

**MONITORING THE EXECUTION OF
PLANS IN SIPE**

October 11, 1984

Technical Note 341

By: David E. Wilkins

Artificial Intelligence Center
Computer Science and Technology Division

**APPROVED FOR PUBLIC RELEASE:
DISTRIBUTION UNLIMITED**

This research was supported in part by the Air
Force Office of Scientific Research under
Contract F49620-79-C-0188.

ABSTRACT

In real-world domains (a mobile robot is used as a motivating example), things do not always proceed as planned. Therefore it is important to develop better execution-monitoring techniques and replanning capabilities. This paper describes the execution-monitoring and replanning capabilities of the SIPE planning system. (SIPE assumes that new information to the execution monitor is in the form of predicates, thus avoiding the difficult problem of how to generate these predicates from information provided by sensors.) The execution-monitoring module takes advantage of the rich structure of SIPE plans (including a description of the plan rationale), and is intimately connected with the planner, which can be called as a subroutine. The major advantages of embedding the replanner within the planning system itself are: 1) The replanning module can take advantage of the efficient frame reasoning mechanisms in SIPE to quickly discover problems and potential fixes, 2) The deductive capabilities of SIPE are used to provide a reasonable solution to the truth maintenance problem, and 3) The planner can be called as a subroutine to solve problems after the replanning module has inserted new goals in the plan. Another important contribution is the development of a general set of replanning actions that will form the basis for a language capable of specifying error-recovery operators, and a general replanning capability that has been implemented using these actions.

1. Introduction

A principal goal of our research in planning and plan execution is the development of a domain-independent, heuristic system that can plan an activity and then monitor the execution of that plan. Over the last two years we have designed and implemented such a system, SIPE (System for Interactive Planning and Execution Monitoring).¹ The basic approach to planning is to work within the hierarchical-planning paradigm, representing plans in procedural networks – as has been done in NOAH [2] and other systems. Several extensions of previous planning systems have been implemented, including the development of a perspicuous formalism for describing operators and objects, the use of constraints for the partial description of objects, the creation of mechanisms that permit concurrent exploration of alternative plans, the incorporation of heuristics for reasoning about resources, and the creation of mechanisms that make it possible to perform deductions.

Given a description of the world, and a set of operators that it can apply, SIPE can generate a plan to achieve a goal in the given world. However, in real-world domains, things do not always proceed as planned. Therefore, it is desirable to develop better execution-monitoring techniques and better capabilities to replan when things do not go as expected. In complex domains it becomes increasingly important to use as much as possible of the old plan, rather than to start all over when things go wrong.

This paper describes the execution-monitoring and replanning abilities that have recently been incorporated into the SIPE system. The particular advantages that can be obtained by the use of the rich structure in our plan representation are shown, as well as more general problems. The environment of a mobile robot has been used as a motivating domain in the development of some of the abilities here, though implementation has been in a general, domain-independent manner. This document does not describe resources, constraints, plan generation, and other features of SIPE, nor does it attempt to justify the basic assumptions behind the system. The interested reader is referred to [5] for this.

The problem we are addressing is the following: given a plan, a world description, and some appropriate description of an unexpected situation that occurs during execution of the plan,

¹The research reported here is supported by Air Force Office of Scientific Research Contract F49620-79-C-0188.

transform the plan, retaining as much of the old plan as is reasonable, into a plan that will still accomplish the original goal from the current situation. This process can be divided into four steps - 1) discovering or inputting the information about the current situation, 2) determining the problems this causes in the plan, if any, (similarly, determining shortcuts that could be taken in the plan after unexpected but helpful events), 3) creating "fixes" that change the old plan, possibly by deleting part of it and inserting some newly created subplan, and 4) determining whether any changes made by the above fixes will conflict with remaining parts of the old plan. Steps 2 and 4, and possibly 3, involve solving a truth maintenance problem since it will be necessary to determine which aspects of a situation are necessary for later parts of the plan. In SIPE, step 4 becomes part of step 3 as only fixes that are guaranteed to work are produced. In addition, serendipitous effects are used to shorten the original plan in certain cases.

The major contributions of the execution-monitoring and replanning module in SIPE result from taking advantage of the system's rich plan representation and from imbedding it within the planning system itself, rather than implementing it as an independent module. This provides a number of advantages, of which the most important follow: 1) the replanning module can take advantage of the efficient frame reasoning mechanisms in SIPE to quickly discover problems and potential fixes, 2) the deductive capabilities of SIPE are used to provide a reasonable solution to the truth maintenance problem, and 3) the planner can be called as a subroutine to solve problems after the replanning module has inserted new goals in the plan. Another important contribution is the development of a general set of replanning actions that will form the basis for a language capable of specifying error-recovery operators (see Sections 5 and 6). A general replanning capability has been implemented using these actions.

SIPE assumes that information provided about unexpected events is correct and, to a certain extent, complete. This assumption avoids many of the hardest problems involved in getting a planner such as SIPE to control a mobile robot. The difficult problem of how to generate correct predicates from information provided by the sensors is not addressed. We expect the translation of the information from the robot's sensors (e.g., the pixels from the camera or the range information from the ultrasound) into the higher-level predicates used by the planner to be crucial to the

application of a SIPE-like planner to a mobile robot. We hope to address this problem in the near future.

In a mobile robot domain, it may often be important to spend considerable effort in checking for other things that might have gone wrong in addition to the unexpected occurrence already noticed. There is a large tradeoff here as interpreting visual input of unexpected scenes may be expensive. The research described here also does not address this problem as it assumes that the minimum is wrong in accordance with the information that has been given (after taking deductive operators into account). The problem of uncertain or unreliable sensors or information is largely unaddressed, except that some predicates and variables may be specified as unknown. What is discussed here is what to do with new information in the form of predicates (assuming you have somehow discovered such predicates). Replanning appropriately with such information is a necessary part of the overall solution.

