Automated Planning for Generating and Simulating Traffic Signal Strategies

Saumya Bhatnagar[1] , **Rongge Guo**[1] , **Keith McCabe**[2] , **Thomas McCluskey**[1] , **Francesco Percassi**[1] and **Mauro Vallati**[1]
[1]School of Computing and Engineering, University of Huddersfield, UK
[2]Simplifai Systems Ltd., Huddersfield, UK
{s.bhatnagar, r.guo, t.l.mccluskey, f.percassi, m.vallati}@hud.ac.uk,
keith.mccabe@simplifaisystems.com

Abstract

There is a growing interest in the use of AI techniques for urban traffic control, with a particular focus on traffic signal optimisation. Model-based approaches such as planning demonstrated to be capable of dealing in real-time with unexpected or unusual traffic conditions, as well as with the usual traffic patterns. Further, the knowledge models on which such techniques rely to generate traffic signal strategies are in fact simulation models of traffic, hence can be used by traffic authorities to test and compare different approaches.

In this work, we present a framework that relies on automated planning to generate and simulate traffic signal strategies in a urban region. To demonstrate the capabilities of the framework, we consider real-world data collected from sensors deployed in a major corridor of the Kirklees region of the United Kingdom.

1 Introduction

The pandemic has changed the way in which people work, shop, and spend their leisure time. This has resulted in significant changes to travel patterns globally, leading to a greater need for tools to manage unpredictable and previously unseen travel patterns. The use of AI techniques becomes imperative under such circumstances, as traditional approaches to urban traffic control tend not to respond in a timely manner to sudden temporal fluctuations of traffic flows. To deal with time-sensitive conditions, strategies of interventions have to be generated on the fly, based on the current actual conditions of the network, and this is considered to be beyond the capacity of human operators. Within the automated planning field of AI, techniques for dealing with urban traffic control problems, with particular focus given to traffic signal control and optimisation, are gaining increasing attention [Gulić *et al.*, 2016; Antoniou *et al.*, 2019; McCluskey and Vallati, 2017; Pozanco *et al.*, 2021]. The growing interest is also due to the fact that automated planning techniques are well-positioned to deal with these kinds of problems, as they can rely on validated and verified knowledge models of traffic. Thanks to such knowledge models, planning techniques can address unusual situations, hence difficult to be dealt with by using machine learning data-driven approaches [Smith, 2020].

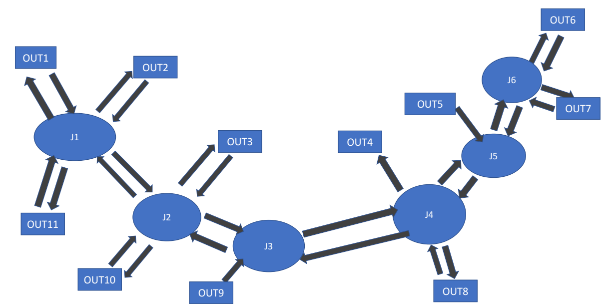

Figure 1: A simplified overview of the modelled urban region, in terms of junctions, links, and boundaries (rectangles). For the sake of readability, the map is not correctly scaled.

The knowledge models needed by automated planning approaches are also ideal to be exploited to simulate traffic condition and evolution [Bhatnagar *et al.*, 2022a]. It is worth noting that, despite their crucial role, traffic simulators are only available for very limited areas of urban regions due to their high cost. To be used, traditional simulators require a traffic authority to commit to significant investments to be generated and calibrated, and require frequent updates and revision. They are also based on the assumption that travel patterns are regular and change little from year to year: in the post-pandemic world, this seems to be no longer the case. On this regards, automated planning lends itself well thanks to its concise and declarative symbolic representation of the dynamics to model, achieved via the standard PDDL+ language [Fox and Long, 2006], and to the modularity of the domain-independent planning paradigm, that allows to swap automated reasoning engines and knowledge models without affecting the overall technical infrastructure.

In this work we present a framework that uses automated planning to generate traffic signal strategies and simulate their impact on an urban region. In particular, we focus on a region where traditional reactive traffic control SCOOT [Taale *et al.*, 1998] is in operation, taking the opportunity to describe how relevant information can be extracted from SCOOT sensors.

