

Explanations and Recommendations for Temporal Inconsistencies

John L. Bresina and Paul H. Morris

NASA Ames Research Center
Moffett Field, CA 94043

John.L.Bresina@nasa.gov and Paul.H.Morris@nasa.gov

Abstract

This paper presents an approach to automatically providing explanations and recommendations for temporal inconsistencies within the context of constraint-based, mixed-initiative planning. In this style of planning, a user operation can cause the planner's temporal constraint network to become inconsistent. For example, if a partial plan has been built and the user is trying to add another activity into the plan, an inconsistency will arise if the activity's associated constraints conflict with the current plan's constraints. In such cases, it is helpful if the system can explain to the user why the user's operation led to the inconsistency, and it is even better if the system can give the user recommendations on how to fix the inconsistency. Our approach to explanation and recommendation is generally applicable to constraint-based reasoning systems; however, the motivation and validation of our approach derive from a deployed planning system for planetary rover missions.

Introduction

The context of this paper is constraint-based, mixed-initiative planning. Inherent in constraint-based planning is the issue of handling inconsistencies that arise in the constraint network underlying the plan. Inconsistencies can be described by a (minimal) *nogood*, which identifies a (minimal) subset of the constraint network that is contributing to the problem. Inherent in mixed-initiative systems is the collaborative problem solving involving humans and machines. In this paper, we address the general issue of collaboratively resolving temporal inconsistencies that arise during the planning process.

We present a general approach to summarizing the essential aspects of nogoods to improve human understanding of temporal inconsistencies. In addition, we present an approach to generating recommendations on how the user can resolve the inconsistency. A successful recommendation restores consistency to the temporal network while achieving what the user was trying to do when the inconsistency arose.

Our motivation derived from the problem of activity planning for the Mars Exploration Rover (MER) mission and the system used to accomplish this task: MAPGEN, Mixed-initiative Activity Plan GENERator (Bresina, et al., 2005a). With the MAPGEN system, the primary role of the user is to direct and focus the plan construction process and to provide qualitative evaluation of plans. The system makes automated planning capabilities available to the user and performs potentially tedious tasks. The planning process is an incremental one in which the user interleaves automatic plan generation and plan editing phases.

The MER mission has been in operation for over a Martian year (more than two Earth years) and MAPGEN is still being employed for activity plan generation for the Spirit and Opportunity rovers. One of the key lessons learned from this experience (Bresina, et al., 2005b) is the need for the automated reasoning component to provide better explanations of its behavior. Especially important is explanation of failure, such as why an activity cannot fit in the plan at a particular time. Our explanation approach addresses this need; in addition, our recommendation approach advises the user on how to work around the failure in order to achieve the user's original intent.

As a starting point for our research, we have used a standalone version of the MAPGEN system. Furthermore, we have modified the MER planning domain model to drastically reduce the number of abstraction levels; in fact, all but a couple of activities have no expansion. However, the examples in this paper are based on real problem instances from the MER mission that have been translated into our modified domain.

In the rest of this paper, we first present further background information on the context of this work, primarily regarding the mixed-initiative, constraint-based planning approach that MAPGEN employs in the MER mission. We then give more details about, and examples of, the inconsistencies that can arise and present a unifying characterization of these inconsistencies. This characterization forms the basis of the explanation approach, which is described next. This is followed by a description of how recommendations are generated and presented to the user. We then briefly describe some related research, and conclude with some remarks on our contributions, current status, and planned next steps.

Background

In the MER mission, the daily commanding cycle proceeds as follows. The engineering and science data from the previous Martian day (sol) are analyzed to determine the status of the rover and its surroundings. Based on this, and on a strategic longer-term plan, the scientists determine a set of scientific objectives for the next sol. At this stage only rough resource guidance is available. Hence, the scientists are encouraged to oversubscribe to ensure that the rover's resources will be fully utilized in the final plan.

In the next step in the commanding process, the science observation requests are merged with the engineering requirements (e.g., testing the thermal profile of an actuator heater) and a detailed plan and schedule of activities is constructed for the upcoming sol. The plan

must obey all applicable *flight rules* that specify how to safely operate the rover and its instrument suite and remain within specified resource limitations. It is in this step that the Tactical Activity Planner (TAP) employs MAPGEN.

Once approved, the activity plan is used as the basis to create sequences of low-level commands, which drive onboard execution. This sequence structure is then validated, packaged, and communicated to the rover. This completes the commanding cycle.

