To appear in Proceedings of IJCAI-93

On the Utility of Systematicity: Understanding Tradeoffs between Redundancy and Commitment in Partial-order Planning

Subbarao Kambhampati*
Department of Computer Science and Engineering
Arizona State University, Tempe, AZ 85287-5406
Email: rao@asuvax.asu.edu

Abstract

Recent work on foundations of partial-order planning has emphasized the importance of systematicity and elimination of redundancy in search space as a way to improve planning performance. In this paper, we investigate the utility of systematicity. Starting with a a rational reconstruction of the motivations for systematicity, we will conclude that it eliminates redundancy in the search space at the expense of increased commitment during planning. This increase in commitment leads to higher backtracking and increased solution depth, both of which can adversely affect the performance of the planner. We will argue that the performance of a planner is correlated more closely with the way it balances the tradeoff between redundancy and commitment, than with the systematicity of its search. We will discuss a spectrum of solutions for dealing with the redundancy-commitment tradeoff and show that systematic planners are at one extreme of this spectrum, with total least commitment planners like TWEAK on the other extreme. Based on empirical studies with seven planners that fall at different points in this spectrum, we will demonstrate that more efficient planners lie in the middle rather than at the extremes of this spectrum. We also identify characteristics of planners and problems that are predictive of planning performance.

1. Introduction

Although the idea of generating plans through partial-order (PO) planning[1] has been around for almost twenty years, it is only recently that the search space characteristics of PO planners have received particular attention. A big thrust in this work has been on reducing the redundancy in the search space of PO planners. This was largely motivated by the belief that redundancy reduction will lead to improvements in planning efficiency[9, 8]. One approach towards redundancy elimination, that turned out to be particularly influential (as evidenced by several closely related extensions [8, 13, 4, 12]), was that of McAllester's [8]. McAllester showed that it is possible to design a PO planner that is *systematic* in the strong sense that it never visits two equivalent plans or plans having overlapping linearizations (see Section 2.). Such systematic planners were claimed to be more efficient than planners that admit redundancy in their search space.

While search space redundancy is an important factor affecting the efficiency of a PO planner, another (perhaps) equally important one is the level of commitment in the planner. After all, avoiding premature commitment was one of the primary motivations for PO planning. There is often a tradeoff between the redundancy elimination and least commitment, in that often the redundancy is eliminated at the expense of increased commitment in the planner. For example, McAllester's planner achieves systematicity by keeping track of the causal structures of the plans generated during search, and ensuring that each branch of the search space commits to and protects mutually exclusive causal structures for the partial plans. We will see that such protection amounts to a strong form of premature commitment, which can increase the amount of backtracking as well as the solution depth, and can have an adverse effect on the performance of the planner.

In this paper we shall argue that the performance of a PO planner is predicted more accurately based on the way it deals with the tradeoff between redundancy and commitment, than on the systematicity of its search. We will start with a rational reconstruction of the motivations behind systematicity in PO planning and show that systematicity is just one extreme solution for the tradeoff between redundancy and commitment. We will then demonstrate that there are a spectrum of solutions to this tradeoff, and identify the dimensions along which they vary. We will explore the relative utility of the different solutions through a comparative study of seven planners that fall at different points on the spectrum. Our studies show that the planners that strike a more judicious balance between redundancy and commitment can outperform both less committed, and more systematic planners.

We start with a rational reconstruction of motivations behind the quest for systematicity in partial-order planning in Section 2., which will help clarify several prevalent misconceptions regarding the notion of systematicity in partial-order planning. Section 3. foregrounds the tradeoff between redundancy and commitment in planning. Section 4. describes seven different planners -- which between them take a spectrum of approaches for dealing with the redundancy-commitment tradeoff -- and analyzes their effectiveness. Section 5. supports the analysis through an empirical comparison of the seven planners. Section 6. summarizes the contributions of this paper. Through out this paper, we shall use the terms *partial plan* and *partially ordered plan* to mean two different things: the former simply refers to any plan which is not yet a complete solution for the problem under consideration. The latter refers to the fact that the steps of the plan (whether partial or complete) are partially, rather than totally ordered with respect to each other.

