Exploring the Trade-off Between Flexible and Deployable Models for PDDL+ Urban Traffic Control (Extended Abstract)*

Sandra Castellanos-Paez, Francesco Percassi, Mauro Vallati

School of Computing and Engineering, University of Huddersfield
{s.castellanos, f.percassi, m.vallati}@hud.ac.uk

Introduction

Traffic signal optimisation in automated urban control has been effectively addressed using the PDDL+ planning formalism (Vallati et al. 2016; McCluskey and Vallati 2017; Kouaiti et al. 2024), which also supports traffic simulation and what-if analysis to evaluate alternative scenarios (Bhatnagar et al. 2022; Percassi and Vallati 2024). This line of research leads to approaches that can efficiently generate high-quality signal plans with significant benefits in terms of congestion and emissions reduction, as demonstrated both in simulations and real-world deployments.

Existing models for automated planning-based traffic signal control can be roughly divided into two classes. (i) Models maximising the flexibility of the traffic controller, see, e.g., Percassi et al. (2023), where the planning system can dynamically adjust the duration of traffic stages without constraints on the overall cycles and on the differences between subsequent cycles. (ii) Models that guarantee the deployability of traffic signal control techniques also on legacy infrastructure, by forcing the AI approach to select the cycle configurations of traffic signals for the controlled junctions from a given pre-defined set (Kouaiti et al. 2024). Of course, both classes offer valuable properties and benefits: the extreme flexibility helps shed light on the potential gains achievable through investment in brand-new infrastructure, whereas the deployable approaches ensure the immediate usability of tools to maximise short-term impact.

To bridge the gap between different model classes and explore the trade-off between flexibility and deployability, we present the TRADE model. It enables the enforcement of key constraints required for deployability while preserving a level of flexibility that surpasses the capabilities of traditional traffic control infrastructure.

The TRADE Model

The proposed TRADE model focuses on adjusting, during a traffic signal cycle, the duration of stages by redistributing green time allocation within the cycle. This model can therefore be seen as a more constrained version of the Extend Reduce (EXRE) model introduced in Percassi et al. (2023). The cornerstone of the TRADE model is the notion of trading time between those stages of a junction that have not yet been executed during the current cycle. Trading time means that one stage reduces its allocated green time to allow another stage to extend its own by the same amount. This behaviour ensures that the overall duration of the cycle remains the same, as green time is moved between stages, while preserving constraints on minimum and maximum green time of each involved stage. It is worth noting that the trade is allowed only between future stages to (i) reduce the options available to the planning engine, and (ii) to avoid trades that would lead to an increase/decrease of the overall duration of the current cycle, as it would be the case if a stage already completed were to be part of the trade.

Further, the TRADE model supports the definition of a constraint (bound) B on the maximum number of time trades between stages at each cycle's junction. This can be used to control how much two subsequent cycles differ from each other, allowing the system's behaviour to be closer to that of traditional traffic control approaches, such as SCOOT, or to remain more flexible. Traditional SCOOT systems are generally bound to change a cycle by up to 5–8 seconds compared to the previous one. Finally, the parameter Gr defines the granularity, i.e., the amount of time traded with each action, to understand the trade-off between multiple small changes and fewer changes involving larger amounts.

The PDDL+ model and an example problem instance are available at https://github.com/smcastellanosp/TradeModel.

Experimental Analysis

Following previous works (Percassi et al. 2023; Kouaiti et al. 2024), we consider six scenarios from a major corridor of Huddersfield, UK. The scenarios were collected on January 26th (Wednesday) and January 30th (Sunday), 2022, using real-world sensors at different time slots: 8:30 am for the morning peak (*morn*), 12:30 pm for noon (*noon*), and 4:30 pm for the evening peak (*eve*).

This variation aimed to assess diverse traffic volumes and conditions. Here we focus on the goal formulation that works best for all approaches, i.e., the one where the goal is to move 350 vehicles as quickly as possible through the five main links of the corridor.

*We report results from our forthcoming work (Castellanos, Percassi, and Vallati 2025).

Model	Scenario						Σ
	26morn	*26noon*	*26eve*	*30morn*	*30noon*	*30eve*	
ExRe	**1.0**	0.96	**1.0**	**1.0**	**1.0**	0.99	**5.95**
Trade($Gr = 1,\ B = 5$)	0.98	**1.0**	0.95	0.99	0.99	**1.0**	5.91
Trade($Gr = 5,\ B = 1$)	0.97	**1.0**	0.94	0.99	0.99	0.99	5.88
Trade($Gr = 5,\ B = 5$)	0.93	0.97	0.98	0.98	0.96	0.95	5.77
Trade($Gr = 5,\ B = \infty$)	0.93	0.96	0.99	0.99	0.95	0.94	5.76
Trade($Gr = 10,\ B = 5$)	0.95	0.96	0.9	0.96	0.99	0.99	5.75
Trade($Gr = 10,\ B = 1$)	0.95	0.96	0.9	0.96	0.99	0.99	5.75
Trade($Gr = 10,\ B = \infty$)	0.95	0.96	0.9	0.96	0.99	0.99	5.75
Trade($Gr = 1,\ B = \infty$)	0.92	0.96	0.99	0.98	0.94	0.93	5.72
$\mathcal{H}$	0.8	0.99	0.92	0.46	0.85	0.88	4.9

Table 1: IPC Score of different TRADE instantiations, EXRE, and historical data $\mathcal{H}$ w.r.t. the ability to move vehicles through the corridor in 900 seconds. Higher values indicate better performance. B denotes the maximum number of trades per cycle, and Gr represents the granularity of trades.