The MAPGEN system is the result of integrating two legacy systems: APGEN, Activity Plan GENERator (Maldague, et al., 1998), and EUROPA, Extendable Uniform Remote Operations Planning Architecture (Frank & Jónsson, 2003). APGEN is an interactive activity plan editor developed by Jet Propulsion Laboratories and used as MAPGEN's graphical user interface. EUROPA is a constraint-based planning framework, developed at NASA Ames Research Center and used to provide MAPGEN's plan representation and reasoning capabilities. A joint JPL-Ames team developed and deployed MAPGEN.

In constraint-based planning, domain rules are specified in terms of activity/state patterns and constraint schemas. A given constraint schema is applied to any instance matching the associated pattern. The EUROPA framework performs sound constraint reasoning and provides a mechanism for firing applicable domain rules. Search methods and other techniques for manipulating partial plans then build on this framework. Partial plans have an underlying simple temporal constraint network, or STN, (Dechter, Meiri, & Pearl, 1991). The consistency of STNs can be determined by checking for arc consistency. Furthermore, each value in an arc-consistent temporal variable domain appears in at least one legal solution of the STN. The set of such values defines a temporal interval that can be represented by its bounds. It is not necessary to immediately ground the variables. Plans with temporal variables left ungrounded are called *flexible plans*; these are the type of plans that MAPGEN constructs. However, the plan that is displayed to the user is a *grounded plan*; i.e., a specific instantiation of the underlying flexible plan. This is chosen to be closest to a *reference schedule* that is initially set by the user and thereafter set to the previous grounded plan; the following is an outline of MAPGEN's solution grounding algorithm (Bresina, et al., 2005a).

For each timepoint x with reference position t do the following:

- (i) If t is within the STN bounds for x ,
then add a grounding constraint that sets x to t .
Else if t is less than the lower bound (lb) for x ,
then add a grounding constraint that sets x to lb.
Else if t is greater than the upper bound (ub) for x ,
then add a grounding constraint that sets x to ub.
- (ii) Propagate the effect of the new constraint.

Explicit temporal constraints fall into three main categories: *model* constraints, *problem-specific* constraints, and *expedient* constraints. The model constraints encompass definitional constraints and mutual-exclusion

flight rules. For example, "you cannot move the arm while the rover is moving" or "you cannot have more than one activity trying to simultaneously point the rover's mast".

The problem-specific constraints comprise relations between specific activities in a planning problem instance. In MER, these constraints were used to ensure that science objectives are satisfied and that the data collected were scientifically useful; thus, for this domain, we also refer to these constraints as *science constraints*. The scientists used two types of problem-specific constraints: temporal bounds and temporal ordering relations. The temporal bounds used were typically constraints on when an activity could start due to, for example, lighting conditions or temperature. The typical ordering relations used were constraints between the end of one activity and the start of another; for example, activity B must immediately follow activity A, or there must be a gap of at least 30 minutes between activities A and B. Temporal bounds are represented in the planner's constraint network as a temporal relation between a start (or end) point of an activity and a distinguished time point called the *origin*. Another type of problem-specific constraints, which are added by the TAP during planning are the *pin constraints*; they are described later.

Expedient constraints are (typically) added during search in automated planning; they result from decisions made to guarantee compliance with higher-level constraints that cannot be directly expressed in an STN. For example, a model constraint might specify that two activities, A and B, are mutually exclusive. This is really the disjunctive constraint that either A must precede B or B must precede A. The choice between the two is made based on heuristics, and is subject to backtracking during the original search.

Planning with MAPGEN

The role of mixed-initiative planning in MAPGEN is to support collaboration to build a high quality activity plan. The primary role of the operator (i.e., the TAP) is to direct and focus the plan construction process and to provide qualitative evaluation of plans. The intended collaboration is that the system handles constraint enforcement constantly in the background, while automated plan construction operations are user invoked.

One of the key design characteristics of MAPGEN is adjustable autonomy: MAPGEN provides a spectrum of automated planning services with different degrees of automation and human guidance. The TAPs tend to build plans in an incremental fashion, checking the energy resource usage by invoking an external power-thermal detailed modeler every now and then. The science requests that are not currently in the plan are kept in a separate display window, called the *hopper*. Due to this incremental approach, the TAPs often apply the *plan-selected* operation. With this operation, the user selects a set of observation requests not in the plan and request that these be inserted anywhere into the current partial plan,

such that all constraints were satisfied. The user can exercise even more control over the planning process via the *place-selected* operation, which is applicable only to individual activities. This operation allows the user to select an activity in the hopper and then choose an approximate temporal placement for it in the plan. The planning algorithm treats the user-chosen time as heuristic guidance and searches for a plan where the selected activity is as close to the desired time as possible.

