In: Proceedings of the ICAPS-04 Workshop *Connecting Planning Theory with Practice*

Toward a Theory of Qualitative Reasoning about Plans

Karen L. Myers
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, California 94025
myers@ai.sri.com

Abstract

Automated planning algorithms embody a theory of causality grounded in linking enabling effects of actions or initial world conditions to preconditions of subsequent actions. This model has several drawbacks when applied in the context of mixed-initiative planning. First, it requires comprehensive causal models that describe for every action its full set of preconditions and postconditions; in many application domains, such models will be unavailable. Second, it does not cover several forms of intraplan relations that a human may wish to document and reason with. To address these shortcomings, we define a *qualitative* framework for reasoning about plan structure. Specifically, we define a set of plan relations that characterize key interactions among plan components and a calculus for reasoning qualitatively about the effects of plan changes. We argue that our qualitative approach is much better suited to mixed-initiative planning than the causal link method.

Introduction

As the AI planning community turns increasingly to realistic problem domains, there has been a growing recognition of the need for mixed-initiative planning techniques. This need is driven by two concerns. First, full automation is inappropriate for many application domains, as users often want control of the planning process. Second, the cost of formulating the correct and comprehensive background theories required by automated planners is prohibitive. These theories consist of causal models that describe for every action its *preconditions* (i.e., the conditions under which it could be applied) and its *postconditions* (i.e., the conditions that result from execution of the action). Most planning knowledge bases to date provide only poor approximations to realistic models of activity. Furthermore, formulation of even these simplified models requires highly trained specialists in AI planning and knowledge representation. Mixed-initiative planning systems can reduce knowledge requirements by relying on the user to provide information when necessary.

Most work to date on mixed-initiative planning assumes a comprehensive store of background knowledge that the system can apply to validate interactions with the user (e.g., (Allen and Ferguson, 2002; Kim and Blythe, 2003)). We are interested in a model of mixed-initiative planning that does not depend on comprehensive background knowledge. Rather, we view plan development as a cooperative problem-solving process to which the system and the user both contribute knowledge.

One consequence of this collaborative model is that the system can no longer generate a comprehensive causal link structure for a plan (Weld, 1999), as the system does not store all required knowledge. Causal structures enable proof of the *correctness* of a plan, meaning that simulated execution of the actions in the plan from the initial state yields a state that satisfies stated objectives. Additionally, they provide the means to answer the following questions (Kambhampati and Hendler, 1992), which are important for both automated and collaborative planning:

A. *What role does a given action, constraint, or assumption play in a plan?*
B. *What impact would a given change have on a plan?*

For these reasons, the planning community has been reluctant to embrace approaches that do not generate causal structures as a by-product.

When background models are incomplete or unavailable, a human planner would have to annotate a plan manually in order to obtain a full causal structure. Supplying a complete set of causal annotations would be a time-consuming and laborious task. Furthermore, the formulation of causal link justifications for activities is a highly technical skill that is beyond most users.[1]

We believe that much of the value of these complex causal models can be attained through simpler models that capture commonsense notions of intraplan relationships. Furthermore, as argued below, the models of intraplan relationships embraced by the AI planning community are overly narrow because of their evolution from the structures required for automated planning. In particular, they cannot capture certain notions of dependency that are important for mixed-initiative planning.

In this paper, we propose a user-centric model for plan structure that is grounded in *qualitative* relationships among plan objects. The basic idea is to trade some of the

[1] Part of the appeal of hierarchical task network (HTN) planning (Erol et al., 1994) is that conditions and effects need not be specified for individual tasks within a network, only for the network as a whole. In particular, detailed causal models are not required for all actions in a domain. This reduction in knowledge modeling costs and the associated improvement in computational properties have helped to make HTN planning the preferred paradigm for large-scale planning problems.

precision of the formal causal models for a simpler approach that is easier to formulate and apply. Our approach still enables answers to questions (A) and (B) above, but in qualitative rather than quantitative terms.