Section 2 of this paper describes how plans are represented in SIPE. To describe unexpected situations, a user (currently a human, but eventually this may be a program controlling and interpreting the robot's sensors) can enter arbitrary predicates at any point in the execution or can specify certain things as unknown. Section 3 describes the details of this process. Once the description of the unexpected situation has been accumulated, the execution monitor calls a problem recognizer, described in Section 4, which returns a list of all the problems it detects in the plan.

In general, recovering from an arbitrary error is a very hard problem. Often very little of the existing plan can be reused. One can always fall back on solving the original problem in the new situation, ignoring the plan that was being executed. The replanning part of SIPE, however, tries to change the old plan, keeping as much of it as possible. Since the general problem is so difficult, one would not expect very impressive performance from a replanner that did not have domain-specific information for dealing with errors. For this reason, we have implemented a number of general replanning actions (i.e., actions that modify a plan in ways that are useful for handling unexpected situations) in SIPE that can be referenced in a language for providing domain-specific error-recovery instructions. In many domains, the types of errors that are commonly encountered

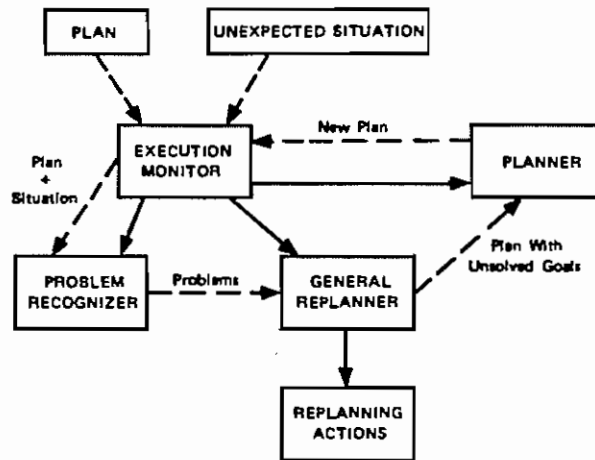


Figure 1
Control and Data Flow in SIPE Modules

can be predicted (e.g., the robot arm dropping something it was holding, or missing something it was trying to grasp). The user can then specify error-recovery operators for these errors, using SIPE'S replanning actions, to take appropriate action after expected errors.

In addition, SIPE provides a general replanning ability that can be applied in the general case and when there are no specific instructions. It is given the list of problems found by the problem recognizer, and tries certain replanning actions in various cases, but will not always find a solution. The replanning actions are described in Section 5 and the general replanner in Section 6. The general replanner changes the plan so that it will look like an unsolved problem to the standard planner in SIPE (e.g., by inserting new goals). After the replanner has dealt with all the problems that were found, then the planner is called on the plan (which now includes unsolved goals) and if it produces a new plan, this new plan should correctly solve all the problems that were found. Section 7 shows examples of the general replanner in operation.

Figure 1 shows the various modules in the SIPE execution-monitoring system. The solid arrows show which modules call which others. The broken arrows show the flow of data and information through the system as it replans for an unexpected situation. These arrows are labeled with a description of the data being passed.

2. Plans in SIPE

Plans in SIPE are represented as procedural networks [2] with temporal information encoded in the predecessor and successor links between nodes. The plan rationale is of primary importance to the execution monitor and is encoded in the network by MAINSTEP links between nodes, and by the use of PRECONDITION nodes (described below). MAINSTEP links describe how long each condition that has been achieved must be maintained. A context must also be given to completely specify a plan, as the network contains choice points from which alternative plans branch. The types of nodes that occur in plans are described below to the extent necessary for understanding the execution-monitoring capabilities.

SPLIT and JOIN nodes provide for parallel actions. SPLITs have multiple successors and JOINs have multiple predecessors so that partially ordered plans can be produced. JOIN nodes have a *Parallel-Postcondition* slot, which specifies the predicates that must all be true in the situation represented by the JOIN node. If a JOIN node originally has N predecessors, then there will be N conjunctions of predicates that must all be true at the JOIN node. (Some branches may have been linearized so there may be fewer than N predecessors after planning.) It is easier to record this at the JOIN node (than by having previous nodes point to the JOIN as their purpose), since a failed parallel-postcondition can more easily be retried during execution monitoring if there is easy access to all parallel-postconditions. The Parallel-Postcondition slot is filled only when the JOIN is first introduced into the plan - it is not updated as more detailed levels of the hierarchy are expanded. As long as the highest level predicates are as desired, it is assumed that the lower level predicates are irrelevant.

COND, ENDCOND, and CONDPATTERN nodes implement conditional plans. COND and ENDCOND are similar to SPLIT and JOIN, but each successor of the COND begins with a CONDPATTERN node that determines which successor will be executed.

CHOICE nodes denote branching points in the search space. They have multiple successors, but the context selects one of these as being in the current plan. Constraints on variables may be posted relative to this choicepoint. Thus, if the part of a plan after a CHOICE node is removed, the corresponding choicepoint in the context should also be removed from the context so that

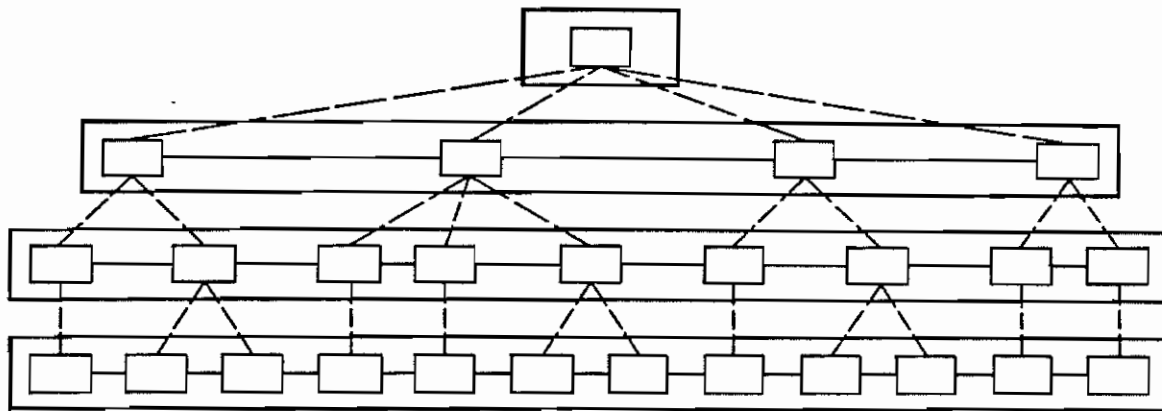
constraints that are no longer valid will be ignored.

