Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

Partial Satisfaction Planning Under Time Uncertainty with Control on When Objectives Can Be Aborted

Sylvain Labranche and **Éric Beaudry**
Université du Québec à Montréal
320 Sainte-Catherine Est Street, Montreal, QC H2X 1L7
(514) 987-3000 # 7890
http://gdac2.uqam.ca/~sylvain/
labranche.sylvain@courrier.uqam.ca, eric.beaudry@uqam.ca

Introduction

Classical planners consider a goal as a non-separable set of objectives, i.e. a goal can be either satisfied or not satisfied in a particular state. However, situations may arise where it is impossible to satisfy every single objective, either because of limited resources, limited time or even mutual exclusions. For these situations, classical planners simply return no plan.

This limitation of traditional planners motivated a branch in planning, which is often called 'over-subscription planning' or 'partial satisfaction planning'(PSP), where a goal is considered as a set of soft objectives, i.e. objectives that may be unsatisfied in a final state.

A typical example of an application which requires PSP is the planning activity for Mars rovers that must gather scientific data. Since such an expedition represents important costs, time must be spent as effectively as possible. Therefore, it makes sense to input the rover more objectives than it is possible to satisfy and leave to automated planning the task of finding the best subset.

When planning under uncertainty, the decision to select or not an objective must sometimes be postponed at execution time, because the effective duration of actions may dynamically influence the optimal achievable subset of objectives. A solution would be to allow dynamic abortion of objectives.

As well as allowing the abortion of objectives, it can also be required to force a deadline on the abortion time. In some situations, it is practical to know that past a certain time, an objective cannot be aborted anymore and will be satisfied. This is the case when planning and execution overlap, like for Mars rovers missions where during execution, planning for the next day needs to already be started. Is it also desirable in situations where commitment before a certain time is needed, like in the delivery business. With deadlines on abortion, it would be possible to know sooner which objectives will be satisfied, because past abortion deadlines soft objectives would become hard.

We present a new approach to solve PSP problems under uncertainty on the duration of actions. Our approach is related to the compilation of soft objectives proposed in (Keyder and Geffner 2009), but gives the possibility to plan the abortion of objectives by introducing special actions that explicitly abort objectives. Our approach also enables to set deadlines on when the abortion can be done. We experiment our method in the Action Contingency Time Uncertainty (ActuPlan) planner (Beaudry, Kabanza, and Michaud 2010; 2012). The current abstract is a short version of a paper appearing in the proceedings of the 27^{th} Canadian Conference on Artificial Intelligence (Labranche and Beaudry 2014).

State of the art

A planning problem is given by $\mathcal{P} = \langle \mathcal{F}, \mathcal{I}, \mathcal{A}, \mathcal{G} \rangle$, where $\mathcal{F}$ is a finite set of state variables, $\mathcal{I}$ an initial state, $\mathcal{A}$ a finite set of actions and $\mathcal{G} = \{g_1, g_2, .., g_n\}$ a goal consisting of a finite set of objectives. Every action in $\mathcal{A}$ has a cost, which is given here by the statistical distribution of the expected time it will take to apply it. Objectives in $\mathcal{G}$ can have deadlines, after which the objective cannot be satisfied anymore.

With Net-Benefit approaches, utilities are assigned to objectives as a way to represent how important they are compared to each other. Then, the subset of goals yielding the highest net benefit (utilities of objectives satisfied minus the costs to achieve them) is generated.

These methods have been criticized. A recent research note has showed that soft objectives can be compiled away by translating a PSP problem into a classical planning problem (Keyder and Geffner 2009). Empirical results also suggest that this translation based approach yields better results.

However, when planning under action duration uncertainty, some objectives might need to be aborted during execution. For example, deadline satisfaction may become impossible. A way to know when objectives should be aborted is then needed, in order to use time in an optimal way.

Some work has already been done to solve PSP under uncertainty on resource consumption (Coles 2012), where the idea is to compute a pessimistic linear plan and branch it with optimistic plans that increase the expected utility. Their method is somehow dual to ours.

Using Abort Actions with Deadlines

We present an approach that extends ideas from (Beaudry, Kabanza, and Michaud 2012; 2010; Coles 2012; Keyder and Geffner 2009) in order to create a conditional plan with branching points that can be actions that abort objectives.

These branching points are given by intervals where, during a certain time, it is more profitable to apply a certain action while, once the interval is over, it is better to apply another one. Therefore, the conditional plan offers different action choices based on execution time. This behaviour leads to the creation of actions that can abort objectives, giving the possibility to modify the goal dynamically. In order to enable some control, these actions have deadlines.

