

PlanWorks: A Debugging Environment for Constraint Based Planning Systems

Patrick Daley^{* †} and Jeremy Frank and Michael Iatauro[‡] and Conor McGann[§] and Will Taylor

Computational Sciences Division
NASA Ames Research Center, MS 269-3
frank@email.arc.nasa.gov
Moffett Field, CA 94035

Introduction

In recent years, model-based planning has moved from reasoning about models largely described with propositions (e.g. STRIPS) to reasoning about time, resources, and constraints on complex numeric quantities (Smith, Frank, & Jónsson 2000), (Fox & Long 2003). Numerous planning and scheduling systems employ underlying constraint reasoning systems to handle the richness and diversity of such constraints (Frank & Jónsson 2003). Debugging such systems involves the search for errors in model rules, constraint reasoning algorithms, search heuristics and the problem instance (initial state and goals). In order to effectively find such problems, users must see why each state or action is in a plan by tracking causal chains back to part of the initial problem instance. They must be able to visualize complex relationships among many different entities and distinguish between those entities easily. For example, a variable can be shared between several constraints, as well as part of a state or activity in a plan; the activity can arise as a consequence of another activity and a model rule. Finally, they must be able to track each logical inference made during planning.

We have developed *PlanWorks*, a browse-based system for debugging constraint-based planning and scheduling systems. *PlanWorks* assumes a strong transaction model of the entire planning process, including adding and removing parts of the constraint network, variable assignment, and constraint propagation. The planner logs all transactions to a relational database that is tailored to support queries for a variety of components. *Visualization* components consist of specialized views to display different forms of data (e.g. constraints, activities, resources, and causal links). Each view allows user customization in order to display only the most relevant information. Inter-view navigation features allow users to rapidly switch views to examine the trace of the process from different perspectives. *Transaction query* mechanisms allow users access to the logged transactions to visualize activities across the entire planning process. While originally developed for debugging, *PlanWorks* has the potential to serve as a knowledge capture tool and an end-user

operations tool as well.

PlanWorks is implemented in Java and employs a MySQL relational database back-end. *PlanWorks* can be used either online while iterative planning and debugging is performed, or offline after capturing the entire planning process. Furthermore, *PlanWorks* is an open system allowing for extensions to the transaction model to capture new planner algorithms, different classes of entity (e.g. complex resource classes) or novel heuristics. *PlanWorks* was specifically developed for the Extensible Universal Remote Operations Planning Architecture (EUROPA₂) developed at NASA, but the underlying principles behind *PlanWorks* make it useful for many planning systems, a point we address at the end of the paper. *PlanWorks* has been used to visualize logs generated by three different planners; two versions of EUROPA₂ as well as a prototype of the Mission Data Systems (MDS) planner developed at JPL (Dvorak *et al.* 2000).

The paper is organized as follows. We first describe some fundamentals of the EUROPA₂ constraint-based planning system, and describe a simple planning domain used throughout the rest of the paper. We then describe *PlanWorks*' principal components. We discuss each component in detail, and then describe inter-component navigation features. We describe how to configure logging of data from the planner to reduce debugging time, which influences what *PlanWorks* can display. This configuration can be done in a configuration file, or using *PlanWorks*. We discuss how *PlanWorks* is used during debugging. Finally, we describe applicability of *PlanWorks* beyond EUROPA₂, discuss system requirements and deployments, and discuss future work.

EUROPA₂

EUROPA₂ provides efficient, customizable *Plan Database Services* that enable the integration of automated planning into a wide variety of applications. These services are based on some simple building blocks. *Plans* are composed of *predicates*, each of which has a name, start time, end time, duration, and a (possibly empty) set of parameters. Each instance of a predicate in a plan is represented by a *token*, and each parameter of the predicate is represented by variables. Predicates are associated with either *timelines* that support totally ordered sequences of states, or *resources* that support possibly concurrent actions that do not exceed a maximum capacity. Timeline or resource instances are referred

* Authors listed in alphabetical order.

† Foothill College

‡ QSS

§ QSS

to as *objects*, and during planning each token is assigned to an object in the plan. *Domain rules* are assertions that if a predicate P is in a plan, then other predicates Q_i must also be in a plan, and are related to P by constraints among the variables of the predicates. Domain rules may also assert that resources are impacted by predicates; resource impacts are called *transactions*, and also have variables that represent them. EUROPA₂ does not implement any planning algorithm; rather, it provides services that support different planning algorithms according to the application. As such, it can be used to support progression planners, regression planners, sequential or causal link planners, and so on. To enable this generality, EUROPA₂ distinguishes between *free* tokens (consequences of rules that haven't been inserted into plans), *active tokens* and *merged tokens*. Planners can insert free tokens into plans (making them active) or merge free tokens with active tokens.

