Planning & Scheduling Applications in Urban Traffic Management

Isabel Cenamor
Departamento de Informática
Universidad Carlos III de Madrid

Lukáš Chrpa and **Falilat Jimoh** and **Thomas L. McCluskey** and **Mauro Vallati**
PARK Research group
University of Huddersfield

Abstract

Local authorities that manage traffic-related issues in urban areas have to optimise the use of available resources, in order to minimise congestion and delays. In this context, Automated Planning and Scheduling can be fruitfully exploited, in order to provide dynamic plans that help managing the urban road network.

In this paper we provide a review of existing planning and scheduling approaches that have been designed for dealing with different aspects of traffic management, with the aim of gaining insights on the limits of current applications, and highlighting the open challenges.

Introduction

The global human population has dramatically increased in the last century. It is expected to reach around 10 billion people by the 2100. While human population is increasing, its distribution is changing: most of humans are now living in large urban areas, with significant impact on traffic and congestion. One big challenge we have to face is the area of the efficient movement of people and goods, especially in urban areas. One part of the solution is in the optimisation of the road resource, in order to minimise congestion and congestion-related issues, such as pollution, delays and increase of transport costs. Road traffic management is the discipline that deals with the control of traffic.

In road traffic management there is a need to promote more efficient use and optimisation of current resources (chiefly roads, but also other resources such as car parks), to help minimise the journey time of traffic, lessen road congestion, and limit the scale and impact of vehicle emissions. In relation to motorways and trunk roads, in the UK the Highways Authority has the responsibility of maintaining a safe and efficient network, whereas in urban areas, the local authority has the management role. In this context, Automated Planning can be fruitfully exploited, in particular to provide plans dynamically that help managing the urban road network. Bast et al. proposed a survey of different techniques algorithms in transportation networks, but there is not intrinsically related with the planning-based approaches.

Several different aspects of the traffic management in urban areas have been modelled and addressed through planning and scheduling approaches. In the 2014 edition of the International Planning Competition (IPC) [1], two models that deal with traffic-related aspects have been introduced. RTAM provides plans for the road traffic accident management, i.e., organises the movements of emergency vehicles for minimising the impact of accidents. CityCar is a study in the direction of evaluating the usefulness of topology of a traffic network. Current self-tuning traffic light cluster is urban areas work well in normal or expected conditions, but fail when a major incident occurs sch as a road closure. Jimoh et al. (2013) proposed a PDDL model (McDermott, Drew et al. 1998) that captures an urban micro-simulation (Treiber and Kesting 2013) of traffic, i.e. vehicles –rather than traffic flows– are modelled, in order to plan routes for vehicles. In particular, they addressed the problem of planning for unexpected situations, like roads become blocked due to some unanticipated incident, as a way of increasing the performance of current systems (rather than replacing them). SURTRAC (Xie, Smith, and Barlow 2012), is a scheduling-based approach for optimising the flow of vehicles in urban areas, by controlling traffic signals. Notably, it has been deployed and tested in a real city showing that it can optimise this behaviour of 9 connected traffic light clusters. Finally, we review two works that deal with multi-modal journey planning by considering different perspectives (Botea, Nikolova, and Berlingerio 2013; Caff, Mauro, and Scala 2014). Although multi-modal journey planning does not seem to be closely related with the other traffic management applications, it should be noted that providing effective journey planners can foster the exploitation of public transport, thus reducing the number of cars on roads.

This paper provides a survey of these planning-based approaches that have been proposed for addressing road traffic control, and related aspects, in particular in urban areas. By analysing the state of the art of automated planning in the complex area of road traffic management, we can gain insights on the limits of current applications, and better understand the open challenges. AI techniques in general have been used in many applications in road transportation, for instance in self - tuning traffic lights, automatic number plate recognition, and route navigation, but AI P & S techniques are still to be exploited in this area.

[1] http://helios.hud.ac.uk/scommv/IPC-14/

Urban Traffic Management

In Urban Traffic Management (UTM), clusters of junctions are managed well by current technology such as SCOOT (Hunt et al. 1982) and MOVA (P and F 1993), when conditions are normal. Additionally, events such as large stadium usage or roadworks can be planned for by having fixed-term plans controlling traffic lights in and around the affected areas. On the other hand, unexpected events can cause problems for which fixed plans would not be available by preplanning.