ActuPlan

We extend the ActuPlan planner (Beaudry, Kabanza, and Michaud 2012; 2010) to support partial satisfaction planning. ActuPlan is a planner that handles action contingency and actions with uncertain durations. The planner uses Bayesian networks to represent time continuously. It allows management of time variables independently of state representation, greatly reducing state space complexity compared to a planner that uses a discrete time representation.

Abort Actions

Based on the problem definition previously given, the abort actions are introduced as follows. To the actions set is added a finite set of size $|\mathcal{G}|$ in which $Abort(g_i)$ is an action with precondition that g_i is in $\mathcal{G}$ and with effect of removing g_i from $\mathcal{G}$. To each $Abort(g_i)$ is assigned an abortion cost, representing how costly it is to not satisfy objective i. Compared to Net-Benefit, the cost to abort an objective can be its utility. A deadline is also assigned to each $Abort(g_i)$, representing the maximum time when objective i can be aborted. States are now modelled with an addition of $|\mathcal{G}|$ state variables, each indicating if objective g_i was aborted or not. Therefore, as in classical planning, a final state s_f is a state where every g_i from $\mathcal{G}$ is satisfied and a plan is a partially ordered sequence of actions leading from $\mathcal{I}$ to s_f.

This approach also naturally handles problems where there are both hard and soft objectives. An objective with an abortion cost of ∞ cannot be aborted and therefore is considered hard. Any objective with abortion cost $< \infty$ is soft and can be aborted. Once the deadline of an abort action is past, the corresponding objective becomes hard and must be satisfied. This approach can therefore be seen as a dynamic compilation of soft objectives

Empirical evaluation

We compared up to now two approaches using abort actions: a conditional planner for which there are no deadlines on objective abortion (Beaudry, Kabanza, and Michaud 2010) and one for which there is.

Adding deadlines on abort actions makes the problem more complex, but more information can be extracted from the plans returned. Results from smaller problems show no significant difference in CPU times, but for bigger problems is it indeed more expansive to compute a conditional plan when there are deadlines on objective abortion.

Results also tend to show that conditional plans with and without abortion deadlines may have the same expected cost. This is normal and means that no abortion deadline triggered an objective to become hard, thus forcing its satisfaction. Therefore, problems where this situation arises need to be created and tested. These results remain interesting because they show that is it possible to have the information on the maximum time when objectives will be aborted without necessarily incurring additional execution costs.

In general, the method is not yet scalable to bigger problems, because the abort actions approach in itself makes the state space quite larger. Ways to reduce this state space explosion will be investigated.

Conclusion

In this paper, we presented an extension of PSP planning to deal with domains having actions with uncertain durations. We introduced new actions to explicitly abort objectives, allowing a dynamic satisfied objectives subset creation. These actions are used at branching points in a conditional plan in order to make the best choice based on the actual execution time. The advantage of our approach is that it supports deadlines on when abortion can be done. These deadlines successfully allow control on the maximum time when an objective can be aborted.

Future work includes the application of PSPs modelled with our approach in a serious real estate broking game (Labranche et al. 2014). PSPs will be used to model non playable brokers behaviours (NPB) who are in competition with the learning player. Therefore, every NPB will have its actions planned individually, as opposed to common behaviour with rules-based scripts We hope that this approach will make the NPBs more realistic and responsive.

References

Beaudry, E.; Kabanza, F.; and Michaud, F. 2010. Planning for concurrent action executions under action duration uncertainty using dynamically generated bayesian networks. In *Proceedings of the 20th International Conference on Automated Planning and Scheduling*, 10–17.

Beaudry, E.; Kabanza, F.; and Michaud, F. 2012. Using a classical forward search to solve temporal planning problems under uncertainty. In *Proceedings of the Problem Solving Using Classical Planners Workshop of the 26th AAAI Conference*, 2–8.

Coles, A. J. 2012. Opportunistic branched plans to maximise utility in the presence of resource uncertainty. In *Proceedings of the 20th European Conference on Artificial Intelligence*, 252–257.

Keyder, E., and Geffner, H. 2009. Soft goals can be compiled away. *Journal of Artificial Intelligence Research* 36:547–556.

Labranche, S., and Beaudry, E. 2014. Partial satisfaction planning under time uncertainty with control on when objectives can be aborted. In *To appear in the Proceedings of the 27th Canadian Conference on Artificial Intelligence*. Springer.

Labranche, S.; Sola, N.; Callies, S.; and Beaudry, E. 2014. Using partial satisfaction planning to automatically select npcs' goals and generate plans in a simulation game. Submitted at the IEEE Conference on Computationnal Intelligence in Games.