A Sample Planning Domain

EUROPA₂ domains are written in a domain description language called NDDL. To illustrate the fundamentals of PlanWorks, we use a planning domain loosely based on a planetary surface robot named *Rover*. *Rover* is a mobile robot that can move from location to location. A *Rover* has a battery on board, and can replenish its energy levels using solar power. Locations are described as follows:

```
class Location {
  int x; int y;
  Location(int _x, int _y){
    x = _x; y = _y;
  }
}
```

The properties of the Rover are described as follows:

```
class Rover {
  predicate At{Location l;}
  predicate Going{Location from;
                  Location to;}
  Resource battery;
  battery = new Battery(10, 3, 30);
}
```

A domain rule in EUROPA₂ describing rover movement is:

```
Rover::Going{
  neq(to, from); // to != from
  meets(object.At a0);
  eq(a0.l, to);
  met_by(object.At a1);
  eq(a1.l, from);
  subgoal(object.battery.transaction tx);
  calcConsumption(tx.quantity, from, to);
  // Consume at the beginning
  eq(tx.time, start);
}
```

Finally, a problem instance for the Rover is:

```
Rover spirit = new Rover();
Location rock = new Location(1, 1);
```

```
Location hill = new Location(2, 3);
Location lander = new Location(5, 8);
goal(Rover.At A);
eq(A.l, rock); eq(A.object, spirit);
leq(A.start, 0); leq(0, A.end);
goal(Rover.At B);
eq(B.l, lander); eq(B.object, spirit);
leq(B.start, 0); leq(0, B.end);
```

Getting PlanWorks the Goods

During planning, EUROPA₂ reads the domain rules to determine if any of them are applicable given the current state of the plan. If so, new tokens, resource transactions, variables, and constraints are created, and the domain rule application is recorded. As the planner makes decisions, tokens can be assigned to timelines, transactions can be assigned to resource instances, variables can be assigned, and constraints can be enforced, leading to reductions in the domains of variables. Each of these activities is logged, and each entity is assigned a unique key that allows for the tracking of entities and their relationships during planning. This information is passed down to PlanWorks to enable users to uncover the relationships between entities in the plan.

PlanWorks can be used in one of two modes. Planners can generate logs for PlanWorks offline, after which PlanWorks is invoked to view the logs. When planning takes a long time, this is impractical. Alternatively, PlanWorks can be started and provided a pointer to a planner. PlanWorks then allows the user to interleave planning and debugging. The user can run the planner for a fixed number of steps, investigate, then continue running the planner or terminate planning.

The basic data integration point is through the database. At each “step” (where step is defined by the planner, but is logically the end of propagation after a planner decision), the current state of the plan graph, the set of transactions that describe the transformation from the previous plan graph, as well as statistics about the graph are logged in a format optimized for importation into the database. The set of files for a single step are collected into a directory, which is itself in a directory representing the particular instance of planner execution. These data are imported and queried upon request, i.e. the statistics of the plan run are loaded when the user opens the Sequence Steps View, transactions are loaded as part of a transaction query, and the plan graph data is loaded when the user opens a step view.

Figure 1 shows a fragment of the Entity Relationship Diagram for the SQL table structure describing the format PlanWorks accepts. The fragment captures the logging of the most important core components of the planning process: Tokens, Objects, Variables, Constraints, Rules and Rule Instances. The ERD indicates key relationships between entities that are logged to ensure efficient database queries; these are later marshalled into the views that we describe in the next section.

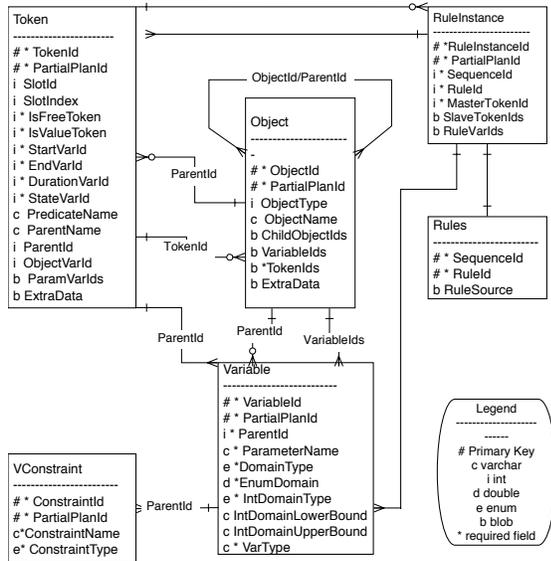


Figure 1: Part of the Entity Relationship Diagram for PlanWorks. This diagram shows the SQL table structure for Token, Variable, Constraint, Object, Rule and RuleInstance tables as well as the relationships between them.

