USING ARTIFICIAL INTELLIGENCE TECHNIQUES FOR AUTOMATED PLANNING AND SCHEDULING

Raymond E. Levitt[1], John C. Kunz[2], and Nabil A. Kartam[3]

1. Abstract

This paper develops a philosophy for the use of Artificial Intelligence (AI) techniques as aids in engineering project management. We start by classifying the subtasks associated with project management as a taxonomy of separate functions and levels. We then assess the cognitive requirements for each project management subtask. Recognizing the cognitive requirements of each subtask and the limitations of existing computer tools for project management decision support, we propose guidelines for using AI and procedural programming techniques to support decision making in each phase and at each level of project management.

First, we propose that traditional domain-independent, "means-end" planners, may be valuable aids for planning detailed subtasks on projects, but that domain-specific planning tools are needed for work package or executive level project planning. Next, we propose that hybrid computer systems, using knowledge processing techniques in conjunction with procedural techniques such as decision analysis and network-based scheduling, can provide valuable new kinds of decision support for project objective-setting and project control, respectively. Finally we suggest that knowledge-based interactive graphics, developed for providing graphical explanations and user control in advanced knowledge processing environments, can provide powerful new kinds of decision support for project management.

The first claim is supported by a review and analysis of previous work in the area of automated AI planning techniques. Our experience with PLATFORM I, II and III, a series of prototype AI-leveraged project management systems built using the IntelliCorp Knowledge Engineering Environment (KEE™), provides the justification for the latter two claims.

---

1 Associate Professor, Dept. of Civil Engineering, Stanford University, Stanford, CA

2 Chief Knowledge Systems Engineer, IntelliCorp, Mountain View, CA

3 Ph.D. Candidate and Research Assistant, Dept. of Civil Engineering, Stanford University, Stanford, CA

2. Introduction

Traditional project management tools, invented during the late 1950's have become indispensable as aids in planning and scheduling construction projects. Although advances in computer technology have greatly facilitated their use on current projects, the fundamental concepts employed by these decision-support tools for project management have not changed significantly since the mid 1960's. This paper highlights the limitations of existing computer-aided project management tools, and sets out guidelines for employing proven and emerging Artificial Intelligence techniques to extend their range and power for each phase of project management.

3. A Taxonomy of Project Management Tasks

We will classify the cognitive tasks associated with project management in two dimensions: first, we will define a set of functions which comprise project management; second, we will define three levels of abstraction at which project management is carried out. We propose the use of different kinds of knowledge processing and procedural computer techniques to aid decision making for each of these subtasks of project management.

3.1 PROJECT MANAGEMENT FUNCTIONS

Project Management consists of project objective-setting, project planning, project scheduling and project control. The words objective-setting, planning, scheduling and control are used in different ways in the project management and AI literature. We will, therefore, start by defining each of these components of project management for the purposes of this paper:

- **Objective Setting** -- consists of the decisions that are made during the early conceptual stages of a project, when trade-offs among the project's location, size, cost, overall duration and performance levels are being firmed up as a set of project goals. This phase is sometimes called conceptual design, and its outputs are typically referred to project goals or project objectives.

- **Project Planning** -- consists of generating a list of the activities and their sequential relationships. These activities must be executed to achieve the specified set of project objectives. The output of the planning process we will term a project plan.

- o **Project Scheduling** -- consists of assigning durations and resources to activities in a project plan, computing the timing of activity starts and finishes, and computing the utilization rates for resources. We will call the output of the scheduling process a project schedule.

- o **Project Control** -- consists of measuring actual activity progress and resource consumption; interpreting past performance data; forecasting the duration and resource consumption of each remaining activity; comparing forecasted durations and resource consumptions against the project schedule; and taking necessary actions to eliminate or reduce unfavorable deviations, and to maintain or increase favorable deviations. These actions may include rescheduling, replanning or even revising project objectives, in the light of large variances from expected plans and schedules.

3.2 PROJECT MANAGEMENT LEVELS

There is a continuum of detail ranging from one-sentence project objective statements to project plans consisting of thousands of activities. We will identify three levels of project plans corresponding to discrete organizational and management responsibility levels for many projects.

- o **Executive Level:** In many contractor and construction buyer organizations, a high level corporate officer has overall responsibility for each project. Activities at this level of abstraction tend to be unique, and involve little or no repetition.

- o **Work Package Level:** There is often a close relationship between what we will call work packages, and the activities in the executive level plan. Each activity in the executive plan can constitute work package plan under the responsibility of a single manager. At this level, there may be some repetition of activities, e.g., the work package *construct control building* on a process plant project might include activities such as: *excavate basement, install drains, form and pour foundations, form and pour 1st floor, form and pour second floor, hang exterior wall north side, hang exterior wall east side, etc.*

- o **Task Level:** Stepping down to the final level, each of the activities in a work package plan can constitute a task level plan under the direction of a first level supervisor such as a foreman. At this task level the number of different kinds of activities can be quite small, with significant amounts of repetition, e.g., the task *form and pour foundations* might consist of the following activities: *form footing$_1$, place rebar in footing$_1$, pour footing$_1$, cure footing$_1$, strip forms from footing$_1$, form footing$_2$, ...cure footing$_{20}$, etc.*

3.3 COMBINING FUNCTIONS AND LEVELS OF PROJECT MANAGEMENT

Plotting the levels of project management as rows and the functions of project management as columns gives us a three by four matrix of project management tasks. Figure 1 illustrates this taxonomy for the project management subtasks involved in construction of a house.

4. Guidelines for Using AI in Project Management

In this section, we will present a set of guidelines for using proven and emerging AI techniques to support each subtask of project management. We propose the use of *knowledge-based interactive graphics* along with specific AI techniques for each of the subtasks of project management.

4.1 USING AI TECHNIQUES TO SUPPORT PROJECT OBJECTIVE-SETTING

At the early conceptual stages of a project, its planners must make gross trade-offs between scope, time and cost, with considerable uncertainty about the detailed parameters of the project. We suggest that the most challenging cognitive task for this function of project management is to comprehend and analyze the multiple implications that possible assumptions or choices about the project's configuration, location, and timing will have on the various dimensions of project performance that must be balanced against each other. Moreover, this cause and effect modeling and trade-off evaluation must often be performed in the presence of considerable uncertainty about states of nature, about unproven technologies, and about market conditions.

Decision support tools for objective-setting should provide the ability to explore the impact of a variety of assumptions on the various dimensions of project performance in an interactive manner. *Spreadsheets* provide one interactive way to analyze implications of choices deterministically, and in a single dimension, e.g., dollars. *Decision analysis* represents an alternative analytical approach to scenario analysis that can capture and