The system also supports an activity movement operation, called *constrained-move*, which takes advantage of the flexibility in MAPGEN's plans. As long as an activity is moved only within the flexibility range defined by the domain in the underlying arc-consistent flexible plan, the result is necessarily another consistent instantiation. During a constrained move, the system actively restricts the movements of the selected activity to stay within the permitted range. Then, once the user places the activity, any dependent activity is automatically updated as necessary to yield a new valid plan instance. Note that the enforcement takes into account *all* the constraints, including expedient constraints that arbitrarily order mutually exclusive activities.

When extending the plan, the planner tries to minimally perturb the previous plan, i.e., move activities as little as possible from their previous positions. If the TAP wants a guarantee that a particular activity will not be moved at all, then the activity can be *pinned*. Each application of a planning operation is done in the context of the current plan and its constraints; thus, previous planning decisions affect what future operations are possible and what additional activities can be fit into the plan. This incremental commitment approach helps achieve a fast response time; however, it can lead to planning failures due to inconsistencies resulting from prior commitments.

Mixed-initiative planning systems must respond and return control quickly to the user. For an automated planning operation, which involves a cascading decision process, MAPGEN relaxes completeness in favor of responsiveness. This has to be done carefully to maximize chances of finding near-optimal solutions within limited time. We developed a backtracking algorithm that noted the difficulty of planning activities, and when the effort to plan an activity exceeded an allowance determined by its priority, the activity was rejected from the plan.

Temporal Inconsistencies

In the MER application, all the constraints, and, hence, all the inconsistencies, are temporal. There are several contexts in which inconsistencies can arise during planning. First, when an activity is considered for insertion, it may be inconsistent with the current plan even before any location is examined. When MAPGEN attempts to insert an additional activity into the plan, it first starts enforcing the constraints associated with that activity. Since the existing plan was formulated without those constraints, they may be inconsistent with previous

ordering decisions made to prevent forbidden overlaps. Second, the activity may be inconsistent with the *specific* location chosen in a place-selected operation. Third, it may be inconsistent with *each one* of the possible locations identified during a plan-selected operation.

The first context generates an immediate nogood in the underlying STN as soon as the constraints associated with the new activity are added. For example, the added constraints might require two activities in the existing plan to be in a different order. In the second context, a nogood can be extracted from the extended STN obtained by temporarily placing the activity in the infeasible location. The third context gives rise to nogoods in every one of the STN extensions arising from the different location choices. This can happen even though there is no immediate nogood, because distinct subsets of the added constraints may rule out the different choices. However, it may be possible to combine these to form a compound nogood that includes the disjunctive choice. In this paper we focus on the first context, which seemed to be a common source of TAP puzzlement. Moreover, within this context, we address a core set of issues also inherent in the other two.

When an inconsistency occurs, MAPGEN does extract a minimal nogood; however, the nogood typically involves complex chains of activities and constraints that could not easily be grasped by the user. For example, during MER, nogoods encountered during planning could involve hundreds of constraints. It is obviously impractical to expect a time-pressured TAP to read, let alone grasp the significance of, a lengthy nogood. Thus, our explanation technique attempts to summarize the essential aspects of nogoods for presentation to the user.

In the following description, a *time point* represents a start or an end of an activity, the origin (defined earlier), or a *sol event*. A sol event is a distinguished absolute time point; for example, the start or end time of the plan.

A *temporal nogood* can be defined as a cycle of time points such that the sum of upper bounds on the links between adjacent time points is negative (Dechter, Meiri, & Pearl, 1991). An equivalent, complementary¹, definition is as a cycle such that the sum of *lower* bounds is *positive*. This latter view turns out to be more useful for our application because most of the constraints that are active in inconsistencies are more naturally viewed as lower bounds. Additionally, this view yields an intuitive "no room" explanation for why activities cannot fit in the plan.

Given a nogood, we outline the summarization process as follows:

1. Categorize each constraint in the nogood.
2. Compress the nogood using the categories.
3. Separate into old and new constraint chains.
4. Present a digest of the two chains.
5. Generate a corrective recommendation.

¹ An upper bound of x between A and B is equivalent to a lower bound of $-x$ between B and A.