PlanWorks Components

Initial Views

Upon startup, PlanWorks presents users with a menu bar offering features for project creation and management. A project contains numerous *planning sequences* corresponding to the execution of a planner on a problem instance. The menu allows users to create new projects, add and delete sequences, and open a sequence for viewing.

When a sequence is opened for viewing, PlanWorks displays two views: the *Sequence Step View* and the *Sequence Query View*. The Sequence Step View, shown in Figure 2, is a broad overview of the planning process. The view is presented as an inverted histogram, with the number of steps of the process along the bottom. The histogram bars are broken into components representing the number of variables, constraints and tokens in the plan at a given point in the planning process. Moving the mouse over a bar of the histogram shows the step number and number of entities of each type in the plan at that point in the process. At a glance, the user sees how the plan's size evolved throughout planning, and can see patterns (such as thrashing in a chronological backtracking algorithm, or local optimal in a local search planner). An indicator above each bar of the histogram shows whether the data for that step has been loaded into PlanWorks.

The Sequence Query View allows the user to request detailed information about the underlying transactions over the entire planning sequence. The scope of logging is crucial to

support these queries. Each time an entity is created or destroyed during planning, this information is logged; at creation time, each entity is given a unique key. These keys make it possible to track entities over the course of planning. Constraints can be tracked when they execute, tokens can be tracked as their state changes (e.g. from creation to insertion on an object), planner decisions can be tracked, and so on.

Examples of supported transaction queries include entity creation, assignments and unassignments of tokens to objects, assignments and unassignments of values to variables, constraint enforcement, checking for variables with only one domain value remaining, and more.

The Sequence Step View is also used to launch numerous other views of the plans generated at each step of the sequence. These views fall into one of three categories: *Plan Views*, *Entity Relationship Views* and *Transaction View*. These views are described further below.

Plan Views

Plan Views are holistic views of the entire plan. Plans are sequences of states or actions over time, so by their nature, the Plan Views are meant to convey a sense of what the plan looks like overall. However, EUROPA₂ can represent plans that are more complex than time-stamped sequences of actions. Plans can be *temporally flexible*; that is, states may have start times or durations that are unknown until plan execution. Further, plans may involve resources whose quantities change over time. Thus, PlanWorks requires visual representations of timelines or resources that are temporally flexible. For this reason, three distinct Plan Views are provided.

The *Timeline View* is designed to show the sequence of predicates on a timeline. Since tokens can be unified, the Timeline View shows the number of unified tokens supporting each predicate; moving the mouse over the predicate shows the keys of the unified tokens and indicates which of these is the active token. The Timeline View shows the possible values of the start and end times of each predicate on the timeline. Finally, the Timeline View shows any free tokens. This is shown in Figure 2. If we also refer back to Figure 1, we can see how queries to the database return the data that is rendered as a TimelineView. PlanWorks asks for the Tokens on a particular Object. In turn, the Tokens are affiliated with Variable objects representing the start and end times; data elements inside the Variables indicate the lower and upper bound of the timepoints, names, etc. This information is marshalled into the correct format to call the rendering and layout routines.

The *Temporal Extent View* is designed to show more temporal information about tokens than the Timeline View. Each token has a series of icons representing the possible values of the start time (downward pointing triangles), end time (upward pointing triangles) and duration (horizontal line bracketed by the triangles). This is shown in Figure 2. Moving the mouse over each of these shows the values, and an absolute time scale at the bottom is used for reference.