GOAL nodes do not occur in final plans as they represent open problems that have not been solved yet. A GOAL node specifies a predicate that is a goal to achieve but which is not true in the situation represented by its location in the procedural network. Replanning actions will insert GOAL nodes in the plan. Each GOAL node has a MAINSTEP slot, which denotes a point later in the plan that depends on the GOAL. (This describes the rationale for having the GOAL in the plan.) Each goal must be maintained as true until the node which is its MAINSTEP is executed. A MAINSTEP slot can have the atom PURPOSE as its value, denoting that the given predicate is the main purpose of the plan, and not preparation for some later action.


PHANTOM nodes are similar to GOAL nodes except that they are already true in the situation represented by their location in the procedural network. They are part of the plan because their truth must be monitored as the plan is being executed. They also contain MAINSTEP slots.

PROCESS nodes represent actions to be performed during execution of the plan, and also have MAINSTEP slots as do PHANTOM and GOAL nodes. In a final plan, all PROCESS nodes will denote primitive actions. (There are also CHOICEPROCESS nodes, which are like PROCESS nodes except that they have a list of actions, one of which must be performed.)

PRECONDITION nodes provide a list of predicates that must be true in the situation represented by their location in the procedural network. Operators may specify preconditions that must obtain in the world state before the operator can be applied. The concept of precondition here differs from its counterpart in some planners, since the system will make no effort to make the precondition true. A false precondition simply means that the operator is not appropriate. Conditions that the planner should make true (and therefore backward chain on) can be expressed as goal or process nodes. This effectively encodes metaknowledge about how to achieve goals as SIPE will use any means to solve a goal node, only the operators listed to solve a process node, and no operators will be tried to solve a PRECONDITION node. Thus, a precondition becoming false does not mean that it should be made into a goal; rather it means that the whole part of the plan produced by the operator producing this precondition is invalid. Such nodes help encode the rationale of a plan since they effectively mean that the part of the plan associated with them (see



(a) plans at different levels



(b) wedges used by the execution monitor

Figure 2
SIPE Plan Viewed from Different Perspectives

below) was produced on the assumption that the predicates in the precondition are true.

In addition to the “horizontal” MAINSTEP, predecessor, and successor links within one level of a plan, there are “vertical” links between different levels of the hierarchy. Each node that is expanded by the application of an operator has descendant links to each node so produced. The descendant nodes in turn have ancestor links back to the original node one level higher in the hierarchy. Starting with a node that was expanded by an operator application, a *wedge* of the plan is determined by following all its descendant links (in the current context) repeatedly (i.e., including descendants of descendants, etc.) to the lowest level. (This definition of wedges is the

same as that used by Sacerdoti in [2].) Figure 2 depicts this graphically, with the large boxes in part (b) representing wedges. The node originally expanded by an operator application is called the *top* of the wedge. A wedge with its top at a high level in the hierarchy will generally contain many lower level wedges within it, and the only nodes that can be the tops of wedges are GOAL, PROCESS, and CHOICEPROCESS nodes.

Since PRECONDITION nodes are created only when an operator is applied, the part of a plan associated with a PRECONDITION node can be found by traversing up the ancestor links to the point where the precondition first became part of the plan (once inserted, PRECONDITION nodes are copied down from level to level). The node that was expanded by an operator to create this precondition is one level higher than where the first PRECONDITION node appears and is the top of the wedge associated with each of the PRECONDITION nodes that are copied from this first one.

3. The Input of Unexpected Situations

During execution of a plan in SIPE, some person or computer system monitoring the execution can specify what actions have been performed and what changes have occurred in the domain being modeled. SIPE permanently changes its original world model to show the effects of actions that have already been performed. At any point during execution, the system will accept two types of information about the domain: 1) an arbitrary predicate whose arguments are ground instances that is now true, false or unknown, and 2) a local variable name that is now unknown. SIPE first checks whether the truth values for the new predicates are different from its expectations, and, if they are, it applies its deductive operators to them to deduce more changed predicates.

It is important to note that the inputting of predicates does not solve the "pixels to predicates" problem, which is the crucial problem in using a planner such as SIPE to control the actions of a robot. This problem involves translating the input of the robot's sensors (e.g., the pixels from the camera or the range information from the ultrasound) into the higher level predicates used by the planner. The research described here involves what to do with the predicates once they have been determined but does not address the question of how to determine them automatically. We hope

to address this latter problem in the near future.

3.1 Unknowns

Unknowns are a new addition to SIPE as it previously assumed complete knowledge of the world. Having unknown quantities constitutes a fundamental modification since even the method of determining whether a predicate is true must be changed. If the truth values of critical predicates are unknown, the planner will quickly fail since none of the operators will be applicable. (Neither a negated or an unnegated predicate in a precondition will match an unknown predicate.) Operators can require predicates to be unknown as part of their precondition, in case there are appropriate actions to take when things are uncertain. Conditional plans have also been implemented as part of the execution-monitoring package in SIPE, so an operator might produce a plan with an action to perceive the unknown value, followed by a conditional plan that specifies the correct course of action for each possible outcome of the perception action. The deductive capabilities have also been enhanced so that operators can deduce that something is unknown.

