Urban Traffic Control via Planning with Global State Constraints (Extended Abstract)

Franc Ivankovic[1], Mauro Vallati[2], Lukáš Chrpa[3], Marco Roveri[1]
[1]University of Trento
[2]University of Huddersfield
[3]Czech Technical University in Prague
franc.ivankovic@unitn.it, M.Vallati@hud.ac.uk, chrpaluk@fel.cvut.cz, marco.roveri@unitn.it

Abstract

Planning with global state constraints is an extension of classical planning such that some properties of each state are derived via a set of rules common to all states. This approach is important for the application of planning techniques in manipulating cyber-physical systems, and has been shown to be effective in practice. Urban Traffic Control (UTC) deals with the control and management of traffic in urban regions, and includes the optimisation of traffic signals configuration to minimise traffic congestion and travel delays. In this paper, we briefly introduce how to cast the UTC problem into the formalism of planning with global state constraints, and we perform a preliminary experimental evaluation considering significant scenarios taken from the literature, and a new one based on real-world data. The results show that the approach is feasible, and the quality of generated solutions has been confirmed in simulation using existing symbolic models.

Introduction and Background

Urban Traffic Control (UTC) is a wide research area that is attracting a growing interest from the AI community given the potential impacts in Smart Cities and the Green Deal as promoted by the European Union. UTC aims at optimising traffic flows in urban areas by reducing travel time delays and avoiding congestion of road links. One possibility, considered here, is optimising traffic light configurations in junctions to deal with either recurrent or unexpected traffic conditions.

A *traffic signal configuration* of an junction is defined by a sequence of green light phases, each with its specified duration, that, in consequence, imposes constraints on the traffic movements. Traffic signal configurations operate in cycles, i.e. the sequence of green phases is being repeated (until the configuration changes).

To deal with the UTC problem, we employ planning with Global State Constraints (GSC) (Ivankovic et al. 2014; Haslum et al. 2018). This refers to an extension of classical planning in which some properties of states are determined by a set of rules that are common to all states. This formalism is well suited for applying classical planning methods on domains that involve a network of interconnected physical systems (Ivankovic, Gordon, and Haslum 2019) controlled by discrete controllable variables (for example power networks). The crucial feature of such domains is that a single discrete action (to change controllable variables), such as opening or closing a switch somewhere, affects the entire network in a way that is dependent on the global state of the system. For a detailed formal definition of this approach see (Ivankovic et al. 2014; Haslum et al. 2018). Here we simply remark that we make a distinction between *primary variables* – which have discrete and finite domains and function exactly as in classical planning; *secondary variables* – which can be of any type and determine the properties that depend on the secondary model (in this case a linear program); and *switched constraints*, which is the mechanism through which primary and secondary variables interact.

We make the following contributions: i) we extend the GSC formalism to allow for a subset of secondary variables to *persist* (to be used in computation of the secondary model in the subsequent state). ii) we provide an encoding of the UTC problem into the extended GSC framework. iii) we consider several UTC reference scenarios taken from the literature and new real-world scenarios for an experimental evaluation comparing the performance and quality of the GSC approach w.r.t. the state-of-the-art approach based on PDDL+ (Vallati et al. 2016; McCluskey and Vallati 2017).

Extending the GSC Framework

In previous work on planning with GSC, values of secondary variables are computed in each state and discarded whenever we move to the next state. In contrast, here we allow for values of a subset of secondary variables, which we refer to as *persistent secondary variables*, from the preceding state to be used in computation of the secondary model of the current state. In the UTC domain, we use persistent secondary variables to track the numbers of cars on the roads and the numbers of cars passing some given points.

Proposed Approach

The traffic network is modelled as a directed graph, with roads $\mathcal{R}$ as edges. Nodes represent either: i) junctions, $\mathcal{I}$ or ii) points through which vehicles enter or exit the network $\mathcal{E}$. Each junction has a set of incoming and a set of outgoing roads (denoted *in*(q) and *out*(q), respectively).