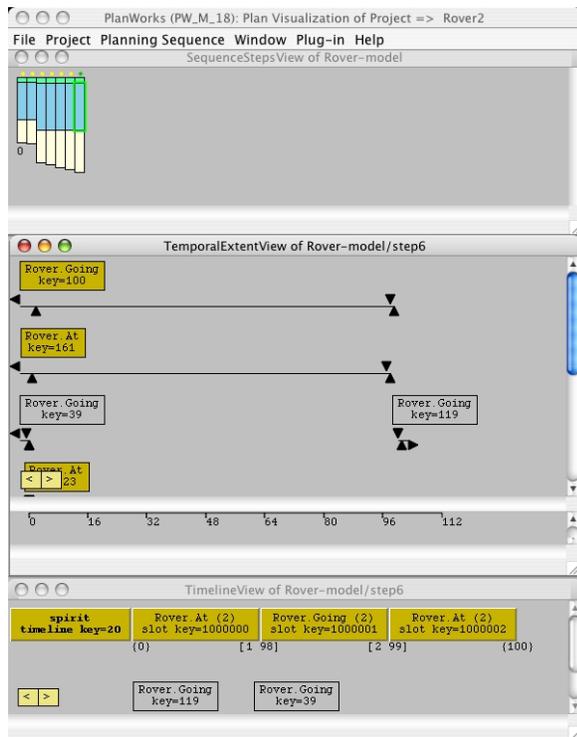


Figure 2: PlanWorks Plan Views. The Sequence Step View is shown at the top, the Temporal Extent View is shown in the middle, and the Timeline View is shown at the bottom.

By moving back and forth between the Timeline View and the Temporal Extent View, users can see how constraints on individual tokens lead to bounds and orderings on predicates in timelines. The Temporal Extent View also includes each resource transaction in the plan. Moving the mouse over the resource transaction shows the impact that transaction has on the resource. The *Resource Transaction View* restricts the Temporal Extent view so that only the resource transactions are shown.

The *Resource Profile View* shows the minimum and maximum quantities of a resource available as a function of time stemming from the transactions in the plan. Again, moving back and forth between the Temporal Extent View (or the Resource Transaction View) and the Resource Profile View, users can see how resource transactions lead to bounds on resource availability.

Entity Relationship Views

EUROPA₂ generates a large number of entities during the course of planning. These entities range from individual tokens, variables representing their parameters and constraints on those variables to object instances and domain rule invocation instances. The *Entity Relationship Views* are graphical views that show each of these entities and how they are related to each other. Under the Help menu, the Shapes option provides a handy guide to the shape each entity takes on in these views.

The *Navigator View* is an entity relationship graph capa-

ble of showing every entity in an individual plan. The Navigator View is launched by selecting an entity from any other view, and initially shows only a small number of entities and relationships. Each entity in the Navigator can be “opened” to show its relationships to other (currently hidden) entities, and subsequently “closed” to hide those relationships. Entities that can be closed are outlined in bold, and those that can be closed are not. The entity graph is directed; the description of the problem defines the initial set of entities, and all entities are derived from them via actions taken by the planner and the rules of the domain. Users can explore the entities and their relationships and find out how various parts of a plan are related to each other. They may do so by hand, manually opening or closing various entities. The Navigator Window also supports a “Find Entity Path” feature that discovers paths between entities (whether they have been opened by the user or not). Navigator View also supports a “Find by Key” feature that highlights an entity with the given key. Each Rule Instance entity can be expanded to show the domain rule text that led to the new token or transaction, as well as the tokens involved in that part of the view.

The *Token Network View* restricts the Navigator view to only tokens, transactions and rule instances. This allows the user to focus exclusively on the “causal chain” that explains why particular tokens were generated. The resulting graph is a directed tree. As with the Navigator View, each Rule Instance entity can be expanded to show the domain rule text that led to the new token or transaction, as well as the tokens involved in that part of the view. This is shown in Figure 3. The Token Network View also supports the same “Find Entity Path” and “Find by Key” features supported by the Navigator.

The *Constraint Network View* restricts the Navigator View to constraints, variables, tokens, transactions and model constants. Initially, the view shows all model constants, predicates, transactions and rules. Model constants may be the values of variables, and may be complex structures; for example, in our simple Rover domain, *paths* consisting of an initial location, final location and cost in terms of energy consumption are constants. Each model constant can be opened to show its underlying structure. Tokens and transactions are associated with sets of variables, which in turn are in the scope of constraints. Domain rules may also have “local variables” to reduce the number of parameters of predicates. The user can incrementally explore the Constraint Network by opening tokens, transactions, or rules, and subsequently opening the variables or constraints. The Constraint Network View also supports the “Find Entity Path” and “Find by Key” features.

