The SimplyfAI Project: Using AI Planning in Urban Traffic Management *or* If at First the Representation Does not Work, Try, Try and Try and Again

Thomas L. McCluskey and Mauro Vallati
School of Computing and Engineering
University of Huddersfield
UK

Abstract

This paper is an experience report on the results of a collaborative one year feasibility study called "SimplyfAI" funded by Innovate UK. This concerned sourcing and enriching urban traffic data, and using this data as inputs to a system to generate urban traffic strategies in order (primarily) to improve air quality. This paper reports on the development surrounding the AI planning component of that work: the engineering and configuration issues that were found in this application. It discusses a range of issues and lessons we learned through the experience of collaborating with end users and technology developers.

Introduction

Conventional road traffic signal management techniques (e.g. traffic-responsive systems such as SCOOT (Taale, Fransen, and Dibbits 1998) or fixed time light changes optimised using historical data) work reasonably well in normal or expected conditions. In exceptional or unexpected conditions, however, these established methods work less well. In these cases Transport Operators may struggle to find a strategy tailored to solve the unexpected situation. Creating such strategies is a manual task that may take several days or weeks, and is therefore infeasible to be done in real-time.

This paper describes an attempt to utilise AI planning technology in the regional management of urban road traffic flows. The long term aim of the work is to provide a tool for urban transport operators that can generate, in real time, strategies to deal with exceptional or emergency situations. For example, transport operators may want to reduce traffic concentrations in a targeted urban area to ameliorate effects of predicted road traffic pollution; or optimise the flow of saturated road links due to an emergency road closure; or produce a strategy to deal with a forthcoming complex situation (e.g. optimising the light timings to deal with the combination of a concert, a football match and some emergency roadworks).

The work reported was carried out during 2016, and formed one of the deliverables of an Innovate/NERC - funded 1 year feasibility study[1]. We choose one traffic region for testing the feasibility - a very busy urban area of Salford, Greater Manchester - where real time and historical data sources could be readily obtained. Though, in the end, we had to focus on a relatively small road network within this region of Salford, the project demonstrated the effectiveness of the auto-generation of goal-directed strategies. The planner used in the project was the well known domain independent planner UPMurphi (Della Penna et al. 2009), which inputs models in PDDL+ (Fox and Long 2006). To produce a working executable of UPMurphi, we had to perform several cycles of iterations over the engineering of the PDDL+ models.

The quality of the strategies output from the planner was evaluated firstly by hand inspecting the strategies to check that they were "sensible". In this case the strategies were inspected to check they embodied common sense plans in them. Secondly, their effect was compared against optimised strategies derived from historical data by simulating their execution using both the AIMSUN micro-modelling software[2], and the off-the-shelf SUMO (Krajzewicz et al. 2012) micro-modelling software. In each case members of the consortium compared the results of simulations using both automated planning generated strategies, and optimised strategies derived from historical data. In both these simulators, run by different members of the consortium, the planner-generated strategies produced sufficient savings to convince the consortium to aim to adopt AI planning within a product for use generally in busy urban areas. On the other hand, the study highlighted several challenges to be overcome before a fielded implementation could be achieved, in particular the ever present problem of scale-up.

Context

This section provides some background on urban transport, and some key advances in configuring planning applications to work in the urban traffic management domain.

Urban Traffic Management (UTM)

Over the years, as traveller safety was built into road traffic infrastructure (e.g. using intergreen delays in traffic signals), road traffic management in urban areas was left with

---

[1] The grant was funded from the call "Solving Urban Challenges with Data", with a consortium consisting of The University of Huddersfield, KAMfutures, InfoHub, British Telecom (BT) and Transport for Greater Manchester (TfGM)

[2] https://www.aimsun.com/

one other goal: to minimise delay to the traveller. Transport support infrastructure –particularly traffic signals– are nowadays optimised to meet this one goal, and in normal circumstances they tend to work reasonably well. There is a need, however, for traffic control support systems to help in the management of traffic by being able to achieve more focussed or detailed goals than simply minimising delay. For example, we may want to prioritize air quality goals explicitly: this is pivotal importance in congested urban areas. At the same time, modern traffic management has the opportunity to take advantage of the huge amount of sensor data available. This sensor data provides real time information on the state of traffic in a network. In particular, the data could alert a Transport Operator to the fact that a region is predicted to exceed an air quality limit in the near future. Generating a detailed strategy of interventions (such as changes to traffic signal timings over a period of time) to avoid this *in real time* is considered to be beyond the capacity of human operators.