	Objective Setting	Planning	Scheduling	Control
Executive Level	**Build House :** * choose location * total area, # BR, # BA * style, consider additional structures (garage, swimming pool) * budget * overall duration	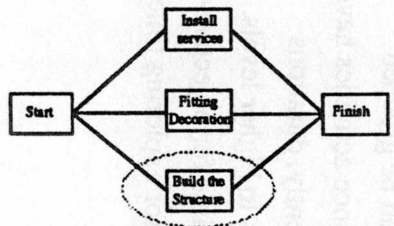	* estimate the overall duration and cost using parameter estimates based on similar, past projects with adjustment.	* monitor completion date, total cost, quality control vs. plans, schedules, and objectives.
Work Package Level	**Build structure :** * choose materials (quality vs. cost, time) * select construction methods (time, cost, quality, environ--mental impact)	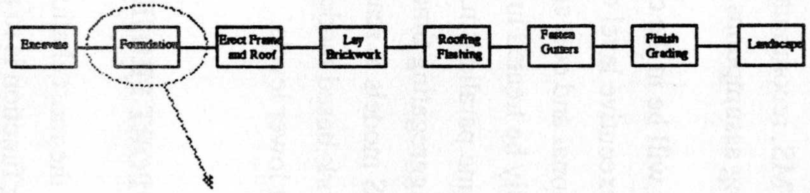	* estimate activities durations and resource requirements based on crew productivity, equipment spreads, and/or subcontractors quotes.	* monitor productivity and production rate data for work package (geotechnical data, concrete cylinder tests, ...etc).
Task Level	**Pour Foundations :** * choose type of foundation (allowable settlement, cost, time) * construction method (crew size, equipment, cost, time)	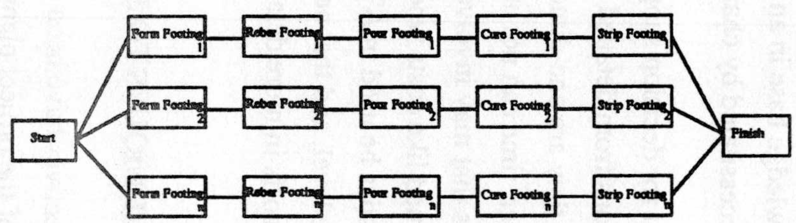	* estimate by analogy with past projects data and/or by modeling the crews and equipment spreads.	* monitor daily labor, material, equipment costs vs. schedules and budgets. * monitor quality in detail (placement of rebar, concrete finishes, ...etc).

FIGURE 1. A Taxonomy of Project Management Tasks for House Construction.
The rows represent levels of management, from executive through task; the columns represent distinct management functions associated with different phases of managing a house construction project. We assert that each cell in the matrix comprises a distinct type of cognitive task with its own set of requirements for decision support.

combine multi-attribute implications of different choices under uncertainty. Yet few project managers currently use either of these tools routinely in exploring project objectives.

Assumption-Based Truth Maintenance Systems (ATMS) [DeKleer86] represent an AI technique that permits an experienced project manager to manipulate knowledge about the implications of choices in the objective-setting stage of a project. Moreover, ATMS techniques facilitate analysis of the effect of changing assumptions in multiple scenarios or "worlds." By maintaining knowledge about which implications depend upon which assumptions, they permit a project team to explore and modify a host of interacting assumptions about a project and to infer the implications of each set of assumptions rapidly. By interfacing analysis programs such as spreadsheets or decision analysis programs with the knowledge base in an ATMS, recomputations of implications can be triggered only when necessitated by changing assumptions.

A key requirement for decision support will be in the computation and aggregation of cost and time implications from detailed to executive level objectives. Costs are especially easy to deal with since they are one-dimensional and cost categories can be grouped hierarchically. However, time can not usually be treated in this way since activities have sequential relationships that may involve some parallelism. Consequently, durations cannot simply be summed like costs when aggregating time objectives to higher levels. We propose that heuristics be used in ATMS models to reason about time for objective-setting at the executive level, and that network-based models be used for exploring time-cost-performance tradeoffs in more detail at lower levels.

4.2 USING AI TECHNIQUES TO SUPPORT PROJECT PLANNING

Once project objectives have been set, the most challenging cognitive task that must be performed as part of the project planning function is to analyze the activities that need to be carried out to reach the project's objectives and synthesize them into a logically correct and efficient plan.

Research by AI scientists to develop systems that can generate project plans and schedules falls into two broad areas: (1) general purpose, domain-independent planners; and (2) mixed systems that employ domain-independent search techniques with some domain-specific knowledge processing. General purpose, *means-end* planners have been widely studied by AI scientists since the early 1970's for a range of planning tasks, typically involving robots. Planners that depend upon more domain-specific knowledge have been developed more recently in a few domains. In this section, we briefly review work in all of these areas and provide some analysis of the capabilities of systems developed to date.

4.2.1 Domain-Independent, "Means-End" Planners

There exists a large body of artificial intelligence research on generating plans -- linear or nonlinear sequences of activities to be executed by some agent (typically a robot) in order to transform an initial world state to some desired goal state. A thorough review of this literature has recently been compiled [Tate85]. We will present a brief overview of the most relevant work on project planning here.

AI planning systems generally use a *means-end* approach. They define the set of simple facts or *literals* that represent the initial state of the world, and the set of *literals* that represent the goal state. A series of potential *actions* are then defined with: *preconditions* of each action (literals that must be true before executing the action); and *effects* of the action on the world state (described by addition and/or deletion of literals). *Means-end* planning systems search through available actions and select an action to be executed which will reduce the difference between the current state and the goal state. This procedure is repeated to generate a sequence of such actions, or a plan.

Since the number of possible plans can be very large, and since actions can interfere with each other's goals, the various planners use different heuristics and procedures to guide their "means-end" search for feasible solutions [Fikes71] [Sacerdoti73] [Sacerdoti75] [Waldinger75] [Tate76]. The classic AI trade-off applies to this problem: reduction in the time taken by the system to search for a solution comes only at the cost of representing and manipulating more knowledge about the particular problem.

NONLIN [Tate76] is a means-end planner that can generate a "partially ordered" plan of parallel and sequential activities, (i.e., a network plan) and it has been applied to generating the plan for a house construction project. Work is currently under way to rewrite and generalize NONLIN as O-PLAN [Currie85] and to enhance its abilities in the areas of project scheduling and resource management [Bell85]. DEVISER, another extension of NONLIN, can constrain activity schedules with the timing of external events, e.g. times during a mission at which the spacecraft's position enables it to take particular photographs [Vere83].

Much of the means-end planning literature deals with robots stacking blocks on a table or moving objects through a series of rooms connected by doors. AI planning systems can have significant value as plan generators in such robot domains or in other domains where the number of possible actions of the agent is limited with actions repeated many times and the amount of knowledge required to choose an appropriate action is limited.

In contrast, most construction projects contain many distinct tasks with relatively little repetition and require significant amounts of knowledge to determine which task should be performed next. When developing high level plans for such projects, the precondition for each unique activity ends up being specified as the completion of another unique activity. The list of activities with their preconditions and effects that must be provided as input to a planner such as NONLIN may be thought of as an *unlinked list* of activity successors. The work of the planner is to "compile" a *linked list* of activities, or to find an explicit sequence of activities which meets the plan objectives. In contrast the list of activities and their successors which is provided to a CPM or PERT program is already an explicit *linked list* which needs no further compiling -- the network is generated directly [Levitt87].

We are asserting here that the *decision-support* utility of the *means-end* planners developed to date appears to be small for overall project planning where large amounts of case-specific knowledge are required to select the next action in a plan from a large number of possible, unique activities. We propose that *means-end* planners are best suited for planning work operations within projects at a fine level of detail or grainsize, in which there is a small number of possible actions that can be performed and there are many repetitions of like actions.

4.2.2 Mixed AI Planners

Some of the more recent AI planning research involves the use of significant amounts of problem-specific knowledge to augment the general search capabilities of means-end planners described in the previous section. We believe that such planners may provide a good compromise between input requirement, power and flexibilty for project planning.

