



How to Plan Roadworks in Urban Regions? A Principled Approach Based on AI Planning

Mauro Vallati¹ , Lukáš Chrpa^{2,3} , and Diane Kitchin¹ 

¹ University of Huddersfield, Huddersfield, UK
{m.vallati,d.kitchin}@hud.ac.uk

² Artificial Intelligence Center, Czech Technical University in Prague, Prague, Czech Republic

³ Faculty of Mathematics and Physics, Charles University in Prague, Prague, Czech Republic
chrpaluk@fel.cvut.cz

Abstract. Roadworks are required to keep roads in acceptable condition, and to perform maintenance of essential infrastructure. Road agencies are facing the problem of how to effectively plan frequent roadworks. In this paper, we exploit Automated Planning for roadworks planning. We introduce a planning domain model that allows us to plan a set of required roadworks, over a period of time, in a large urban region, by specifying constraints to be satisfied and suitable quality metrics. Our empirical analysis shows the suitability of the proposed approach.

Keywords: AI planning · Roadworks · Urban traffic management

1 Introduction

Despite the pressing need for efficiency, exacerbated by the current growth of urbanisation, there is a lack of approaches designed for supporting agencies in the decision-making process of creating a plan of roadworks for the controlled urban area. Research in the area is mainly focused on providing formal frameworks for the analysis of the impact of works [5], or on the evaluation of the process used by local authorities [6]. For this reason, roadwork plans are usually generated manually, without any explicit notion of “quality”, and a limited control on constraints and assessment of the impact of works on the managed network traffic.

Automated planning, which deals with the problem of finding a plan (a sequence of actions) that transforms the environment from an initial state to some desired goal state [7], is an effective tool for decision making. Plans consist of sequences of actions to be performed in order to achieve desired goals. Domain-independent planning requires only to specify a *planning domain model* and a *planning problem description* in a standard language such as PDDL [2], and then

to use one of the available generic planning engines to solve the problem and find a plan. This independence between reasoning and domain knowledge gives engineers a high level of flexibility since planning domain models and planning engines can be used as “black-boxes” that can be easily embedded into larger systems. Examples of embedding planning into larger frameworks include recent works in road traffic accident management [1], and urban traffic control [8].

In this paper, we propose a principled approach to use Automated Planning as an effective Decision Support tool in planning and scheduling roadworks in urban areas. We provide a formalisation of the roadworks planning problem, and we specify and develop a planning domain model that allows the representation of the constraints related to roadworks planning within large urban regions. The domain model is encoded in PDDL [2] that is widely supported by the large number of existing domain-independent planning engines. We empirically validate our approach by considering a urban region of Yorkshire (UK), and by modelling scenarios based on actual roadworks performed in the region.

2 Problem Specification

Here we consider the case in which a number of roadworks need to be performed in a controlled urban region. The controlled region, for which works have to be planned, is represented in terms of roads and areas. Roads r are connected to each other via junctions, and are grouped into areas a . The notion of area allows the effective and efficient encoding of the proximity of roadworks, as discussed later in this section. Each work w is described by a quadruple $\langle R, t_i, t_f, d \rangle$. R indicates the set of roads affected by the work w . The set can include a single road or more, according to the typology and size of w . For the sake of simplicity, here we consider that, if multiple roads are affected by the same work, they are grouped into the same area. d indicates the duration of the work, expressed in time units. Finally, the couple t_i, t_f represents the time window in which the work must be completed such that t_i is the earliest time unit in which the can start, and t_f represents the latest time unit in which the work can be completed, and thus $t_f > t_i$ must hold.

The time window must be at least as long as the duration of the considered work, i.e. $d \leq t_f - t_i$. Time is discretised, following a suitable discretisation for the considered region and type of works. In this paper we consider that each time unit corresponds to one week, but that can be easily adjusted (e.g. to one day) without changes to the model.