2. A Rational Reconstruction of motivations for Systematicity in PO Planning

Whatever the exact nature of the planner, the ultimate aim of (classical) planning is to find a *ground operator sequence*, which is a **solution** to the given problem (i.e., when executed in a given initial state will take the agent to a desired goal state). The ground operator sequences thus constitute the space of **potential solutions**. From a first principles perspective, the objective of planning is to navigate this space, armed with the problem specification, and find the operator sequences that are solutions for the given problem. Figure 1 illustrates how **refinement planners** navigate the space of potential solutions. When a refinement planner reaches a given operator sequence that is not a solution, it may refine the sequence further in the hopes of making it a solution.

Completeness and Systematicity in Base Space: Given this first principles view of planning, a planner is said to be **complete** if it finds *all* the operator sequences that are solutions of any given problem. The planner is said to be **systematic** if it never examines and refines any operator sequence (potential solution) more than once. The primary motivation behind systematicity is to avoid redundancy in the search space generated by the planner. An important corollary of systematicity is thus that *in the worst case, a planner which is*

---

*This research is supported in part by National Science Foundation under grant IRI-9210997.

[1] Throughout this paper, we shall use the term *PO planning* rather than *nonlinear planning*, as the former avoids confusion with *linearity* assumption.

Figure 2: Examples showing redundancy and looping in TWEAK search space. In all the examples, the operators are shown with preconditions on the right side, and effects on the left (with ‘‘+’’ sign for add list literals, and ‘‘−’’ sign for delete list literals). The Init and Goal lists specify the problem.

Further more, since TWEAK does not keep track of which goals it already achieved and which it is yet to achieve, it may achieve and clobber a goal/subgoal arbitrarily many times within the same refinement branch. This causes unnecessary looping and generates many useless operator sequences. Example (b) in Figure 2 shows a scenario where TWEAK repeatedly achieves and clobbers the conditions Q and R at action $a1$. In fact, in this example, TWEAK would loop infinitely if the action $a4$ is not available.

Book-keeping through Causal Links: The looping behavior in the above example could have been avoided had TWEAK kept track of the reasons for introducing different steps into the partial plan in the first place. In particular, had it remembered that $a2$ is introduced to give Q at $a1$, then while introducing $a3$, into the plan, it would have realized that there is no way the interaction between $a2$ and $a1$ can be resolved, and would thus have abandoned that path. This would have obviated the possibility of repeated establishment and undoing of the same goal (and hence the looping).[3]

It is this realization that served as the first impetus for protection intervals/causal links in planning. Protection intervals were first introduced in Sussman’s Hacker [15] as a book-keeping device for keeping track of the goals that have already been worked on, and thereby avoiding undoing them or working on them multiple times. They were first used in partial-order planning in Tate’s Nonlin [16]. A protection interval, denoted as $\{s\} \xrightarrow{p} w$, is best seen as a commitment by the planner that the condition p at step w, first established by the step s will henceforth be protected from being violated before w. Specifically, if there is a possibility of p being not true before w, the planner tries to see if the steps that are deleting p could be ordered to come either before s or after w.

Redudancy elimination through Exhaustive causal links: McAllester [8] was the first to point out that a natural extension of the commitment to a protection interval, $\{s\} \xrightarrow{p} w$, viz., s will be the effective contributor (i.e., last step before w that provides p) of p for w in every topological sort of the current partially ordered plan -- can be used to eliminate redundancy in the planner’s search space. Notice that with this extension, any step v that can possibly come between s and w and either delete or *add* p is going to threaten the commitment promised by the causal link $\{s\} \xrightarrow{p} w$, and thus is considered a *threat*. Specifically, v is considered a **-ve threat** if it deletes p and a **+ve threat** if it adds p. Causal links with these semantics are called **exhaustive causal links** (c.f. [4]). To make an exhaustive causal link $\{s\} \xrightarrow{p} w$ **safe**, every threat must either be *promoted* to come after w or be *demoted* to come before s. (Notice that white-knight declobbering [1] is inconsistent with the semantics of an exhaustive causal link).