The **BB1** blackboard approach represents an opportunistic or open approach to planning and design problems [Hayes-Roth85, 86]. In this approach, partial plans consisting of one or more activities can be assembled and combined into overall plans or designs using multiple sources of knowledge. **BB1** employs knowledge at three levels: strategic, to achieve high level goals such as "Close Building In"; tactical (called "focus"), such as "Schedule West Wing Concrete Activities"; and implementation" (called "heuristic"), corresponding to individual activities in a plan. In **BB1**, multiple knowledge sources whose trigger conditions are satisfied by the evolving state of a plan are activated. When the preconditions of a possible action in one of these knowledge sources are satisfied, the knowledge source uses heuristics, means-end planning, embedded constraint propagation and resolution, or other problem-solving techniques to propose Knowledge Source Activation Records (KSAR's) for execution. These are all posted to a control blackboard where they can be compared. Using heuristically derived weights, the scheduler selects one KSAR for execution on each cycle. Execution of the KSAR modifies attributes of objects on the control blackboard, triggering new KSAR's for consideration. In this manner, **BB1** constructs a dynamic control plan which uses its knowledge to change tactics or even strategy at any time in response to the evolving plan.

COMTRAC-O [Koo87] employs a decentralized knowledge-based approach to planning. Its design is based upon multiple "intelligent agents" planning their own work (using means-end planning techniques) and then using case-specific knowledge with which they have been provided to communicate selectively with one another by sending messages asking for commitments to perform required predecessor activities. Ongoing research in this area is addressing issues involving assymetry between the agents (analogous to hierarchy in human project teams) and conflict resolution strategies for sharing global resources (such as space or scaffolding).

The Knowledge-Intensive Planner (**KIP**) [Luria87] employs the notion of concerns, goals whose probability of failure is inferred to be significant using specific knowledge about the problem context. Concerns are only injected into the plan as subgoals to be achieved when their failure probability exceeds a specified threshold, and are ignored in planning where they are very likely to succeed. **KIP** has been implemented only in one domain -- user assistance in the UNIX operating system with tasks such as writing or deleting files -- but this approach to knowledge-based planning appears to offer some new ways of limiting the planner's search space in other, more complex planning domains.

[Wilkins84] designed the architecture of SIPE as a combination of a means-end search procedure with constraint propagation and resolution of domain-specific knowledge. A least commitment approach is used to control the search for a valid plan. Objects or operators in SIPE are described in terms of constraints whose variables bind to the domain-specific knowledge in the frames as late as possible. This work is a generalization of concepts first proposed in MOLGEN for planning experiments in a particular domain, molecular biology [Stefik81]. Wilkins claims that SIPE achieves efficiencies over previous planners by: the use of a frame-based hierarchy to represent all parts of the domain in a uniform way; the use of constraints to postpone instantiation of objects; and the use of "resources" to resolve harmful interactions between parallel actions.

Systems employing virtually all domain-specific knowledge, such as traditional knowledge-based expert systems, have been developed to generate plans in very limited domains [Hendrickson86b]. However, we suggest that systems like the four discussed here, which combine domain-specific and domain-independent planning approaches, may provide a compromise between depth and breadth of knowledge that will be better suited for executive or work package level project planning than either means-end planners or expert systems.

4.3 USING AI TECHNIQUES TO SUPPORT PROJECT SCHEDULING

In preparing detailed activity durations and resource requirements for a project at the *work package* level, schedulers may devote considerable effort to modeling the crews and equipment spreads planned for a project, including detailed consideration of work methods and the layout of temporary facilities. For *executive* level scheduling, the process is more likely to involve interpretation of overall project conditions in comparison with a group of similar, past projects, to estimate directly the duration, labor and other resources required to complete each major activity.

At the *task* level, schedulers tend to use a combination of the above approaches. For some tasks, especially those to be subcontracted and for which they may have less detailed knowledge, they typically adjust unit resource consumption rates and production rates from past projects using attributes of the project's design and working environment. For tasks to be executed with their own forces, they may design hypothetical crews and build up detailed estimates using qualitative or even quantitative simulation models of the construction or manufacturing process.

Once activity durations and resource consumption levels have been estimated, most existing project scheduling algorithms then carry out an analysis of the network and assign activities to their early start times, assuming resources are infinite. Rescheduling to smooth resource peaks, or to satisfy hard resource constraints, involves subsequent cycles of analysis and synthesis.

Advanced AI programming environments such as *IntelliCorp's* Knowledge Engineering Environment (KEE) provide facilities for graphical input and output to and from models embedded in frame/rule systems. Building on the ideas proposed by [Paulson71], we propose the use of such *knowledge-based interactive graphics* systems in which users can perform the synthesis by modifying schedules incrementally, and can then receive analysis results immediately to show the implications of their most recent changes. Knowledge Processing System (KPS) approaches have been explored and found viable for generating executive level activity durations and resource needs from overall project data, and for generating work package level schedules from more detailed project attributes [Hendrickson86a]. Task level scheduling involves a combination of the types of decision making involved in executive and work package level scheduling. Consequently, KPS should provide valuable leverage for scheduling at this level, too.

We, therefore, propose that KPS techniques be used in combination with knowledge-based interactive graphics and procedural analysis tools to support project scheduling at all levels. At the executive level the KPS will primarily interpret data from past projects and make duration and resource estimates for activities based upon global characteristics of the current project; whereas at lower levels it can use simulation or other modeling techniques to generate durations and resource requirements for activities constructively from more detailed information about the current project.

4.4 USING AI TECHNIQUES TO SUPPORT PROJECT CONTROL

Existing project management packages employ database management systems extensively as aids to project control. They generate reports of task, work package and executive level project performance measures against original or revised plans. Such reports provide potentially useful data for project managers at all levels to exercise control of their projects. Since the volume of project status information can be overwhelming, exception reports or highlights are widely used to filter project status information. Two construction firms (Hitachi in Japan, and Guy F. Atkinson company in the US) have taken this to considerable lengths, and have implemented filtering systems which use several different attributes of each activity or cost account to determine whether it should be flagged for exception reporting [Niwa83] [Teicholz86]. However, most companies use a simple variance of actual cost against budgeted cost, or actual duration against planned duration, as a filter.

Meaningful project control, especially at the executive level, involves high level cognitive skills of:

- **Interpretation** -- "Which among all of these performance data are significant and what do they imply?"
- **Prediction** -- "What will be the durations and resource requirements for remaining activities if present trends continue?"
- **Diagnosis** -- "What caused the significant variances from expectations?" and
- **Prescription** -- "What remedial actions should be taken?"

When they are used by knowledgeable project managers, techniques such as CPM or PERT network scheduling and discrete event simulation, which create deterministic or stochastic models of resource and time consumption on projects, are valuable analysis tools in support of the interpretation, diagnosis and prediction tasks associated with project control. However, we have pointed out in a previous paper [Levitt85] that the constant requirement for this type of knowledgeable input during a project makes the cost of using existing network-based programs for real-time project control too high in many cases.

KPS techniques have been successfully used to store and use knowledge required for interpretation, prediction, diagnosis and prescription tasks in many domains, including project management [Levitt85]. We, therefore, propose that hybrid project control systems, which leverage traditional project analysis tools with knowledge processing capabilities for database access, interpretation, diagnosis and prediction, will permit effective capture and use of much of the expert knowledge currently involved in project control at all levels of detail.

Our guidelines for the use of AI techniques to support each of the project management functions at the various levels of management are summarized in **Figure 2.**

5. The PLATFORM Experiments

Using the taxonomy of project management tasks developed above, we will briefly describe the PLATFORM AI-leveraged project management systems. Experience from the development and subsequent use of these systems provides strong support for many of the guidelines presented above.