The ability to specify variables as unknown is simply a tool provided by the system that will hopefully be useful in some domains, and particularly in a mobile robot domain. The idea behind this tool is that the location of an object may become unknown during execution. Rather than make predicates unknown, which may cause the application of operators to fail, we simply say that the variable representing the location is instantiated to the atom UNKNOWN rather than to its original location. All predicates with have this variable as an argument may then still match as if they were true. Thus the system can continue planning as if the location were known. The only restriction is that no action can be executed that uses an unknown variable as an argument. When such an action is to be executed (e.g., go to LOCATION1) then the actual instantiation of the variable must be determined before executing the action (possibly through a perception action). Note that it would be incorrect to continue planning if the truth values of important predicates depended on the instantiation of the location variable. It is the responsibility of the user not to use this tool (i.e., the unknown variable) if predicates depend on the variable's value. This tool may or may not prove useful in practice.

3.2 Interpreting the input

SIPE assumes that the minimum is wrong in accordance with the information that has been given (after taking deductive operators into account). Alternatively, we could decide on some basis (which would have to be provided as part of the domain-specific description) how much effort to spend with perception actions to see if more than the minimum has gone wrong. For example, if we are told that (ON A B) is not true when we expected it to be, we might want to look if B is where we thought it was. As it is, SIPE will just deduce that B is clear (if no other block is on B) and will not try to execute actions to make further checks about the world. Doing the latter could be very expensive for a mobile robot without good domain-specific knowledge about what was worth checking.

The user need not report all predicates that have changed since many of these may be deduced by SIPE's deductive operators. The system's deductive power has been increased recently (see next section) so many effects can be deduced from certain critical predicates. There is a problem in deciding how the unexpected effects interact with the effects of the action that was currently being executed (e.g., did they happen before, during, or after the expected effects?). Our solution to this problem is to assume the action took place as expected and to simply insert a "Mother Nature" action after it that is assumed to bring about the unexpected effects (and things deduced from them). The system assumes that any effects of the action being executed that did not actually take place are either provided or can be deduced from the information that is provided. This solution interfaces cleanly and elegantly with the rest of the planner and avoids having to model how the unexpected effects might interact with the expected effects.

4. Finding Problems in a Plan

Having just inserted a MOTHER-NATURE node (MN node) in a plan being executed, SIPE must now determine how the effects of this node affect the remainder of the plan. This involves solving the truth maintenance problem, since it is necessary to know on which facts the remainder of the plan depended. This is discussed later in this section. Because of the rich information content in the plan representation (including the plan rationale), there are only six problems that

must be checked. These are discussed below. All problems in the remainder of the plan are found, and this list is later given to the general replanner, which attempts to change the plan to solve these problems.

1 - *Purpose not achieved.* If the MN node negates any of the main effects of the action just executed, then there is a problem. The main effects must be reacheived.

2 - *Previous phantoms not maintained.* SIPE keeps a list of phantom nodes that occur before the current execution point, and whose MAINSTEP slot specifies a point in the plan that has not been executed yet. These are phantoms that must be maintained. If the MN node negates any of these, then there is a problem. The phantoms that are no longer true must be reacheived.

3 - *Process node using unknown variable as argument.* If a variable has been declared as unknown, then the first action using it as an argument must be preceded by a perception action for determining the value of the variable (see Section 3).

4 - *Future phantoms no longer true.* A phantom node after the current execution point may no longer be true. It must be changed to a GOAL node so that the planner will try to achieve it.

5 - *Future precondition no longer true.* A PRECONDITION node after the current execution point may no longer be true. In this case, we do not want to reacheive it, but rather pop up the hierarchy and perform some alternative action to achieve the goal at that level of the hierarchy.

6 - *Parallel-postcondition not true.* All the parallel-postconditions may no longer be true at a JOIN node. (This could be handled by maintaining phantoms, but is more convenient to handle separately.) In this case, we must insert a set of parallel goals after the JOIN, one for each untrue parallel-postcondition. The parallel-postconditions of the new JOIN will be the same as those on the old JOIN.

Note that only the last three problems below require truth maintenance since only they must know the truth value of predicates in situations after the current execution point. In addition to the above problems, possible serendipitous effects are also noticed and included in the list of problems. If the main effect of some action later in the plan is true before the action is executed, then that is noted as a possible place to shorten the plan (this is discussed in more detail in the next section).

Because of the way plans are encoded in SIPE, these are the only things that need be checked for determining if a MN node affects the remainder of a plan. It should be noted that this depends upon the fact that processes (actions) are assumed to work whenever their precondition is true and when all phantoms whose MAINSTEP slot points to the process are true. (All such necessary conditions should be encoded as either preconditions or goals in any case.) There is currently no check for loops in case the same error happens repeatedly with the same fix proposed each time. Various simple checks could easily be added if this were a problem.

4.1 Solution to the Truth Maintenance Problem

SIPE's solution to the truth maintenance problem is based on the efficiency of its deductive capability. Since it is assumed that processes work as expected whenever their precondition is true and all phantoms whose MAINSTEP slot points to the process are true, only the deduced effects need to be checked for their dependence on unexpected effects. (The execution monitor will solve problems having to do with preconditions and phantoms that are not true).

SIPE's deductive capability was designed to find a good balance between expressiveness and efficiency. While providing the power of many useful deductions, it nevertheless keeps deduction under control by severely restricting the deductions that can be made, and by having triggers to control the application of deductive operators. All deductions that can be made are performed at the time a node is inserted into the plan. Since deduction is not expensive, the truth maintenance problem is solved simply by redoing the deductions at each node in the plan after a MN node. Even this can be avoided in simple cases, because SIPE carries a list of changed predicates as it goes through the plan, and, if they all become true later in the plan (without any deduced effects changing in the interim), then the execution monitor need not look at the remainder of the plan (either for redoing deductions or for finding problems).

5. Replanning Actions

The six replanning actions described below, REINSTANTIATE, INSERT, INSERT-CONDITIONAL, RETRY, REDO, and POP-REDO, have all been implemented in SIPE. These actions

provide sufficient power to alter plans in a way that often retains much of the original plan. These are domain-independent actions, and the intention is to use them as a basis for domain-specific error-recovery operators in SIPE. They are also used in the general replanner. Both of these uses are described in more detail in the next section. In all actions below, the context argument merely specifies the context of the current plan.

