but highlighted for illustration how traffic enters/exits the controlled part of the network.

As it is apparent, the four maps pose different challenges for the effective distribution of traffic: for instance, in the typical European city and in the Huddersfield town centre map, the ring road has three times higher bandwidth than the roads passing through the centre. In contrast, Map2 contain a backbone road (CW to CE and vice-versa) that has three times higher bandwidth than the other roads. Finally, the Manchester example considers only the city centre, as the ring road is not usable due to floods. Specifically, the *heavy*, *medium* and *light* thresholds are set to 10, 6, and 3 vehicles over 30 seconds, respectively, for minor roads (e.g. town centre roads), and 30, 20, and 10 respectively for major roads (e.g. ring roads) in Maps 1–3. In Map4, the thresholds are 20, 13, 7, respectively, for Spinningfields road (north-west); 15, 10, 5, respectively,

for the road around Deansgate station (centre of the map); and 30, 20, 10, respectively, for the rest of the roads. The bandwidth thresholds are provided in consultation with UK traffic engineers, but different values can be exploited straightforwardly. The road lengths that represent time needed by vehicles for moving through the road (from its tail to its head) is specified proportionally according to the above mentioned maps.

For each map, we generated problems consisting of 10, 30, ..., 150 vehicles entering the map over a 30-seconds time window. With increasing the number of considered vehicles we simulate rush hour conditions, where effective distribution of traffic is essential for keeping the roads congestion free. For the sake of comparison, in critical urban regions of the Manchester (UK) metropolitan area, there is a flow of traffic among main directions of approximately 2,000 vehicle over a 1-hour period of morning rush hour, corresponding to an average of 16.7 vehicle every 30 seconds per an entry point. Therefore, the higher numbers of vehicles considered in our experimental analysis correspond to exceptional conditions, where traffic is ten times a typical rush hour.

Entry and exit points of vehicles were distributed randomly (entry and exit point must be different for a particular vehicle), in order to simulate different traffic flows. In Map1 and Map3 every entry or exit point has the same chance to be selected, in Map2 CE and CW have three as much change for being selected as the other entry/exit points. In Map4, traffic enters the network at south and exits it at north. The distribution of vehicles is even, i.e., every entry point (at south) has the same chance to be selected as well as every exit point (at north) has the same chance to be selected.

As a baseline comparison we consider situations in which drivers always take the shortest (and quickest) route, as it simulates the exploitation of standard SATNAV systems with no updated information about traffic conditions. Two state-of-the-art planning engines have been considered, Yahsp [11] and Mercury [3], according to their good performance in a small subset of considered problems, and to their good performance in the recent International Planning Competition [9]. Both planning engines were given 10 seconds time limit, in order to simulate scenarios where plans are needed relatively quickly to assign routes to vehicles before they actually enter the controlled road network. Yahsp was run in the agile setting, i.e., after a plan was found the planning engine terminated. Mercury, on the other hand, was run in the satisficing settings, i.e., once a plan has been found, the planning engine keeps running in order to find a better quality solution. Both the standard and congested models have been used. The experiments were performed on an i7 2.5 GHz CPU, 8GB RAM, Ubuntu 16.04.

Table 1 provides a comparison between average travel times of the planning approaches and the baseline while considering delays for congested road segments, in particular 10% delay for the Medium level, 50% for the Heavy level and 200% for the Congested level, respectively. These values are optimistic delay estimations derived after discussing with expert traffic engineers. In particular, value 0.90 for the travel time change means that the average travel time is by 10% shorter than in the baseline case (for the corresponding problem and planning approach). Value 0.47 for traffic flow decrease means that the traffic flow on the most congested road segment with respect to its bandwidth (for the corresponding problem and planning approach) is 47% of the traffic flow on

Table 1: Travel (T)ime change with respect to the baseline (the shortest route) and maximum relative Traffic (F)low decrease with respect to the baseline. The suffix “-c” stands for the model supporting the drive-congested actions. “-” stands for no plan has been generated in the given time limit (10 seconds).