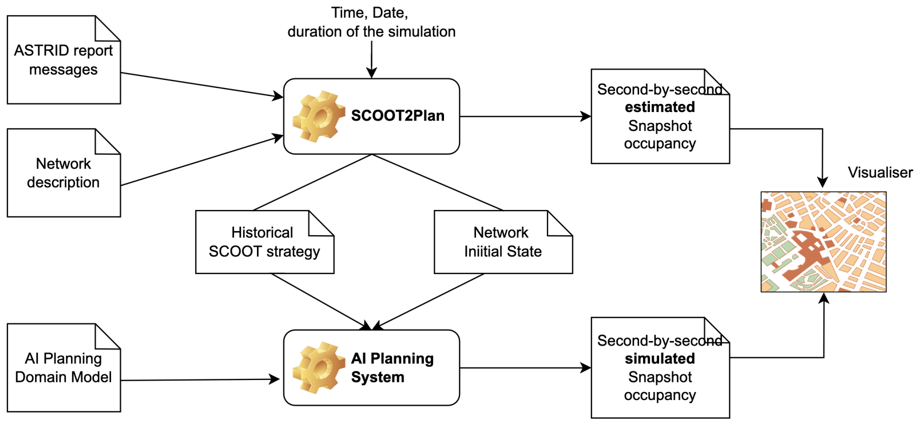

Figure 2: An overview of the proposed architecture.

2 Case Study

We focus on an urban region currently controlled via SCOOT system. SCOOT is a demand-driven, traffic-responsive control aimed at handling cycle-to-cycle changes in demand. In response to changes in traffic flows, SCOOT would gradually adapt and adjust the traffic signal timings of a set of managed neighbouring junctions. SCOOT is dependent on its own local data sensors, usually inductive loops embedded in the road surface. Such loops can cover one or more lanes and are usually calibrated accordingly. It stores the data coming from its sensors, and its internal behaviour, into a dedicated database called ASTRID [Hounsell and McDonald, 1990].

The modelled region is situated in West Yorkshire, United Kingdom, specifically within the Kirklees council. It consists of a major corridor that links the Huddersfield ring road with the M1 highway and the southern part of the Kirklees council. It is heavily used by commuters and by delivery vans to get to the centre of Huddersfield town or to move between the M62 and the M1 highways. The corridor is approximately 1.3 kilometres long and consists of 6 junctions and 34 road links. Each junction has between 4 and 6 stages, and between 10 and 17 valid traffic movements. A simplified schema of the considered urban region is shown in Figure 1, in terms of links, junctions, and connections with the outside region.

3 Architecture

The architecture we designed to support traffic control and simulation in a urban region by means of automated planning is presented in Figure 2, and consists of 2 main modules. SCOOT2Plan is in charge of processing the messages produced by ASTRID to generate a snapshot of the network conditions at a give time and date (Network initial state) and to provide the strategy of intervention implemented by SCOOT over a considered period of time (Historical SCOOT strategy), if available. These two outputs of SCOOT2Plan can be used by the planning system, together with a dedicated AI Planning Domain Model, to simulate the evolution of traffic conditions on the considered urban network. We assume that the topology of the network is encoded as part of the automated planning system – since it is static and will not change over time. The cases of traffic accidents or road works that impact the structure of the network are dealt with by modifying the characteristics of the links and of the expected traffic movements, with no changes to the topology. The planning system can provide as output, beside the traffic light strategy, the second-by-second simulated status of all the links of the network over the length of the plan. A similar output can be provided also by the SCOOT2Plan module, and outcomes can be compared for validation and verification purposes, via a visualiser.

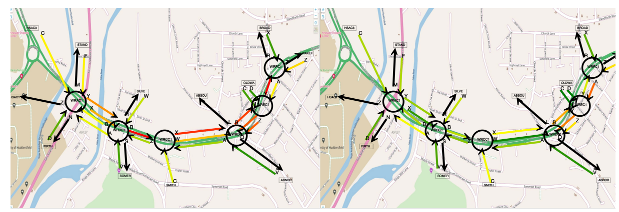

Figure 3: An example of a comparison of different traffic light strategies using the visualiser.

Notably, the proposed architecture can also be used to simulate new traffic signal strategy plans that were not in operation when historical data was collected. That is possibly the most interesting aspect of the proposed approach, and can be done by leveraging on the planning system to generate a new traffic light control strategy. This can be achieved by providing to the planning system the generated network initial state, the domain model, and a goal to be achieved. In the proposed architecture, a set of goals have been defined by the traffic authority in charge of the region, and can be used according to the current traffic conditions. Newly generated plans can then be simulated as per historical plans, and presented via the visualiser. An example use of the visualiser for comparing traffic light strategies is provided in Figure 3.

