A Hill-Climbing Approach for Planning with Temporal Uncertainty

Janae N. Foss **Nilufer Onder**
Michigan Technological University, Department of Computer Science
1400 Townsend Drive, Houghton, MI 49931
{jnfoss,nilufer}@mtu.edu

Abstract

We present a hill-climbing algorithm to solve planning problems with temporal uncertainty. First an optimistic plan that is valid when all actions complete quickly is found. Then temporal reasoning techniques are used to determine when the plan may fail. At time points that cause an unsafe situation, contingency branches are inserted. We describe our planner PHOCUS-HC, give preliminary results, and discuss future work.

Introduction

Constructing optimal plans for domains with uncertainty is a challenging problem. Such domains often include uncertain discrete or continuous effects, oversubscribed goals, and possibly parallel actions. Many planners have been built that prepare contingency plans when actions may affect the world in uncertain ways. However, the problem of generating plans where there is uncertainty in action duration has been studied less. We approach this problem by finding temporally contingent plans, i.e., plans with branches that are based on the duration of actions at execution time, using a hill-climbing algorithm. We take an optimistic approach by first finding a plan that is valid when all actions complete quickly, and then insert temporal contingency branches at time points that cause unsafe situations.

As an example, consider the problem of traveling from home to a conference. One solution plan is to drive to the airport, fly to the destination city, take a shuttle to the conference venue, and finally register for the conference. Another solution plan could involve taking a taxi instead of a shuttle to the venue. Assuming the objective is to minimize money spent, the plan with the shuttle action would be preferred. However, the taxi may be faster than the shuttle and due to constraints on conference registration time, if the flight takes too long there may only be enough time for the more expensive taxi option. To always have a safe plan, and be able to save money when possible, our approach would generate a temporally contingent plan: drive to the airport, fly to the destination, take the shuttle if there is enough time, otherwise take the taxi, and register for the conference.

We have several contributions: (1) we define the notion of temporally contingent plans, (2) we provide a hill-climbing algorithm that uses efficient temporal reasoning techniques to insert branches based on time rather than world conditions, (3) we show that plans with maximal expected utility can be generated in this framework by using our implemented planner PHOCUS-HC, (4) we provide example domains including a disaster evacuation domain which can benefit from our approach.

Planning with Temporal Uncertainty

We define temporal uncertainty by assigning each action a closed interval duration [*min-d*, *max-d*], where *min-d* and *max-d* denote the minimum and maximum possible durations for the action. It is assumed that the agent has no control over the duration of the action and the actual duration is only known after the action completes execution. Also, actions can be defined with temporal constraints to restrict their execution to given times or during given time windows. Respecting these bounds, a *planning problem* is defined as a quadruple <**D**, **I**, **G**, **M**>, where **D** is a domain description that lists the available actions (including interval durations and temporal constraints), **I** is a description of the initial state, **G** is a description of the goals, and **M** is a plan metric that represents the objective function for ranking plans.

PHOCUS-HC uses a Just-In-Case style algorithm (Drummond, Bresina, & Swanson 1994) where a seed plan is generated, the points where it is likely to fail are located, and then contingency branches are inserted (when available) at those points (Fig. 1). To generate the seed plan (line 0 in Fig. 1), *min-d* is assigned as the duration of each action to remove uncertainty at planning time. This allows generation of plans using any planner that can handle durative actions, timed initial literals, and optimize based on an objective function. A plan *P* returned by such a planner will be temporally deterministic. Our algorithm factors temporal uncertainty back in by converting *P* to a directed, edge-weighted graph called a *distance graph DG*, thus expressing *P* as a simple temporal network (STN) (Dechter, Meiri, & Pearl 1991). This conversion is described in detail in our earlier work (Foss & Onder 2005).

Since *DG* contains all temporal constraints given in the domain, it can be used to determine when *P* becomes unsafe (line 10) (Dechter, Meiri, & Pearl 1991). In the loop

---

```
PHOCUS-HC (D, I, G, M)
 1: $P_0 \leftarrow$ GENERATE-SEED-PLAN (D, I, G, M)
 2: $P_{current} \leftarrow P_0$
 3: loop do
 4:     $DG \leftarrow$ CONSTRUCT-DISTANCE-GRAPH($P_{current}$,D,I)
 5:     if SAFE-PLAN ($P_{current}$, DG, D, I, G, M) return $P_{current}$
 6:     $P_{next} \leftarrow$ MAKE-PLAN-SAFE ($P_{current}$, DG, D, I, G, M)
 7:     if $P_{next}$ is null return failure
 8:     $P_{current} \leftarrow P_{next}$

MAKE-PLAN-SAFE (Plan P, DistanceGraph DG, D, I, G, M)
 9: for i = downto 1 in P
10:     maxAllowedDuration $\leftarrow$ SHORTEST-PATH-DISTANCE($s_i$, $e_i$, DG)
11:     if maxAllowedDuration $\geq$ max-d of i
12:         DG, D $\leftarrow$ DG, D updated to constrain i to always require max-d of i
13:         DG, D $\leftarrow$ DG, D updated to constrain i to always start at latest
                possible time that allows max-d of i
14:     else
15:         newMinDuration $\leftarrow$ maxAllowedDuration + 1
16:         $D_{mod} \leftarrow D$ modified so that action i requires newMinDuration
17:         $P_{new} \leftarrow$ generate plan with $D_{mod}$
18:         if P and $P_{new}$ have the same steps through step i
19:             return a contingency plan created out of P and $P_{new}$
20:         else
21:             return $P_{new}$
```