The first step in the summarization process categorizes each active constraint in the nogood based on its source; e.g., whether it is an expedient constraint imposed by the planner, a problem-specific constraint imposed by the scientists, an expansion constraint arising from the model, or a pin constraint (typically) imposed by the TAP.

Part of the reason the nogoods from MER were so lengthy was because the durations of high-level activities were derived from numerous low-level constraints that expanded them into more primitive sub-activities with fixed durations. If a high-level duration is active in an inconsistency, then the nogood involves all these low-level constraints. Similarly, the mutual exclusion rules are implemented by means of a hierarchy of planner subgoals that are ordered on timelines, and the low-level constraints between such subgoals are also included in the nogoods. This is illustrated in the following example. (In the examples, $[x] \rightarrow$ indicates an edge with a lower bound of x , and the constraint edges are labeled as follows: expedient (Planner), problem-specific (Science), activity expansion (Expand), duration (Dur), and pinned activity (Pin).)

```

Origin [181196592]→ Plan_Start          (Pin)
Plan_Start [0]→ Start APXS_1           (Science)
Start APXS_1 [0]→ Start ARM_MOVE_1     (Expand)
Start ARM_MOVE_1 [100]→ End ARM_MOVE_1 (Dur)
End ARM_MOVE_1 [0]→ Start APXS_ON_1    (Expand)
Start APXS_ON_1 [480]→ End APXS_ON_1   (Duration)
End APXS_ON_1 [0]→ Start APXS_ACQ_1    (Expand)
Start APXS_ACQ_1 [28800]→ End APXS_ACQ_1 (Dur)
End APXS_ACQ_1 [0]→ Start APXS_OFF_1   (Expand)
Start APXS_OFF_1 [60]→ End APXS_OFF_1  (Dur)
End APXS_OFF_1 [1719]→ End APXS_1     (Expand)
End APXS_1 [0]→ Start APXS_2          (Science)
Start APXS_2 [0]→ Start ARM_MOVE_2     (Expand)
Start ARM_MOVE_2 [100]→ End ARM_MOVE_2 (Dur)
End ARM_MOVE_2 [0]→ Start APXS_ON_2    (Expand)
Start APXS_ON_2 [480]→ End APXS_ON_2   (Dur)
End APXS_ON_2 [0]→ Start APXS_ACQ_2    (Expand)
Start APXS_ACQ_2 [28800]→ End APXS_ACQ_2 (Dur)
End APXS_ACQ_2 [0]→ Start APXS_OFF_2   (Expand)
Start APXS_OFF_2 [60]→ End APXS_OFF_2  (Dur)
End APXS_OFF_2 [0]→ End APXS_2        (Expand)
End APXS_2 [0]→ Start MB              (Science)
Start MB [0]→ Start ARM_MOVE_3         (Expand)
Start ARM_MOVE_3 [100]→ End ARM_MOVE_3 (Dur)
End ARM_MOVE_3 [0]→ Start MB_ON        (Expand)
Start MB_ON [120]→ End MB_ON           (Dur)
End MB_ON [420]→ End UHF_Mutex        (Planner)
End UHF_Mutex [0]→ Start UHF_Mutex     (Planner)
Start UHF_Mutex [0]→ Start UHF         (Planner)
Start UHF [-181241466]→ Origin         (Pin)

```

The second step in our inconsistency summarization process compresses the nogood by combining low-level sub-chains in the cycle that have the same source. For example, a chain of low-level edges that determine an expansion is combined into a single expansion edge whose lower bound is the sum of the component lower bounds. In

the following pseudo-code, “similar” category means both are “Planner” or both are “Expand” or “Dur.” The latter combines into “Expand.” The combination of $x[a] \rightarrow y$ and $y[b] \rightarrow z$ edges yields a single $x[a+b] \rightarrow z$ edge.

```

DO
  WHILE current and next edge have similar category
    DO combine into single edge;
    Advance to next edge;
  UNTIL current edge is first edge

```

For raw MER nogoods, the compression step typically achieves close to an order of magnitude reduction in size. The reduction is not as great for our modified model because the expansions have already been significantly flattened, but compression is, nevertheless, a useful step. For the example we consider to illustrate our approach, the original MER nogood contains 49 edges (i.e., active constraints). In our modified model, it contains 30 edges, shown above. The nogood compresses down to 9 edges, shown below.

```

Origin [181196592]→ Plan_Start          (Pin)
Plan_Start [0]→ Start APXS_1           (Science)
Start APXS_1 [31159]→ End APXS_1       (Expand)
End APXS_1 [0]→ Start APXS_2          (Science)
Start APXS_2 [29440]→ End APXS_2       (Expand)
End APXS_2 [0]→ Start MB              (Science)
Start MB [220]→ End MB_ON              (Expand)
End MB_ON [420]→ Start UHF            (Planner)
Start UHF [-181241466]→ Origin         (Pin)