Hence there is a need to be able to produce strategies in real time which deal with abnormal or unexpected events (such as capacity losses through road closures). These cause huge delays and decreased air quality because of excessive congestion and stationary traffic. Looking to the future, there is a need for technology which has the capability of taking into account a range of controls regionally: not just traditional traffic lights, but real-time support in forming strategies combining traffic signals and other controls e.g. variable speed limits, variable message signs and extra lane introduction.

AI Planning applied to UTM

While there are many examples of the application of general AI techniques to road traffic monitoring and management (Various 2007; Miles 2006), the application of AI Planning to help in the management of road traffic is fairly novel (McCluskey and Vallati 2014; Cenamor et al. 2014). The most mature work coming out of the ICAPS community appears to be SURTRAC, a distributed scheduling system which controls traffic signals in urban areas (Xie, Smith, and Barlow 2012). Run by a schedule-driven intersection control algorithm, the system is intended for use in grid-based town centres. It is currently being trialled in Pittsburgh, USA, with its distributed approach suggestion better scale-up but less flexibility than a centralised AI planner.

One of the technologies seen as important in the EU's COST Action 1102 "Autonomic Road Transport Support Systems" (ARTS) (McCluskey et al. 2016) was AI Planning, and through work in that Network a particular role for planning emerged: to help manage operations during exceptional events (Jimoh et al. 2013). This work, and subsequent work (Chrpa et al. 2015), took a microsimulation approach to planning, meaning that vehicles were specified individually in the domain and problem model, and durative actions were used to describe the movement of traffic (at a microscopic level) between junctions. While this line of work looked promising, the problem of scaling up to larger urban regions (involving thousands of vehicles) was still beyond reach.

It was with this background that the SimplyfAI proposal was formed, under the assumption that, once we had achieved automation of interpretation of a wide range of sensor data, some general planning technology would be available to generate strategies to help Transport Operators deal with excessive traffic and consequent extra pollution in exceptional situations. Whether this would be feasible was still unclear when the consortium was formed. Two achievements in 2015, however, both initiated by the COST ARTS Network, gave us more confidence that AI planning would scale sufficiently[3].

The first work (chronologically) was initiated by Matija Gulic while on a COST-funded short term scientific mission (STSM) visit to the University Carlos III of Madrid, with hosts Ricardo Olivares and Daniel Borrajo (Gulić, Olivares, and Borrajo 2016). This work involved joining together a SUMO simulator (Krajzewicz et al. 2012) to an AI Planner, via a monitoring and execution module called the "Intelligent Autonomic System". The planning representation was done using PDDL 2.1 (Fox and Long 2003), with no explicit representation of vehicles in the planner. Instead, traffic concentrations on road links are represented by relative descriptors, such as very-low, low, medium and high. Light change actions are enumerated to cover all the ways that a particular configuration would effect the arrangements of road links: for example, an action to change the density of vehicles on a road link from medium to low, by green lighting ways out of that road link, is given in (Gulić, Olivares, and Borrajo 2016). By abstracting away from explicit counts of vehicles, the system can deal with regions containing thousands of vehicles. In their work vehicles are not explicitly modelled, but considered as density. Also, the close coupling with SUMO demonstrates the use of monitoring and replanning very effectively, and allows exhaustive testing of the system under sets of disturbances (vehicle influx, road closures).

The second was initiated by another COST STSM, the visit of Mauro Vallati to the Delft University of Technology, where he worked with Bart De Schutter, member of the Delft Centre for Systems and Control. In that research centre, modelling traffic "flows" and analysing them using Model Predictive Control, was a well established practice (Lin 2011; van den Berg et al. 2004). They approached Daniele Magazzeni, a co-author of the AI planning engine UPMurphi (Della Penna et al. 2009), and started a collaboration which resulted in the encoding of a flow model of vehicles through junctions in PDDL+, and the configuring of UPMurphi to solve goals (in terms of numbers of vehicles on road links) by changing traffic signals over a period of time. With a particular PDDL+ representation and heuristics, they demonstrated that UPMurphi could solve traffic problems containing thousands of vehicles, in response to exceptional conditions (Vallati et al. 2016). They showed the efficacy of the resulting strategy by comparing its execution with a "fixed time" plan using SUMO. The largest scenario used

---

[3] these two achievements were recognised in October 2015 as they were joint winners of "The Second COST ARTS Competition: Increasing the Resilience of Road Traffic Support Systems by the Use of Autonomics" https://helios.hud.ac.uk/cost/comp2.php

was a hand simulation of a real problem that had occurred in Manchester in 2015. There had been an emergency closure of a link in the inner ring road causing all the traffic to re-route through the City Centre. The strategy generated by UPMurphi was shown by SUMO to be approximately twice as efficient as a fixed time strategy.