An interesting parallel can be drawn between plan development and other synthesis tasks such as mechanical design. Commercial design systems support human designers by providing automation for low-level, tedious tasks. A designer does not document every step or component within a design. Indeed, studies have shown that humans are resistant to providing full rationale information as part of the design process, due both to the substantial time commitment required and the changes in work style that result (Carroll and Moran, 1991). However, designers do track key dependencies and requirements that they believe will be important in understanding and maintaining a design downstream. We believe that this type of user-driven approach should be the aim for documenting design rationale within mixed-initiative planning technologies.

We begin by defining a set of qualitative plan relations. We then present an example that contrasts the use of standard causal link structures with our qualitative approach. This example draws on a simplified version of a noncombatant evacuation operation (NEO) domain that was developed for the PASSAT mixed-initiative planning system (Myers et al., 2002). Finally, we define a calculus for reasoning qualitatively about plan changes.

Qualitative Plan Model

Our framework for reasoning qualitatively about plans builds on both causal links and qualitative information about plan relationships. Here, we define our model of a plan and the qualitative and causal link relations.

Plan Elements

Our model of a plan contains three types of element: *actions* (activities that can be undertaken), *effects* (conditions, either to be achieved, the expected result of executing an action, or a property of the initial world state), and *parameters* (arguments to an action or condition). We use the symbol *obj* (i.e., plan object) to denote an arbitrary plan element from any of the above types. A plan relation is defined between a *source object* and a *target object* and is represented as *Reln: source-obj → target-obj*.

Causal Link Relation

The *causal-link* relation (see Figure 1) is the primary relation within most automated planning systems. It indicates that the source effect is a necessary condition for the target effect.

Causal-link: {Effect} → {Effect}

Figure 1. Causal Link Relation

Qualitative Relations

Figure 2 summarizes our qualitative plan relations. These relations can be separated into two categories: *temporal* (QR1 – QR3) and *logical* (QR4 – QR5).

QR1 *Precedes:{Action,Effect} → {Action, Effect}*
QR2 *Necessary-for:{Action,Effect} → {Action,Effect}*
QR3 *Supports:{Action,Effect} → {Action,Effect}*
QR4 *Parameter-dependence:{Param} →{Param}*
QR5 *Condition-dependence:{Effect} → {Action,Effect,Param}*

Figure 2. Qualitative Plan Relations

Qualitative Temporal Relations

The qualitative temporal relations capture the notion that a given action or effect in a plan must precede some other action or effect. We consider three types: *precedes*, *necessary-for*, and *supports*.

The *precedes* relation captures the notion that the source action or effect should occur before the target action or effect, without providing any indication of why. This type of relation can be used to capture a preference for performing activities in some designated order when there is no necessary reason for that order. (Task-ordering links in the HTN paradigm provide a comparable capability.) For example, consider the actions of preparing an evacuation site and flying evacuees to the evacuation site. Although it would be possible to perform those actions in parallel, an individual planner may have a preference for completing the preparation prior to the start of the airlift of the evacuees, possibly to enable a delay of the airlift in the event of problems with the preparation.

The *necessary-for* and *supports* relations specialize *precedes* to capture semantic motivations for the ordering. *Necessary-for* indicates that a given action or effect must occur before a designated action or effect in order to enable the target plan element. For example, it would be *necessary-for* evacuees to be marshaled to an assembly point before they could be loaded onto an evacuation aircraft. In essence, the *necessary-for* relation constitutes a qualitative abstraction of the *causal-link* relation.

The *supports* relation indicates that the source action or effect contributes to the target action or effect in some noncritical way. For example, a patrol mission may provide additional support to an evacuation activity, without being essential to its undertaking. Hence, if the aircraft performing the patrol were redirected, the evacuation process would not be jeopardized. Source objects for *supports* relations correspond to 'redundant' actions or effects that, while unnecessary, lead to improved plan robustness or quality. While redundancy is considered good practice in human-authored plans, causal link planners prohibit redundancy by imposing conditions of minimality on a plan's causal structure.

Qualitative Logical Relations

The qualitative logical relations (QR4 -- QR5) can be used to declare a dependency between the source and target elements when the dependency cannot be defined precisely in terms of a deductive specification or mathematical formula. Such relationships arise frequently, given that many factors that influence decision-making are problematic to formalize. These factors can include conditions that are too complex to codify (i.e., the *qualification problem* (McCarthy, 1977)) or subjective preferences that vary among human planners.