```

APXS and MB refer to different spectrometers on the rover’s arm, and UHF refers to communication with an orbiter. Note that the final edge of this compressed nogood has a negative lower bound; this is equivalent to a positive upper bound in the opposite direction (i.e., from the origin to the start of UHF). Observe also that the nogood tends to alternate between Start/End activity-expansion constraints and End/Start activity-separation constraints, as is typical.

This nogood arose when the TAP tried to bring APXS_1 into the plan. Its science constraints conflicted with constraints associated with the other activities in the plan. We note that a previous planner decision ordered MB_ON (which is a sub-activity of the high-level MB) before the UHF communication, because they cannot overlap. This has the effect that there is not enough room for the APXS_1, which must come before the sequence of APXS_2 followed by MB.

Explanations

Given the types of temporal constraints, there are many types of nogoods that could theoretically arise; Figure 1 illustrates six abstract examples. In these examples, the TAP is trying to insert activity B into the plan, which causes the nogood; we refer to B as the goal activity. Activity B is constrained to be placed “between” A and C; more precisely, $\text{start}(B)$ is constrained by $\text{end}(A)$ and $\text{end}(B)$ is constrained by $\text{start}(C)$. Both A and C are




Figure 1: Abstract nogood examples

activities, except in case 4 where A is a sol event. Thick vertical lines indicate the temporal boundaries of the nogood. The arrows represent edges in the nogood; the dotted arrows represent edge-pairs through the origin.

The edges adjacent to B may have non-zero lower bounds; we refer to these as *buffers*. In case 4, the line on the arrow from end(B) to start(C) indicates a positive buffer; this means that C cannot start earlier than m time units *after* the end of B, where the buffer is $+m$. In case 6, there is a negative buffer on the edge from end(A) to start(B); this means that B cannot start earlier than m time units *before* the end of A, where the buffer is $-m$.

In case 5, the middle unlabelled block indicates a *backwards* chain of activities; in general, there could be several of these in a nogood, causing it to zigzag in direction. To understand this, observe that if we “cut” a nogood by removing the goal activity and its adjacent edges, the remaining edges fall into an ordered sequence. However, there can be subsequences that intersperse negative and zero lower bounds. These typically involve activities traversed in the reverse direction and give rise to backward blocks in the nogood. A negative path through the origin, as in cases 1, 2, and 4, may also be viewed as a

backward block. Because of backward blocks, A may precede C in the plan, as seen in several of the examples. However, except for paths through the origin, backward blocks are rare in MER.

Although the examples look quite different, it is remarkable that they can all be explained in a unified way. The intuition behind the explanation is that there is not enough “room between” A and C to fit B. Note that both the room needed and the room available can be either positive or negative depending on the buffers and on backward blocks. In order to make room, either C has to be moved to the right or A has to be moved to the left, or both. The amount of temporal distance that A and C have to be moved “apart” is determined by the duration of B as well as by the buffers (if any) in the edges from A to B and from B to C. This movement of A and/or C will cause one or more activities within the nogood’s boundaries to be moved outside of the boundaries.

We can thus divide the explanation into two parts. The first part describes why A needs to move relative to C, which involves only the constraints associated with the edges adjacent to B. For the cases we consider, the duration of the goal (i.e., B) is active in the inconsistency;

thus, one of the nogood edges represents the lower-bound constraint on the goal duration. The following pseudo-code generates this first part of the explanation.

Set dur = lower-bound (lb) of goal duration edge.
Set gap1 = lb of edge preceding goal duration edge.
Set tp1 = left time point of preceding edge.
Set gap2 = lb of edge following goal duration edge.
Set tp2 = right time point of following edge.

Output: "For <goal> to fit in the plan, the planner would need to slide <tp2> to no earlier than (gap1+dur+gap2) after <tp1> because of science constraints requiring <preceding-edge constraint> and <following-edge constraint>"

Cosmetic changes may be made to this output; for example, "-x after" would be replaced by "x before." The second part summarizes why A cannot be moved relative to C in the way desired because of the remaining portion of the nogood, which comes from constraints in the existing plan. The *nogood span* is defined as the sum of all the lower bounds in the nogood.