A possible advantage in using the continuous PDDL+ -based approach over the classical PDDL approach was that the representation contained exact counts of vehicles, and modelled continuous change of vehicle numbers on road links during green times. In other words, the use of PDDL+ was semantically closer to the problem conceptualisation than using a discrete PDDL. than On the other hand, there are very few PDDL+ planners apart from UPMurphi, and the tests of the PDDL+ method generated strategies in SUMO did not involve monitoring and re-planning. In the end, however, the use of UPMurphi appeared to be the best choice for SimplyfAI, given the accuracy of the continuous models it uses, and the demonstration of its potential to scale up.

The Application

The initial focus of the one-year SimplyfAI project (September 2015 - spring 2016) concentrated on the Semantic Enrichment of the data in a collaboration involving BT and members of the PARK research group at Huddersfield. The raw data was taken from transport and environment sources and integrated into a BT "Data Hub", using semantic technologies such as the universal RDF triple format and a data ontology. The method was to take real time feeds and process them until they produced logical facts about a traffic scenario, which could serve at part of an initial sate of a AI planner. Figure 1 indicates the anticipated architecture of the final system, built around a sense-control loop typical of autonomous systems. To work towards that, however, the data enrichment and strategy generation have to be tested in a real scenario, hence rather than taking in real-time current data, we adjusted the system so that what would be translated into the current state would be from *historical* data. This would allow checking the performance of the system against the observed performance from historical data, in order to assess its feasibility, before real trials of the system in a future project. Also, although environmental data was gathered into the data hub, for simplicity we decided to concentrate on the targeted reduction of congestion, and investigate effects on air quality after that had been achieved.

The Main Data Sources

As a basis for exploring exceptional or emergency traffic conditions, we chose to use historically averaged traffic data from a time/day when the road links were most congested: morning rush hour, between 8am and 9am on a non-holiday weekday. The main data source was the "Saturn" system[4]. From this and other Transport Engineer documentation records our partners BT and TfGM extracted, for the Salford region:

1. the topology of the road links (a link is a uni-directional part of a road between two junctions);
2. the vehicle capacity of all the road links (in numbers of "passenger car units" –pcu– which takes into account the differing size of vehicles) ;
3. the average traffic flows between links in number of pcu's per second. This number represented the number of vehicles flowing through a particular junction at a certain time of day, when the corresponding traffic signal phase is green. A special case of this were flows in and out of boundary junctions.
4. the traffic signal position, stages of signals, minimum and maximum time that a signal stage can be set for;
5. inter-green timings between each of the stages of the signals,
6. the data at a certain instance that the plan is expected to start from: number of vehicles on each link, and the settings of each light stage.

These data items made up the "initial state" of a problem file in planning terms. The "goal" language of the planner is what the actions in the domain model can effect - in this case the goals are made up of numerical expressions denoting predicates on the occupancy levels of road links.

Initial Trials

As is the case in fielded trials, as academics we had no overall control over many of the system parameters: in this case the region chosen, the nature of the data feeds and the validation of the end result was largely in the hands of the rest of the consortium. After the data describing an initial region of Salford was used to populate a PDDL+ initial state, UPMurphi was initiated with a simple test goal, and using a similar configuration to our previously published earlier work(Vallati et al. 2016). The initial region contained over 200 road links, 70 junctions, and over 300 vehicle flows. Whilst the initial region contained a similar number of vehicles to the hand crafted region that had been used in the successful trials with with UPMurphi, the number of road links and junctions used in the earlier work were approximately 10% of the size of the real scenario. Also, the previous tests used simple junctions - in the real situation some of the junctions were complex. Junction 1202 in Figure 2, for example, contains a cycle of seven stages, where each stage is defined by a different set of flows being active. Hence, we failed to produce a runnable executive, and it became clear that scenarios with hundreds of road links and flows were not feasible.

Exacerbating the problem was the fact that our representation was making certain assumptions that were not realistic –notably we were not modelling the time between consecutive light stages ("inter-greens"). The duration of inter-greens often varied –it could be as much as 25 seconds for pedestrian crossings across a busy junction, or 0 seconds if the difference between 2 stages was the green lighting of a filter arrow.