A cornerstone of the architecture shown in Figure 2 is the planning domain model, that encodes in the PDDL+ language the dynamics of urban traffic control that we are aiming at simulate. Building on top of [Vallati *et al.*, 2016;

McCluskey and Vallati, 2017] and of a corresponding patent[1], which proposed an approach based on PDDL+ to generate traffic signal strategies to reduce the impact of traffic accidents on a controlled urban region, here we describe the model that has been designed to accurately simulate the traffic conditions of a urban network.

We use a software tool to generate a PDDL+ model of the SCOOT region under consideration [Bhatnagar *et al.*, 2022b] which allows the method to be efficiently applied to any pre-existing SCOOT-controlled region. The road network represented as a directed graph, where edges stand for road links and vertices stand for junctions. One additional vertex is used for representing the outside of the modelled region. Each link is characterised by its *static capacity*, i.e. the maximum number of PCUs (Passenger Car Unit) that can be at a single time in the link, and by the *snapshot occupancy* that indicates the number of vehicles currently estimated to be in the link.

Junctions are described in terms of the corresponding traffic signal stages. For each stage, the *next* valid stages are defined, and for each possible next stage, it is also defined the length of the corresponding *intergreen*. This representation allows to model junctions where one or more stages are on-demand only, and/or can be skipped. For each stage, the corresponding active phases are described in terms of the active traffic movements. The status of a junction is defined by the stage (or intergreen) currently active, and by the time spent in the current stage (or intergreen). Traffic movements are represented as *turn rates*. A turn rate represents the expected traffic flow between two links connected to a considered junction, when a corresponding traffic signal phase is active. Flow rates are expressed in PCUs per time unit that, on average, move from the incoming link to the outgoing link of the junction. Flow rates can vary over time, and depends also on the signal stage. In other words, the expected flow between two links can be different depending on which stage is currently active. Flow rates can also be associated with intergreen stages, to model for instance the fact that a specific phase is activated before the others.

PDDL+ processes are used for modelling the flows of vehicles described by turn rates, that are activated when a corresponding traffic signal phase is on green. Dedicated processes are also used to measure the time spent on green by the traffic signal stages (or intergreens) on the considered junctions. PDDL+ events are used to stop flows of vehicles when the receiving link is completely full or the discharging link is empty. Finally, a dedicated action *SwitchStage* is used to model the fact that one stage is stopped and the junction is transitioning, after the intergreen, to a new stage.

With the above-mentioned constructs, it is possible to fully describe the initial status of a network and, given a plan representing a traffic signal strategy for a specified period of time, simulating the evolution of the network conditions over time. An excerpt of the initial state description of a traffic network is provided in Figure 4, where the status of a junction J1 and of some corresponding links are defined.

The SCOOT2Plan system generates the network conditions for a particular time and duration from available SCOOT data. It is important to emphasise that the state information extracted from SCOOT can be used in one of two modes: (a) for use with the AI Planner (PDDL+) simulation of a SCOOT plan (extracted from the corresponding ASTRID messages). This mode is used to validate the simulation, as it can be compared to a traffic distribution generated only the SCOOT sensor values; (b) for use with the AI Planner (PDDL+) simulation of a newly generated traffic strategy, to test out its effectiveness before being implemented.

```
(active J1_stage2)
(= (greentime J1) 9)
(= (intergreen J1_stage1 J1_stage2) 5)
(= (intergreen J1_stage2 J1_stage3) 6)
(= (snapshot_occupancy J1:linkX) 35)
(= (static_capacity J1:linkX) 90)
(= (turnrate J1_stage2 J1:linkW J1:linkX) 0.33)
```

Figure 4: Excerpt of PDDL+ initial state description.

The planning engine used for generating traffic signal strategies is the well-known ENHSP [Scala *et al.*, 2016b; Scala *et al.*, 2016a], and the same engine can also be used to simulate SCOOT strategy plans or plans generated using different techniques [Bhatnagar *et al.*, 2022a].

4 Demonstration

The demonstration will focus on the use of the proposed architecture to generate traffic light strategies, and to show how generated strategies can be compared with the SCOOT-implemented plans on historical data.