The last three actions mentioned below all change the plan so that it will contain unsolved problems. The intention is that the plan will then later be given to the normal planning module of SIPE (possibly after a number of these replanning actions have changed the plan). The planner will then attempt to find a solution which solves all the problems that have been corrected in the plan. The planner automatically checks whether things it splices into the middle of the plan cause problems later in the plan so any solution found will be correct. (It does this when copying nodes down to the next lower level during planning.)

REINSTANTIATE (predicate node context)

The action attempts to instantiate a variable differently in order to make the given predicate true in the situation specified by the given node. This appears to be a commonly useful replanning action as it might correspond to using a different resource if something has gone wrong with the one originally employed in the plan, or deciding to return to the screw hopper for another screw rather than trying to find the one that has just been dropped.

This is done by looping through the arguments of the given predicate and, for each one, checking if there is another instantiation for it that will make the predicate true. This is cheap and efficient in SIPE since it merely involves removing the INSTAN constraint on the variable from the current context (and also from all variables constrained to be the same as this one), and then calling the normal matcher to determine if the predicate is now true (which will return possible instantiations). If new instantiations are found, the REINSTANTIATE action checks the remainder of the plan to see if any parts of it might be affected by the new instantiation. This is done by a routine similar to the problem detector described in Section 4 (in fact, the two share much of their code). REINSTANTIATE currently accepts new instantiations only if they cause

no new problems (see discussion below on tradeoffs). If all new instantiations are rejected, the old INSTAN constraint is simply replaced.

There are many tradeoffs in writing a replanning action such as this. There are also tradeoffs in deciding when to apply REINstantiate as it exists, but these are discussed later in the paper. The implementation described above opts for re instantiation only when it is likely to be the correct solution. For example, new instantiations could be accepted even though they cause problems if these problems are less severe than the problems entailed by keeping the old instantiation. Since SIPE has no way of comparing the difficulty of two sets of problems, REINstantiate does not do this.

One could also expend more effort in finding new instantiations. As implemented, this replanning action will find re instantiations when only one variable is changed. Some problems could be solved by re instantiating a whole set of variables, but this would be more expensive and perhaps involve a search problem to decide which variables to include in the set. The decision to try only one variable was made because it appeared powerful enough to be useful. If the ability to re instantiate sets of variables appeared useful, it would certainly be tractable to implement.

INSERT (node1 node2)

This action inserts the subplan beginning with node1 (which has been constructed) into the current plan after node2. All links between the new subplan and the old plan are correctly inserted. This is used as a subroutine by many of the actions below.

INSERT-CONDITIONAL (variable node context)

This action is not very interesting, but complements the unknown variable feature, which may be useful. It simply inserts a conditional around the given node that tests whether the given variable is known. If it is, the given node is executed next; otherwise a failure node is executed.

RETRY (node)

This replanning action is very simple. The given node is assumed to be a phantom node and it is changed to a goal node so that the planner will see it as unsolved.

REDO (pred node context)

This action creates a GOAL node whose goal is the given predicate. It then calls INSERT to put this new node after the given node in the plan. The planner will see the new node as an unsolved goal.

POP-REDO (node predicates context)

This is the most complicated of the replanning actions; it is used to remove a hierarchical wedge from the plan and, in some cases, replace it with a node at the lowest level. This is used both when a PRECONDITION node is no longer true and another action must be applied at a higher level, and when there may be a serendipitous unexpected effect. POP-REDO could also be used to find higher-level goals from which to replan when there are widespread problems causing the replanning to fail (this is not currently implemented).

In the case of redoing a precondition failure, it is easy to determine the wedge to be removed since PRECONDITION nodes are copied down from level to level. The top of the wedge to be removed is the node that was expanded to initially place the given PRECONDITION node (or one of its ancestors that is a PRECONDITION node) in the plan. Actually only the bottom of the wedge is spliced out of the plan, as planning will continue only from the lowest level. The subplan that is removed at the lowest level is replaced by a copy of the GOAL or CHOICEPROCESS node that was at the top of the wedge. (The INSERT replanning action is used for this.) This is seen as an unsolved goal by the planner, which automatically checks whether expansions of this node cause problems later in the plan.

There is one further complication involved. Various constraints may have been posted on the planning variables because of decisions made in the wedge of the plan that has been effectively removed. Because of SIPE's use of alternative contexts, this is easily solved. A context is a list of choicepoints, and constraints are posted relative to the choicepoint that forced them to be posted. Therefore, our problem is solved by removing from the current context all the choicepoints that occurred in the wedge of the plan that was effectively removed. This new context is given as the context argument to future planning actions, and no further action need be taken. This results in




Figure 3
Hierarchical Wedges with a Common Last Action

ignoring exactly those constraints that should be ignored.

The case of serendipitous effects possibly shortening the plan is similar, except that, after the wedge is removed, no node is inserted. However, in this case it is nontrivial to decide which wedge to remove. There may be various wedges that are candidates, and, as with REINstantiate, these candidates may cause problems later in the plan when they are removed. SIPE currently handles this case in the same way in handles the REINstantiate case. Namely, it removes a wedge, checks if this causes any problems, and replaces the wedge if there are any problems. Thus serendipitous effects are taken advantage of only if doing so does not change the rest of the plan. This is a tradeoff like the one discussed previously.

SIPE also generates only one candidate wedge, and gives up taking advantage of the serendipitous effect if this one does not work. This candidate is generated by following ancestor links from the given node as long as some main effect of the candidate node is made true by one of the predicates in the list of given predicates. The candidate node found in this manner determines the candidate wedge. The wedge is rejected immediately unless all its main effects are true in the given list of predicates.