After several iterations of reducing the region's size, without success, we attempted to use other AI Planners capable of inputting a form of PDDL+, such as DReal (Bryce et al. 2015) and Popf (Coles and Coles 2014). Like UPMurphi, none of these options worked and the reason appeared to be

---

[4] http://www.saturnsoftware.co.uk/saturnmanual

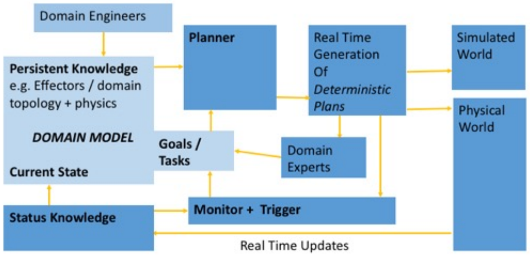

Figure 1: The Target Architecture

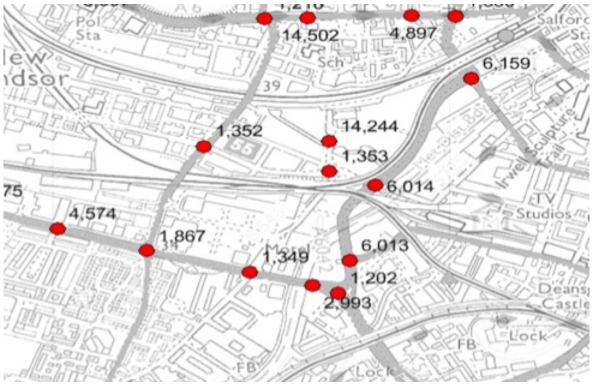

Figure 2: The Modelled Area

the same: the size of the problem was prohibitive. We also experimented with larger memories courtesy of the HPC lab that had been used for IPC 2014 (Vallati et al. 2015). This provided no significant scale up. The project was at this stage in its last few weeks, hence another option of changing to a simpler PDDL representation, and following the work of Gulic (Gulić, Olivares, and Borrajo 2016), was infeasible in the time left.

A successful approach

Our final course of action was motivated by the fact that the project was a *feasibility study*, and hence showing any working system with the real data was better than none. We agreed with the Transport Engineers (Figure 2) a reduced region to focus on, and adopted a systematic approach, starting from the simple, to encode it, as follows: (i) create a model of a very small portion of the network (with only 2 junctions - 1202 and 1349, in figure 2 and 3), and successfully apply the planner (UPMurphi) to solve a simple goal. This involved stripping the initial state file provided to us by the Consortium of all details apart from surrounding area of the 2 junctions. (ii) extend the modelled network until the limits of the planner were reached. We added new junctions until the configuration would no longer compile. (iii) re-represent the model to improve efficiency until the configuration compiled. We iterated three times over steps (ii) to (iii). We concentrated on engineering the PDDL+ model in a very efficient way, so to minimise the size of groundings produced by the planner. Focusing on this, we were also able to represent features that made the model more realistic, such as the addition of inter-green processes to all the junctions, and the introduction of processes representing roadworks.

The final modelled region within the Salford district of Manchester abstracted in Figure 3, and identified by its "Saturn" numbering for junctions. Directed links are identified by the concatenation of the names of their start and end junctions.

- The model consists of 15 junctions and 34 road links: 7 junctions are controllable junctions (in red) and the 8 outer junctions are not modelled as controllable, but act as a boundary to the region.
- The controllable junctions have 69 controllable flows in total between them. For instance, junction "1349" has 12 flows, as a vehicle has a choice of 3 exits when entering from any of the 4 directions.
- Each controllable junction has between 2 and 7 variable-time light stages. A stage determines which of the flows through that junction are on and have traffic flowing. When the set of flows that are on at any instant changes, this is identified as a stage change.
- Each light stage has a fixed inter-green period between its end of green and the start of the next green light stage. The duration of each inter-green is dependent on the stage and junction, and may be 0 (e.g. when a new stage consists of the previous stage augmented with a filter arrow turning on) or over 20 seconds (e.g. when a pedestrian crossing is in use between stages). All flows are considered off during the inter-green.
- Roadworks could be placed in links as follows: they were modelled as simple junctions with 2 flows, one off and one on at any point in time. The intergreen would vary in size depending on the size of the roadworks. A similar model could be used for pedestrian crossings. In both cases, however, the introduction would add two extra links and two extra process flows to the total.
- Each boundary road link going into or out of the region was modelled as a single flow process.

