Automated Planning in Temporal Domains: Some Recent Advances and Current Research Topics

Alfonso Gerevini
Dipartimento di Elettronica per l'Automazione
Università degli Studi di Brescia
Via Branze 38, 25123 Brescia, Italy
gerevini@ing.unibs.it

Automated planning is a central area of artificial intelligence, involving the design of languages and computational models for reasoning about actions, change and time. In domain-independent planning, a planning problem is specified by the description of an initial world state, a set of desired goals to achieve and a set of possible actions or action schemata (domain operators). A solution of a planning problem is a (partially) ordered set of actions forming a valid plan, whose execution in the initial state transforms it into a world state where the problem goals are satisfied.

Many computational models (algorithms, data structures, heuristics) and languages for planning have been developed since the beginning of AI. Current planning systems support languages that are much more expressive than the simple STRIPS language largely studied in the past. In this talk, we focus on planning in temporal domains specified by the Planning Domain Description Language (PDDL) [2, 3, 5, 9], which is an important standard reference for the automated planning community. In particular, we will consider the last two versions of this language, PDDL2.2 [2] and PDDL3 [5].

The temporal information in a PDDL2.2 planning problem concerns: the static or dynamic (state-dependent) duration of the plan actions; the relative time when the action conditions have to be satisfied and the effects occur (at the beginning, at the end, or throughout the execution of the action); possible predictable exogenous events happening at known times, on which certain action conditions may depend; possible deadlines for the problem goals or plan quality metrics requiring minimization of the plan makespan. Moreover, the plan actions are partially ordered by the causal structure of the plan and some constraints that prevent the overlapping of interfering actions.

A planner has to take this temporal information into account both to guarantee the correctness of the synthesized plan, to generate a plan of good or optimal quality (makespan), and to use effective search heuristics for fast planning.

After a brief description of the relevant features of PDDL2.2, we will give an overview of a recent approach [7, 8] for solving the class of temporal planning problems that can be expressed in this language. In this approach, the causal structure of the plan is managed by a graph-based representation, temporal information is managed by efficient constraint-based reasoning, and plan search is guided by stochastic local search. These techniques have connections to other prominent approaches to AI planning and automated reasoning, such as Blum and Furst's planning graphs [1], disjunctive temporal reasoning problems [13, 14], and the Walksat procedure for satisfiability checking [12]. Our methods are implemented in a fully-automated domain-independent planner, called LPG, that was awarded at the 2002 and 2004 international planning competitions [10, 11].

The last part of the talk will be dedicated to a brief description of the temporal features in the new version of PDDL, PDDL3 [5, 6], which was used to develop the benchmark problems of the 2006 planning competition [4]. PDDL3 extends the previous versions of the language with some new constructs that can better characterize plan quality by allowing the user to express strong and "soft" constraints about the structure of the desired plans, as well as strong and soft problem goals. Plan trajectory constraints are particular linear temporal logic formulae expressing constraints on possible actions in the plans and on intermediate states reached by the plans. Soft goals and constraints are "preferences" that we desire to satisfy in order to generate a good-quality plan, but that do not necessarily have to be achieved for the correctness of the plan. Dealing with plan (strong or soft) trajectory constraints poses a new challenge to fully automated planning. Some of the planners that entered the 2006 planning competition address this research topic and are capable of solving a large set of the competition benchmark problems.

References

[1] A. Blum and M. Furst. Fast planning through planning graph analysis. *Artificial Intelligence*, 90:pp. 281–300, 1997.

[2] S. Edelkamp and J. Hoffmann. PDDL2.2: The language for the classic part of the 4th international planning competition. Technical Report 195, Institut für Informatik, Freiburg, Germany, 2004.

[3] M. Fox and D. Long. PDDL2.1: An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research (JAIR)*, 20:pp. 61–124, 2003.

[4] A. Gerevini, Y. Dimopoulos, P. Haslum, and A. Saetti. The fifth International Planning Competition, 2005. http://ipc5.ing.unibs.it/index.html.

[5] A. Gerevini and D. Long. Plan constraints and preferences in PDDL3. Technical Report 2005-08-47, Università degli Studi di Brescia, Dipartimento di Elettronica per l'Automazione, 2005.

[6] A. Gerevini and D. Long. Preferences and soft constraints in PDDL3. In *Proc. of the Workshop on Preferences and Soft Constraint in Planning*, pages 46–53, Cumbria, UK, 2006. ICAPS 2006.

[7] A. Gerevini, A. Saetti, and I. Serina. Planning through stochastic local search and temporal action graphs. *Journal of Artificial Intelligence Research (JAIR)*, 20:pp. 239–290, 2003.

[8] A. Gerevini, A. Saetti, and I. Serina. An approach to temporal planning and scheduling in domains with predictable exogenous events. *Journal of Artificial Intelligence Research (JAIR)*, 25:187–231, 2006.

[9] M. Ghallab, A. Howe, C. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL - the planning domain definition language. Technical Report CVC TR98-003/DCS TR-1165, Available at http://cs-www.cs.yale.edu/homes/dvm/, 1998.

[10] J. Hoffmann and S. Edelkamp. The deterministic part of IPC-4: An overview. *Journal of Artificial Intelligence Research (JAIR)*, 24:519–579, 2005.

[11] D. Long and M. Fox. The 3rd international planning competition: Results and analysis. *Journal of Artificial Intelligence Research (JAIR)*, 20:pp. 1–59, 2003.

[12] B. Selman, H. Kautz, and B. Cohen. Noise strategies for improving local search. In *Proceedings of the Twelfth National Conference of the American Association for Artificial Intelligence (AAAI-94)*, pages 337–343, Seattle, WA, 1994. Morgan Kaufmann.

[13] K. Stergiou and M. Koubarakis. Backtracking algorithms for disjunctions of temporal constraints. *Artificial Intelligence*, 120(1):pp. 81–117, 2000.

[14] I. Tsamardinos and M. E. Pollack. Efficient solution techniques for disjunctive temporal reasoning problems. *Artificial Intelligence*, 151(1-2):pp. 43–89, 2003.