For example, the choice of assembly point in an evacuation plan will impact the type of aircraft that can be used for transporting evacuees (e.g., a small helicopter may be necessary for evacuation from an embassy while a larger aircraft could be used at a football stadium). However, there is no hard-and-fast rule for determining the type of aircraft to be used for a particular location.

Qualitative logical relations can be designated between plan parameters (QR4), or between effects and any type of plan object (QR5). The connection between the choice of assembly location and transport aircraft in the example above corresponds to a *parameter-dependence* relation; the connection between security level and choice of assembly point corresponds to a *condition-dependence* relation.

The qualitative logical relations can be made 'quantitative' by associating definite constraints with them. For the *parameter-dependence* relation, these constraints would be in the form of a set of equations linking the two parameters. We introduce the term *parameter-constraint* relation to refer to this specialization of the *parameter-dependence* relation. One of the most valuable forms of *parameter-constraint* relation is an equality constraint indicating that two planning variables be instantiated to the same value. (Such constraints are sometimes referred to as *codesignation* constraints.) A comparable *condition-constraint* relation could similarly be defined.

Properties of the Model

The qualitative model trades the precision of exhaustive causal links for simplicity and ease of specification. Indeed, there is a natural abstraction from an exhaustive causal link structure to corresponding qualitative models that replaces every *causal-link* relation with a *precedes* or *necessary-for* relation.

While the qualitative relations provide less precision than a full set of causal links, they offer two key advantages. First, they are simpler and more intuitive to specify, thus making them better suited for use in a mixed-initiative planning environment. Second, there are relationships that can be modeled in the qualitative framework that are not supported by causal links or logical constraints associated with operator preconditions. In particular, the *supports* relation enables the description of a connection between plan objects that is not essential for plan correctness. As well, the *precedence* relation allows the expression of ordering information independent of causality. Furthermore, the *condition-dependence* and *parameter-dependence* relations enable ill-defined connections between plan objects to be expressed; this ability is useful when the precise logical or mathematical relationship is either not known or not easily formalizable, yet it is desirable to document that some dependency exists.

Sources

Our motivation for defining the qualitative framework is to enable reasoning about plan changes within a mixed-initiative planning environment. Within this context, several sources would contribute to the set of causal relations for a given plan. First, background knowledge could capture both conditions and effects for individual actions, and causal link and qualitative relations within 'standard operating procedures' encoded as task networks. Second, the human planner would contribute additional relations at planning time. In the future, learning mechanisms could be applied to hypothesize relations from a user-authored plan (e.g., along the lines proposed by (El Fattah 2003)), yielding a baseline that a user could then modify.

Example

To illustrate the use of the qualitative relations, consider the plans for a simple evacuation mission in Figures 3 and 4. These plans are designed to achieve the goal conditions *Prepared(Camp1)* and *At(Evacuees Camp1)*, based on the operators defined in Figure 5. To simplify reference, each plan action is labeled with a unique identifier (e.g., *N1*).

The plan in Figure 3 corresponds to a solution that an automated causal link planner might produce for this problem. It includes a full set of causal links that match action preconditions to earlier effects in the plan.

The plan in Figure 4 represents a solution that a human planner might construct using a plan authoring tool. It contains all the actions in Figure 3 plus a *Patrol* action (*N7*) that provides additional security for the sector to which the evacuees will be moved. Given the logic of the domain operators, this action is redundant because it does not establish any effects that are required in the plan. However, human planners typically build such redundancy into plans to provide additional safeguards in the face of unexpected events. The plan in Figure 4 also contains a set of causal link and qualitative relations that document what the user might view as the key dependencies.

The main differences with the plan in Figure 4 compared to that in Figure 3 are as follows.

- Causal-link relations are limited to dependencies on initial state conditions that might be expected to change (e.g., the security of key locations, the position of the vehicle to be used for transporting the evacuees), and important intermediate effects of actions (e.g., the evacuees remain at the embassy until they are loaded onto the transport vehicle). Static initial conditions and unimportant intermediate effects have been omitted.