This configuration was at the edge of the limit for the version of UPMurphi that we were using –to add roadworks, for example, we needed to abstract outgoing process flows (this would abstract any limit on the volume of traffic leaving via an out link) leading from the region. Where the abstracted outgoing process flows were not near road links involved in goals, this had little or no effect on the result.

As a comparison, beside all the extensions needed for modelling intergreens, in/out-flows, etc., the largest test presented in (Vallati et al. 2016) involved 9 junctions with between 2 - 3 light stages, 21 links, and 27 controllable flows.

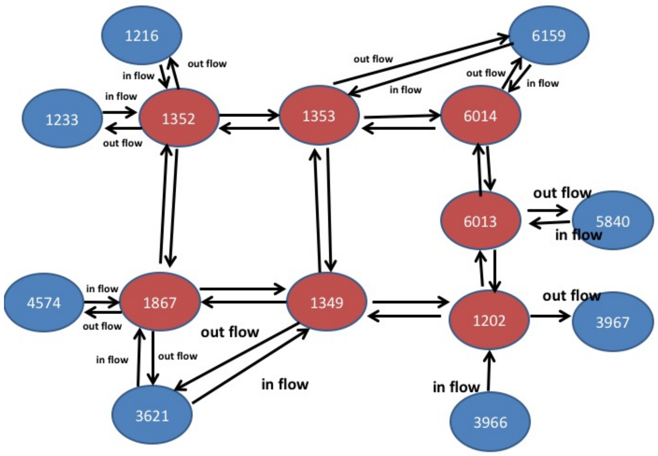

Figure 3: Modelled Flows and Links in the Salford Area: Abstracted View

Evaluation

Test Set-Up

We tested the software configuration on a range of classes of problems (i) to clear a saturated road link as soon as possible; (ii) to clear several saturated road links as soon as possible; (iii) to clear a region as soon as possible; and (iv) to clear a saturated road link with nearby road works.

The idea behind the tests was that, when a problem was spotted, the normal fixed time strategy would be turned off, and replaced by the planner-generated strategy. When the plan achieved the goal, the fixed time strategy would be turned back on.

All the goals in the tests below have the format:

```
X1 < N1 & X2 < N2 ...
```

where `Xi` is the road link occupancy, and `Ni` is the desired occupancy level. Hence, in this context, clearing road links equates to lowering the occupancy to less than a certain –predefined– value.

UPMurphi is configurable using two types of heuristics (as described in (Vallati et al. 2016)). One allows certain preconditions to be put on action choices, and another specifies the goal heuristic. In a nutshell, the first heuristic can provide some guidance about when it is more promising to apply available actions, while the goal heuristic provides an estimation of the distance of the current state from a goal state.

For all the tests with the configuration below, we specified only the goal heuristic, as the other heuristic did not seem to play a clear role in the success of the tests. The goal heuristic amounted to minimising the values of link occupancy in the goal expression (`X1`, `X2`, ...). The tests that completed generated strategies in less than 30 seconds, on a standard Linux PC with 2GB of memory. The strategies (plans) were composed of sequences of the instantiation of the single operator schema in the PDDL+ model: to move on a traffic a signal stage on to the next stage (respecting intergreen intervals, of course).

Validation of the strategies generated was carried out in several steps:

1. Comparison with what would be expected in a "common sense" solution.
2. Comparison of the effects of the generated strategies with a fixed strategy which had been optimised for the time of day by Transport Engineers, using simulation software (SUMO and AIMSUN). Clearly this fixed strategy was not generated to deal with the exceptional event, but nevertheless this was assumed a good comparison as that strategy would be operational when an event occurred.
3. Estimates of savings in terms of tail-pipe emissions.

Results

The first test was in part intended to investigate the connection between the planner's internal traffic model (based on flow values), the microsimulation model SUMO being utilised by Infohub, and the AIMSUN microsimulation package used by TfGM. We were aware that if the PDDL+ model was correct/sufficiently accurate, then its generated strategy was guaranteed to solve the goal when executed; and if the independent simulation tests showed that it does not, then we would conclude that the planner's PDDL+ model was not correct or sufficiently accurate.