We consider historical data collected in January 2022, as COVID-related movement restrictions were lifted in the United Kingdom. We identify a week day and a weekend day where no major disruptions were recorded for the considered region, no major event happened, and no faults were recorded on the SCOOT infrastructure – to reduce the probability of noisy sensors readings. According to the aforementioned criteria, we select Wednesday the 26th of January and Saturday 30th, 2022. We run a 15-minutes simulation of the modelled region during the morning peak hour, at lunch time, and at evening peak hour for the weekday, and at noon for the Saturday scenario. A video demonstration can be found at https://tinyurl.com/yc5vyvar.

5 Conclusion

In this paper we presented an architecture that allows to leverage on automated planning to generate and simulate traffic light strategies in urban areas. While in this demonstration we focused on a single corridor, it is worth noting that the approach can scale to larger and combined regions, including multiple corridors and a large number of junctions and links. Our experience and preliminary experimental analysis indicate that scaling up generally results in a near-linear increase in processing time.

The corresponding system is currently being trialled in the Kirklees region of the United Kingdom, and plans are undergoing to test it in other areas of England.

[1]Traffic strategy system and method of implementing the same. U.S. Patent Application 17/422,996, 2022.

Acknowledgements

This work is partially supported by the UKRI Future Leaders Fellowship [grant number MR/T041196/1].

References

[Antoniou *et al.*, 2019] Grigoris Antoniou, Sotiris Batsakis, John Davies, Alistair Duke, Thomas L. McCluskey, Evtim Peytchev, Ilias Tachmazidis, and Mauro Vallati. Enabling the use of a planning agent for urban traffic management via enriched and integrated urban data. *Transportation Research Part C: Emerging Technologies*, 98:284–297, 2019.

[Bhatnagar *et al.*, 2022a] Saumya Bhatnagar, Rongge Guo, Keith McCabe, Lee McCluskey, Enrico Scala, and Mauro Vallati. Leveraging artificial intelligence for simulating traffic signal strategies. In *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2022.

[Bhatnagar *et al.*, 2022b] Saumya Bhatnagar, Sumit Mund, Enrico Scala, Keith McCabe, Thomas Leo McCluskey, and Mauro Vallati. On-the-fly knowledge acquisition for automated planning applications: Challenges and lessons learnt. In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence, ICAART*, pages 387–397, 2022.

[Fox and Long, 2006] Maria Fox and Derek Long. Modelling mixed discrete-continuous domains for planning. *J. Artif. Intell. Res.*, 27:235–297, 2006.

[Gulić *et al.*, 2016] Matija Gulić, Ricardo Olivares, and Daniel Borrajo. Using automated planning for traffic signals control. *PROMET-Traffic&Transportation*, 28(4):383–391, 2016.

[Hounsell and McDonald, 1990] N. Hounsell and M. McDonald. Astrid: Automatic scoot traffic information database. Project report. transport research laboratory, 1990.

[McCluskey and Vallati, 2017] Thomas Leo McCluskey and Mauro Vallati. Embedding automated planning within urban traffic management operations. In *Proc. of ICAPS*, pages 391–399, 2017.

[Pozanco *et al.*, 2021] Alberto Pozanco, Susana Fernández, and Daniel Borrajo. On-line modelling and planning for urban traffic control. *Expert Systems*, 2021.

[Scala *et al.*, 2016a] Enrico Scala, Patrik Haslum, and Sylvie Thiébaux. Heuristics for numeric planning via subgoaling. In *Proc. of IJCAI*, pages 3228–3234, 2016.

[Scala *et al.*, 2016b] Enrico Scala, Patrik Haslum, Sylvie Thiébaux, and Miquel Ramírez. Interval-based relaxation for general numeric planning. In *Proc. of ECAI*, pages 655–663, 2016.

[Smith, 2020] Stephen F. Smith. Smart infrastructure for future urban mobility. *AI Mag.*, 41(1):5–18, 2020.

[Taale *et al.*, 1998] H Taale, W.C.M Fransen, and J. Dibbits. The second assessment of the SCOOT system in Nijmegen. In *IEEE Road Transport Information and Control*, number 21-23, April 1998.

[Vallati *et al.*, 2016] Mauro Vallati, Daniele Magazzeni, Bart De Schutter, Lukás Chrpa, and Thomas Leo McCluskey. Efficient macroscopic urban traffic models for reducing congestion: A PDDL+ planning approach. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 3188–3194, 2016.