Transaction View

EUROPA₂ was designed to support multiple planning paradigms, from heuristically driven chronological backtracking planners to local searching planners to iterative sampling planners. Consequently, logging of information

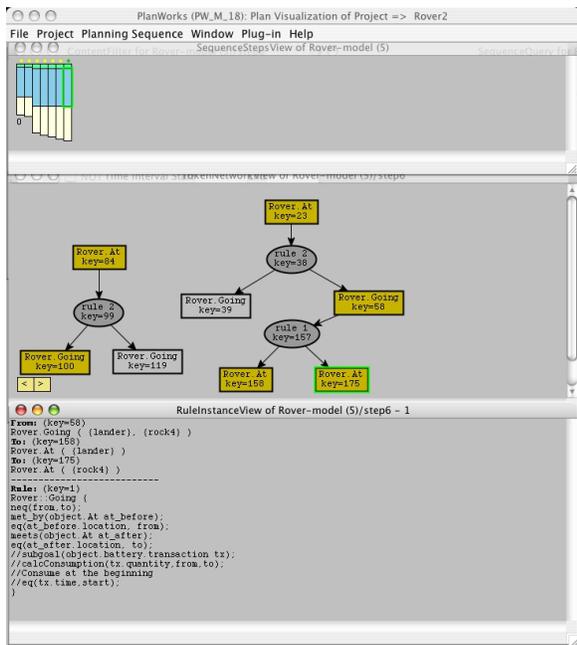


Figure 3: PlanWorks Token Network View and Rule Instance for `Rover::Going`.

about how the planner makes decisions is the responsibility of the planner, while logging the consequences of planner decisions is the responsibility of EUROPA₂. The *Transaction View* shows every transaction EUROPA₂ performed this step. This includes checking domain rule applicability, entity creation and destruction, variable assignment, token state manipulation, constraint enforcement, and so on. Figure 4 shows the Constraint Network View and the Transaction View together.

Navigating PlanWorks

An early decision was made in PlanWorks to create separate Views that contain information that users typically want grouped. However, PlanWorks contains numerous features that allow users to efficiently navigate between Views. These features allow users of PlanWorks to rapidly move from View to View when debugging planning domains and planners.

Launching Views

Almost all Views can be launched from any other View. The Navigator View can be launched when the mouse is over an entity such as a token, variable variable, constraint, constant or rule in a View (except the Transaction View). All other views can be launched when the mouse is over the background of a view.

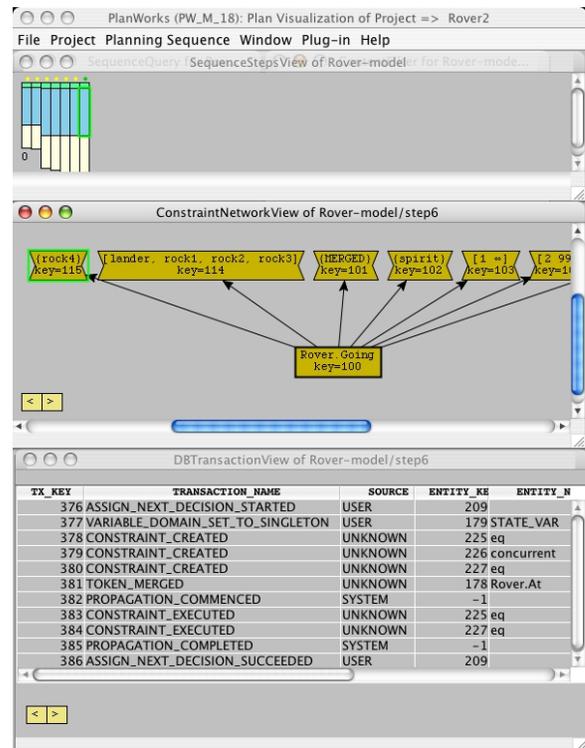


Figure 4: PlanWorks Constraint Network View and Transaction View.

Tracking Objects by Key

All objects have keys, and every view has a method to search for entities by key. Furthermore, the Views opened on starting PlanWorks have facilities for querying transactions by key. Finally, moving the mouse over entities will show the keys of entities. The entity relationship views can be used to provide keys used for querying the transaction database.

Snapping to Entities

PlanWorks provides additional features to navigate between Views. An entity can be made *active* in one view, and the user can then “Snap to Active Entity” in a second view.

Filtering Views

In addition to manually opening and closing entity relationships to incrementally explore views, PlanWorks provides a custom filter for each View. This filter allows rapid reduction of the View to exclude designated predicates, classes (time-lines or resources), or predicates in particular time ranges.

Overviews

Every View can have an associated Overview window that shows the view at a maximum zoom. This allows users to simultaneously examine a small number of related entities in a View, while also seeing the “Big Picture”.

Stepping Forward and Backwards

All Views pertain to one step of the planning sequence. Each View has buttons that allow the user to advance or retreat the step the View shows. This permits a primitive “animation” feature that shows how a View changes during planning. As the View changes, the window is updated. Furthermore, the mouse allows users to either advance or retreat all Views simultaneously.