The proposed planning approach (Jimoh et al. 2013) addresses the problem of effective navigation of cars through a given Road Network from entry points to exit points during unforeseen situations such as road incident, road re-construction, car breakdown and when traffic demand changes rapidly within a short time interval. The introduced model, is divided into two main parts – static and dynamic. The static part represents road network topology, i.e., roads, their capacity, length and junctions connecting the roads. The dynamic part stands for how many cars are on each road (and where) and whether the road is operational. The term 'operational' means that the road is available and accessible within the road network system. Clearly, the dynamic part is changing through the time as cars are moving through the road network. The model architecture is divided into four parts:

- The properties of the road represented by road layout segment.
- Controlled intersections are represented by the road-junctions relationship within the road network.
- The road users properties represented by the road capacity and queue length.
- Route properties that span several adjoining road segments.

Jimoh et al. tested the feasibility of their approach in a real-world scenario, by considering the urban area depicted in Figure 1. Though complexity of a model at the level of the individual vehicle is likely to cause problems as the model increases, for models of the scale illustrated in the diagram, current planning technology is able to provide acceptable planned routes for vehicles in a reasonable time.

RTAM

Road Traffic Accident Management is based on the managing road accidents model previously presented in (Shah et al. 2013). The road network in this model is represented by a graph where vertices are locations, and edges are roads between different locations. The graph represents a local area or a region. In order to manage and clear traffic accidents, a number of different figures are represented: police, ambulances, fire brigades and two trucks. The stakeholders are distributed on the considered area, according to the position of corresponding buildings. For instance, ambulances are at hospitals, police cars at police stations, etc.

When an accident occurs, police has to go there in order to confirm and describe the situation. After that, ambulances are required for dealing with accident victims, and

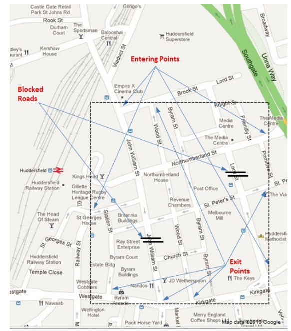

Figure 1: Map showing the network entry and exit points and the blocked roads in a part of town centre of Huddersfield, West Yorkshire, United Kingdom. It is used for our empirical analysis.

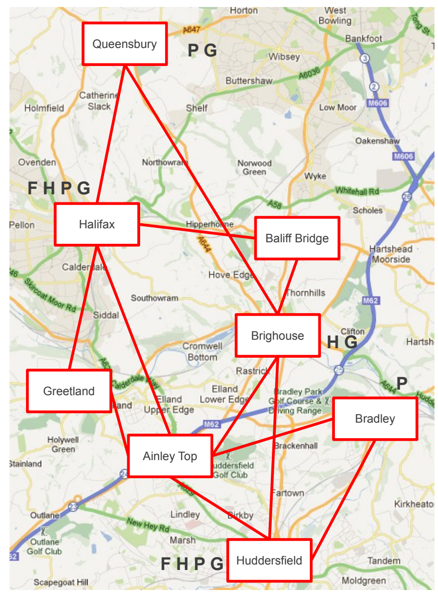

Figure 2: The Road Traffic Domain Model used for empirical analysis. It consists of a portion of the county of West Yorkshire. H, F, P and G respectively stand for Hospital, Fire station, Police station and Garage locations.

tow trucks should remove damaged vehicles for clearing the area. Moreover, in case of fire or trapped victims, fire brigades are involved. For an accident to be completely cleared, victims have to be taken to hospitals, and damaged vehicles removed.

This domain model has been used in the IPC-14, but Shah et al. tested it in a previous work (Shah et al. 2013). In their empirical analysis, Shah et al. demonstrated the effectiveness of the approach by modelling an inter-urban area of Yorkshire (see Figure 2), and by considering real data in terms of distribution of emergency vehicles, speed and probability of accidents.