- The hybrid annotation replaces certain of the *causal-link* relations with qualitative *necessary-for* relations, indicating that it is essential for the source activity to precede the target activity in the plan but without documenting the effects that link the actions. From the user's perspective, these effects are obvious (e.g., the aircraft has to be loaded before it can be unloaded), and so documenting them explicitly is of little value.
- The qualitative annotations include a *precedes* relation from node *N6* to node *N4*, indicating a (noncausal) preference for ordering those two actions. A *condition-dependence* relation has been added from the predicate *#Evacuees(25)* in the initial world state to the parameter *?place* in *N1* where the evacuees are to be assembled. This link indicates that the choice of assembly location depends on the number of evacuees; if the number changes, the choice may need to be revisited.

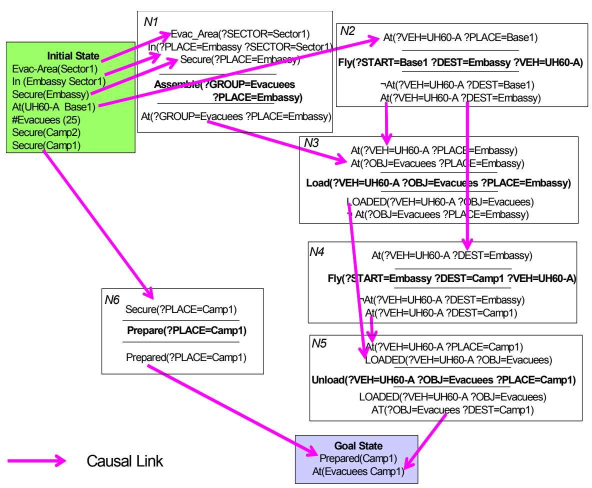

Figure 3. Plan with Complete Causal Link Relations

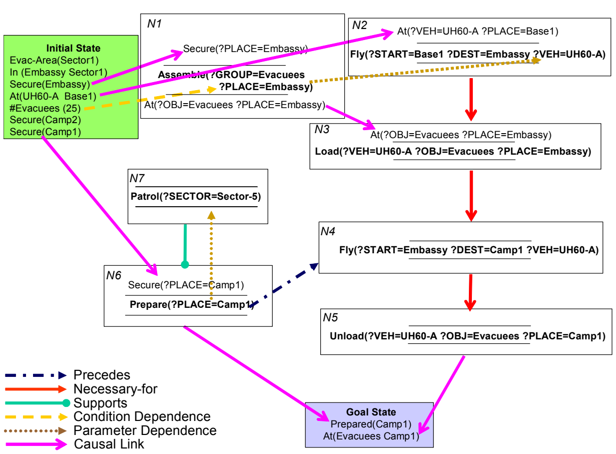

Figure 4. Plan with Key Qualitative and Causal Link Relations

- A *parameter-dependence* relation has been added from *?place* in *N1* to *?veh* in *N2* and a *condition-dependence* relation added from *#Evacuees(25)* in the initial world state to *?veh* in *N2*. These relations show that the choice of vehicle depends on both the number of evacuees and their assembly location, although the precise nature of the dependency is unknown.
- There is a *supports* relation from *N7* to *N6*, documenting that the *Patrol* action is being performed in service of the *Prepare* action. No comparable link is possible in the causal link view because there is no enabling relationship between effects produced by *N7* and required by *N6.* A *parameter-dependence* relation has been added from the *?place* variable in *N6* to the *?sector* variable in *N7* indicating that the choice of patrol area depends on the evacuation site. In this case, the relationship could be expressed algebraically by adding the constraint *?sector=Sector-of (?place)* to the *parameter-dependence* relation, yielding a *parameter-constraint* relation.