Primary variables are used to represent traffic light configurations. To a junction $q \in \mathcal{I}$ we assign variable L_q, whose domain is $\mathcal{D}(L_q) = [0 \ldots N_q]$, where N_q is the number of possible configurations. Additionally, a Boolean variable *time* is set to $\top$ by the action that advances time and $\bot$ by every other action.

Secondary variables include total numbers of vehicles, vehicle inflows and vehicle outflows over a cycle (represented by n_i, $f_{i,in}$ and $f_{i,out}$, for road i, respectively). Additionally f_{ij} is a secondary variable representing the flow from road i to road j over a cycle. **Persistent secondary variables** include i) n_{ij}, which is the number of vehicles currently in i intending to turn to j at the intersection and ii) C_i, which tracks the number of vehicles that left i. The values of n_{ij} and C_i in the previous state are denoted $n_{ij,s-1}$ and $C_{i,s-1}$, respectively.

The following set of constraints is created for each road i. We denote i's origin with a and the destination with b.

$$\begin{array}{ll} n_i = \sum_{j \in out(b)} n_{ij} & n_i \leq n_{i,max} \\ \textit{time} \rightarrow C_i = C_{i,s-1} + f_{i,out} & \neg\textit{time} \rightarrow C_i = C_{i,s-1} \\ f_{i,in} = \sum_{k \in in(a)} f_{ki} & f_{i,out} = \sum_{j \in out(b)} f_{ij} \\ f_{i,out} \leq f_{i,out,max} & f_{i,in} \leq f_{i,in,max} \end{array}$$

If $a \in \mathcal{E}$, then we have an additional constraint $f_{i,in} = a_{enter}$, where a_{enter} is a constant representing the inflow into the network at node a (over a cycle).

At each junction, we have a pair of switched constraints corresponding to each pair of roads between which movement of vehicles is possible:

$$\begin{array}{l} \textit{time} \rightarrow n_{ij} = n_{ij,s-1} + \alpha_{ij} f_{i,in} - f_{ij} \\ \neg\textit{time} \rightarrow n_{ij} = n_{ij,s-1} \end{array}$$

where α_{ij} is a constant denoting the portion of the cars on road i intending to turn to j at b (this value is obtained from historical data). We also have a set of constraints related to traffic light configurations. At each junction q, for each configuration $l \in \mathcal{D}(L_q)$, we need one switched constraint for each incoming-outgoing road pair ij.

$$L_q = l \rightarrow f_{ij} \leq c_{ij} t_{ij,l}$$

where c_{ij} is the flow capacity and $t_{ij,l}$ is the time within a cycle during which the vehicles are allowed to move from i to j. $t_{ij,l}$ must respect the rules regarding green light phase lengths, such as maximum and minimum green time. The objective function (active in every state) is: max $\sum_{f \in \mathcal{F}} f$ where $\mathcal{F}$ contains all flows f_{ij}.

We have two types of **actions**: i) actions that change configurations of traffic lights and ii) an action that advances time for the duration of one traffic light cycle.

The **goals** include *decongestion goals* where we need to reduce the number of cars on a given set of roads below the specified values and *flow goals* where we need to get a specified number of vehicles to pass through a given road.

Experimental Evaluation

Benchmark set. We consider three different networks, and 7 different scenarios, all of which were previously modelled in PDDL+. Scenarios (A-B3) correspond to those used in (Vallati et al. 2016). Scenario A focuses on a simple network composed of three junctions, while scenarios B-B3 are based on a large section of Manchester urban area, composed by 11 junctions, and corresponding to a huge traffic situation. Scenarios C1 and C2 consider a real heavily congested situation from a corridor (1.3 kilometers long, with 6 junctions) situated in West Yorkshire, UK (see (Bhatnagar et al. 2022) for details). In scenarios A-B3, the goals are to decongest a given set of links, while in scenarios C1 and C2 goals are given in terms of vehicles that need to transit through the corridor (flow goals).