CityCar

This CityCar PDDL domain model aims at evaluating the usefulness of roads in traffic networks. The city is represented by a graph where vertices are junctions, and the edges are "potential" roads. All the junctions are connected in this graph, but the roads between junctions are not built yet. In problems of this domain, cars are entering the network from some random positions, and have to reach a specific goal destination, that is identified as a garage. There is a maximum number of roads that can be built, and that can be active on the network at the same time; it is possible to build and remove roads from the network. The overall objective is to minimise the total cost. Building and removing roads as a cost, as well as moving cars around.

The figure 3 shows a simple problem example. There is one car, that is initially at J_1, and should reach its final destination at J_5. This map has seven different junctions (J_1 – J_7), junctions are connected by edges. In the initial state, three different roads are in place: a diagonal road between J_1 and J_4, and two straight roads between J_2 and J_4 and between J_4 and J_6. Intuitively, this problem can be optimally solved by building the road that connects J_6 and J_5. On the other hand, if the total number of roads is limited at three, finding the optimal solution is not straightforward.

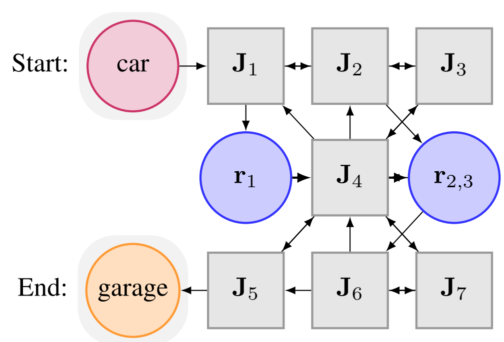

Figure 3: An example problem in the CiryCar domain. The car is initially at the pink circle (J_1) and the goal is the garage (orange circle) at the location (J_5). The built roads are the blue circles with the direction indicate for the arrow. Junctions are represented as squares.

This domain is encoded in PDDL , and includes negative preconditions, action cost and conditional effects. There is another version of the same domain with temporal representation (Fox and Long 2003). In this domain(Map-Analyser), cars have different speed and the time needed for building roads depend on their length.

SURTRAC

The SURTRAC system features a decentralised schedule-driven approach to provide network-wide optimisation of the dynamic flow of vehicles in a urban area (Xie, Smith, and Barlow 2012). Flows are controlled by managing traffic signals at junctions. At the network level, many junctions are tightly-coupled, since travel times between them can be short and the demands can be high. The main issue related to this problem is the scalability; while managing the single junction, or a small neighbourhood, can be easily done, it is hard to effectively control the whole network of an urban area.

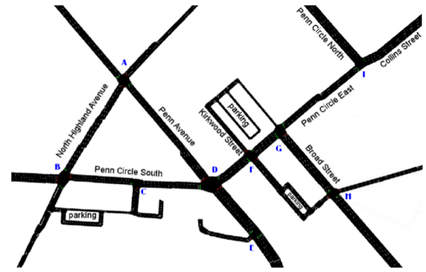

Figure 4: The real-world map used for testing the SURTRAC approach.

In the Xie, Smith, and Barlow model, each junction is controlled by a different agent, which exploits look-ahead scheduling strategy to operate with a limited predictions of incoming traffic. In the basic coordination mechanism, each agent uses the planned output flows of vehicles from its upstream neighbours to generate an optimistic observation, which includes look-ahead information from neighbours. In addition to optimistic non-local observation, two other coordination mechanisms help to prevent "nervousness" –this happens in downstream agents because of the uncertainty and disruption associated with the predictions of upstream agents that are using on-line control strategies with finite horizons–, and dynamic instability in the network (Kumar and Meyn 1995). Authors tested their approach in a number of scenarios, which include also a real-world network (see Figure 4). They compared SURTRAC with an isolated control strategy (where every agent ignores the rest of the network), a coordination strategy that uses a moving average forecast, and a coordination strategy that uses adaptive offset calculation via simulation on a 5X5 grid network designed to present a challenge for decentralised network-wide optimisation. The results obtained demonstrated the ability of SURTRAC to establish traffic flow coordination with lower average wait times than other methods. On the other hand, a few issues have been raised. One of them concerns the design of more effective coordination mechanisms: in their approach sometimes a downstream agent has to sacrifice its own interest for the sake of an upstream agent. Such one-sided sacrifice is pointless if the upstream agent changes its schedule quickly.