Action: *Assemble(?group ?place)*
Preconditions: *Secure(?place), Evac-Area(?sector), In(?place ?sector)*
Effects: *At(?group ?place)*

Action: *Fly(?start ?dest ?veh)*
Preconditions: *At(?veh ?start)*
Effects: *At(?veh ?dest), ¬At(?veh ?start)*

Action: *Load(?veh ?obj ?place)*
Preconditions: *At(?veh ?place), At(?obj ?place)*
Effects: *Loaded(?veh ?obj), ¬At(?obj ?place)*

Action: *Unload(?veh ?obj ?place)*
Preconditions: *At(?veh ?place), Loaded(?veh ?obj)*
Effects: *At(?obj ?place), ¬Loaded(?veh ?obj)*

Action: *Prepare(?place)*
Preconditions: *Secure(?place)*
Effects: *Prepared(?place)*

Action: *Patrol(?place)*
Effects: *Prepared (?place)*

Figure 5. Evacuation Planning Operators

Qualitative Calculus

The qualitative calculus determines the impact of changes on a plan. Here, we consider only changes to plan objects (parameters, actions, effects); more generally, changes to plan relations should also be considered.

We consider the set of plan changes listed below. Each entry lists a type of plan change along with the class of plan elements to which the change applies.

Change-Parameter: {Parameter}
Cancel: {Action | Effect}
Delay: {Action | Effect}
Move-forward: {Action | Effect}
Change-Truth-Value: {Effect}

Direct Effects of Plan Changes

Plan changes can have a range of effects on plan elements linked by qualitative and casual link relations. Our qualitative model considers a set that includes temporal, correctness, and minimality effects.

In terms of temporal effects, *delay* indicates that a plan element may be moved further back in time, while *move-up* indicates that a plan element may be moved forward in time. In terms of correctness, *disable* indicates that the plan element is no longer 'supported' within the plan. In terms of minimality, *obviate* indicates that the plan element may no longer be needed in the plan (e.g., it no longer plays a role within the plan). The effect *affect* indicates an impact about which no conclusions can be determined.

In general, qualitative reasoning will enable inference of potential rather than definite impact of a change. For this reason, we categorize qualitative effects according to the modalities of *necessity* and *possibility*, indicating whether the effect necessarily occurs or only may occur. Because of the loss of precision inherent to the qualitative models, it is generally difficult to establish necessary effects. However, consideration of explicit quantitative information enables inference of necessary effects in some cases.

Plan changes can have impacts that flow along both 'directions' of a plan relation. In particular, changing the source element in a relation could impact the target element (i.e., the *forward* direction) while changing the target could impact the source (i.e., the *backward* direction).

Below, we describe the direct impact of each type of change on objects related by the qualitative relations QR1 – QR5, as well as the *causal-link* and *parameter-constraint* relations. Figure 6 summarizes the results.

Cancel The *cancel* operation deletes an action or effect from a plan. Equivalently, it can be interpreted as having an action that fails at execution time or an effect that is violated. A *cancel* operation impacts only *necessary-for* and *causal-link* relations in the forward direction. If the *source-obj* of a *necessary-for* or *causal-link* relation is canceled, the *target-obj* would have a necessary enabling effect violated and hence would be disabled. Cancellation of the *target-obj* could potentially obviate the need for the *source-obj* for *necessary-for*, *supports*, and *causal-link* relations. However, that would be the case only if the *source-obj* was not also the source for a *necessary-for*, *supports*, *condition-dependence*, or *causal-link* relation for some other plan element.

Change-Parameter *Change-parameter* modifies an argument to an action or effect; it can impact elements related by *parameter-dependence* or *parameter-constraint* relations. If the source is changed and the target is an unbound parameter (or, the target is changed and the source is an unbound parameter), there is no qualitative impact. Below, we consider only instantiated parameters.

For the *parameter-dependence* relation, a change to *source-obj* could necessitate a change to *target-obj*, although the precise nature of the change could not be determined without additional information. Thus, the impact of a change in the forward direction is *affect*. In contrast, the *source-obj* is independent of the value of *target-obj*, so changes to the latter will not affect the former. To illustrate, consider a *parameter-dependence* relation from the number of evacuees in a NEO operation (*source-obj*) to the choice of assembly site (*target-obj*). If there is an increase in the number of evacuees, a planner may need to reconsider the chosen assembly site (i.e., it may be too small to accommodate the increased number of people). However, changing the assembly site would not require reconsideration of the number of evacuees involved (i.e., that number is beyond the control of the planner).

