

Systematic Nonlinear Planning

Commentary by Daniel S. Weld¹

To appear in AI Magazine 2010

The field of automated planning seeks to develop fast methods for generating programs of action that allow an agent to achieve goals or maximize rewards. Initially, researchers made many simplifying assumptions, creating the *classical* planning problem: produce a sequence of atomic actions which will achieve a logically specified goal in a completely known world without uncertainty. McAllester and Rosenblitt's (M&R) paper "Systematic Nonlinear Planning,"¹ presented nineteen years ago at AAAI-91, had two major impacts on the field: 1) an elegant algorithm and 2) endorsement of the *lifting* technique. The paper's biggest impact stems from its extremely clear and simple presentation of a sound and complete algorithm (known as SNLP or POP) for classical planning. While it is easy to define such an algorithm as search through the space of world states, SNLP is a "partial-order" planner, meaning it searches the space of partially-specified plans, where only partial constraints on action arguments and ordering decisions are maintained. Here, M&R benefited from Chapman's elegant TWEAK planner (IJCAI-85), which greatly clarified previous partial-order algorithms. SNLP's key feature is the use of a data structure, called a *causal link*, to record the planner's commitment to establish a precondition of one action with the postcondition of another. While the notion of causal link wasn't novel, originating much earlier in Tate's NONLIN planner (IJCAI-77), SNLP was so much simpler than NONLIN that M&R could prove it was sound and complete. A bonanza of research results followed. Penberthy & Weld's UCPOP planner (KR-92) handled actions modeled in an expressive language and was widely distributed. Other extensions included temporal planning, planning with metric resources, contingent planning in response to uncertain effects, methods for guiding the planner's non-deterministic search, *etc.*

A second contribution of M&R's paper was its clear realization that partial order planning was *not* aimed at finding partially ordered plans, but rather *lifting* the underlying search from operator sequences to *equivalence classes* over operator sequences (represented by a set of ordering constraints). This insight allowed the authors to 1) steer away from "necessary and sufficient conditions for correctness of partially ordered plans" as formalized by Chapman as the *modal truth criterion*, and 2) provide conditions under which the search is *systematic*, *i.e.* that the equivalence classes don't overlap.

The desire to speed planning led many to abandon partial-order planning in the late 1990s. The impetus for change came from three alternative approaches to planning, each of which *appeared* to dominate the SNLP family. Blum and Furst's Graphplan (IJCAI-95) was the first on the scene. Graphplan operates in two phases – first, in polynomial time it constructs a leveled graph whose nodes are either actions or propositions and whose edges are akin to causal links. In the second phase, Graphplan searches backwards from the graph's final level, looking for a consistent set of actions to achieve the goal. The following year, two more planning algorithms diverted attention from partial-order planners. Using a Graphplan-like, leveled-graph model,

¹ <https://www.aaai.org/Papers/AAAI/1991/AAAI91-099.pdf>

SATPLAN (Kautz & Selman, AAAI-96) showed that planning problems could be compiled into SAT problems, which modern SAT solvers could solve to create plans, blazingly quickly. Finally, McDermott's UNPOP (AIPS-96) introduced a method for automatically-generating domain-independent heuristics, and extensions of the method allowed simple state-space planners to solve larger problems than previously believed possible. Taken together, these exciting new directions led many researchers away from the SNLP lineage of planners.

But were researchers correct to abandon partial-order planning? While the idea of representing a plan as a partially-specified and partially-ordered set of actions embodied the popular ideas of least commitment and constraint posting, it had become apparent that the representation harbored an inherent weakness – there was no fast and easy way to compute the world-state at a “point” in the plan. Heuristically estimating the distance (*i.e.* number of plan modification operations) from one partial plan to another, that solved the goal, was extremely tricky. In contrast, if one represents a plan with a totally-ordered sequence of concrete actions (as with McDermott's UNPOP), then it is easy to calculate the resulting state and thus heuristically estimate the distance to a desired goal state. In time researchers would recognize that a major component of Graphplan's speed could be attributed to the planning graph's role as a fast way to estimate *reachability* in the state space. This realization led to even faster planners which used a planning graph to quickly generate heuristic values for use in regression and progression planning.

Eventually, researchers took the new understanding of automated heuristic construction back to the space of partially-ordered plans, completing the cycle. Nguyen & Kambhampati's RePOP (IJCAI-01) sparked renewed interest in SNLP-style planning by introducing domain-independent heuristics that rendered an SNLP-like algorithm competitive with the state of the art. Interestingly, some of these heuristics were computed from a Graphplan-style planning graph, while others were derived from graph cuts through the partial plan's causal link structure, affording the equivalent to a view of intermediate states during execution of a potential linearization of the partially-ordered steps. Then, in the 2002 International Planning Competition, two planners using partial-order representations surprised the community with their performance. Indeed, one of them, Gerevini's LPG (AIPS-02), won the competition.

Up until now discussion has been restricted to classical planning, but partial-order planning has many advantages in two currently-active research areas: temporal planning and integrated planning and execution. Temporal planning abandons the classical assumption that actions execute atomically, instead modeling the expected time duration required for execution. In this framework a world-state representation creates many more irrelevant distinctions (e.g., depending on the order in which overlapping actions are started); in contrast, the causal link representation extends naturally. Indeed, most work on temporal planning originally used this representation. While, recent temporal competitions have been dominated by so called “decision epoch planners,” work by Cushing *et al.* (ICAPS-07) shows that they suffer from fundamental problems of expressiveness. Furthermore, when a planner is connected to physical effectors and must recover from execution errors and unexpected outcomes, a partially-ordered, causal-link representation allows more flexibility to the executive, and allows

for quicker recovery if replanning is required. These reasons likely explain why NASA researchers chose a partial-order representation for EUROPA, which planned actions for the Mars Rovers Spirit and Opportunity (Bresina *et al.* ICAPS-05).

ⁱ This work was supported by Office of Naval Research grant N00014-06-1-0147 and the WRF / TJ Cable Professorship. Space constraints have necessitated the omission of a number of important planning innovations for which I apologize. I like to acknowledge many helpful comments and suggestions from Peng Dai, Oren Etzioni, Hector Geffner, Rao Kambhampati, Andrey Kolobov, Mausam, and David Smith.