Multi-Modal Journey Planning

In it simplest form, when a single modality in considered, journey planning is a path planning problem: given a graph representation of a transportation network, find a route between an origin and a destination location. Clearly, in more realistic models of the real world, uncertainty must be taken into account. Multi-modal journey planning allows combining different transportation types, such as bus riding and cycling, in a single journey. In real-world multi-modal journey planning, it is of critical importance to consider uncertainty: for instance, small variations of bus arrival times can result in a missed connection, with a great negative impact on the actual arrival time at the destination. It is also worthy noting that multi-modal journey planning can optimise very different quality criteria, such as the total travel time, the number of interchanges, and the cost.

Botea, Nikolova, and Berlingerio encoded the multi-modal journey planning problem as a probabilistic planning problem, and developed a planner that performs heuristic search in an and/or tree, computing probabilistic plans (i.e., contingent plans where branches have probabilities). The planner is an incremental planner that optimises the robustness of provided plans. In their experimental analysis, authors used data built from real historical GPS traces of buses across Dublin, Ireland. A total of 3617 locations considered include 3559 bus stops, 44 bike stations, and 14 points of interest. The results [2] demonstrate that in many cases, plans of good quality are provided fast.

Similarly to the approach exploited in (Botea, Nikolova, and Berlingerio 2013), Caff, Mauro, and Scala (2014) addressed the multi-modal journey problem by exploiting planning techniques. In doing so, they relaxed the uncertainty of the problem, and presented a PDDL formulation representing temporally expressive actions encompassing the multiple schedule for an action homogeneously. Their work shows how the majority of the temporal conditions can be translated by exploiting numeric fluents and constraints. They tested the proposed approach with existing temporal planners, and evaluated different possible quality metrics.

Observations

While authorities have manually designed very good plans for expected situations, like football matches and other major events that strongly affect traffic flows, these fixed plan take several weeks to create and validate. AI Planning can be effectively exploited when unexpected events happen, and a new plan is required in close to real-time. In these cases, having a system that quickly provide a reasonable plan for deviating traffic, or for dealing with related issues, can provide a strong support to local authorities.

All the existing planning-based approaches work at the micro-simulation level, i.e. by planning the route of every single vehicle within the same urban area. This approach has a number of drawbacks: i) can limit the scalability; ii) exact position of each vehicle is required; iii) it assumes that authorities can communicate with each vehicle and that each vehicle respects indications. On the other hand, car navigators are now getting information by accessing the open data provided by authorities, thus a level of communication is already in place. It should also be noted that micro-simulations are usually not feasible for real-time planning, and for all the vehicles on the network; they can usually be applied on small number of vehicles or planned well-ahead. Nevertheless, planning approaches that manage traffic flows at higher level of abstractions are needed, in order to evaluate their effectiveness and compare them with current traffic plan adopted by local authorities.

In order to plan with traffic flows, and also to further improve existing PDDL models, the exploitation of different description languages is advisable. A more realistic modelling, with regards to current PDDL ones, can be achieved by using PDDL+ (Fox and Long 2006). Also, considering languages not propositionally-based, like ANML (Smith, Frank, and Cushing 2008) or NDDL (Frank and Jónsson 2003) can possibly provide better support for modelling traffic-related problems. Albeit such languages provide more expressivity for modeling such problems, their use will not allow to exploit the large number of high-performance domain-independent planners that supports PDDL.

On the side of performance of domain-independent planners, the organisers of the 2014 editions of the International Planning Competitions introduced the aforementioned domain models (RTAM, CityCar) in order to assess the performance of the state of the art. Results showed that high-performance planners are able to efficiently solve problems from such domains, but usually on instances that involve a rather small number of vehicles. It would be interesting to evaluate if additional knowledge can help state-of-the-art planners in significantly improve their applicability. Moreover, the area seems to be feasible for exploiting case-based planning approaches (Borrajo, Roubíkoví, and Serina 2015). For instance, a large number of problems, which consider different situations that have been observed, can be generated and solved, to be used as a library of ready-to-use plans, that can be quickly adapted for online use.