Output: "Currently, <tp2> is barred from going later than (nogood_span - (gap1+dur+gap2)) after <tp1> because of ..."

The output goes on to refer to the constraints in the remaining portion of the nogood. The following is the explanation generated for the example presented earlier.

For APXS_1 to fit in the plan, the planner would need to slide Start of APXS_2 to no earlier than 31159 after Plan_Start (because of science constraints requiring Start of APXS_1 to be no earlier than Plan_Start and Start of APXS_2 to be no earlier than End of APXS_1).

Currently, Start of APXS_2 is barred from going later than 14794 after Plan_Start because of planner orderings involving End of MB_ON before Start of UHF together with science constraints or pins.

This style of summarization obeys certain principles that we believe facilitate understanding. The first part of the explanation (why C needs to be moved) is given in detail. This follows the principles of *locality*, i.e., adjacent to the goal, and *recency*, i.e., newly activated constraints. The second part (why C cannot be so moved) spells out in detail only the planner constraints. This follows a principle of *visibility*. We can expect the TAP to have an internal mental model of the constraints in the unfolding plan; however, that picture usually does not include the expedient constraints, which are installed "behind the scenes" by the planner and are not graphically shown in MAPGEN's interface.

Recommendations

Following these explanation principles and presenting the user with a concise summary of the inconsistency is very useful, and it may be all the user needs to work around the problem. However, the system can offer further assistance to the user by recommending how to resolve the inconsistency and fit the goal activity into the plan.

When the science constraints are input to MAPGEN, they are checked for consistency, and if inconsistent, the user must correct them before using the planner. Hence, the recommendation approach focuses on the expedient constraints. Recall that an expedient constraint imposes an arbitrarily chosen order between activities in order to ensure that they do not overlap; thus, it may be possible to impose the reverse order instead. The idea is that by reversing one or more of the expedient orderings in the nogood, we might be able to facilitate the necessary activity movements, as specified in the explanation. This is the idea behind the recommendation; however, it must be expressed in user operations supported by MAPGEN. There is no user operation that directly deletes an expedient constraint; instead, such constraints are deleted as side effects of other operations.

One way the user can cause an expedient constraint to be deleted is by unplanning one of the (mutually exclusive) activities involved. However, as a side effect of the unplan operation, *re-linking* will occur, i.e., new expedient constraints will be imposed to maintain ordering decisions previously made for activities still in the plan. These new constraints may again frustrate the achievement of our intent, which is to get the goal activity into the plan. Hence, the only sure way to forestall such blockage is to eliminate all current and possible expedient constraints between planned activities in the nogood. This can be achieved by unplanning a subset of the activities from the nogood such that no mutual exclusion rules apply to the remaining planned activities in the nogood. Ideally, we want to unplan the minimal number of activities necessary.

The recommendation technique determines this minimal subset of activities to unplan as follows. The first step is to identify all the pairs of mutually exclusive activities in the nogood. Note that the goal activity is not planned at this point, so it is not considered in this step. The second step is to replace any lower-level activity in a pair with its top-level activity. The reason for this step is that, in MAPGEN, it is not possible to unplan only part of an activity expansion structure; i.e., the top-level activity and its descendants are either all in the plan or all in the hopper. The third step is to determine a minimal set cover of these pairs; i.e., the smallest set that contains a member from each pair. This third step is equivalent to the vertex-covering problem. Finding an optimal solution is an NP-hard problem; for tractability, we use a standard greedy approach, based on (Cormen, et al., 1990), which guarantees the result to be no worse than twice the optimal. The recommendation presented to the user is the following:

1. Unplan the activities in the computed set cover; this places them back in the hopper.
2. Apply plan-selected to the goal activity.
3. Replan the previously unplanned activities, using the plan-selected operation on the set.

However, this recommendation is not guaranteed to succeed; it may fail due to conflicts beyond the scope of the nogood or due to aspects of the planning process, e.g., the order in which a set of activities is (re)planned. If the recommendation is unsuccessful, the user can try the following alternative recommendation:

1. *Set relaxed enforcement mode.*
2. Move A left and/or move C right until they are at the recommended relative positions.
3. *Set full enforcement mode.*
4. Apply the plan-selected operation to the goal activity.