Scenario	GSC solver	PDDL+&UPMurphi
Scenario A	48	125
Scenario B	9,045	19,210
Scenario B1	2,295	1,550
Scenario B2	6,120	3,430
Scenario B3	810	950
Scenario C1	400	unsolved
Scenario C2	400*	unsolved

Table 1: Quality of the plans, in terms of simulated time (in seconds) required to reach the goal. * denotes that the solution could not be validated by the PDDL+ simulator.

Experimental environment. We used a modified version of GSC planner (Haslum et al. 2018) with best-first search and a domain-specific heuristic. The experiments were run on a Intel x86 2.1Ghz laptop equipped with 32Gb of RAM. We considered a timeout of 30 seconds (s typical timing under which UTC has to operate), and a memory limit of 10 Gb. As a comparison, we consider the plans generated by the approach presented by Vallati et al. (2016), consisting of UPMurphi (Penna et al. 2009). We validated the execution of the GSC solutions against the corresponding PDDL+ domain model and problem instance. Although our formulation of the problems is more abstract than the PDDL+ model, we can, still convert our solutions to PDDL+ plans, and validate them accordingly.

Results. Table 1 shows the plan quality, in terms of simulated time needed to reach the goal, of the solutions generated by the GSC approach and the PDDL+ approach. Considering the ability to generate a solution, it is worth noting that the GSC approach is capable of providing a solution for all the considered scenarios, while UPMurphi fails to solve the most challenging scenarios. The solution generated for scenario C2 could not be validated though, as one of the goals was not achieved. This is likely due to the level of abstraction of the proposed model.

Conclusion and Future Works

We presented an approach based on GSC planning to perform UTC via traffic signal optimisation. For future work we plan to explore more accurate traffic modelling, and investigate the reasons the GSC solutions obtained with our abstracted model fail to reach the goal when simulated in PDDL+. Finally, we plan to interface the proposed approach with traffic simulators, such as SUMO (Lopez et al. 2018), to better understand the performance of the proposed approach.

Acknowledgments

The work of the first and the last authors is partially funded by the grant MOSES, Bando interno 2020 Università di Trento "Covid 19". Lukáš Chrpa is supported by the OP VVV funded project CZ.02.1.01/0.0/0.0/16_019/0000765 "Research Center for Informatics". Mauro Vallati was supported by a UKRI Future Leaders Fellowship [grant number MR/T041196/1].

References

Bhatnagar, S.; Mund, S.; Scala, E.; McCabe, K.; McCluskey, L.; and Vallati, M. 2022. 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 2022*, 387–397.

Haslum, P.; Ivankovic, F.; Ramírez, M.; Gordon, D.; Thiébaux, S.; Shivashankar, V.; and Nau, D. S. 2018. Extending Classical Planning with State Constraints: Heuristics and Search for Optimal Planning. *J. Artif. Intell. Res.*, 62: 373–431.

Ivankovic, F.; Gordon, D.; and Haslum, P. 2019. Planning with Global State Constraints and State-Dependent Action Costs. In *ICAPS 2018*, 232–236. AAAI Press.

Ivankovic, F.; Haslum, P.; Thiébaux, S.; Shivashankar, V.; and Nau, D. S. 2014. Optimal Planning with Global Numerical State Constraints. In *ICAPS 2014*.

Lopez, P. A.; Behrisch, M.; Bieker-Walz, L.; Erdmann, J.; Flötteröd, Y.-P.; Hilbrich, R.; Lücken, L.; Rummel, J.; Wagner, P.; and Wießner, E. 2018. Microscopic Traffic Simulation using SUMO. In *Proceedings of ITSC*.

McCluskey, T. L.; and Vallati, M. 2017. Embedding Automated Planning within Urban Traffic Management Operations. In *ICAPS 2017*, 391–399. AAAI Press.

Penna, G. D.; 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*. AAAI.

Vallati, M.; Magazzeni, D.; Schutter, B. D.; Chrpa, L.; and McCluskey, T. L. 2016. Efficient Macroscopic Urban Traffic Models for Reducing Congestion: A PDDL+ Planning Approach. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 3188–3194.