The PLATFORM systems represent a series of prototype project management decision support systems developed as hybrids of traditional, procedural analysis packages and KPS capabilities. They were implemented over a period of about two years by the authors as a collaborative research effort to explore the potential applications of AI programming environments such as **KEE** in the domain of project management.

5.1 PLATFORM I: USING PLANNERS' KNOWLEDGE FOR PROJECT CONTROL

PLATFORM I is a KPS, developed using the *IntelliCorp* Knowledge Engineering Environment (**KEE**) system, which employs model-based reasoning for project control [Levitt85]. It uses a PERT project network to model the sequential and temporal connections between the activities required to build a concrete gravity type of offshore oil platform in the North Sea. It integrates this PERT modeling capability with KPS techniques for reasoning with heuristic knowledge provided by the planner about risk factors that could impact each activity's duration. PLATFORM I represents and reasons with planners' knowledge in two ways:

5.1.1 Contingent Subnetworks

A project may have some activities whose construction method -- and hence duration -- will depend significantly upon the value of some variable whose value cannot be determined a priori. Such activities can be treated as a series of *Contingent Subnetworks* which the planner defines, each containing the tasks that would have to be performed for one state of the unknown variable. The appropriate contingent subnetwork will subsequently be activated when the uncertainty is resolved, manually by the user, or automatically by

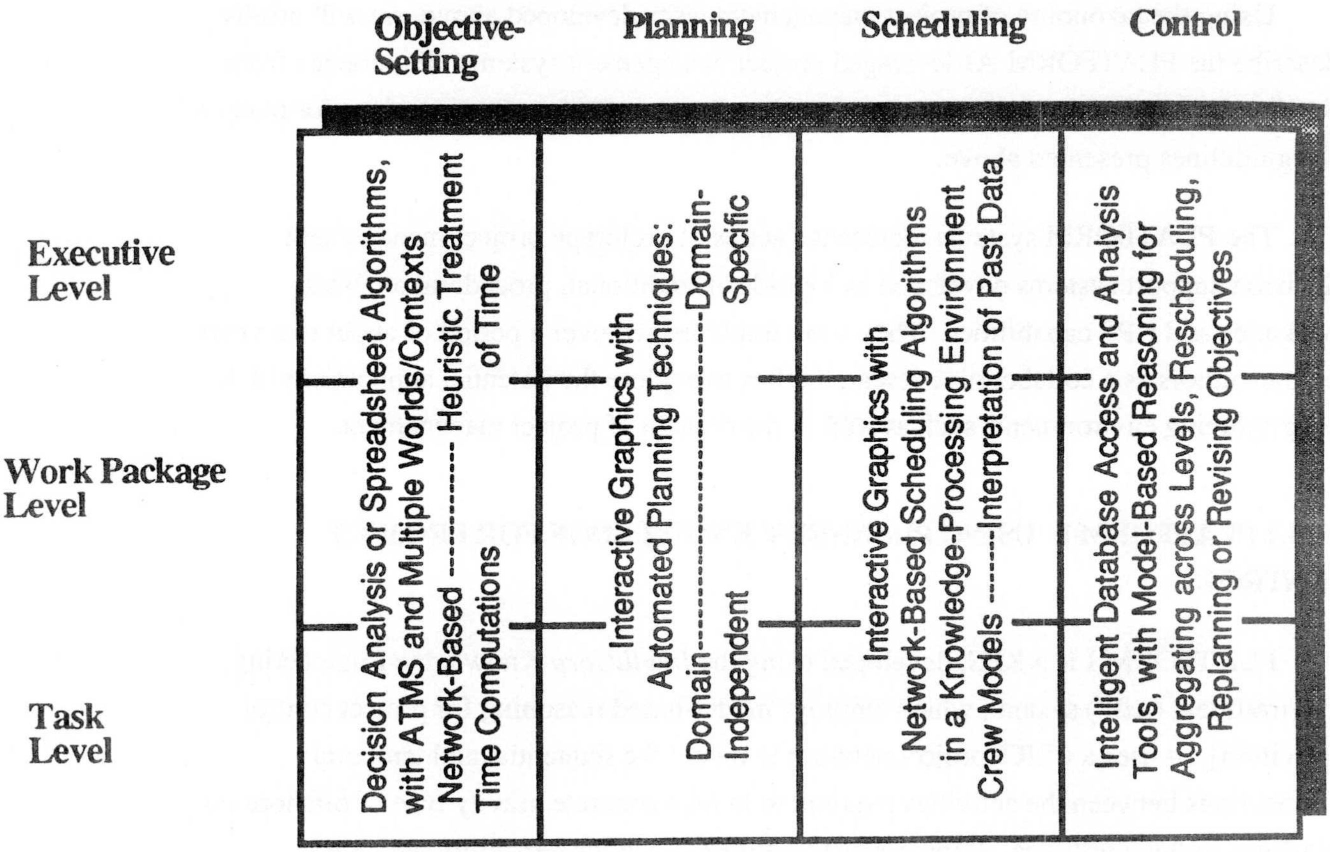

FIGURE 2. Guidelines for the use of AI Techniques in Support of Project Management Functions at Various Levels of Decision Making. The entries in each cell of the matrix are the AI techniques that we propose as being most valuable for decision support of that function at that level.

PLATFORM I according to prespecified heuristics.

The example provided in [Levitt85] illustrates three different sets of tasks required for construction of a temporary drydock in different types of soil -- sand, silt, or clay -- according to a set of rules for classifying the *in-situ* soil.

5.1.2 Knights and Villains

Traditional project management approaches assume that activity durations are either deterministic (CPM) or vary independently of one another (PERT). The actual durations of the activities on any project will vary and may be closely correlated with a series of risk factors such as weather, and hence with each other. Previous researchers have attempted to use statistical techniques to capture this corellation but the problem becomes computationally and cognitively intractable for large networks [Carr71] [Woolery83].

In PLATFORM I, we use frames to record the planner's knowledge about risk factors which might influence each activity's duration, and then use a small number of heuristics to infer from interim performance data on a project which risk factors have become *knights* (i.e., they are having a positive effect on durations) or *villains* (i.e., they are delaying activities). The durations of future activities are then revised downwards to a more optimistic value if the activity is impacted by one or more *knights*, and upwards to a more pessimistic duration if the activity is impacted by one or more *villains*. Since any *villain* can delay an otherwise productive activity, we assume that *villains* will override *knights*.

With this modest amount of domain-specific knowledge and a few simple heuristics, PLATFORM can interpret schedule performance on a partially completed project, diagnose potential risk factors as *knights* or *villains*, and provide revised forecasts for the durations of remaining activities. The forecasts generated by PLATFORM I for an offshore oil platform project have been judged by experienced project managers to be reasonable.

This knowledge-based updating has several important benefits. First, it assists a project manager to sort through vast amounts of project status information and identify meaningful patterns of variance. Second, it provides project personnel with more realistic forecasts of project completion time than either assuming that future activity durations will be as planned (the default assumption in existing packages) or revising them by simple trending of average durations. This can help to develop appropriate levels of concern early in a project when there is still a chance for corrective actions. Finally, since *knights* and *villains* are identified explicitly in PLATFORM's analysis, it helps to direct management attention on the project towards maintaining the status of any *knights*, and "vanquishing" any *villains*.