The first step, in MAPGEN, deletes all the expedient constraints, allowing the user to move activities much more freely than the constrained-move operation; however, the planner still enforces the science constraints in this mode. The third step re-establishes whatever new expedient constraints are necessary to enforce the flight rules. The A and C mentioned in step 2 are defined as in Figure 1, i.e., the goal is intended to fit “between” A and C, and the recommended distance is presented in the explanation. Our recommendation process currently presents the user with both of these alternatives in case the first one does not succeed. The first alternative is preferred, if it works, because it involves less potential change to the existing plan. If the alternate recommendation also fails, then the user will most likely have to sacrifice some of the previously planned activities in order to get the goal activity into the plan. This tradeoff involves human-level scientific judgment.

We now illustrate this recommendation process with our previous example nogood. The first step finds two mutual exclusive pairs of planned activities: (APXS_2, MB) and (MB_ON, UHF). The second step converts the pairs to (APXS_2, MB) and (MB, UHF). The third step computes a minimal set cover of {MB}. The recommendation generated is then the following:

1. Unplan the following activities: MB.
2. Plan the goal.
3. Try to replan the activities.

If that doesn't succeed, try this alternative:

1. Go into relaxed enforcement mode.
2. Move Plan_Start and/or APXS_2 so that Start of APXS_2 is no earlier than 31159 after Plan_Start.
3. Go into full enforcement mode.
4. Plan the goal activity.

In this case, the first part of the recommendation does work; that is, as a result of following the prescribed steps, the goal activity APXS_1 fits into the plan without having

to sacrifice any activity that had been in the plan. The order between MB_ON and UHF is reversed, i.e., MB_ON is moved to the right of UHF, thus creating enough room on the left of the UHF for both of the APXS activities.

Related Research

There has been work related to explanations in a wide variety of research areas, including expert systems (e.g., Wick & Slagle, 1989), case-based reasoning (e.g., Sørmo, Cassen, & Aamodt, 2005), web-based systems (e.g., McGuinness & da Silva, 2004), and constraint-based reasoning. Within this last research area, many earlier systems tended to indicate or “explain” a failure, exhibited as an inconsistency in a constraint network, by outputting a complex, low-level nogood. However, as mixed-initiative systems are becoming more common, more attention is being paid to helping the user understand and resolve such failures. Work on improved explanation facilities has been carried out within different constraint-based research areas, including constraint programming (e.g. Ouis, Jussien, & Boizumault, 2002), multi-agent planning (e.g. Burstein, Ferguson, & Allen, 2000), and scheduling (e.g. Smith, et al., 2004).

This last referenced effort by Smith, et al. is probably the most similar to the work reported in this paper. Their approach employs a *filtering* step to reduce the constraints in the nogood. This step bears some similarity to our compression step, in that both eliminate the details of the expansion (or decomposition) structure in the domain model. Whereas our recommendation approach focuses exclusively on expedient constraints, their approach considers both *search constraints* (which would include expedient constraints) as well as constraints resulting from user decisions. The greatest difference between the two approaches is in the types of resolution methods recommended. Their resolution options all seem to involve user-alteration of specific domain constraints; e.g., changing task duration, due-date, or release-date, which relax the problem specification. In contrast, our generated recommendations are in terms of high-level planning operations that the user can perform; such recommendations change the solution rather than the problem specification.

Concluding Remarks

Our current implementation is integrated with our version of MAPGEN such that our explanations and recommendations pop up in an APGEN dialog window when the user's planning operation causes an inconsistency in EUROPA's constraint network. We are still in the process of mining the examples of nogoods generated by MAPGEN during MER operations in order to evaluate the generality of the implementation and the quality of the recommendations in terms of how often do they solve the current problem. We would also like to evaluate the

quality of the explanations, which would require a careful study of MAPGEN users. The overall objective of this work is to make the TAPs more effective users of MAPGEN. Hence, carrying out a user study that would determine whether we have achieved this objective would be ideal, but non-trivial. We are currently involved with the design of a new ground operations system that is baselined for the Phoenix and Mars Science Laboratory (MSL) missions. The possibility exists to import the explanation and recommendation techniques into this new system and to collect feedback from mission operations engineers during readiness exercises and during mission operations.

In addition to pursuing evaluations of the techniques, there are also some future research directions that we'd like to explore. One of the extensions we are investigating is to combine related inconsistencies into one explanation and recommendation. Consider the example of planning three goal activities that are constrained to be in sequence, and there is no room in the current plan for any of them to fit. Currently, the implementation would explain to the user how to make room for each goal activity; whereas, it would be better to generate a single recommendation that made room for all three goal activities.