The impact on a *parameter-constraint* relation is similar in the forward direction but differs in the backward direction. The *parameter-constraint* relation is 'bidirectional', being grounded in an actual algebraic constraint. For this reason, changes to the *target-obj* of such a relation would yield a possible change to the *source-obj*; thus, the qualitative impact of such a change would be *affect.* To illustrate, consider a *parameter-constraint* relation between the parameters *?X* and *?R* defined by the equation $?X=Times(\pi,\ ?R^2)$. In this case, changing *?X* or *?R* will clearly impact the value of the other.

Change-Truth-Value The *change-truth-value* operation modifies the truth-value of an effect, from *true* to *false* or *false* to *true.* Such changes could impact plan elements linked by a *condition-dependence* relation. In the forward direction, change to the source effect of such a relation threatens the viability of the target, hence the impact is *disable*. Change to the target element could eliminate the need for the source effect, specifically when the target element is an effect and the change establishes the truth of the effect. In this case, the source object may no longer be required in the plan, provided that no other plan element depends on it (i.e., the effect is not the source object for some other qualitative relation). For this reason, the qualitative effect is listed as *obviate*. The same qualitative effects apply for the *causal-link* relation, which is a specialization of the *condition-dependence* relation.

Delay A *delay* operation impacts the three temporal qualitative relations *precedes, necessary-for,* and *supports*, as well as *causal-link* relations. In each case, *target-obj* will possibly be delayed when *source-obj* is delayed, but a delay to *target-obj* will have no impact on *source-obj*.

Move-Up A *move-up* operation impacts only the temporal qualitative relations (*precedes, necessary-for, supports*) and the *causal-link* relation, and only in the backward direction. The qualitative impact will be a possible *move-up* of *source-obj* for such relations.

FORWARD IMPACT	Delay	Move-up	Cancel	Change Param	Change TV
Precedes	*delay*	*	*	*	*
Necessary-for	*delay*	*	*disable*	*	*
Supports	*delay*	*	*	*	*
Parameter-dependence	*	*	*	*affect*	*
Condition-dependence	*	*	*	*	*disable*
Causal-link	*delay*	*	*disable*	*	*disable*
Parameter-constraint	*	*	*	*affect*	*

BACKWARD IMPACT	Delay	Move-up	Cancel	Change Param	Change TV
Precedes	*	*move-up*	*	*	*
Necessary-for	*	*move-up*	*obviate*	*	*
Supports	*	*move-up*	*obviate*	*	*
Parameter-dependence	*	*	*	*	*
Condition-dependence	*	*	*	*	*obviate*
Causal-link	*	*move-up*	*obviate*	*	*obviate*
Parameter-constraint	*	*	*	*affect*	*

Figure 6. Summary of Direct Qualitative Effects on Plan Objects

Change Propagation

Figure 6 summarizes the *direct* qualitative effects of plan changes. Such direct effects would need to be propagated across relations to establish the full qualitative impact of a change on a given plan. Propagation involves repeatedly treating each direct effect as a change, and then determining the new effects triggered by those changes. Although beyond the scope of this paper, it is straightforward to show that the propagation algorithm terminates in time polynomial to the length of the plan.

Example: Reasoning with Qualitative Effects

For the plan in Figure 4, we consider the qualitative impact of three changes.

1. Move up by 1 hour the start of the FLY action in *N4*.

There is a path of qualitative temporal relations from every action node in the plan to *N4*, except from *N5*. Thus, the propagated qualitative effects of this change consists of a *move-up* effect for each node except *N5*.

2. Revise the initial state predicate *#Evacuees(25)* to *#Evacuees(50)*

The direct qualitative effect of this change is an *affect* inference for the assembly site (i.e., *?place* in *N1*); as a result, the value for this parameter may need to change. Propagation of this direct effect would lead to an *affect* inference for the vehicle used to transport the evacuees (i.e., *?veh* in node *N2*) by virtue of the *parameter-dependence* link from *?place* in *N1*to *?veh* in node *N2.*

3. Eliminate *N7* from the plan.

No effects result (i.e., no other part of the plan critically depends on this node).