Contingent subnetworks and *knights and villains* represent two kinds of domain-specific planning knowledge that are not usable by traditional project management systems. By integrating them in a knowledge processing system together with a traditional network analysis algorithm, the advantages of each are exploited. Readers interested in additional details about the architecture used for representation, reasoning and explanation in PLATFORM I are referred to [Levitt85].

5.2 PLATFORM II: DECISION SUPPORT FOR PLANNING AND SCHEDULING

The goal of PLATFORM II was to develop a highly interactive graphical environment that could be used to create and modify project plans and schedules. In addition to this, however, the system should provide KPS capabilities for critiquing plans, generating duration and resource estimates, interpreting performance and other intelligent functions. Accordingly it was developed as an extension of PLATFORM I described below, in the KEE environment.

SIMKIT™ is a knowledge-based, discrete event simulation package developed by *IntelliCorp* to work in conjunction with KEE. We specialized SIMKIT's graphical model editor to develop a graphical network editing tool, and used the *KEEPICTURES* knowledge-based interactive graphics facility of KEE to create customized graphics for an interactive Gantt Chart. KEE provides the environment under which PLATFORM II runs and affords access to its rule system and active images. The frame and inheritance mechanisms in KEE permit us to create templates for different kinds of projects (e.g., oil platforms, or software development efforts) which can be conveniently specialized for a particular project by providing specific project data or overriding the local default values from the template.

Schedule logic is defined and edited interactively with a mouse pointer by selecting icons representing different types of activities or milestones from a library, placing them on the *canvas*, and connecting them with precedence lines to appropriate predecessors and successors. Duration and resource information can be provided in this view if desired.

Figure 3 shows the interactive planning canvas from PLATFORM II.

When a model's logic has been entered, activities can be provided with durations and resources via menus and table entries. Templates are helpful here in determining which resources are likely to be required by each activity, e.g., a concreting activity selected from the menu, will cause the activity to inherit *cement masons* as a resource, with unknown quantity. All the user has to do is fill in the blank quantity.

Once activity durations have been provided, the user can select a Gantt chart representation of the project. What makes this Gantt chart different than those provided in other project management packages is that it is alive. That is, the user can mouse on bars representing activities and move them back and forth within the *float windows* that define their earliest start and latest finish times. Successors or predecessors will be bumped backwards or forwards as necessary, cascading through the entire network, to preserve precedence relationships defined in the network model. Modifying the display results in equivalent changes to the slots in activity frames, termed *units* in **KEE**, in which the planned start and finish times are stored, and where they can be read by rules, error-checked by value class settings, and monitored as *active values* (demons) to trigger other computations or alarm messages.

Simultaneously, the user can view one or more live histograms showing the project's demand for a given resource in each time period against the availability constraint for that resource. As activities are moved around, these resource histograms are updated and displayed immediately. This permits the user to carry out resource leveling interactively for multiple resources with unprecedented ease. A *GanttAlive* chart with several resource histograms is shown in **Figure 4**.

5.3 PLATFORM III: DECISION SUPPORT FOR OBJECTIVE-SETTING

Decision Analysis is a technique that has frequently been used for considering alternative scenarios and selecting optimum choices for aerospace projects, and has occasionally been used for construction projects. The authors and some colleagues developed PLATFORM III, a prototype KPS to assist a project manager in exploring alternatives at the early objective-setting phase of a project, using the *Multiple Worlds* feature of *IntelliCorp's* **KEE** system [Kunz86].

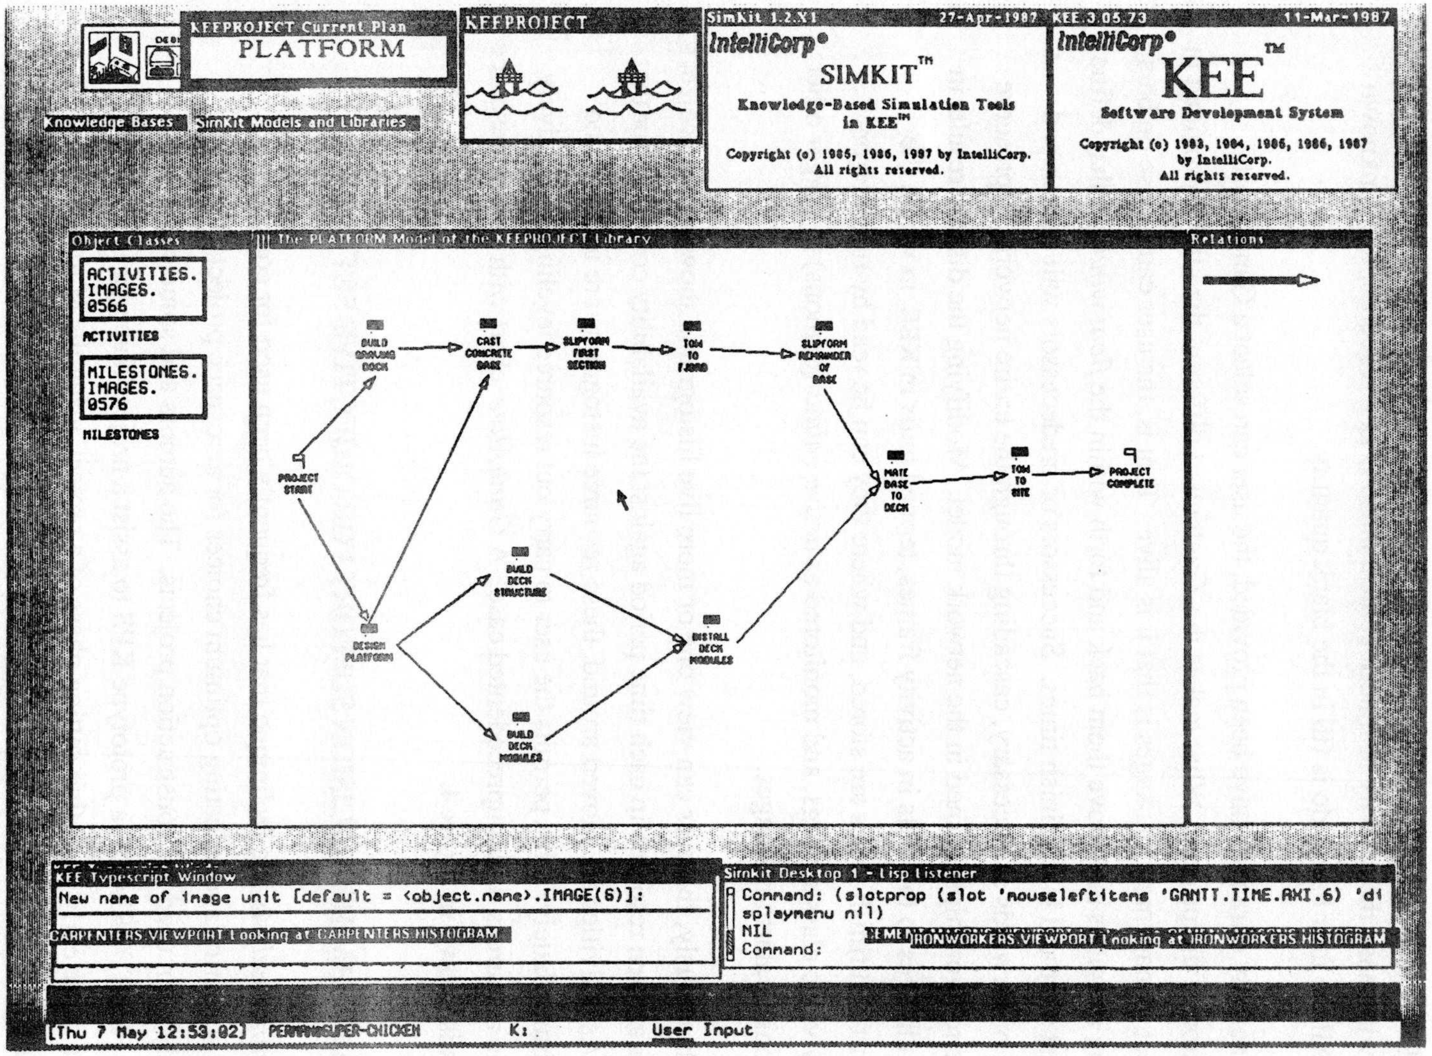

FIGURE 3. PLATFORM II: An Interactive Project Planning System. Activities are placed on the canvas and connected to predecessors and successors using a mouse pointing device with menu driven commands for setting durations and resources. Data defined via this interface are automatically stored in activity frames where they can be read by rules or demons.