The first test was inspired by a possible scenario. Assume there was an extreme vehicle build upon a link (in our case 3966_1202) entering into the region, and the consequent air quality implications around the link were unacceptable. This problem would be to clear the link as soon as possible. It is formalised by assuming the link contains at the initial state an unexpectedly large number of vehicles (in this case, 300), and the goal state is to reduce the number to less than 10. In the test scenarios, this was the only time that we introduced our own data into the problem formulation, in order to simulate the occurrence of some exceptional event.

The common sense, approximate strategy to solve this kind of problem would be as follows. At the junction that the link leads to (in this case 1202) called the "primary junction": give maximum green time to those light stages which allow vehicles to leave the link, and minimise those stages which do not, so that the lights will quickly cycle back to the stages letting out traffic. At the junctions that lead off from the primary junction (in this case 6013 and 1349): give at least enough green time to the links leading in from the primary junction to make sure that the links do not get congested and the increased level of traffic can go through them smoothly. This strategy may have to be repeated through junctions further away if necessary. To visually inspect the quality of the strategy, we checked that it was indeed close to this common sense solution.

Considering the simulation, the traffic models (AIMSUN and SUMO) were run independently by the transport authority and InfoHub, respectively, using the planner-output strategy and the optimised strategy. In the first test, after validating that the simulations were fairly consistent, the reduction in time to clear a junction using the planner-output strategy

	Original	**Planned**
run 1	730	700
run 2	907	800
run 3	733	710
run 4	907	690
run 5	763	700

Table 1: Seconds required for achieving the goal set for the first test by the fixed strategy, and by the strategy generated using the proposed planning-based approach. Simulation has been performed using SUMO. In order to take into account the variability of traffic flows, five independent runs have been considered.

was approximately 12% using all the simulations. Table 1 presents the results of SUMO simulation in terms of time needed for achieving the set goals using the fixed strategy and the planned strategy. In order to take into account the variability of traffic (e.g., acceleration, brake, bus stops, etc.) five independent runs have been executed.

The next batch of tests was run to test the idea of clearing congestion from several links at the same time, which lead into the region from a particular direction. For example, 3 links may be in an air quality management zone, and some exceptional event is forcing traffic into the region from this direction. A common sense, approximate strategy to solve these kinds of problems is much more difficult to formulate than for the single link examples (and hence one of the reasons for automation). Clearly, we need to give a great deal of green time to let the vehicles out of the links from which they are emanating. However, given the congestion, other measures need to be put in place such as keeping other traffic flows into the region down, and maximising the flows away from the congestion. This is further complicated by the fact that, when moving the green light stages from one stage to another, several flows may be affected (some which might help and some not). However, a sensible pattern appeared to exist in the planner-generated strategies, to green light the correct junctions. Similarly, simulations with the planner-generated strategy versus the optimised strategy proved that the former achieved the goal in much less time, at least 20%. Tests dealing with larger subregions of the targeted region, and the introduction of traffic lights, produced similar savings.

Reductions in Tail-pipe emissions

The assumption we use is that clearing a junction (in particular, reducing it from a level of saturation as quickly as possible) will lead to a reduction in tail-pipe emissions, and hence overall pollution. We illustrate this by deriving the expected emission reduction along the strategic link 3966_1202 of test 1. The potential emission reduction achieved by the strategy has been calculated, very approximately, as follows:

$$\text{Emissions Reduction} = (Y - X) * (E1 - E2)$$

where:

- Y = Time taken for the goal to be reached by the normal strategy provided that the link is congested;
- X = Time taken for the goal to be reached by the planned strategy provided that the link is congested;
- E1 = The Emission expected given that the model is congested and the normal strategy is being used;
- E2 = The Emission expected given that the model is not congested and the normal strategy is being used.

E1 and E2 emissions have been provided from a “capacity” case (E1) and a normal case (E2). For both, default fixed timings were used in the AINSUM model. In the ’normal’ case, SATURN demand for the L3966_1202 link was used.

The overall effect of applying the planner-generated strategy was measured using the TRACI[5] impact assessment tool built into SUMO.

As well as estimating the emissions reduction in the link referred to in the goal, the emissions reduction from the overall effect of applying the strategy to the model given that certain links carry more weight (e.g. those that are in an air quality management zone) was calculated. The emissions around the link to be cleared were calculated to drop by 5%, whereas the overall drop over the region was 2.5%. It must be stressed that these results are preliminary, however, with still much more testing to be done.