Multiple Views

PlanWorks can display multiple windows with the same view (e.g. two Token Network Views). These views can display data from the same model at different steps; the step number is indicated in each View’s title bar. Also, it is possible to display data from different versions of the model, so that users can see the impact of debugging. The name of the project has a version number appended to it, indicating revisions and re-running of the planner.

Managing Views

All Views are labeled at the top with View name, Project, Sequence number and Step. Thus, at a glance, a user can automatically tell what information they are seeing in a View. In addition, there is a drop-down menu named Window that allows users to see at a glance what windows are open, as well as either Tile or Cascade all open windows. Finally, using the mouse, users may automatically close, hide or open all Views.

Configuring PlanWorks Logging

EUROPA₂ is capable of logging large amounts of information during planning. For large plans with many constituent parts or for large searches, this can be both time consuming and produce more information than the user desires. For this reason, PlanWorks allows users to configure logging to reduce the amount of information logged. This is especially useful as debugging proceeds and the user is able to focus on parts of the search space where information about possible bugs is most likely. Users can specify where plan steps should be logged and where to look for model files. Users can specify that every step is logged, only the final plan is logged, or the frequency of steps to log (i.e. every 5th step). Finally, users can specify which transactions to log.

PlanWorks Planner Control

The PlanWorks Planner Control window allows the user to control EUROPA₂ planner execution on a per step basis. Specifically, PlanWorks permits fine control of the planner’s data logging, which greatly speeds planner execution and reduces planner data storage requirements. Debugging with integrated Planner Control involves creating and configuring a PlanWorks project followed by cycles of running, viewing, changing, and rerunning the plan.

Configuring a PlanWorks Project

PlanWorks provides the user the capability to create, configure and save multiple Planner Control projects. The setup of

a project includes specifying the path to the planner dynamic library, the path to the user’s model library and initial state file, and the path to the destination directory for the planner’s output. Each of these configurations is saved in a file describing the configuration as discussed previously.

Running the Planner

Once a new project has been configured or an existing project has been opened, the user is ready to start debugging. The Create New Sequence window starts the planner control process. First it allows the user to modify the project configuration. After PlanWorks finds all of the specified files it launches a Transaction Types to Log window that enables the user to limit the transaction types the user would like EUROPA₂ to log. PlanWorks then launches a Planner Control window that provides the controls for stepping the planner. The Planner Control window includes buttons for:

- running the planner up to step N and logging data for only that step.
- running the planner for the next N steps logging data for each step.
- running the planner to completion, logging data for only the last step in the plan.
- early termination of the run.

Once a plan has run to completion or terminated, it can be rerun without exiting PlanWorks.

Debugging with Planner Control

When debugging a model, the user may first want to skip to the last step in the plan. This may be the final plan step of a completed plan or the last step before a planner time out. This function is provided with the WriteFinalStep button. If the planner does not find a solution, the user can then use the Sequence Step View to display histograms of variables, constraints, and tokens used in each step. If, for instance, excessive backtracking is observed, the user can rerun the model, skipping to the step where backtracking occurs. From this point the user can proceed to debug the problem. Identifying steps of interest before enabling data logging is crucial for debugging large models.

Once a problem has been identified, the user can modify and rebuild the model library and initial state file in a terminal window outside of PlanWorks and then initiate a New Sequence Window from PlanWorks. The planner can be directed to skip to the step of interest before logging any data. In most cases, where only the model or initial state has been modified this will be a quick and easy process.

Viewing Multiple Sequences

PlanWorks is capable of displaying multiple sequences from different runs of the planner. This allows users to compare runs and determine whether they have fixed a problem. An example of this is shown in Figure 5.

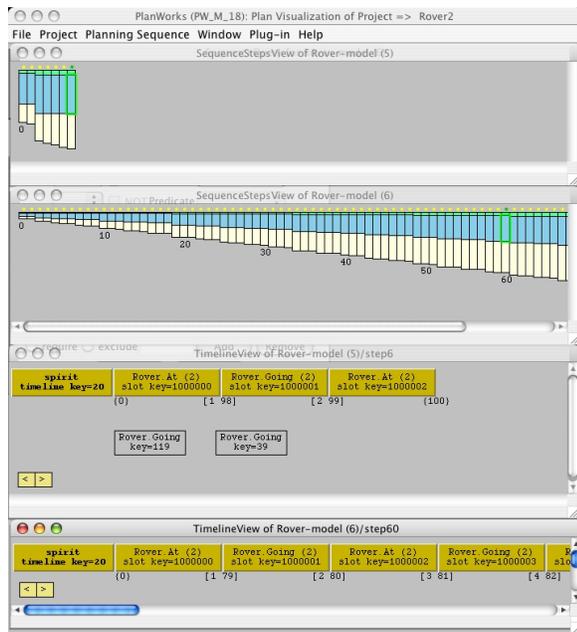


Figure 5: PlanWorks displaying data from two runs of the planner after modification of the model. The Sequence Step Views are shown at the top, and Timeline Views of the last plan step are shown at the bottom.