We compare the plans generated by using the TRADE model with the strategies historically implemented by SCOOT in the reference region. Additionally, we consider plans obtained by using the EXRE model. All planning problems were solved using the same search scheme as in Percassi et al. (2023), specifically ENHSP with Greedy Best-First Search combined with a heuristic tailored for EXRE (Scala et al. 2020), which also works natively for TRADE.

All plans generated have been validated and simulated on historical data via the architecture designed by Bhatnagar et al. (2022). Experiments were conducted on a multi-core processor running at 2.30 GHz and 8 GB of RAM.

The TRADE model was instantiated with different parameter settings: $Gr \in \{1, 5, 10\}$ and $B \in \{1, 5, \infty\}$, excluding the variant with $Gr = 1$ and $B = 1$, as it would permit only minimal modifications between cycles, effectively limiting the system's ability to rapidly adapt to changing traffic conditions. In this context, a higher value of B increases the model's flexibility, while different values of Gr affect the ability to perform rapid, large changes. Notably, the variant with $B = \infty$ does not impose an explicit upper limit on the number of trades per cycle.

Table 1 compares the performance of the TRADE variants with the EXRE and historical data $\mathcal{H}$. Results are presented as quality IPC scores based on the total number of vehicles that moved through the corridor within a 900-second time window. This is used as a proxy to measure the ability of the generated strategies to facilitate the efficient movement of vehicles through the corridor. At first glance, it is noticeable that the TRADE generates strategies that are generally better than the historical ones, i.e., those implemented by the SCOOT system in operation.

The results indicate that the TRADE($Gr = 1, B = 5$) configuration achieves the highest overall performance among TRADE variants, suggesting that fine-grained time exchanges ($Gr = 1$) with moderate trade flexibility ($B = 5$) lead to more effective signal plans. This configuration, in some cases, can also outperform the more flexible EXRE.

Higher values of B do not always improve performance, and in some cases, excessive flexibility ($B = \infty$) results in lower scores. Our intuition here is that the bounds may provide valuable implicit knowledge regarding the suitable number of changes to perform.

While it might seem that large, sudden changes to traffic light cycles would improve traffic flow, our analysis shows the opposite. Larger Gr values, which allow for these drastic changes, actually led to slightly worse results. This is likely because the existing SCOOT system already provides reasonable cycle timings. In this regard, it is worth highlighting that a Gr value of 10 for a 90-second cycle means a single trade action alters the cycle by 20 seconds, which is more than 20% of the overall cycle duration.

The experimental analysis indicates that the TRADE model strikes a good balance between flexibility and deployability.

Conclusion

In this paper, we investigated the capabilities of TRADE, a PDDL+ knowledge model designed to explore the trade-off between flexibility and deployability in traffic signal optimisation. By integrating practical constraints into an extremely flexible model, TRADE achieves performance comparable to EXRE while ensuring a higher degree of compatibility with existing urban traffic infrastructure. Future work will focus on refining the TRADE model through the development of tailored heuristics for improved scalability and the investigation of junction-specific constraints.

Acknowledgments

Francesco Percassi and Mauro Vallati were supported by a UKRI Future Leaders Fellowship [grant number MR/Z00005X/1].

References

Bhatnagar, S.; Guo, R.; McCabe, K.; McCluskey, T. L.; Scala, E.; and Vallati, M. 2022. Leveraging Artificial Intelligence for Simulating Traffic Signal Strategies. In *ITSC*, 607–612. IEEE.

Castellanos, S.; Percassi, F.; and Vallati, M. 2025. Improved Models for Automated Planning-based Urban Traffic Control. In *MT-ITS*. IEEE.

Kouaiti, A. E.; Percassi, F.; Saetti, A.; McCluskey, T. L.; and Vallati, M. 2024. PDDL+ Models for Deployable yet Effective Traffic Signal Optimisation. In *ICAPS*, 168–177. AAAI Press.

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

Percassi, F.; Bhatnagar, S.; Guo, R.; McCabe, K.; McCluskey, T. L.; and Vallati, M. 2023. An Efficient Heuristic for AI-based Urban Traffic Control. In *MT-ITS*, 1–6. IEEE.

Percassi, F.; and Vallati, M. 2024. Leveraging AI Planning in a What-If Analysis Framework for Assessing Traffic Signal Strategies. In *ITSC*, 1330–1335. IEEE.

Scala, E.; Haslum, P.; Thiébaux, S.; and Ramírez, M. 2020. Subgoaling Techniques for Satisficing and Optimal Numeric Planning. *J. Artif. Intell. Res.*, 68: 691–752.

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 *AAAI*, 3188–3194. AAAI Press.