Systematicity of SNLP: McAllester described a planner based on the idea of exhaustive causal links (which later became known as

---

[3] Planners using causal links do not *eliminate* looping; they merely reduce it. In particular, although no causal-link based planner will loop on Example (b) of Figure 2, there are other examples on which they will.

Figure 3: Example demonstrating that arbitrary threat deferment does not preserve systematicity

in multiple search branches. It is however solution-systematic since all the redundant search branches will be ultimately pruned when the planner attempts to resolve the +ve threats. In other words, a planner that is solution systematic may do arbitrary amount of redundant work in the branches that will ultimately get pruned. Furthermore, a planner that is solution-systematic may have a *larger* search space than a corresponding planner that is not solution-systematic. In the current example, a planner which completely ignores +ve threats (such as McNonlin to be described in the next section), and is thus not solution systematic, would stop before resolving the positive threats, and thus would generate a smaller search space than SNLP$'$. Thus, from the point of view of redundancy elimination, which was the primary reason for systematicity, solution-systematicity is of dubious utility.

3. The Redundancy-Commitment Tradeoff

In the previous section, we saw that SNLP avoids redundancy in the search space through a strong form of commitment to causal links. TWEAK, on the other hand, does not commit to causal links, and as a result introduces redundancy into its search space. We shall now focus on how this affects their performance in practice.

Clearly, worst case search space size will have a strong correlation with performance only when the planner is forced to explore a significant portion of its search space in solving the problem. Since the goal of planning typically is to find *one*, rather than *all*, solutions of a problem, unless the problem is unsolvable; or the solution density is low and the planner's initial choices lead it towards non-solution branches, systematicity may not necessarily lead to improvements in planning performance (c.f. [6]).

If systematicity may not always improve performance, the next question is: does it *hurt* performance? While redundancy in the search space is an important factor affecting the performance of a PO planner, another (perhaps equally) important factor is the level of commitment in the planner. Indeed, as mentioned earlier, the original impetus for PO planning was the realization that total ordering planners such as STRIPS [10] commit too soon to specific action orderings.

The use of protection intervals and causal links, while systematizing the search, also results in increased commitment to particular establishment structures. Increased commitment leads to higher backtracking as well as increased depth of the solution, which in turn adversely affect the performance. This is in contrast to TWEAK, which avoids commitment to causal links, but at the expense of increased redundancy.

This leads us to the hypothesis that performance depends not solely on the level of redundancy or commitment in the planner, but rather on the tradeoff between them. This also suggests that neither SNLP which eliminates redundancy without worrying about commitment, nor TWEAK which guarantees least commitment without worrying about redundancy, is guaranteed to be efficient all the time. In the next two sections, we will explore this tradeoff further by comparing a spectrum of solutions that fall in the middle of SNLP and TWEAK, and strike a better balance in this tradeoff than either.

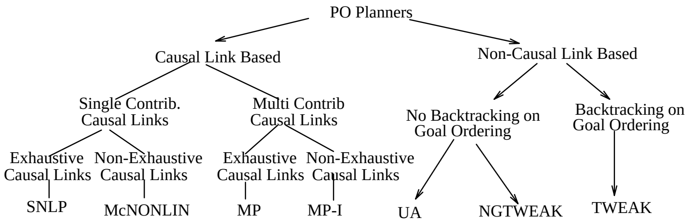

Figure 4: A spectrum of solutions to the tradeoff between redundancy and commitment in PO planning

4. A Spectrum of solutions for the tradeoff between redundancy and commitment

Figure 4 shows other *sound* and *complete* PO planners that fall in the middle of these two extremes. In this section, we will characterize the dimensions along which these planners vary, with particular emphasis on the way they deal with redundancy-commitment tradeoff.