Figure 1: The PHOCUS-HC algorithm.

that contains line 10 the plan is analyzed one step at a time to find the last action which makes the rest of the plan unsafe. If an action is found to be safe in line 11 the domain and the corresponding distance graph are updated to provide topmost flexibility to the earlier actions (lines 12,13). Otherwise, the domain and problem are modified so that the action minimally uses the duration that causes the plan to fail and a new plan meeting the new constraints is sought for at lines 15 through 17. If the new plan shares a head with the current plan, a contingency plan is formed (step 19). Otherwise, the new plan is returned and becomes a new seed plan (step 21). The process is repeated until all steps have been analyzed.

Experimental Results

To the best of our knowledge, there are no planners that prepare contingency branches based on time. We therefore designed our experiments to show that our algorithm works and to help identify the ways in which it can be improved. We used LPG-*TD* (Gerevini *et al.* 2004) for generating seed plans and branches. Tests were run with both sequential and parallel versions of the conference domain. In the parallel version the agent must grade exams and read papers while traveling. Tests were also run with an evacuation domain having one bus and one helicopter available to evacuate a school and a hospital with the objective of maximizing lives saved. Multiples trips with the helicopter and bus were required and if early trips took too long, some lives were lost.

For comparison, we generated plans using LPG-*TD* assuming minimum duration, maximum duration and average duration for actions. Expected utility (EU) was calculated for each plan as $\sum_b probability(b) \times utility(b)$ where b is a complete branch or path that can be taken in a plan. Table 1 shows a comparison of the EU for the three different plans generated by LPG-*TD* and the conditional plan generated by PHOCUS-HC for six problems. More results and discussion can be found in (Foss & Onder 2006).

problem	min	avg	max	cond
conf-2	-467.33 *(0.03)*	450.00	450.00	451.00
conf-3	-348.51 *(0.02)*	161.11 *(0.55)*	280.00	541.66
p-conf-2	-437.33 *(0.03)*	480.00	480.00	481.00
p-conf-3	-318.51 *(0.02)*	191.11 *(0.55)*	310.00	571.66
p-evac-2	84.89 *(0.31)*	84.89 *(0.31)*	80.00	98.67
p-evac-3	130.50 *(0.05)*	136.90 *(0.09)*	137.00	137.96

Table 1: EU for min, avg, max, and conditional plans for each problem. When probability of success is less than 1 it is given in parenthesis after EU. Number in problem name denotes number of branches in conditional plan. Problems beginning with p contain parallel steps.

Conclusions and Future Work

We have presented a framework for characterizing and directly dealing with temporal uncertainty and have implemented our hill-climbing approach in a planner called PHOCUS-HC. We have tested PHOCUS-HC on many problems from two domains and provided preliminary results. In addition to generating contingency branches, our hill-climbing approach has the advantage of being able to replace the entire seed plan when adding a contingency branch is not possible, or when starting with a new seed plan yields a better expected utility. In the current version of PHOCUS-HC, a uniform distribution is assumed over all uncertain action durations. In the future we plan to further develop the implementation to allow user specified distributions. Also, the current implementation always searches until a plan with 100% safety is found. We plan to improve PHOCUS-HC so that the user can choose the level of safety that is required.

Acknowledgements. We would like to thank Ryan S. Richards for helping implement the parser for PHOCUS-HC. Janae N. Foss' research was supported by the Harriett G. Jenkins Predoctoral Fellowship Program and a grant from the Michigan Council of Women in Technology.

References

Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artificial Intelligence* 49:61–95.

Drummond, M.; Bresina, J.; and Swanson, K. 1994. Just-in-case scheduling. In *Proc. 12th National Conf. on Artificial Intelligence*, 1098–1104.

Foss, J., and Onder, N. 2005. Generating temporally contingent plans. In *IJCAI 2005 Workshop on Planning and Learning in A Priori Unknown or Dynamic Domains*.

Foss, J., and Onder, N. 2006. A hill-climbing approach for planning with temporal uncertainty. Technical Report CS-TR 06-02, Computer Science Dept., Michigan Technological University.

Gerevini, A.; Saetti, A.; Serina, I.; and Toninelli, P. 2004. Planning in PDDL2.2 domains with LPG-*TD*. In *International Planning Competition booklet (ICAPS-04)*.