Figure 3 helps show the idea behind this selection process. Frequently, the last action at one level of a wedge will achieve the main effect of every level above that. The above selection process requires that all goals generated at a higher level than the candidate wedge be achieved before the wedge becomes a candidate, while goals generated at a lower level than the top of the candidate

wedge need not have been serendipitously achieved. Thus, for wedge 2 to be selected in Figure 3, the serendipitous effects must include the main effects of the top of wedges 1 and 2, but need say nothing about the main effects of wedge 3. (It is assumed that, as long as the highest-level goal is achieved, we do not care about the lower-level goals that were necessary to bring this about.) The main effects of wedge 1 must be true, because they will be copied down to be effects of the top node of wedge 2 in the case when this is the node that achieves these effects.

6. Guiding the Replanning

The replanning actions of the previous section form the basis for a general replanning capability that has been implemented and for a language capable of specifying domain-specific error-recovery instructions that has been designed but not implemented. This section describes the automatic replanner and briefly outlines the error-recovery operators.

6.1 The General Replanner

The general replanner takes a list of problems from the problem recognizer described in Section 4 as well as possible serendipitous effects, and calls one or more of the replanning actions in Section 5 in an attempt to solve the problem. It first checks that a listed problem is still a problem since the REINstantiate action may solve many problems at once.

If the problem is a purpose not being achieved, the system tries a REDO, which inserts the unachieved purpose as a GOAL node after the Mother-Nature node. If the problem is a previous phantom not being maintained, SIPE first tries REINstantiate and, if that fails, it calls RETRY. The idea is that, if there is another object around with all the desired properties, then it would be easier to use that object than to re achieve the desired state with the original object. If a PROCESS node has an unknown variable as an argument, then INSERT-CONDITIONAL is called. If a future phantom is no longer true, then RETRY is called. As with maintaining phantoms, REINstantiate may be more appropriate, but, in both cases, this depends entirely on the domain so the selection here is arbitrary. For preconditions that are not true, the general replanner first calls REINstantiate and, if that fails, it calls POP-REDO. If parallel-postconditions are

not true, the general replanner constructs the appropriate parallel goals to reach them and calls INSERT to place them after the original JOIN node.

One cannot expect very impressive performance from a replanner that does not have domain-specific information for dealing with errors. For example, whether or not REINstantiate is likely to succeed will be dependent on the domain. The automatic replanner makes reasonable guesses at what might be a good choice in the domains on which SIPE has been tested. Since it simply chooses a replanning action for each type of problem that is found, it is very simple and could easily be rewritten for different domains.

6.2 Error Recovery Language

We also plan to extend the operator description language so that instructions for handling foreseeable errors can be included in operators. The error recovery operators will be in the same syntax as all other SIPE operators, with some new additions made to this language as described below. The plots of these operators will include references to the replanning actions in Section 5. SIPE's ability to specify conditional plans in operators can be used to try a second replanning action only if the first fails.

The error-recovery operators will match their argument list to the arguments of the node being executed so original problem variables can be bound. There are two ways to invoke these operators, one for general operators that solve problems that have been recognized, and one for more specific operators that act directly on unexpected predicates. The latter ability seems attractive since it can avoid a lot of effort when there is good domain-dependent error-recovery information available. If one of these latter operators matches an unexpected predicate, it may be possible to simply apply the operator and assume that it will solve any problems caused by the unexpected predicate, thus shortcircuiting the normal problem detection mechanism.

The general operators will be applied after a MN node is added and problems have been found. Preconditions of these operators will be matched in the situation specified by the MN node. The general replanner will apply any general error-recovery operator that matches a given problem (the matching process is described below) instead of using its default replanning action.

Nodes in plots of regular operators will be able to specify an ERROR slot that gives names of error-recovery operators. The specific error-recovery instructions will be applied immediately after the input of an unexpected predicate (this will assume the predicate is problematical), but only when they are specified in the ERROR slot of the action being executed and match the predicate that was input (see below).

Like deductive operators, error-recovery operators will have a TRIGGER slot [5] to determine when they should be applied. The trigger will be a predicate (for specific operators) or a combination of keywords and predicates, where the keywords refer to the six types of problems. Specific operators are applied when their trigger matches the predicate that is input in the situation represented by the node being executed. General operators will have triggers that match when their keyword matches the problem being tried and any predicate in the trigger matches the appropriate predicates given in the problem.

7. Examples

This section presents two examples of SIPE actually monitoring the execution of a simple plan, and replanning when things do not go as expected. SIPE has been tested on larger and more complex problems than those presented here. These examples are simple to facilitate comprehension and to save space. Everything typed by the user is in boldface - nearly everything below is generated automatically by the system.

This first problem was constructed to show the successful use of the REINstantiate replanning action. The problem is to get A on C in parallel with getting any blue block on any red block. In the initial world B1 and B2 are the only blue blocks (they are both on the table) and R1 and R2 are the only red blocks (R1 is on B1 and R2 is on the table and clear). Since A and C are both clear initially, the planner quickly finds a two-action plan of putting A on C in parallel with putting B2 on R2. While executing the moving of A to C, the SIPE is told that D has suddenly appeared on top of R2. It notices the problem in the parallel branch of the plan and the general replanner tries REINstantiate, which succeeds. The original plan is retained in its entirety and B2 is placed on R1 instead of R2, thus achieving the original goal.