While most of the planning approaches deal with traffic management within the same urban area, it will be interesting to design more models that allow to manage vehicle flows at different levels, such as inter-urban. Currently, only in RTAM (Shah et al. 2013) a larger area is considered, although only behaviour of emergency vehicles is planned. A different area of investigation would be the highway traffic management, that deals with huge traffic flows on motorways, and needs to manage the use of accessing ramps.

Finally, multi-modal journey planning is becoming more and more popular (e.g. public transport option on Google maps). However, availability of information on public transport, such as reliability timetables, or impossibility to automatically process them, is a major issue that might hinder deploying multi-modal journey planning in some cities.

Conclusion

In this paper we presented a state of the art of Automated Planning and Scheduling approaches for managing traffic and related aspects in urban areas. We observed that existing techniques are covering different and various aspects of traffic management, and we believe that the planning community interest in the topic has significantly grown in recent years. We provided a discussion on possible future avenues, for further improvements in planning and scheduling approaches. This appears an application area that could benefit a great deal from P & S, though the potential of the community's technology within the Transport Studies area is not well studied.

[2] http://helios.hud.ac.uk/scommv/IPC-14/resDoc.html

References

Bast, H.; Delling, D.; Goldberg, A.; Müller-Hannemann, M.; Pajor, T.; Sanders, P.; Wagner, D.; and Werneck, R. 2014. Route planning in transportation networks. *Tech. Rep. MSR-TR-2014-4, Microsoft Reseach, Microsoft Corporation*.

Borrajo, D.; Roubíkova, A.; and Serina, I. 2015. Progress in case-based planning. *ACM Computing Surveys*.

Botea, A.; Nikolova, E.; and Berlingerio, M. 2013. Multi-modal journey planning in the presence of uncertainty. In *ICAPS*.

Caff, S. C. M.; Mauro, F. D.; and Scala, E. 2014. A numeric pddl based approach for temporally constrained journey problems. In *ICTAI*.

Fox, M., and Long, D. 2003. PDDL2.1: an extension to PDDL for expressing temporal planning domains. *J. Artif. Intell. Res. (JAIR)* 20:61–124.

Fox, M., and Long, D. 2006. Modelling mixed discrete-continuous domains for planning. *J. Artif. Intell. Res. (JAIR)* 27:235–297.

Frank, J., and Jónsson, A. 2003. Constraint-based attribute and interval planning. *Constraints* 8(4):339–364.

Hunt, P.; Robertson, D.; Bretherton, R.; and Royle, M. 1982. The scoot on-line traffic signal optimisation technique. *Traffic Engineering & Control* 23(4).

Jimoh, F.; Chrpa, L.; McCluskey, T. L.; and Shah, S. 2013. Towards application of automated planning in urban traffic control. In *Intelligent Transportation Systems-(ITSC), 2013 16th International IEEE Conference on*, 985–990. IEEE.

Kumar, P., and Meyn, S. P. 1995. Stability of queueing networks and scheduling policies. *Automatic Control, IEEE Transactions on* 40(2):251–260.

McDermott, Drew et al. 1998. PDDL–the planning domain definition language. Technical report, Available at: `www.cs.yale.edu/homes/dvm`.

P, K., and F, D. 1993. Mova and lhovra: Traffic signal control for isolated intersections. *Traffic engineering and control* 34(4):195–200.

Shah, M. M.; Chrpa, L.; Kitchin, D.; McCluskey, T. L.; and Vallati, M. 2013. Exploring knowledge engineering strategies in designing and modelling a road traffic accident management domain. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, 2373–2379. AAAI Press.

Smith, D. E.; Frank, J.; and Cushing, W. 2008. The anml language. *Proceedings of ICAPS-08*.

Treiber, M., and Kesting, A. 2013. Traffic flow dynamics. *Springer*.

Xie, X.-F.; Smith, S. F.; and Barlow, G. J. 2012. Schedule-driven coordination for real-time traffic network control. In *ICAPS*.