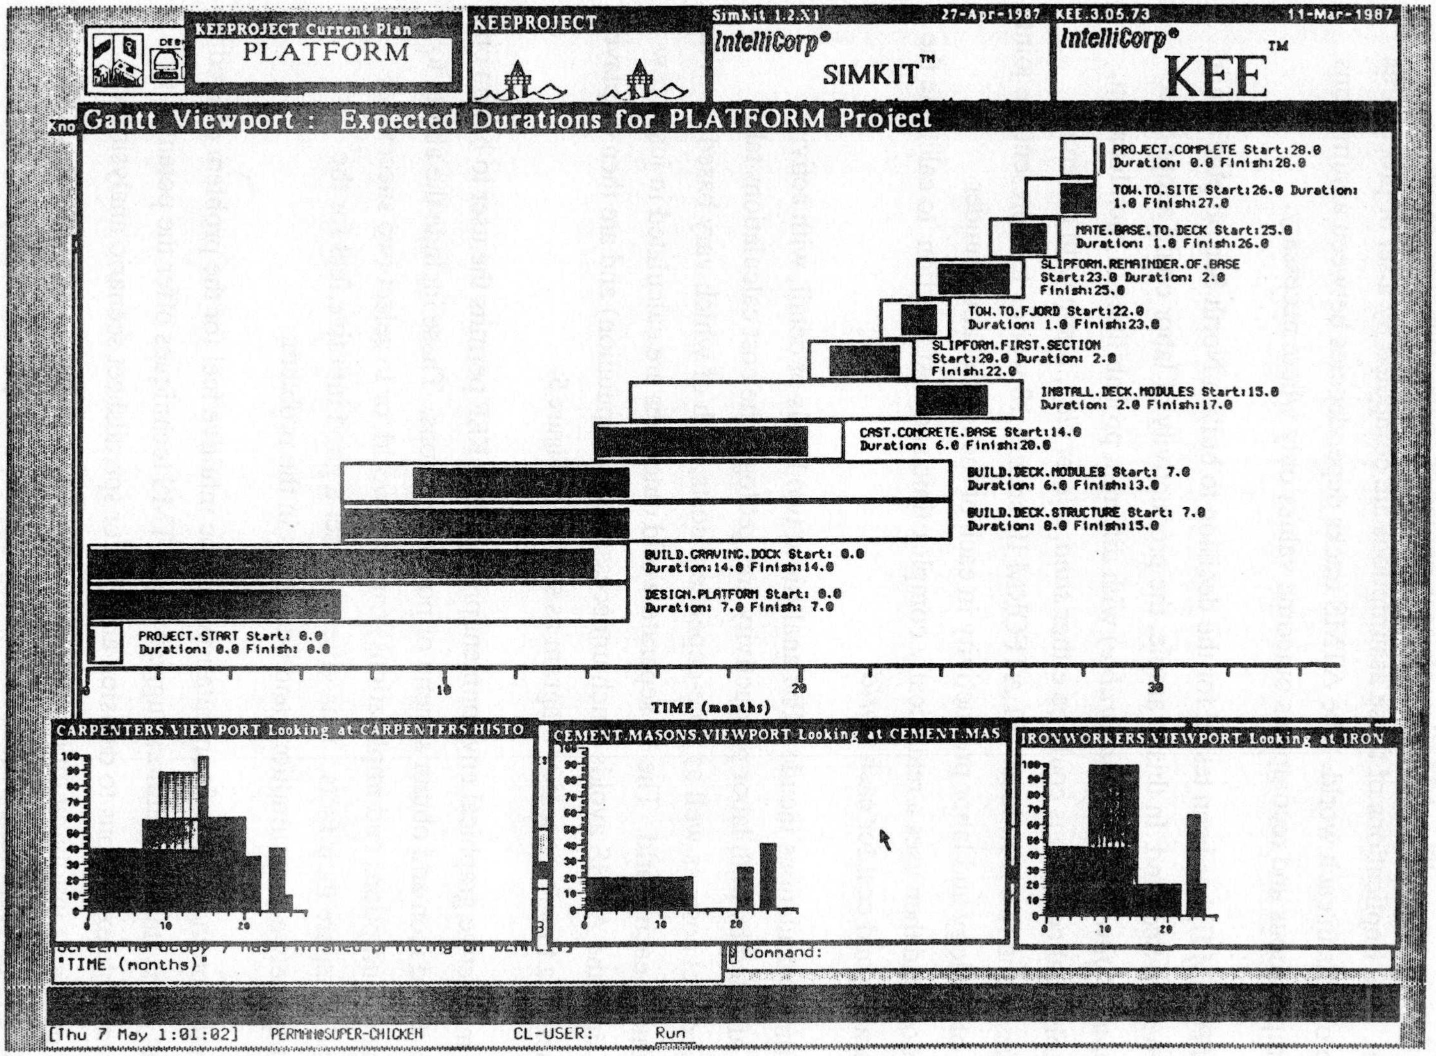

FIGURE 4. The GanttAlive Interface in PLATFORM II: An Interactive Scheduling System with Real-time Display of Resource Implications. This interface permits a user to experiment with different planned start and fininsh times for activities by moving them around within their float windows, or extending the project's duration to obtain more float for activities. Planned start and finish times are posted to slots in activity frames where they can be read by demons or rules in KPS applications that support scheduling.

PLATFORM III uses the KEE system's multiple context mechanism (termed *Multiple Worlds*) and its assumption-based truth maintenance system (ATMS) to create a series of parallel scenarios or *Worlds*, each of which corresponds to a single leaf node in a traditional decision tree. The ATMS treats each choice or chance outcome as an assumption. The implications of those assumptions are computed by a set of procedures attached to objects in each world. The ATMS tracks dependencies between assumptions and their implications and recomputes outcome values only where necessary.

PLATFORM III has been tested for the decision to build a North Sea oil platform in either Norway or Scotland. In this example, the productivity of labor could adopt three values (*nominal, favorable or unfavorable*) with different probabilities in each location, and the geological conditions could be either *sand, silt,* or *clay*, again with different probabilities in the two locations. PLATFORM III creates 18 worlds to represent the joint outcomes of geology and labor productivity in each of the two locations under consideration, and then uses a realistically complex time and cost function to calculate the project duration and cost for each *world.*

The time calculations incude CPM analysis at two levels of detail, with activity durations affected by both labor productivity and geology; the cost calculations take account of direct costs as well as time-dependent costs, both of which vary based upon the assumptions in each world. The 18 separate world outcomes are calculated in just a few seconds (since the ATMS avoids much unnecessary recomputation) and are then displayed for the user on a time/cost scatter diagram as shown in **Figure 5.**

The interactive graphical environment provided by KEE permits the user to button on a particular world icon and obtain a menu of possible actions. These include the ability to view facts (assumptions and implications) true in a world, or to select two successive worlds and compare their facts. This gives the user a good intuitive feel for the relationships between assumptions and outcomes in the problem.