PROBLEM: PROB6

PARALLEL

BRANCH 1:

GOALS: (ON A C);

BRANCH 2:

GOALS: (ON BLUEBLOCK1 REDBLOCK1);

END PARALLEL

END PROBLEM

Plan being executed:

PLANHEAD: P0171

PLANHEAD: World model: (ON A B)(ON B TABLE)(ON C TABLE)(ON D TABLE)(ON E TABLE)
(ON R1 B1)(ON B1 TABLE)(ON B2 TABLE)(ON R2 TABLE)(CLEAR R1)(CLEAR R2)(CLEAR B2)
(CLEAR TABLE)(CLEAR A)(CLEAR D)(CLEAR E)(CLEAR C)

SPLIT: C0170

Parallel branch:

SPLIT: C0189

Parallel branch:

PHANTOM: P0194

Goals: (CLEAR C);

Mainstep: P0197;

Parallel branch:

PHANTOM: P0191

Goals: (CLEAR A);

Mainstep: P0197;

JOIN: C0190

PROCESS: P0197

Action: PUTON.PRIM;

Effects: (ON A C);

Deduce: (CLEAR B), -(ON A B), -(ON A OBJECT3) OBJECT3 UNIVERSAL, -(CLEAR C);

Mainstep: PURPOSE;

Parallel branch:

SPLIT: C0160

Parallel branch:

PHANTOM: P0165

Goals: (CLEAR R2);

Mainstep: P0168;

Parallel branch:

PHANTOM: P0162

Goals: (CLEAR B2);

Mainstep: P0168;

JOIN: C0161

PROCESS: P0168

Action: PUTON.PRIM;
Effects: (ON B2 R2);
Deduce: -(ON B2 TABLE), -(ON B2 OBJECT3) OBJECT3 UNIVERSAL, -(CLEAR R2);
Mainstep: PURPOSE;

JOIN: C0172

P0197 P0168

PICK ONE OF THE ABOVE NODES TO EXECUTE NEXT (? FOR HELP): P0197

Executing action P0197

Unexpected effect (? for help): **(ON D R2)**

Unexpected effect (? for help): **NIL**

Adding deduced predicate: -(ON D TABLE)

Adding deduced predicate: -(ON D OBJECT3) OBJECT3 UNIVERSAL

Adding deduced predicate: -(CLEAR R2)

Problem:

The following predicate negates predicate following it: -(CLEAR R2)

(CLEAR R2)

Causing purpose of following node not to be achieved.

EXECUTED: P0165

Effects: (CLEAR R2);

Mainstep: P0168;

Trying to reinstantiate to make this predicate true: (CLEAR R2)

matching condition: (CLEAR OBJECT1)

(collected 1 possibilities)

Adding INSTAN constraint: OBJECT1 BOUND TO R1

Success, new instantiation: OBJECT1 BOUND TO R1

New plan produced for continuing execution:

PLANHEAD: P0171

SPLIT: C0170

Parallel branch:

EXECUTED: P0197

Action: PUTON.PRIM;

Effects: (ON A C);

Deduce: (CLEAR B), -(ON A B), -(ON A OBJECT3) OBJECT3 UNIVERSAL, -(CLEAR C);

EXECUTED: P0350

Action: MOTHER.NATURE;

Effects: (ON D R2);

Deduce: -(ON D TABLE), -(ON D OBJECT3) OBJECT3 UNIVERSAL;

Parallel branch:

PROCESS: P0168

Action: PUTON.PRIM;

Effects: (ON B2 R1);

Deduce: -(ON B2 TABLE), -(ON B2 OBJECT3) OBJECT3 UNIVERSAL, -(CLEAR R1);

JOIN: C0172

Executing action P0168

Unexpected effect (? for help): **NIL**

Goal achieved.

The second problem, shown below, is the same as the first except that the red block is constrained not to be R1, which will cause REINstantiate to fail. The original plan is the same and the unexpected situation is the same. This time SIPE tries REINstantiate and it fails, so it calls RETRY, which causes the (CLEAR R2) phantom to be made into a goal. The planner solves this to produce a plan that puts D back on the table before B1 is placed on R2. The original plan is not shown below and the other plans are abbreviated by removing MAINSTEPS, deductions, and some phantoms.

PROBLEM: PROB7

PARALLEL

BRANCH 1:

GOALS: (ON A C); BRANCH 2:

GOALS: (ON BLUEBLOCK1 REDBLOCK1); END PARALLEL

GOAL

ARGUMENTS: BLUEBLOCK2, REDBLOCK2 IS NOT R1;

GOALS: (ON BLUEBLOCK2 REDBLOCK2);

END PROBLEM

P0254 P0224

PICK ONE OF THE ABOVE NODES TO EXECUTE NEXT (? FOR HELP): P0254

Executing action P0254

Unexpected effect (? for help): **(ON D R2)**

Unexpected effect (? for help): **NIL**

Adding deduced predicate: **-(ON D TABLE)**

Adding deduced predicate: **-(ON D OBJECT3) OBJECT3 UNIVERSAL**

Adding deduced predicate: **-(CLEAR R2)**

Problem:

The following predicate negates predicate following it: **-(CLEAR R2)**

(CLEAR R2)

Causing purpose of following node not to be achieved.

EXECUTED: P0221

Effects: **(CLEAR R2);**

Trying to reInstantiate to make this predicate true:

(CLEAR R2)

matching condition: **(CLEAR OBJECT1)**

Condition fails.

P0221 being reset to GOAL for replanning.

initial problem:

PLANHEAD: P0227

PLANHEAD: World model: (ON D R2)(ON A C)(CLEAR B)(ON B TABLE)(ON C TABLE)
(ON E TABLE)(ON R1 B1)(ON B1 TABLE)(ON B2 TABLE)(ON R2 TABLE)(CLEAR R1)
(CLEAR B2)(CLEAR TABLE)(CLEAR A)(CLEAR D)(CLEAR E)

SPLIT: C0226

Parallel branch:

SPLIT: C0246

Parallel branch:

EXECUTED: P0251 Goals: (CLEAR C);

Parallel branch:

EXECUTED: P0248 Goals: (CLEAR A);

JOIN: C0247

EXECUTED: P0254

Action: PUTON.PRIM;

Effects: (ON A C);

EXECUTED: P0366

Action: MOTHER.NATURE;

Effects: (ON D R2);

Parallel branch:

SPLIT: C0216

Parallel branch:

GOAL: P0221 Goals: (CLEAR R2);

Parallel branch:

EXECUTED: P0218 Effects: (CLEAR B2);

JOIN: C0217

PROCESS: P0224

Action: PUTON.PRIM;

Effects: (ON B2 R2);

JOIN: C0228

planner succeeds

Plan being executed:

SPLIT: C0448

Parallel branch:

SPLIT: C0483

Parallel branch:

EXECUTED: P0488 Goals: (CLEAR C);

Parallel branch:

EXECUTED: P0482 Goals: (CLEAR A);

JOIN: C0485

EXECUTED: P0486 Action: PUTON.PRIM;
Effects: (ON A C);
EXECUTED: P0487 Action: MOTHER.NATURE;
Effects: (ON D R2);

Parallel branch:
SPLIT: C0446

Parallel branch:
EXECUTED: P0489 Goals: (CLEAR B2);

Parallel branch:
SPLIT: C0435
Parallel branch:
PHANTOM: P0440 Goals: (CLEAR TABLE);

Parallel branch:
PHANTOM: P0437 Goals: (CLEAR D);
JOIN: C0436

PROCESS: P0443 Action: PUTON.PRIM;
Effects: (ON D TABLE);
Deduce: (CLEAR R2), -(ON D R2), -(ON D OBJECT3) OBJECT3 UNIVERSAL;

PHANTOM: P0450 Goals: (CLEAR R2);
JOIN: C0451

PROCESS: P0452 Action: PUTON.PRIM;
Effects: (ON B2 R2);
JOIN: C0453

PHANTOM: P0454 Goals: (ON B2 R2);

Executing action P0443
Unexpected effect (? for help): NIL
Executing action P0452
Unexpected effect (? for help): NIL
Goal achieved.

8. Comparison to other systems

There is very little previous work in this area since most domain-independent planning systems do not address the questions of execution monitoring and replanning (e.g., NONLIN [3] and DEVISER [4]). Hayes [1] and Sacerdoti [2] have addressed this problem. The approaches used in both these systems were considerably simpler and less powerful than SIPE. NOAH did not even allow the input of arbitrary predicates, so the general replanning problem never arose. It did

permit the user to specify that whole wedges had been executed at once, and did allow the node just executed to be planned for again if it failed. This essentially provides one limited replanning action that is useful only in very specific situations.

Hayes's system does allow the input of some information about unexpected situations. It is not clear what types of information can be provided, but it appears less general than the arbitrary predicates accepted by SIPE. The system's only replanning action is to delete part of the plan. This permits the planner to reach higher-level goals, but they must be the same higher-level goals that were already present in the plan. The system deletes everything that depended on any effect of a decision that is no longer valid. This will, in general, be wasteful since much of the plan may be unnecessarily removed. If only one of many effects of an action has failed, subplans depending on the effects that did not fail do not need to be deleted. SIPE would find problems with such subplans.

SIPE provides a much more powerful replanning capability than either of these systems. It allows input of arbitrary predicates, computes how these affect the rest of the plan (and only finds problems that really are problematical), and uses a large number of replanning actions (including REINSTANTIATE) to fix problems in ways that allow much of the original plan to be maintained.

9. Summary and Limitations

Given correct information about unexpected events, SIPE is able to determine how this affects the plan being executed, and, in many cases, is able to retain most of the original plan by making changes in it to avoid problems caused by these unexpected events. It also is capable of shortening the original plan when serendipitous events occur. It cannot solve difficult problems involving drastic changes to the expected state of the world, but it does handle many types of small errors that may happen frequently in a mobile robot domain. The execution-monitoring package does this without the necessity of planning originally to check for these errors.

The major contributions of this work center around taking advantage of the rich structure of SIPE's planner and its plans, and the development of a general set of replanning actions that are used as the basis of an automatic replanner and can be used as the basis of a language for specifying

domain-dependent error-recovery information. The replanner calls the standard planning system so that it can take advantage of the efficient frame-reasoning mechanisms in SIPE to quickly discover problems and potential fixes, and use the deductive capabilities to provide a reasonable solution to the truth maintenance problem. The fixes need involve only inserting new goals in the plan, since calling the planner as a subroutine will solve these goals in a manner that assures there will be no conflicts with the rest of the plan. SIPE's execution-monitoring capabilities make extensive use of the explicit representation of plan rationale in plans. The problem detector makes use of the information encoded in MAINSTEP slots, phantoms, and preconditions to quickly find all the problems with a plan. Furthermore, it does not remove parts of the original plan unless the parts are actually problematical. SIPE's deductive capability is instrumental in the solution of the truth maintenance problem. The replanning actions make use of constraints, alternative contexts, and wedges in SIPE whenever they consider removing part of the plan.

Another important contribution is the development of a general set of replanning actions that will form the basis for a language capable of specifying error-recovery operators, and a general replanning capability that has been implemented using these actions. These actions provide sufficient power to alter plans in a way that often retains much of the original plan, (e.g., the REINSTANTIATE action). The general replanner attempts to solve all problems that are found. It is unlikely to be very successful without being tuned for particular domains. The design of the language for error-recovery operators allows for both operators that will handle very specific situations and operators that will give more general advice to the replanner.

The major limitations of this research result from the assumption of correct information about unexpected events. This avoids the hard problems of generating predicates from information provided by the sensors, deciding how much effort to expend checking facts that may be suspect, and modeling uncertain or unreliable sensors. These problems are all crucial to providing execution-monitoring capabilities to a mobile robot and must yet be addressed.

ACKNOWLEDGMENTS

Many people influenced the ideas expressed in this paper. Special thanks go to Michael Georgeff

for many enlightening discussions.

REFERENCES

1. Hayes, Philip J., "A Representation for Robot Plans", *Proceedings IJCAI-75*, Tbilisi, USSR, 1975, pp. 181-188.
2. Sacerdoti, E., *A Structure for Plans and Behavior*, Elsevier, North-Holland, New York, 1977.
3. Tate, A., "Generating Project Networks", *Proceedings IJCAI-77*, Cambridge, Massachusetts, 1977, pp. 888-893.
4. Vere, S., "Planning in Time: Windows and Durations for Activities and Goals", Jet Propulsion Lab, Pasadena, California, November 1981.
5. Wilkins, D., "Domain-independent Planning: Representation and Plan Generation", *Artificial Intelligence* 22, April 1984, pp. 269-301.