The problem is constrained by the maximum number of works that can be planned at the same time in a specific area. More formally, for each time unit t the following must hold:

$$\max W_{a_i} \geq |\{w_j \mid w_j \in W, (active(w_j, t) \wedge in_area(w_j, a_i))\}| \quad (1)$$

Where $\max W_{a_i}$ is the maximum number of works that can be active during the same time unit in the a_i area of the controlled region, $active(w_j)$ indicates whether the work w_j is planned or not for the considered time unit, and

$in_area(w_j, a_i)$ is used to specify if the work w_j has to be performed in the area a_i . A similarly-structured constraint is used to encode cases where, due to the fact a number of roadworks have been assigned to the same company, they can not be executed concurrently due to limited resources:

$$maxPW_{c_i} \geq |\{w_j \mid w_j \in W, (active(w_j, t) \wedge assigned_{t_o}(w_j, c_i))\}| \quad (2)$$

Where $maxPW_{c_i}$ is the maximum number of works that can be performed at the same time by a company c_i , $active(w_j)$ indicates whether the work w_j is planned or not for the considered time unit, and $assigned_{t_o}(w_j, c_i)$ represents the fact that w_j has been assigned to the company c_i .

The goal is achieved by planning all the required roadworks so that they can be completed within the corresponding defined time windows, while respecting the constraints presented in Eqs. 1 and 2. Planning is done by allocating works to subsequent time units, according to the expected duration.

3 Domain Model Specification

The conceptualisation of the problem requires to specify *object types*, which stand for classes of objects considered in the planning process, *predicates and numeric variables*, which describe the context, and *actions*, which allows the state to be modified. In our model we consider five main object types: *works*, *roads*, *time units*, *companies*, and *areas*. Works consist of the number of roadworks that have to be performed. Roads are the links of the considered urban region. Time units consist of the slots available for planning the roadworks. Companies represent the construction companies involved in the planning task. Areas indicate the various areas in which the controlled region has been divided.

The static part of the model, i.e. the aspects that do not change during the planning process, describes the road network and to some properties of the objects. We consider an abstract road network, where only roads affected by works, and main roads of the network are represented. A predicate *assigned* is used to denote that a work has been assigned to a company. A predicate *inArea* encodes that a road is included in an area of the controlled region. Similarly, a predicate *onRoad* denotes that the corresponding roadwork has to be performed on a road; multiple instances of this predicate can be used to indicate that a roadwork is affecting multiple roads. A numeric variable *maxWorks* is used to encode the maximum number of works that can be planned in the same time unit for a given area of the network. Similarly, another numeric variable *maxCompany* is used to encode the maximum number of parallel works a given company can deal with. Note that although the *maxWorks* and *maxCompany* values do not change during the planning episode, they might be initially set differently for different planning episodes. Time units objects are ordered using a *subsequent* predicate, which indicates that a time unit $t_i + 1$ is after a corresponding time unit t_i . Each time unit is also assigned a numeric *value*, which corresponds to its position in the problem. For instance, t_1 has a value of 1, and t_2 has a value of 2, etc. Values are exploited for optimisation purposes, and for encoding the time

horizon of each roadwork. Finally, numeric variables *startTime* and *deadline* represent the earliest initial and latest final time unit of each work.

The dynamic part of the model encodes the aspects which are under the control of the planning engine. The main task of the planning process is to allocate roadworks to suitable time units so that constraints are satisfied. The *toStart* predicate is used to indicate that the corresponding roadwork has yet to be started, i.e. the planning engine did not start allocating it yet. The numeric variables *duration* denotes the “remaining” duration of a roadwork, in terms of time units. Starting from the actual duration of the roadwork in object, the variable is reduced every time the work is allocated to a time unit. A duration value of 0 indicates that the corresponding work has been completely allocated. The numeric variables *companyWork* and *parallelWork* are used to count, respectively, the number of works assigned to a company and the number of works in the same area, that are planned in the same time unit. Finally, *initiated* is a general numeric variable which is exploited for optimisation purposes. It stores the sum of the differences between the planned starting time and the earliest possible starting time, of all the considered roadworks.

One roadwork cannot be allocated twice to the same time unit; this is a very common constraint for numerical planning and has been encoded using well-known PDDL constructs. It is therefore omitted from the following description of actions.