4. Change the goals to *Prepared(Camp2)* and *At(Evacuees Camp2)*.

No qualitative effects would be deduced in this case. The ideal, however, would be to infer that *?place* in nodes *N5* and *N6* should both change to *Camp2*, with further impacts then assessed.

Complete causal links support this form of reasoning, as they connect the changed parameters to the impacted variables. This inference can be attained within the qualitative framework provided that the user has specified the appropriate *parameter-dependence* and *parameter-constraint* relations with equality constraints among such codesignating parameters. However, humans are unlikely to document their plans to that level of detail.

A better solution is to augment the qualitative calculus with a mechanism that responds to the changing of a parameter value from *v1* to *v2* by marking all other occurrences of *v1* in the plan (as well as in the goal or initial conditions linked to any portion of the plan) as having a possible effect *affect* as a result of the change. In this example, the mechanism would then identify the *?place* parameters in nodes *N5* and *N6* as having the effect *affect,* as well as the initial state condition *Secure(Camp1).*

By binding those *?place* parameters to *Camp2,* the user would in turn trigger recalculation of the parameter *?sector* in node *N7* via the *parameter-dependence* relation with constraint *?sector=Sector-Of(Camp2).* The marking of the initial state condition *Secure(Camp1)* as having the effect *affect* would alert the user to the need to ensure the revised condition *Secure(Camp2)*. Because that condition holds in the initial state, no further actions would be needed.

Conclusions

Our framework for qualitative reasoning about the effects of changes on plans has several advantages over standard logical/deductive approaches. First, it does not require comprehensive and correct causal theories. While qualitative inferences can be drawn from incomplete models, however, more complete models will yield more informative results. Second, the qualitative relations are simpler and more intuitive to define, making it possible for users to annotate plans with qualitative relations that reflect their specific needs and interests. In contrast, traditional deductive approaches require sophisticated models that have proven to be difficult for users to formulate. Third, the qualitative models include relationships that do not require complete formalization of concepts, making them usable in situations where precise dependencies among plan elements cannot be articulated.

One shortcoming of our qualitative framework is that it is derived from commonsense notions of plan relations and interactions. It would be valuable to define a formal semantic model that grounds the qualitative relations in abstractions of plan structures. Such a model would enable us to show that the calculus 'does the right thing', rather than simply matching intuitions as to what should happen.

Acknowledgments This work was supported by DARPA under Air Force Research Laboratory contract F30602-00-C-0058. The author thanks Peter Jarvis, Mabry Tyson, and Michael Wolverton for their help in developing these ideas.

References

Allen, J. and Ferguson, G. (2002). "Human-Machine Collaborative Planning". In *Proc. of the 3rd Intl. NASA Workshop on Planning and Scheduling for Space*.

Carroll, J. M. and Moran, T. P. (1991). Special issue on design rationale. *Human-Computer Interaction*, 6(3-4).

Erol, K., Hendler, J., and Nau, D. (1994). Semantics for Hierarchical Task-Network Planning. Technical Report CS-TR-3239, Computer Science Department, University of Maryland.

El Fattah, Y. (2003). "Heuristic Causal Modelling for Mixed Initiative Planning". *Submitted for Publication.*

Kambhampati, S. and Hendler, J. (1992). "A Validation-structure Based Theory of Plan Modification and Reuse", *Artificial Intelligence*, 53(2).

Kim, J. and Blythe, J. (2003). "Supporting Plan Authoring and Analysis". In *Proc. of the 7th Intl. Conf. on Intelligent User Interfaces*.

McCarthy, J. Epistemological Problems of Artificial Intelligence. (1977). In *Proc. of the Intl. Joint Conf. on Artificial Intelligence*.

Myers, K. L., Tyson, W. M., Wolverton, M. J., Jarvis, P. A., Lee, T. J., and desJardins, M. (2002). "PASSAT: A User-centric Planning Framework". In *Proc. of the 3rd Intl. NASA Workshop on Planning and Scheduling for Space*.