SNLP, McNonlin, MP and MP-I , the four planners on the bottom left in Figure 4 use causal links to organize their search. They introduce causal links to support each open condition of the plan, and ensure that none of the causal links are threatened. All of them do threat resolution before link establishment. The planners differ only in their definition and treatment of threatened causal links. We say that $\mathcal{S} \xrightarrow{p} w$ is a generalized (multi-contributor) causal link of plan $\mathcal{P}$ if the step w and the set of steps $\mathcal{S}$ (called the *contributor set*) belong to $\mathcal{P}$, such that all the steps in $\mathcal{S}$ are ordered to precede w and each of them is capable of giving the precondition p to w. Given a step v, and causal link $\mathcal{S} \xrightarrow{p} w$ of the plan $\mathcal{P}$ such that v can possibly come between every step of $\mathcal{S}$ and w, v is called a *negative threat* for the causal link if v deletes p. v is called a *positive threat* for the causal link if v adds p. A causal link $\mathcal{S} \xrightarrow{p} w$ of a plan $\mathcal{P}$ is said to be *exhaustive* (c.f. [4]), if for every ground linearization of $\mathcal{P}$, some step $s' \in \mathcal{S}$ will be the last step preceding w that gives p in that linearization. It is easy to see that a correct plan with exhaustive causal links, none of the causal links will have positive or negative threats. If every causal link of the plan has a singleton contributor set (i.e., $|\mathcal{S}| = 1$), then the plan is said to have *single contributor* causal structure.

Given this formulation of causal links, we can now describe the four causal link planners as four points along a spectrum. Two of them, SNLP and McNonlin, use the traditional *single contributor* causal links to guide planning. SNLP uses exhaustive causal links (and thus resolves +ve as well as -ve threats). In contrast, McNonlin[7] uses non-exhaustive causal links, and thus ignores +ve threats. In both planners, a threat v to a causal link $\{s\} \xrightarrow{p} w$ is resolved by ordering it to come either before s or after w.

The other two planners, MP and MP-I , use multi-contributor causal links [4] to guide planning. MP maintains exhaustive causal links. In MP , a threat v to a link $\mathcal{S} \xrightarrow{p} w$ is resolved by either promoting v to come after w or demoting v to come before some step $s \in \mathcal{S}$ (a choice point). In addition, if v is adding the condition p, we can also make the link safe by ordering v to come before w and *merging* v with $\mathcal{S}$. MP-I on the other hand maintains non-exhaustive causal links and thus ignores +ve threats. It however introduces an additional declobbering resolution operation for -ve threats: Suppose there is a negative threat v to a causal link $\mathcal{S} \xrightarrow{p} w$ and there exists step s' such that it follows v and adds p. In this case, the threat can be resolved by ordering s' to come before w and merging s' with $\mathcal{S}$. The overall effect is that MP-I is less eager than MP in accumulating new contributors to a causal link.

Clearly, the four planners balance the tradeoff between redundancy and commitment in different ways. In particular, SNLP and McNonlin reduce redundancy in the search space by increasing the commitment to specific contributors, while MP and MP-I reduce commitment to individual contributors at the expense of increased redundancy. MP and SNLP maintain exhaustive causal structures which allow them to split the possible partial plans into sets of plans with mutually exclusive establishment structures, there by controlling the redundancy in the search space. (Specifically, although MP is not systematic, it can be shown to be causal structure systematic; Section 2.). However, exhaustiveness also increases effective solution depth, since unlike McNonlin and MP-I , MP and SNLP have to deal with positive as well as negative threats to the causal links.

UA, TWEAK and NGTWEAK, the three planners on the bottom right in Figure 4, do not use causal links (or protection intervals) in their search. Thus, they completely eliminate the commitment to contributors. Of these three, TWEAK closely follows the idea of inverting Chapman's MTC (with the minor difference that it doesn't use external white-knights for declobbering). It is the only planner among the seven that backtracks on goal orderings. NGTWEAK reduces some of the redundancy in the TWEAK search space by avoiding backtracking on goal ordering decisions (similar to the implementation described in [17]). UA [9], reduces the redundancy further by maintaining partial plans that are *unambiguous* in the sense that each prerequisite in the plan is either necessarily true or necessarily false. UA achieves this by ordering every newly introduced step with respect to all possibly interacting steps. This policy allows UA to eliminate plans with overlapping linearizations from its search space (unlike TWEAK). However, like TWEAK, UA still allows multiple paths to the same partial-order, as well as establishing and violating the same condition multiple times within the same search branch.