Debugging in PlanWorks

The capabilities of PlanWorks to visualize a specific plan are its main strength. This provides confidence in both the final result (primarily through Resource and Timeline Views) and the means by which that result was obtained (primarily through the Token Network View). The Sequence Step View can also provide a rapid means of assessing “thrashing” behavior in backtracking style planners by exhibiting obvious patterns (assigning 50 variables then backtracking).

Should the user suspect a problem because a plan does not look right, the user can visualize the network of constraints and obtain a view of the propagation in that constraint network using the Transaction Query. Badly constructed models usually succumb to this type of investigation, as missing rules or extra or wrong constraints become obvious. We now present two small examples of how PlanWorks can be used to find bugs in EUROPA₂.

Missing Model Rule

Suppose that the rule governing rover movement was missing a part:

```
Rover::Going{
  neq(to, from); // to != from
  meets(object.At a0);
  eq(a0.l, to);
  met-by missing
  ...
}
```

The resulting plan could then have two consecutive

Rover::Going tokens; this would appear in the Timeline View. From here, the user could proceed in several ways. One option is to launch the Token Network View and see that only one Rover::At results as a subgoal from Rover::Going. Upon opening the Rule Instance View, the user would see the rule text and the context in which the rule was invoked, and be able to revise the rule. Alternatively, the user could launch the Navigator View and discover the problem.

The Wrong Constraint

As another example, suppose that the user used the wrong constraint in a rule:

```
Rover::Going{
  eq(to, from); // to == from? silly!
  ...
}
```

In this case, the user might see an unexpectedly long plan, as is shown in Figure 5. Again, there are several candidate debugging scenarios. One option is for the user to open the Token Network View, and observe that Rover::Going begets Rover::At, which begets another Rover::Going ad-infinitum. Mousing over a Rover::Going reveals that its parameters from and to are equal. Opening a Rule Instance View on the Rover::Going, the user will notice the incorrect use of the eq constraint. This is shown in Figure 6. Another possibility is that the user immediately suspects a problem with constraint reasoning, and opens the Constraint Network View. After finding the key of a parameter of a Rover::Going, the user then can check the transactions enforced on this variable using the Transaction View. Upon realizing that no neq constraints are enforced, the user can then check the Rule Instance View, and realize that the wrong constraint was used in the rule.

Applicability of PlanWorks

PlanWorks was specifically designed with EUROPA₂ in mind, but planners using modeling languages such as PDDL can (with some work) make use of PlanWorks. The rules for the simple Rover domain can easily be translated into PDDL, either with or without temporal annotations. A plain STRIPS encoding of Rover::Going looks like this (omitting the static predicate and object descriptions):

```
(:actiondrive
  :parameters (?v ?l1 ?l2)
  :precondition(and(
    (location ?l1) (location ?l2)
    (vehicle ?v) (at ?v ?l1)))
  :effect (and (at ?v ?l2) not(at ?v ?l1))
)
```

The notions of state and action as sub-classes of tokens can be added to the ERD of Figure 1 and logged appropriately. Variables are straightforward. Preconditions enable actions to occur; thus, connections between states and rules that are enabled is also easily logged. Capturing the

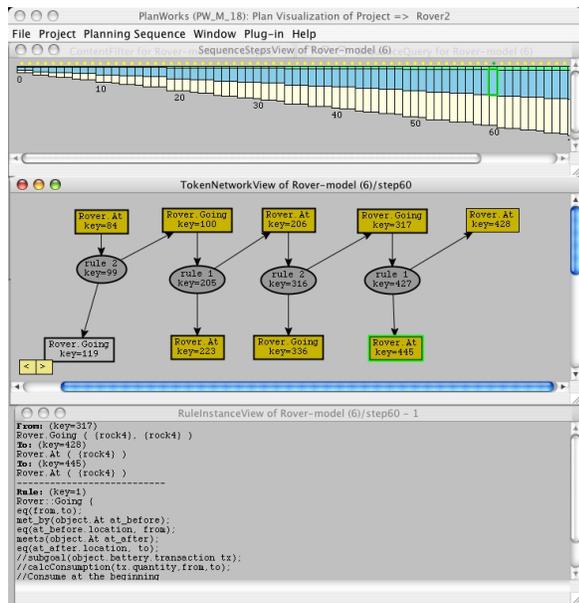


Figure 6: PlanWorks displaying a suspicious Sequence Step View (top), the long chain of rules (middle) and faulty rule in the Rule Instance View (bottom).