Based upon the speed of computation and the intuitive feel for the problem offered by this prototype system, the authors suggest that ATMS techniques offer the potential to provide valuable extensions to decision analysis or spreadsheet scenario analysis for complex decisions involving multiple uncertainties and complex computations. The prototype and its preliminary results are described more fully in [Kunz86].

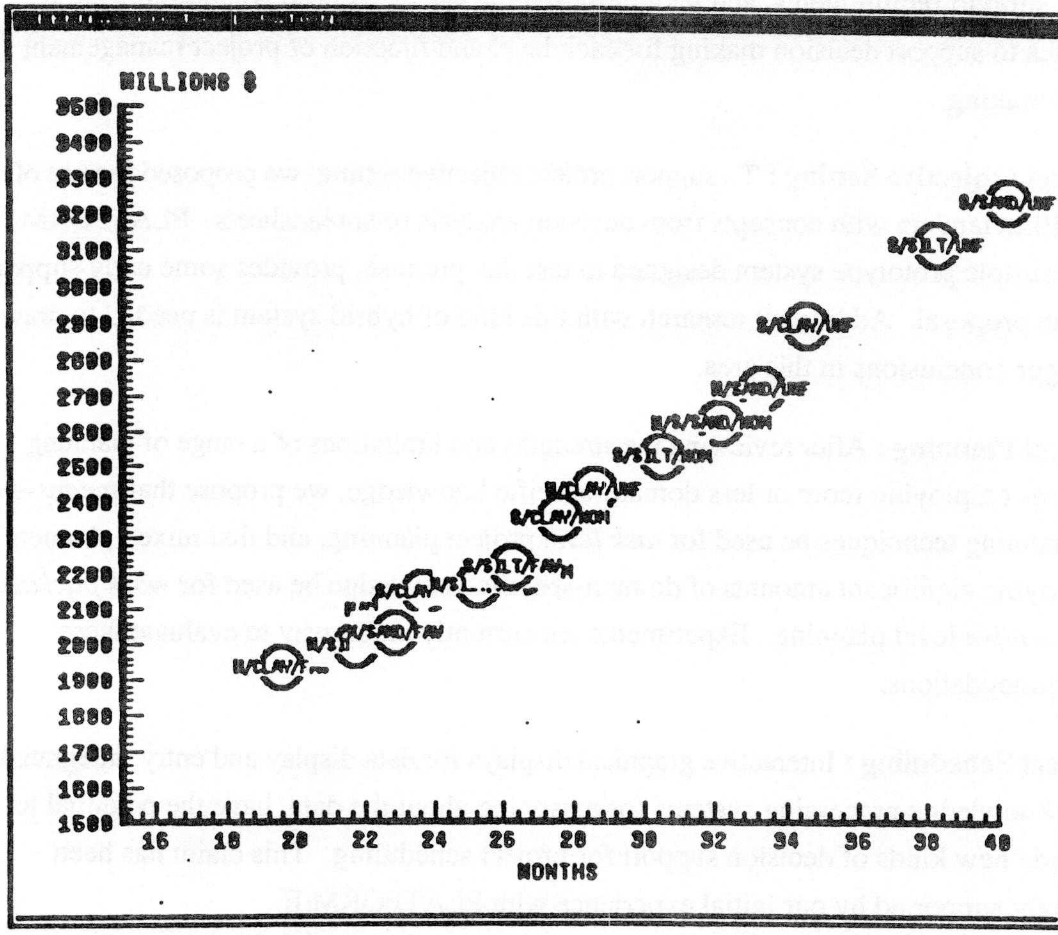

FIGURE 5. PLATFORM III: Analysis of Alternatives in Objective-setting Phase Using ATMS. This time-cost scatter diagram for constructing an offshore oil platform is rapidly computed by PLATFORM III for 18 Possible scenarios ("worlds") based upon nine different sets of assumptions about labor productivity and geological conditions in Norway versus Scotland. The icons representing each possible world are active and can be selected with a mouse pointing device to show facts true in each world. Two or more worlds can likewise be compared to highlight their differences. New worlds with slightly different assumptions can be "sprouted" from previously generated worlds.

6. Conclusions

We defined a taxonomy of project management tasks with specific cognitive and decision support requirements, and then proposed the use of specific AI and procedural techniques to support decision making for each level and function of project management decision making.

Project Objective Setting : To support project objective setting, we proposed the use of ATMS in tandem with concepts from decision analysis or spreadsheets. PLATFORM III, a simple prototype system designed to test this premise, provides some early support for the proposal. Additional research with this kind of hybrid system is needed to draw stronger conclusions in this area.

Project Planning : After reviewing the strengths and limitations of a range of planning systems employing more or less domain-specific knowledge, we propose that *means-end* AI planning techniques be used for *task level* project planning, and that mixed planners employing significant amounts of domain-specific knowledge be used for *work package* or *executive* level planning. Experiments are currently under way to evaluate these recommendations.

Project Scheduling : Interactive graphical displays for data display and entry, integrated with knowledge processing systems for reasoning about the data, have the potential to provide new kinds of decision support for project scheduling. This claim has been strongly supported by our initial experience with PLATFORM II.

Project Control : For project control at all levels, we proposed the use of hybrid systems combining traditional network-based algorithms with KPS capabilities. Our initial results from the PLATFORM I system are most encouraging in this regard. The *knights* and *villains* knowledge-based approach to interpretation, diagnosis and prediction of project schedules has proven to be a powerful alternative to statistical techniques. Several researchers are currently experimenting with variations of this approach to project control. We plan to continue this line of research in collaboration with industry experts to develop KPS for project control that will exhibit truly expert levels of performance.

Our analysis of the AI literature and our work with hybrid AI-procedural systems leads us to the conclusion that AI techniques can extend the range and power of conventional computer tools for project management in important ways. Our analysis of research on AI planning has shown that, if used appropriately, both domain independent and domain specific planners have potential for use in project planning. Moreover, our experience in developing hybrid AI-graphics-procedural systems for project objective-setting, scheduling and control has convinced us that small increments of knowledge can be used by such systems to provide significant additional power in decision support over conventional procedural tools for a range of project management tasks.

7. Acknowledgements

The development of the PLATFORM systems was initiated with funding from a sabbatical study grant by *IntelliCorp* to the first author during 1984-85. Several members of *IntelliCorp's* staff have subsequently contributed ideas and code to the PLATFORM project: Thomas Bonura, Ray Carhart, Eric Lazarus, Philip McBride, Michael McNabb, Catherine Perman, Deborah Sharpe and Marilyn Stelzner helped to implement various features of the PLATFORM prototypes described in this paper. Austin Tate and Ken Currie of the *Artificial Intelligence Applications Institute* in Edinburgh shared valuable insights from their research on AI planning systems with us. We gratefully acknowledge the help we have received from all of these colleagues and friends.

8. References

[Bell85]

Bell, C.E. and Tate, A. : *Use and Justification of Algorithms for Managing Temporal Knowledge in O-Plan,* AIAI-TR-6, Artificial Intelligence Applications Institute, University of Edinburgh, UK, Jan. 1985.