The seven planners described in the previous section thus represent a spectrum of solutions for the tradeoff between redundancy and commitment. Informally, in Figure 4. the search space redundancy increases from left to right, while the level of commitment decreases from left to right.

5. Empirical Analysis of the tradeoff between Redundancy and Commitment

To understand how the various solutions to the tradeoff between redundancy and commitment affect the planning efficiency, we implemented all seven planners described in the previous section on a common platform, and performed focused empirical studies. The broad objectives of the studies were to demonstrate that neither systematicity nor least-commitment alone are good predictors of planning performance, and that performance can be better predicted by the way a planner balances the redundancy-commitment tradeoff. This section describes the study and analyzes the results.

5.1. Experimental Setup

Our test domains included classical toy-worlds such as blocks world, as well as the synthetic domains used in Weld et. al.'s work [13]. In this paper, we will concentrate on the results from Weld et. al.'s synthetic domains and our variants of them, as they provide for a more controlled testing of the tradeoffs. Weld et. al.'s original domains include ART-MD and ART-1D (also called D^mS^1, D^1S^1 respectively), which are designed to contain easily serializable and laboriously serializable sub-goals respectively. The domains are defined as:

ART-MD $A_i\ \underline{prec}:\ I_i\ \underline{add}:\ G_i\ \underline{del}:\ \{I_j | j < i\}$
ART-1D $A_i\ \underline{precond}:\ I_i\ \underline{add}:\ G_i\ \underline{del}:\ I_{i-1}$

To these, we also added our own variants: ART-MD-RD and ART-1D-RD, which introduce preconditions achieved and deleted by multiple operators. The variants are produced by making every even numbered action require an additional precondition *he*, delete that precondition and add an additional postcondition *hf*.[8] The odd numbered actions similarly require and delete *hf* and add *he*. The specification of ART-MD-RD is shown below. ART-1D-RD is produced similarly from ART-1D

---

[7] McNonlin is a considerable simplification of Tate's original implementation of NONLIN[16], which was a *hierarchical* nonlinear planner.

[8] *he* and *hf* are supposed to be mnemonics for $handempty$ and $handfull$ conditions in the traditional blocks world domain [10], which are achieved or deleted by many of the actions in the domain.

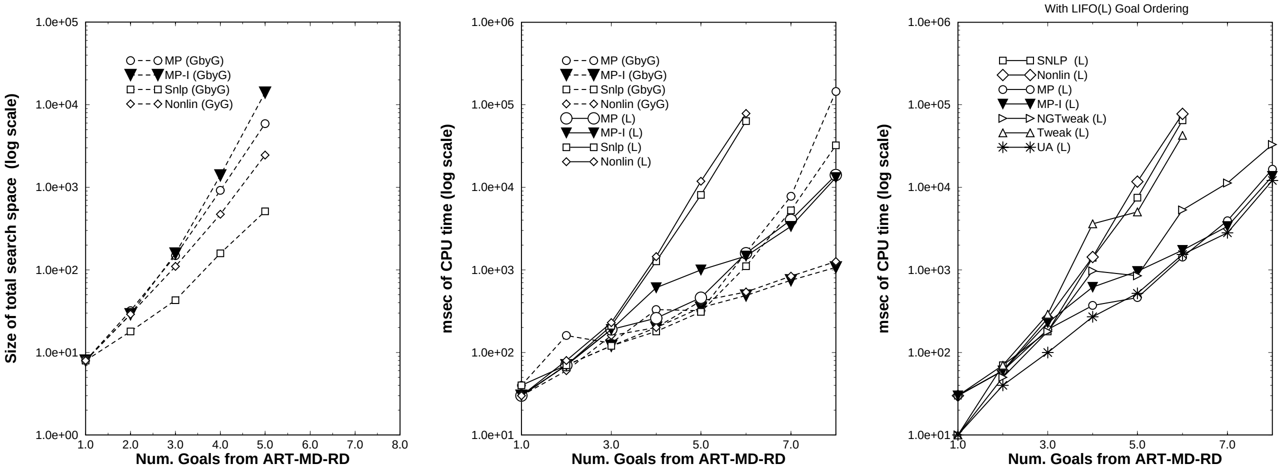

Figure 5: The left plot compares the full search space sizes of the four causal link planners in ART-MD-RD. The middle one compares the performance on solvable problems for two different goal orderings. The last compares the performance of all seven planners.