rule text for rules that are applied in a state can be done by tracking file line number of the model input file (which is how PlanWorks receives the data from EUROPA₂ based planners). Progression and regression planners have an advantage over EUROPA₂ planners in that the token transaction model is much simpler; there is no unification (while EUROPA₂ planners can do progression and regression planning, unification is supported; however, no such transactions will be logged.) Partial order causal link (POCL) planners use a transaction model similar to that of EUROPA₂ in that preconditions can be shared between actions (resulting in unification) and causal links are added to the plan to ensure those preconditions are protected, corresponding to the addition of constraints (see the a discussion of constraints as generalized causal links in (Frank & Jónsson 2003)). Elaborations to handle numeric fluents and temporal constraints of newer versions of PDDL can proceed along similar lines.

A significant drawback to using PlanWorks for PDDL-based planners is that PlanWorks currently has no facility to visualize domain independent heuristics such as the PlanGraph. However, with some work, views such as the TokenNetwork View can be modified to visualize either Blum and Furst's PlanGraph with mutual exclusions (Blum & Furst 1995) (and by extension Smith and Weld's TGP graph (Smith & Weld 1999)), the relaxed plan graph, and cost-based heuristics that are appended to the PlanGraph (for example (Bonet & Geffner 2001)).

Build Requirements

PlanWorks was built largely with COTS technology, and has been fielded on Mac OSX 10.2 and 10.3, Linux and Solaris. PlanWorks requires Java 1.4.1 or higher, MySQL 4.0.13 or

higher, and makes use of Ant (a replacement for Make used to manage build commands, using XML-based build files), Swing (graphics, layout and user interface toolkit bundled with Java), and JGO (additional diagram graphics libraries associated with Swing). Ant and JGO have been bundled with PlanWorks, while numerous public versions of MySQL (up to 4.1.7) have been integrated with PlanWorks. As an aside, it was found that numerous JGO classes rendered desired views quite slowly, so some customization of layouts was done by the team.

Conclusions and Future Work

We have described PlanWorks, a system designed to debug planning domains and planners. While PlanWorks was designed to debug planners built on top of the EUROPA₂ system, it can be used by any planner that obeys a small number of rules about how to log its inner workings. PlanWorks was originally conceived of as an Integrated Development Environment for building and managing projects with EUROPA₂. In the near future, PlanWorks will be extended to handle visual model building, visualizing plan execution and associated constraint reasoning. Furthermore, EUROPA₂ is designed to support many different planning algorithms. We will also extend PlanWorks to enable user customization to visualize different planner algorithms and heuristics.

Acknowledgements

The authors would like to acknowledge Andrew Bachmann's contributions to the NDDL language used to describe EUROPA₂ planning domains, Tania Bedrax Weiss for ongoing work on EUROPA₂, Sailesh Ramakrishnan for contributions as PlanWorks' prototype user, and Bob Kanefsky for the Potato prototype. This project was funded by the NASA Intelligent Systems Program.

References

- Blum, A., and Furst, M. 1995. Fast planning through planning graph analysis. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI 95)*, 1636–1642.
- Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence* 129(1-2):5–33.
- Dvorak, D.; Rasmussen, R.; Reeves, G.; and Sacks, A. 2000. Software architecture themes in JPL's mission data system. In *IEEE Aerospace Conference*.
- Fox, M., and Long, D. 2003. Pddl 2.1: An extension to pddl for expressing temporal planning domains. *Journal of Artificial Intelligence Research* 20.
- Frank, J., and Jónsson, A. 2003. Constraint-based interval and attribute planning. *Journal of Constraints Special Issue on Constraints and Planning*.
- Smith, D. E., and Weld, D. S. 1999. Temporal planning with mutual exclusion reasoning. In *IJCAI*, 326–337.
- Smith, D.; Frank, J.; and Jónsson, A. 2000. Bridging the gap between planning and scheduling. *Knowledge Engineering Review* 15(1).