The actual plan of road works can be generated by a planning engine using the two modelled actions:

- **WorkStart**(w,r,t,a,c) – identifies a time unit t as the starting time of a road work w , that has been assigned to a company c ; the work has to be done on road r , which is part of area a . As preconditions, beside requirements on location of roadworks and areas, *toStart* (w) must hold, and the deadline of the work must be satisfied by starting at time unit t . Effects include the fact that the work is considered as started (i.e., the *toStart* predicate is falsified), duration is reduced by one time unit, and the number of works in the area and works performed by company c in time unit t are increased by one.
- **WorkContinue**(w,r,t_b,t_s,a,c) – a work w that has already been allocated to a time unit t_b is allocated to the subsequent time unit t_s (i.e., $t_s = t_b + 1$). The preconditions and effects are the same as the *WorkStart* action except *toStart* (w) is replaced by *onGoingWork*(w,t_b) in the precondition.

The PDDL encoding of two different actions gives a high degree of flexibility to the model, as works with significantly different durations can be considered. Furthermore, the step-by-step approach which is forced by requiring that a different action is used to allocate a work to each time unit guarantees a higher level of control, as at each step all the constraints can be checked.

In addition, to the described actions, two PDDL constraints on the planning trajectory are enforced. Such constraints are used to ensure that the limits on the maximum number of works per area and per company, allocated to one time slot, are respected. The model is available at <http://bit.do/roadworkplan>.

4 Evaluation of the Approach

The Planning domain model and problems have been encoded in PDDL 2.1. For solving such problems, we selected three planning search techniques implemented by the Metric-FF planner [3]: Fast Forward (FF), Weighted A* (WA*), and Enforced Hill Climbing (EHC). The first technique ignores the quality of the solution found, but focuses on generating a solution as fast as possible, while the considered WA* and EHC approaches do take the quality of plans into account and try to optimise the generated solution. All the techniques are guided by the well-known delete-relaxation heuristic: such heuristic ignores all the negative effects of actions, i.e. it only reasons by considering positive (additive) effects. In this context, quality is measured in terms of the average “delay” of planned starting time of roadworks, with regards to the possible earliest starting time. Better quality corresponds to plans where works, on average, are started as early as possible. This is a key indicator of quality because in a typical urban context, all the roadworks which need to be planned are usually essential, and should be completed as soon as possible.

Experiments have been performed on a system equipped with 2.0 GHz Intel i7-3667U Processors, 8 GB of RAM and Linux operating system. The VAL tool [4] has been used for validating the generated plans, in order to check their correctness with regards to the designed domain model, and also to identify the presence of bugs or flaws in the model.

Our analysis has been focused on the northern part of the Kirklees urban region, that includes 11 different areas.

4.1 Empirical Results

In order to collect informative and realistic data, we analysed the number of roadworks being executed in the considered region over a period of time of 2 months, during 2017. According to the observed data, we synthetically generated a set of instances, with the following main characteristics: number of roadworks to be planned ranging from 5 to 20; plan horizon of 6 months, roadworks duration ranging between 1 and 10 weeks. Start time and deadline values have been randomly generated. For each considered number of roadworks to be planned, we randomly generated three instances with different durations and start times. As a general deadline, we required all the works to be completed by the end of the modelled 6-months period. Presented results are averaged on the three instances, in order to take into account noise and variability due to the randomised aspects.

The results of this first set of experiments, designed for testing the fact that the proposed system is effectively able to plan the roadworks required in the controlled region, are presented in Table 1. Unsurprisingly, the use of the FF planning techniques leads to plans of poor quality, with regards to the quality metric considered. This is due to the fact that this approach does not exploit any cost optimisation for generated plans. FF tends to plan roadworks 17 weeks later, on average, than the earliest starting time. By analysing the generated plans, we noticed that the works are frequently scheduled in order to finish exactly by

Table 1. Runtime (CPU-time seconds) needed by FF, WA*, and EHC for providing a plan for an increasing number of roadworks in the considered area, over a 5-month period of time. Quality of generated plans is measured as the average delay (time units) with regards to the first possible start time of the work. Bold indicates best results.

# works	Runtime			Plan quality		
	5	10	20	5	10	20
FF	0.1	0.3	1.6	18.3	18.9	19.0
WA*	0.1	0.9	7.8	1.4	1.0	1.2
EHC	0.1	0.3	1.5	18.2	18.4	18.5

the deadline, even if it would be possible to complete them before. Surprisingly, plans of poor quality are also generated by the EHC algorithm, which instead exploits cost optimisation. Our intuition is that this behaviour is due to the search space topology, that does not allow the enforced hill climbing approach to identify promising areas to visit in order to increase the quality of an initially identified solution. Instead, the use of a weighted-A* search algorithm allows to find high quality plans: works are started usually at the earliest possible time, except in cases where the constraints on maximum parallel works that can be performed by companies or in specific areas could not be satisfied by starting works earlier.