Discussion

At the end of the project, the consortium was convinced enough by the results of using AI planning as to want to pursue field trials and potentially a software product. Using a domain independent planning engine was, in the end, adequate for showing the proof of concept of a planning-driven approach to the solution of a real problem. While we did not get to the stage of monitoring and re-planning, the plan generation speeds during the trials made re-planning in real time look feasible.

The main advantage of the planning approach appears to be its ability to generate a useful strategy in real time to meet the needs of a new unexpected situation. This relies on the flexibility of the approach, as well as the speed of a planner, in that goals involving different road links, and different initial states, can be generated in real time to suit the kind of problem to be solved. The ability to generate complete initial states and triggered goals in real time (and so be responsive to a detected event) was also a persuasive factor for the consortium in the use of an AI Planning approach. Also, new effectors such as the exploitation of variable speed limits or variable message signs (affecting traffic flows) can be added to the planner’s domain model modularly, meaning that new strategies generated will contain instances of those effectors if they help achieve a goal.

There were many lesson learned, however, and challenges still to be overcome: we summarise the main ones below.

- *Problems with the data:* the “meaning” of the flow values obtained from Saturn were not as we had anticipated them. While we expected these values to be the maximum number of vehicles –expressed in a suitable unit, such as PCU– that could flow assuming a queue was formed in

---

[5] http://link.springer.com/article/10.1007/s10098-010-0338-9

the oncoming link, in fact they denoted the flow averaged for the particular time of day.

- *Problems with adequacy of our representation:* the PDDL+ model embodied several assumptions that made it inaccurate. Firstly, it assumed that as soon as vehicles enter into a link, they are queued at the next. Secondly, there were breaks in links that we did not model, such as roundabouts and pedestrian crossings.
- *Problems in complexity measures:* the field trial demonstrates the crucial importance of estimating accurately measures of the trial (region) size a priori, and acquiring planning machinery which would cope with that. In our case the measure of "number of vehicles in a region" was not as relevant for determining limits as other factors such as the total number of links, and consequently we were over optimistic in our expectations.
- *Problems with understanding the chosen planning engine:* Several classes of scenarios when input to UPMurphi would not yield results - for example in the first class of tests, instead of raising the occupancy of a road link to 300 (well about its maximum value), the normal approach would be to increase greatly the *flow-in* value. In such a scenario we were unable to obtain an output. From extensive tests it appeared that if the goal was one in which actions could make immediate progress towards, then an answer would be extracted. On the other hand, if UPMurphi was initiated in a heuristic "canyon" it was likely that no result would be output. Given a fixed number of vehicles to start, however, the path to the goal heuristic (minimise occupancy on 3966_1202) was monotonic, which seemed to guarantee a resulting plan.
- *Problems with a purely goal directed strategy:* while the effect of a generated plan was successful for solving the goal, other junctions through the region were not optimised. In fact the light signals in other junctions in the region were all left to run to maximum (actions to move them on a stage would not be taken unless it helped towards solving the specific goal). Also, goals such as the *maintenance* of a goal value are desirable in some situations. For a future system, we need to design a richer goal language.
- *Problems in joining up the technology:* When engineering a planning component into a larger application we naturally use the high level interface input language – in this case PDDL+. Components of the initial state are assembled automatically from the data hub. In our application, a different team had responsibility for producing the tool which assembles PDDL+ elements. As this work was chronologically scheduled first, there was an over-commitment to a particular target representation. The work following on from this involved configuring the planner, and required changing the PDDL+ representation of the problem domain many times, during the extensive testing with the simulated scenarios. Hence the coding of a PDDL+ –assembling tool, and hence any work on the end to end effectiveness of the system, would need to be completed only after a final PDDL+ representation had been agreed upon.

Conclusions

In this paper we have described the operation and results of the SimplyfAI project, a collaboration between a transport authority, academics, a large technology provider and two SME's, which included in its remit the use of AI Planning to generate plans of traffic light changes to achieve transport goals. The trials involved using historical data describing the traffic in an area of Salford, Greater Manchester. The plans (timing changes of traffic signals) output were judged to be useful for dealing with exceptional situations, using both hand inspecting the strategies to check that they were *sensible* and simulating their execution using two different traffic modelling software packages AIMSUN and SUMO. We believe that this is the first successful demonstration of AI Planning technology to create useful strategies for UTM where the overall control for the region chosen, the nature of the data feeds and the validation of the end result was largely in the hands of non-academic stakeholders.