[Carr71]

Carr, R.I. : *Synthesis of Uncertainty in Construction Planning,* Ph.D. dissertation, Department of Civil Engineering, Stanford University, 1971.

[Chandra86]

Chandra, N. and Marks, D.H. : Intelligent Use of Constraints for Activity Scheduling, *Proceedings of the First International Conference on Applications of Artificial Intelligence to Engineering Problems,* edited by D. Sriram and R. Adey, Southampton University,UK, pp. 369-382, April 1986.

[Cohen82]

Cohen, P. and Feigenbaum, E. : "Planning and Problem Solving" , Chapter XV in *The Handbook of Artificial Intelligence,* Volume 3, William Kaufmann Inc., Los Altos, California, 1982.

[Currie85]
Currie, K. and Tate, A. : *O-Plan - Control in the Open Planning Architecture*, AIAI-TR-12, Artificial Intelligence Applications Institute, University of Edinburgh, UK, December 1985.

[DeKleer86]
DeKleer, J. : An Assumption-Based TMS, *Artificial Intelligence*, Vol. 28, No. 1, 1986.

[DeLaGarza87]
DeLaGarza, J.M. and Ibbs, W. : Issues in Construction Scheduling Knowledge Representation, *Proceedings of the CIB W-65 Symposium*, London, 1987.

[Fikes71]
Fikes, R.E. and Nilsson, N.J. : STRIPS: a New Approach to the Application of Theorem Proving to Problem Solving, *Artificial Intelligence*, 2, pp 189-208, 1971.

[Fondahl61]
Fondahl, J.W. : *A Noncomputer Approach to the Critical Path Method for the Construction Industry*, Technical Report No. 9, Stanford University, November 1961.

[Hayes-Roth85]
Hayes-Roth, B. : A Blackboard Architecture for Control, *Artificial Intelligence Journal*, Vol. 26, pp. 251-321, 1985.

[Hayes-Roth86]
Hayes-Roth, B., Garvey, A., Johnson, M.V. and Hewett, M. : *A Layered Environment for Reasoning about Action*, Department of Computer Science, KSL-86-38, Stanford University, California, 1986.

[Hendrickson86a]
Hendrickson, C., Martinelli, D. and Rehak, D. : "Hierarchical Rule-Based Activity Duration Estimation", *ASCE Journal of Construction Engineering and Management*, 1986.

[Hendrickson86b]
Hendrickson, C., Zozaya-Gorostiza, C., Rehak, D., Baracco-Miller, E., and Lim, P. : *"An Expert System Architecture For Construction Planning"*, EDRC-12-07-87, Carnegie Mellon University, Dec.1986.

[Koo87]
Koo, C. : Intelligent Planning for Assembly Manufacturing, unpublished paper, 1987.

[Kunz86]
Kunz, J., Bonura, T., Stelzner, M. and Levitt, R. : Contingent Analysis for Project Management Using Multiple Worlds, *Proceedings of the First International Conference on Applications of Artificial Intelligence to Engineering Problems*, edited by D. Sriram and R. Adey, Southampton University,UK, pp. 707-718, April 1986.

[Kunz87]
Kunz, J. : *Model Based Reasoning*, unpublished manuscript, 1987.

[Levitt85]
Levitt, R.E. and Kunz, J.C. : Using Knowledge of Construction and Project Management for Automated Schedule Updating, *Project Management Journal*, Vol. 16, No. 5, December 1985.

[Levitt87]
Levitt, R.E. and Kunz, J.C. : Using Artificial Intelligence Techniques to Support Project Management, *The Journal of Artificial Intelligence in Engineering Design, Analysis and Manufacturing*, Vol 1, No.1, 1987.

[Luria87]
Luria, Marc : Concerns: A Means of Identifying Potential Plan Failures, *In the Proceedings of the third IEEE conference on AI applications*, Orlando, Florida, Feb.1987.

[Moder83]
Moder, J.J., Phillips, C.R. and Davis, E.W. : *Project Management with CPM, PERT and Precedence Diagramming*, Var Nostrand Reinhold, 1983.

[Morris86]
Morris, P. and Nado, R. : Representing Actions with an Assumption-Based Truth Maintenance System, *Proceedings of the AAAI*, 1986.

[Nay85]
Nay, L.B. and Logcher, R.D. : *Proposed Operation of an Expert System for Analyzing Construction Project Risks*, Technical Report R85-5, Order #CCRE-85-Z, M.I.T. Department of Civil Engineering, February 1985.

[Niwa83]
Niwa, K. and Sasaki, K. : A New Project Management System Approach: The "Know-How" Based Project Management System, *Project Management Journal*, Vol. 14, No. 1, March 1983.

[O'Connor86]
O'Connor, M.J., DeLaGarza, J.M. and Ibbs, C.W. : An Expert System for Construction Schedule Analysis, in *"Expert Systems in Civil Engineering"* edited by Kostem and Maher, ASCE, April 1986.

[Paulson71]
Paulson, Boyd : *Man-Computer Concepts For Project Management*, Tech. Rep. #148 , Department of Civil Engineering, Stanford University, August 1971.

[Sacerdoti73]
Sacerdoti, E.D. : Planning in a Hierarchy of Abstraction Spaces, *International Joint Conference on Artificial Intelligence*, pp. 412-422, Palo Alto, California, 1973.

[Sacerdoti75]
Sacerdoti, E.D. : The Non-linear Nature of Plans, *International Joint Conference on Artificial Intelligence*,pp. 206-214, Tbilisi, USSR, 1975.

[Sathi86]
Sathi, A., Morton, T.E. and Roth, S.F. : Callisto: An Intelligent Project Management System, *Artificial Intelligence*, Winter 1986.

[Stefik81]
Stefik, M.J. : Planning with Constraints, *Artificial Intelligence*, 16, pp 111-140, 1981.

[Stefik86]
Stefik, M.J. : The Next Knowledge Medium, *Artificial Intelligence*, Vol. VII, No.1, Spring 1986.

[Tate76]
Tate,A. : *Project Planning Using A Hierarchic Non-linear Planner*, Department of Artificial Intelligence, University of Edinburgh, Research Report No. 25, August 1976.

[Tate77]
Tate,A. : Generating Project Networks, *International Joint Conference on Artificial Intelligence*, pp. 888-893, Boston, Massachusetts, 1977.

[Tate85]
Tate,A. : A Review of Knowledge-Based Planning Techniques, *Knowledge Engineer's Review*, British Computer Society Specialist Group on Expert Systems, Vol 1, No. 2, June 1985.

[Teicholz86]
Teicholz, P. : Personal communication with Dr. Paul Teicholz, Manager of Corporate Information Systems, *Guy F. Atkinson Co.*, South San Francisco, California, 1986.

[Vere83]
Vere, S. : Planning in Time: Windows and Durations for Activities and Goals, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 246-267, May 1983.

[Waldinger75]
Waldinger, R. : *Achieving Several Goals Simultaneously*, Stanford Research Institute-AI Center, Technical Note 107, Menlo Park, California, 1975.

[Wilkins84]
Wilkins, D.E. : "Domain-independent Planning: Representation and Plan Generation", *Artificial Intelligence* 22, pp. 269-301,1984.

[Woolery83]
Woolery, J.C. and Crandall, K.C. : Stochastic Network Model for Planning Scheduling, *Journal of Construction Engineering and Management*, ASCE, Vol. 109, No. 3, September 1983.