A_i $\underline{prec}$: I_i, he $\underline{add}$: G_i, hf $\underline{del}$: $\{I_j | j < i\} \cup \{he\}$ for even i
A_i $\underline{prec}$: I_i, hf $\underline{add}$: G_i, he $\underline{del}$: $\{I_j | j < i\} \cup \{hf\}$ for odd i

In each of the domains, we compared all seven planners, over *solvable* problems with 1 to 8 goals from the set $\{G_1 \cdots G_8\}$. Since the effect of commitment to contributors depends significantly upon the order in which the various goals and subgoals are addressed by the planner, we tested with two types of goal order strategies: strategy L which is a LIFO strategy where a goal and all its recursive subgoals are addressed before the next higher level goal is addressed; and strategy GbyG, which is a FIFO strategy where all the top level goals are addressed before their subgoals are considered by the planner (strategy L corresponds to a depth-first traversal of the goal-subgoal tree, while strategy GbyG corresponds to a breadth-first traversal). An A^* search regime was used for all the planners (see [5]).

5.2. Experimental Results and Discussion

5.2.1. Comparison between the Causal Link Based Planners

We started by comparing the cpu time (in m.sec. on a SUN SPARC-II), taken by the four causal link based planners for each of the goal ordering strategies in ART-MD and ART-1D respectively. We found that there is no appreciable difference in time taken by the planners for solving problems in these domains (plots omitted due to space limitations).

Next, we studied the performance of these planners in ART-MD-RD, and ART-1D-RD domains. We started by comparing the sizes of the overall search spaces of all the planners. The leftmost plot in Figure 5 compares the total search space sizes of these four planners for one of the goal orderings.[9] It shows that they have varying amounts of redundancy in the search space, with SNLP having the smallest search space, and MP-I having the largest. If the search space size were to be sole indicator of performance, we would expect SNLP to perform best, followed by McNonlin, MP and MP-I .

The middle plot in Figure 5 compares the performance of the planners in ART-MD-RD domain which contain the easily achieved and deleted conditions *hf* and *he* (the performance profiles for ART-1D-RD were very similar, but are omitted here due to space limitations). The missing data points on a plot correspond to the problems that couldn't be solved before the time bound, which for us was the time it took for the lisp to run out of memory and fail. These empirical performance relations are surprisingly robust. In particular, we found that short of providing the correct order in which to attack goals, there is no other way to change the qualitative performance relations between SNLP and MP-I. Specifically, static threat delay strategies, such as those described in [12], while improving the overall performance of all planners slightly, leave the qualitative relations between the planners unchanged. Same is true of the heuristic of working only on the goals that are not necessarily true (a la TWEAK and UA).

Analysis: The behavior of the four causal link based planners in our experiments can be explained in terms of the tradeoff between redundancy and commitment that we discussed earlier. In particular, the near identical performance of the four planners in the original ART-MD and ART-1D domains can be explained by the fact that in these domains each precondition is ultimately provided by a single action in the plan, and there is no penalty for committing prematurely to that action as the contributor. Premature commitment does hurt performance In ART-x-RD domains, and we find that:

- MP and MP-I perform better than SNLP and McNonlin in the case of the LIFO goal ordering strategy L (see plots in Figure 5) even though the latter two search in exponentially smaller search spaces (see plots in Figure 5), with SNLP in particular having the smallest search space, as guaranteed by its systematicity property.

- MP-I and McNonlin perform better than SNLP and MP in the case of GbyG strategy, even though the former two don't use exhaustive causal links and thus have more redundancy.

Once again, this behavior can be explained in terms of the way each planner handles the tradeoff between redundancy and commitment. Since both MP and MP-I can change their commitment to establishment structures within the same branch, they may visit a potential solution more than once. They thus have larger search spaces compared to SNLP. Moreover, MP and SNLP, which maintain exhaustive causal structures, have smaller search spaces than MP-I and McNonlin respectively.