On the other hand, the success is limited by several factors discussed above. While the results of the plan generation component seem acceptable to the stakeholders, a certain amount of scale-up is required in terms of traffic area covered, and granularity of representation, before the project can progress further. Whether this can be achieved by a future PDDL+ planner, or by the utilisation of a discrete representation as in the work of (Gulić, Olivares, and Borrajo 2016), remains to be seen.

Acknowledgements

This work was supported by the ESF-funded COST Action 1102, and UK NERC grant NE/N007239/1 issued through Innovate UK. Many previous works made this trial possible. Firstly, we are indebted to the members of the SimplyfAI consortium, and in particular Grigoris Antoniou and Ilias Tachmazidis who designed the semantic enrichment engine for the traffic data, and to Sam Corns and Louis Burrows who conducted the comparative simulations with AIMSUN and SUMO respectively. We would like to thank Bart De Schutter for hosting Mauro Vallati's STSM in 2015, and Daniele Magazzeni for introducing us to the implementation and configuration of UPMurphi. Finally, we acknowledge the help of members of the COST ARTS Action too numerous to name who explained to us the finer points of Transport Engineering.

References

Bryce, D.; Gao, S.; Musliner, D. J.; and Goldman, R. P. 2015. Smt-based nonlinear PDDL+ planning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 3247–3253.

Cenamor, I.; Chrpa, L.; Jimoh, F.; McCluskey, T. L.; and Vallati, M. 2014. Planning & scheduling applications in urban traffic management. In *Proceedings of PLANSIG*.

Chrpa, L.; Magazzeni, D.; McCabe, K.; McCluskey, T.; and Vallati, M. 2015. Automated planning for urban traffic control: Strategic vehicle routing to respect air quality limitations. In *6th Italian Workshop on Planning and Scheduling (IPS)*.

Coles, A. J., and Coles, A. I. 2014. PDDL+ planning with events and linear processes. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, ICAPS*.

Della Penna, G.; Magazzeni, D.; Mercorio, F.; and Intrigila, B. 2009. UPMurphi: A tool for universal planning on PDDL+ problems. In *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS)*.

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. *Journal of Artificial Intelligence Research* 27:235–297.

Gulić, M.; Olivares, R.; and Borrajo, D. 2016. Using automated planning for traffic signals control. *PROMET-Traffic&Transportation* 28(4):383–391.

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

Krajzewicz, D.; Erdmann, J.; Behrisch, M.; and Bieker, L. 2012. Recent development and applications of SUMO - Simulation of Urban MObility. *International Journal On Advances in Systems and Measurements* 5(3&4):128–138.

Lin, S. 2011. *Efficient model predictive control for large-scale urban traffic networks*. TU Delft, Delft University of Technology.

McCluskey, T. L., and Vallati, M. 2014. Extracting information from urban traffic data to improve road traffic management. In *Final Report for Transport Catapult Innovation Voucher*.

McCluskey, T.; Kotsialos, A.; Müller, J.; Klugl, F.; and Schumann, R. 2016. *Autonomic Road Transport Support Systems*. Autonomic Systems. London, UK: Springer.

Miles, J. C., W. A. J. 2006. The potential application of artificial intelligence in transport. *Journal of Intelligent Transport Systems* 153.

Taale, H.; Fransen, W.; and Dibbits, J. 1998. The second assessment of the SCOOT system in Nijmegen. In *IEEE Road Transport Information and Control*, number 21-23.

Vallati, M.; Chrpa, L.; Grzes, M.; McCluskey, T. L.; Roberts, M.; and Sanner, S. 2015. The 2014 international planning competition: Progress and trends. *AI Magazine* 36(3):90–98.

Vallati, M.; Magazzeni, D.; De Schutter, B.; Chrpa, L.; and McCluskey, T. L. 2016. Efficient macroscopic urban traffic models for reducing congestion: a pddl+ planning approach. In *The Thirtieth AAAI Conference on Artificial Intelligence (AAAI)*.

van den Berg, M.; De Schutter, B.; Hegyi, A.; and Hellendoorn, J. 2004. Model predictive control for mixed urban and freeway networks. In *Proceedings of the 83rd Annual Meeting of the Transportation Research Board*, volume 19.

Various. 2007. Artificial intelligence in transportation. *Transportation Research Circular E-C113, Transport Research Board*.

Xie, X.-F.; Smith, S.; and Barlow, G. 2012. Schedule-driven coordination for real-time trafc network control. In *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS)*.