Vehicles	10		30		50		70		90		110		130		150	
	T	F	T	F	T	F	T	F	T	F	T	F	T	F	T	F
map1																
yahsp	1.00	1.00	1.04	0.75	1.07	0.61	1.00	0.57	0.86	0.54	0.88	0.46	0.87	0.47	0.95	0.47
yahsp-c	1.00	1.00	1.04	0.75	1.07	0.61	1.00	0.57	0.86	0.54	0.88	0.46	0.87	0.47	0.95	0.47
mercury	1.00	1.00	1.00	0.75	1.02	0.66	0.96	1.00	0.82	0.90	0.86	0.66	0.87	0.62	0.88	0.58
mercury-c	1.00	1.00	1.00	0.75	1.02	0.66	0.96	1.00	0.85	0.90	0.99	0.93	0.99	0.93	1.00	1.00
map2																
yahsp	1.00	1.00	0.99	1.00	0.96	1.00	0.90	0.90	0.80	0.77	0.71	0.66	-	-	-	-
yahsp-c	1.00	1.00	0.99	1.00	0.96	1.00	0.90	0.90	0.80	0.77	0.74	0.73	0.78	0.72	0.86	0.57
mercury	1.00	1.00	0.97	0.60	0.98	1.00	0.94	1.00	0.81	0.83	0.73	0.66	-	-	-	-
mercury-c	1.00	1.00	0.97	0.60	0.98	1.00	0.94	1.00	0.97	1.00	1.02	1.00	0.90	1.00	0.86	0.80
map3																
yahsp	1.32	0.33	1.38	0.29	1.12	0.25	0.99	0.25	0.93	0.33	0.97	0.24	0.93	0.25	0.92	0.22
yahsp-c	1.32	0.33	1.38	0.29	1.12	0.25	0.99	0.25	0.93	0.33	0.97	0.24	1.09	0.26	1.37	0.27
mercury	1.04	0.75	1.38	0.29	1.15	0.25	0.99	0.31	0.88	0.31	1.01	0.23	0.97	0.25	-	-
mercury-c	1.04	0.75	1.38	0.29	1.15	0.25	1.00	0.29	0.88	0.31	1.01	0.23	1.12	0.25	1.30	0.26
map4																
yahsp	1.11	0.80	1.11	0.83	1.08	0.88	-	-	-	-	-	-	-	-	-	-
yahsp-c	1.11	0.80	1.11	0.83	1.08	0.88	1.20	0.87	1.01	0.94	0.83	0.85	0.92	0.86	0.90	0.83
mercury	1.06	1.00	1.08	0.72	1.09	0.88	0.92	0.76	-	-	-	-	-	-	-	-
mercury-c	1.06	1.00	1.08	0.72	1.09	0.88	1.17	0.84	1.04	0.94	0.85	0.85	1.02	0.86	0.98	0.83

the most congested road segment with respect to its bandwidth for the baseline approach. In the Manhattan-like road network topology, the average travel time decreases with increasing number of vehicles along with congestion decrease. In the RingRoad-type scenarios, the average travel time is affected more by a longer distance the vehicles have to take (in Map 3, using the Ring Road instead of going through the town centre can increase the travel distance by up to 40%). If, however, the congestion is considerably reduced, then despite taking a (much) longer route, the average travel time decreases. The most apparent example is for 90 vehicles for both RingRoad maps (Map1 and Map3). The baseline approach has reasonable results for smaller number of vehicles (up to 50) as the highest congestion level is usually at most Medium. For larger number of vehicles (from 90), the centralized planning approach better utilizes the road network and thus reduces the travel time. Noteworthy, longer term congestion of a few road segments, which our approach aims to mitigate, can propagate across the whole network resulting in larger delays than we (optimistically) assume.