However, the performance is not directly correlated with the search space size. In the LIFO ordering strategy L, SNLP and McNonlin are forced to commit to a specific contributor for the *hf* and *he* subgoals (preconditions) of a top level goal G_i, before the other top level goals are expanded. Since both *hf* and *he* are easily added and deleted by many actions in the domains, such premature commitment has a high probability of being wrong. Since both SNLP and McNonlin protect causal links to eliminate redundancy, the only way they can get rid of a wrong causal link is to backtrack over it (i.e., go over to another branch of the search space). Such backtracking turns out to be very costly in terms of performance. MP and MP-I avoid problematic backtracking as they can deal with

---

[9] The size of the overall search space is found by setting the termination condition for each planner to be uniformly false, thus forcing the planners to visit every node in the search space before giving up.

their initial wrong commitment by merging additional contributors into the contributor list as and when they become available.

Premature commitment turns out to be less of a problem when the FIFO ordering strategy GbyG is used, since in this case G_i and I_i are addressed before hf and he, and each action A_i is capable of giving only one of the goals G_i. Since only the initial state is capable of giving all I_i, the orderings imposed to deal with the deletions of I_i's by individual actions of the plans constrain the plan enough so that by the time hf and he are addressed, the only available contributor choice also happens to be the correct choice.

Although premature commitment is not a problem, SNLP and MP consider positive as well as negative threats for causal links and thus have higher solution depth. In particular, note that plans that are complete for MP-I and McNonlin may still need to be refined further to ensure exhaustiveness of their validation structure and thereby make them complete with respect to MP and SNLP. This explains the higher planning cost of SNLP and MP compared to McNonlin and MP-I respectively.

Note however that ability to avoid premature commitment, and solution density are more predictive of performance than is the presence or absence of +ve threats. In particular, the middle plots in Figure 5 show that McNonlin performs worse than SNLP in the LIFO strategy (this effect was even more pronounced in the case of ART-1D-RD domain [5]). This can be explained by noting that the redundancy that McNonlin introduces into its search space by ignoring +ve threats, adversely affects the performance exactly when the planner's initial commitments all turn out to be wrong, forcing it to look at a significant portion of its search space.

5.2.2. Comparison between Causal Link based and Non-causal Link based Planners

Given that the reduced commitment to contributors provided by multi-contributor causal links allows them to strike a better balance between redundancy and commitment than is the case with single contributor causal links, a natural question arises as to whether planners that completely avoid causal links strike an even better balance. To answer this question, we compared the performance of all seven planners shown in Figure 4. The rightmost plot in Figure 5 shows the results of these comparisons in ART-MD-RD for the LIFO ordering (the plots for GbyG ordering are omitted due to space limitations; see [5]).

As expected, both SNLP, which completely eliminates redundancy at the expense of increased commitment, and TWEAK which avoids all forms of commitment, fare badly in this domain. The former gets penalized for the overcommitment (especially in the LIFO goal ordering strategy), and the latter gets penalized for the excessive redundancy (high branching factor). For the LIFO goal ordering strategy, the planners using multi-contributor causal links, and those using no causal links (except TWEAK), do better than the planners using single contributor causal links, including SNLP and TWEAK. This can be explained by the fact that LIFO goal ordering heavily penalizes planners that commit too early to specific contributors and protect the commitments. The best performance for both goal orderings is shown by UA and MP-I with NGTWEAK coming third. This also shows that the cost of maintaining causal links is not a factor affecting performance. The real issue is the level of commitment forced by them. Although UA, a causal-link-less planner performed as well as the best causal link planner, MP-I in this domain, we believe that this may change as we go to more expressive action representations. In these latter situations, not maintaining causal links leads to an additional form of redundancy which can adversely affect performance: the planner may work on multiple secondary preconditions in aid of the same (sub)goal (c.f. [11]). We are currently testing this hypothesis [5].