One of the issues we are currently debating is validation of recommendations. Currently, the system does not test out the recommendations before presenting them to the user; rather it leaves it up to the user to choose whether to follow the recommendation and, if so, to try it out. There are different degrees of validation that could be carried out by the system. Making progress on resolving one inconsistency may cause another inconsistency; thus, full validation would involve chasing down all the inconsistencies until the plan was completely repaired. Since we want the user to have the initiative in the resolution decision, it does not make sense for the system to spend much effort on plan repair just to generate a recommendation. Furthermore, the user gains understanding of the problem by dealing with the multiple inconsistencies one at a time.

In this paper, we have presented a general approach for a mixed-initiative planning system to explain temporal constraint inconsistencies and recommend to the user how to resolve them. The foundation of our approach is a unified characterization of the varied types of nogood structures. Although we were motivated by the MER mission experience and we validated our approach by incorporating it into a version of MAPGEN, it should be easy to adapt it for many other mixed-initiative systems that employ temporal reasoning. Further, we expect the principles underlying our approach to be widely applicable.

Acknowledgements: This work was funded by NASA, initially under the SISIM program and then the ETDP. The authors would like to acknowledge the rest of the MAPGEN team: Mitch Ai-Chang, Len Charest, Brian Chafin, Adam Chase, Kim Farrell, Jennifer Hsu, Ari Jónsson, Bob Kanefsky, Adans Ko, Pierre Maldague, Kanna Rajan, Richard Springer, and Jeffrey Yglesias.

Secondly, we would like to acknowledge our collaborators, on this research project, for helpful discussions: Lina Khatib, Conor McGann, Karen Myers, and Michael Wolverton. We want to again thank Ari Jónsson, who was an early collaborator on this project and produced the standalone version of MAPGEN that we use.

References

- J. Bresina, A. Jonsson, P. Morris, & K. Rajan (2005a). Activity Planning for the Mars Exploration Rovers. In *Proc. of Fourteenth International Conference on Automated Planning and Scheduling*.
- J. Bresina, A. Jonsson, P. Morris, & K. Rajan (2005b). Mixed-Initiative Planning in MAPGEN: Capabilities and Shortcomings. In *Workshop on Mixed-Initiative Planning and Scheduling, ICAPS05*.
- M. Burstein, G. Ferguson, & J. Allen (2000). Integrating Agent-Based Mixed-Initiative Control with an Existing Multi-Agent Planning System. In *Fourth International Conference on Multi-Agent Systems (ICMAS'00)*.
- T. Cormen, C. Lieserson, & R. Rivest (1990). *Introduction to Algorithms*. MIT Press/McGraw-Hill 1990.
- R. Dechter, I. Meiri, & J. Pearl (1991). Temporal constraint networks. *Artificial Intelligence*, 49:61–95.
- J. Frank & A. Jonsson (2003) Constraint-Based Interval and Attribute Planning, *Journal of Constraints Special Issue on Constraints and Planning*.
- Deborah L. McGuinness & Paulo Pinheiro da Silva (2004). Explaining Answers from the Semantic Web: The Inference Web Approach. *Journal of Web Semantics*. Vol.1 No.4., pages 397-413.
- P. Maldague, A. Ko., D. Page, & T. Starbird (1998). APGEN: A multi-mission semi-automated planning tool. *First International NASA Workshop on Planning and Scheduling*, Oxnard, CA.
- S. Ouis, N. Jussien, & P. Boizumault (2002). COINS: a constraint-based interactive solving system. In Alexandre Tessier (ed), *Proc. of the Twelfth International Workshop on Logic Programming Environments*.
- S.F. Smith, G. Cortellessa, D.W. Hildum, & C.M. Ohler (2004). Using a Scheduling Domain Ontology to Compute User-oriented Explanations. In *Proc. of the Workshop Planning and Scheduling: Bridging Theory to Practice*, ECAI-04, Valencia, Spain, pp 85-94.
- Frode Sørmo, Jörg Cassens, & Agnar Aamodt (2005). Explanation in Case-Based Reasoning – Perspectives and Goals. *Artificial Intelligence Review*; Volume 24, number 2, pp. 109-143.
- M.R. Wick & J.R. Slagle (1989). An explanation facility for today's expert systems. *IEEE Expert*, Volume 4, Issue 1, Spring 1989 Page(s):26 – 36.