4 CONCLUSIONS

In this paper, we focused on the problem of smartly calculating routes for vehicles navigating through urban areas, according to the current and expected traffic conditions, in order to maximise the utilization of the network and reduce congestion of the roads. Specifically, we considered a scenario in which vehicles are equipped with a satellite navigation system that can communicate with a centralized controller, which is in charge of providing such routes.

We formulated the problem as a planning problem, using basic features of the PDDL language, which is supported by a wide range of planning engines. The proposed approach has then been extensively tested on four different urban area networks, with different characteristics and structures, considering an increasing number of vehicles accessing the controlled road network at the same time. The results of the performed analysis demonstrated on a range of scenarios that our centralized approach for navigating vehicles

thought the road network has the capability of their effective distribution, minimizing the risk of having some roads congested and hence reducing vehicles’ average travel time.

Acknowledgements. L. Chrpa was partially funded by the Czech Science Foundation (project no. 17-17125Y). M. Vallati was partially supported by the EPSRC grant EP/R51343X/1 (AI4ME).

REFERENCES

- [1] RD Bretherton. 1989. SCOOT urban traffic control system: Philosophy and evaluation. In *Proceedings of the 6th IFAC/IFIP/IFORS Symposium on Control, Computers, and Communications in Transportation*. 237–239.
- [2] Lukáš Chrpa, Daniele Magazzeni, Keith McCabe, Thomas Leo McCluskey, and Mauro Vallati. 2016. Automated planning for Urban traffic control: Strategic vehicle routing to respect air quality limitations. *Intelligenza Artificiale* 10, 2 (2016), 113–128.
- [3] Carmel Domshlak, Jörg Hoffmann, and Michael Katz. 2015. Red-black planning: A new systematic approach to partial delete relaxation. *Artificial Intelligence* 221 (2015), 73–114.
- [4] Maria Fox and Derek Long. 2003. PDDL2. 1: An Extension to PDDL for Expressing Temporal Planning Domains. *J. Artif. Intell. Res. (JAIR)* 20 (2003), 61–124.
- [5] Matija Gulić, Ricardo Olivares, and Daniel Borrajo. 2016. Using automated planning for traffic signals control. *PROMET-Traffic&Transportation* 28, 4 (2016), 383–391.
- [6] Falilat Jimoh, Lukas Chrpa, T.L. McCluskey, and Mohammad Munshi Shahin Shah. 2013. Towards Application of Automated Planning in Urban Traffic Control. In *2013 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. Institute of Electrical and Electronics Engineers (IEEE), 985–990.
- [7] PR Lowrie. 1982. The Sydney coordinated adaptive traffic system-principles, methodology, algorithms. In *Proceedings of the International Conference on Road Traffic Signalling*.
- [8] Thomas Leo McCluskey and Mauro Vallati. 2017. Embedding Automated Planning within Urban Traffic Management Operations. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*. 391–399.
- [9] Mauro Vallati, Lukáš Chrpa, Marek Grzes, Thomas L. McCluskey, Mark Roberts, and Scott Sanner. 2015. The 2014 International Planning Competition: Progress and Trends. *AI Magazine* 36, 3 (2015), 90–98.
- [10] Mauro Vallati, Daniele Magazzeni, Bart De Schutter, Lukáš Chrpa, and Thomas L. McCluskey. 2016. Efficient Macroscopic Urban Traffic Models for Reducing Congestion: a PDDL+ Planning Approach. In *The Thirtieth AAAI Conference on Artificial Intelligence (AAAI)*.
- [11] Vincent Vidal. 2014. YAHSP3 and YAHSP3-MT in the 8th International Planning Competition. In *Proc. of the 8th International Planning Competition (IPC-2014)*.
- [12] R.A. Vincent and J.R. Pierce. 1988. Self-optimising Signal Control for Isolated Intersections. In *Crowthorne: Transport and Road Research Laboratory Research Report*.