6. Concluding Remarks

In this paper we provided a rational reconstruction of the motivations for systematicity, and argued that the performance of a PO planner is correlated more closely with the way it balances the tradeoff between redundancy and commitment, than on the systematicity of its search. We showed that there are a spectrum of solutions to this tradeoff, and identified the type of causal links used (single *vs.* multi-contributor *vs.* none), and the level of commitment to the causal links (exhaustive vs. non-exhaustive) as the important dimensions of variation. We conducted focused empirical studies to understand how these different dimensions affect the performance of a planner. The main objective of the empirical study has been to establish the presence of tradeoffs. The studies characterize SNLP and TWEAK as two extreme solutions to the redundancy-commitment tradeoff, and demonstrate that planners which strike a more judicious balance in this tradeoff can outperform both less committed, and more systematic planners. Although our empirical studies were mainly done in artificial domains, we believe that the hypotheses regarding tradeoffs between redundancy and commitment will also apply to other domains. Our future plans include a more analytical formulation of these tradeoffs [5].

Acknowledgements: Special thanks to Bulusu Gopi Kumar, Laurie Ihrig, Suresh Katukam and Eric Cohen of ASU planning seminar for helping me clarify my thinking. Thanks are also due to Mark Drummond, Will Harvey, Richard Korf, Pat Langley, David McAllester, Steve Minton, Mark Peot, Dave Smith, Austin Tate and Dan Weld for many helpful discussions.

References

[1] D. Chapman. Planning for conjunctive goals. *Artificial Intelligence*, 32:333--377, 1987.

[2] S. Hanks and D. Weld. Systematic Adaptation for Case-Based Planning. In *Proc. 1st Intl. Conf. on AI Planning Systems*, 1992.

[3] W.D. Harvey and M.L. Ginsberg and D.E. Smith. Deferring Conflict Resolution Retains Systematicity. Submitted to AAAI-93.

[4] S. Kambhampati. Multi-Contributor Causal Structures for Planning: A Formalization and Evaluation. Arizona State University Technical Report, CS TR-92-019, July 1992. (A preliminary version appears in the Proc. of First Intl. Conf. on AI Planning Systems, 1992).

[5] S. Kambhampati. A Comparative analysis of Search Space Size, Systematicity and Performance of Partial-order planners. Arizona State University Technical Report (In preparation).

[6] P. Langley. Systematic and Nonsystematic search strategies. In *Proc. 1st Intl. Conference on AI Planning Systems*, June 1992.

[7] D. McDermott. Regression Planning. *Intl. Jour. Intelligent Systems*, 6:357-416, 1991.

[8] D. McAllester and D. Rosenblitt. Systematic Nonlinear Planning. In *Proc. 9th AAAI*, 1991.

[9] S. Minton, J. Bresina and M. Drummond. Commitment Strategies in Planning: A Comparative Analysis. In *Proc. 12th IJCAI*, 1991.

[10] N.J. Nilsson. *Principles of Artificial Intelligence*. Tioga Publishers, Palo Alto, CA, 1980.

[11] E.P.D. Pednault. Generalizing nonlinear planning to handle complex goals and actions with context dependent effects. In *Proc. IJCAI-91*.

[12] M.A. Peot and D.E. Smith. Threat-Removal Strategies for Nonlinear Planning. In *Proc. AAAI-93*

[13] A. Barrett and D. Weld. Partial Order Planning: Evaluating Possible Efficiency Gains Technical Report 92-05-01, Department of Computer Science and Engineering, University of Washington, Seattle, WA, June 1992.

[14] D.E. Smith and M.A. Peot. Postponing Simple Conflicts in Nonlinear Planning. In Working notes of AAAI Spring Symp. on Foundations of Automatic Planning: The classical approach and Beyond.

[15] G.J. Sussman. *A Computer Model of Skill Acquisition*. American Elsevier, New York, 1975

[16] A. Tate. Generating Project Networks. In *Proceedings of IJCAI-77*, pages 888--893, Boston, MA, 1977.

[17] Q. Yang, J. Tenenberg and S. Woods. Abstraction in nonlinear planning University of Waterloo Technical Report CS-91-65, 1991.