All plans are generated extremely fast by the considered techniques. WA* is the slowest approach, and requires less than 8 CPU-time seconds to solve the most complex scenario considered in this analysis. The other approaches require less than 2 CPU-time seconds.

Generated plans have been validated using the well-known VAL tool. This guarantees that plans are valid with regards to the exploited domain model. However, this sort of validation does not provide any insights into the actual usefulness of the plans in the real-world applications. In order to check this aspect of the validity of generated plans, they have been visually inspected by transport experts, who confirmed their overall quality and that they look similar to plans one would expect from an expert-generated solution.

4.2 Time Horizons

The size of the time window is an extremely important aspect, as it can significantly affect the shape of the search space: here we assess the impact of time windows on the performance of the proposed planning-based approach. Focusing on the 10 roadworks scenario, we generated four planning problems by varying (i) the duration of roadworks and (ii) the size of the time windows. Works can have a duration of 1 or 10; the time window can be exactly the duration of the work, or can cover the whole modelled period. In other words, in the first problem all the works have duration 1, and the time window is 1 for all of them, in the second problem all the works have duration 10, and the time window is 10

for each of them, etc. Results indicate that the duration of the considered works has no impact on the performance of the considered planning approaches. On the other hand, the size of the time window has a remarkable impact on performance. The more constrained the time windows are, the easier it is for the three considered planning approaches to find a solution: in the more constrained case, all the 3 approaches can find a plan in approximately 0.5 CPU-time seconds. Moreover, the more constrained the problem is, the more similar the quality of the generated solutions is. Conversely, larger time windows lead to a higher computational time (up to 10 CPU-time seconds for the WA* approach), and a remarkable variability in terms of quality of the generated solution plans.

5 Conclusion

In order to assist road agencies and authorities in the critical duty of planning essential roadworks, here we investigated the use of automated planning techniques as an effective decision support tool. The main contributions of this work are: (i) a formalisation of the roadworks planning problem; (ii) a PDDL model encoding of the domain; and (iii) a thorough experimental analysis, that considers realistic scenarios and compared plans generated using three different planning approaches, embodied in the same domain-independent planning engine. The performed experimental analysis demonstrates the extent to which the proposed approach is able to efficiently and effectively create a roadworks plan that satisfies the identified constraints, and follows a notion of quality. Future works will focus on including different degrees of uncertainty in the models, and evaluating the approach in different urban regions.

Acknowledgements. L. Chrpa was partially funded by the Czech Science Foundation (project no. 17-17125Y). M. Vallati was partially supported by the EPSRC grant EP/R51343X/1 (AI4ME).

References

1. Chrpa, L., Vallati, M.: On the exploitation of automated planning for efficient decision making in road traffic accident management. In: Proceedings of CDC (2016)
2. Fox, M., Long, D.: PDDL2. 1: an extension to PDDL for expressing temporal planning domains. *J. Artif. Intell. Res.* **20**, 61–124 (2003)
3. Hoffmann, J.: The metric-ff planning system: translating “ignoring delete lists” to numeric state variables. *J. Artif. Intell. Res.* **20**, 291–341 (2003)
4. Howey, R., Long, D., Fox, M.: VAL: automatic plan validation, continuous effects and mixed initiative planning using PDDL. In: Proceedings of ICTAI (2004)
5. Huerne, H.T., Berkum, E.V., Hermelink, W.: A multi objective framework to optimise planning of road works. In: 5th Eurasphalt and Eurobitume Congress (2012)
6. Hussain, R.S., Ruikar, K., Enoch, M.P., Brien, N., Gartside, D.: Process mapping for road works planning and coordination. *Built Environ. Proj. Asset Manag.* **7**(2), 157–172 (2017)

7. Ghallab, M., Nau, D., Traverso, P.: Automated Planning Theory and Practice. Elsevier Science, Amsterdam (2004)
8. McCluskey, T.L., Vallati, M.: Embedding automated planning within urban traffic management operations. In: Proceedings of ICAPS (2017)