

Integrating Planning and Reactive Control

David E. Wilkins Karen L. Myers
 SRI International Artificial Intelligence Center
 333 Ravenswood Ave. Menlo Park, CA 94025
 phone: 415-859-2057 fax: 415-859-3735
 {myers, wilkins}@ai.sri.com

KEY WORDS AND PHRASES

planning, reactive control, intelligent agents, replanning, uncertain reasoning

INTRODUCTION

Our research is developing persistent agents that can achieve complex tasks in dynamic and uncertain environments. We refer to such agents as *taskable, reactive agents*. An agent of this type requires a number of capabilities. The ability to execute complex tasks necessitates the use of strategic plans for accomplishing tasks; hence, the agent must be able to synthesize new plans at run time. The dynamic nature of the environment requires that the agent be able to deal with unpredictable changes in its world. As such, agents must be able to react to unanticipated events by taking appropriate actions in a timely manner, while continuing activities that support current goals. The unpredictability of the world could lead to failure of plans generated for individual tasks. Agents must have the ability to recover from failures by adapting their activities to the new situation, or replanning if the world changes sufficiently. Finally, the agent should be able to perform in the face of uncertainty.

Many domains of interest require problem-solving agents with the capabilities described above. Military and space operations provide good examples. Certainly one would not engage in an undertaking such as Desert

Storm or repairing the Hubble Space Telescope without first formulating a strategic mission plan. Reactive response and failure recovery are necessary because unexpected equipment failures, weather conditions, enemy actions, and other events may require changes to the overall strategic plan.

The Cypress system, described here, provides a framework for creating taskable, reactive agents. Several features distinguish our approach: (1) the generation and execution of complex plans with parallel actions, (2) the integration of goal-driven and event-driven activities during execution, (3) the use of evidential reasoning for dealing with uncertainty, and (4) the use of replanning to handle run-time execution problems.

Our model for a taskable, reactive agent has two main intelligent components, an *executor* and a *planner*. The two components share a library of possible actions that the system can take. The library encompasses a full range of action representations, including *plans, planning operators, and executable procedures* such as predefined standard operating procedures (SOPs). These three classes of actions span multiple levels of abstraction.

The executor is always active, constantly monitoring the world for goals to be achieved or events that require immediate action. In accord with its current beliefs and goals, the executor takes actions in response to these goals and events. Appropriate responses include applying SOPs stored in the action library, invoking the planner to produce a new

plan for achieving a goal, or requesting that the planner modify a previous plan during execution. The planner should be capable of synthesizing sophisticated action sequences that include parallel actions, conditional actions, and resource assignments. The planner plans only to a certain level of detail, with the executor taking that plan and expanding it at run time by applying appropriate library actions at lower levels of abstraction.

CYPRESS

Cypress constitutes a framework in which to define taskable, reactive agents based on the above model. The architecture of Cypress is depicted in Figure 1.

The motivation for Cypress was to build a heuristically adequate system for use in practical applications. To this end, Cypress relies on mature, powerful planning and execution technologies, namely the SIPE-2 generative planner [5] and the PRS-CL reactive execution system [5]. We have applied Cypress to a number of demanding problems, including real-time tracking, fault diagnosis on the Space Shuttle, production-line scheduling, and military operations [5].

PRS-CL is a framework for constructing persistent, real-time controllers that perform complex tasks in dynamic environments while responding in timely fashion to unexpected events. It has been used to monitor the Reaction Control System (RCS) of the Space Shuttle [5]. This application illustrates the use of multiple agents, and has been used to detect and recover from most of the possible malfunctions of the RCS, including sensor faults, leaking components, and regulator and jet failures. The system demonstrated guaranteed response, support for asynchronous inputs, interrupt handling, continuous operation, and handling of noisy data.

SIPE-2 is a partial-order AI planning system that supports planning at multiple levels of abstraction. It has the properties required by our agent model, including the ability to generate plans that include parallel

actions, conditional actions, resource assignments, and the ability to modify previously generated plans. In contrast to most AI planning research, heuristic adequacy has been a primary design goal of SIPE-2.

PRS-CL and SIPE-2 employ their own internal representations for plans and actions for efficiency. For this reason, Cypress supports the use of an *interlingua* called the ACT formalism [5] that enables these two systems to share information. ACT provides a language for specifying actions and plans for both planners and executors. Cypress includes translators that can automatically map Acts onto SIPE-2 and PRS-CL structures, and one that can map SIPE-2 operators and plans into Acts. Using the ACT interlingua, PRS-CL can execute plans produced by SIPE-2 and can invoke SIPE-2 in situations where run-time replanning is required. The ACT-Editor subsystem supports the graphical creation and display of Acts. Gister-CL [5] implements a suite of evidential reasoning techniques that can be used to analyze uncertain information about the world and possible actions. For example, Gister-CL can be used to reason about uncertain information in order to choose among candidate Acts in either the planner or executor.

In contrast to many other agent architectures, planning and execution operate asynchronously in Cypress, in loosely coupled fashion. This approach makes it possible for the two systems to run in parallel, even on different machines, without interfering with the actions of each other. In particular, PRS-CL remains responsive to its environment during plan synthesis. While the subsystems of Cypress can function independently, Cypress is used most advantageously as an integrated framework that supports a wide range of planning and execution activities.

APPLICATIONS

An example from military operations planning [4] is currently the only implemented application that illustrates the use of all sub-

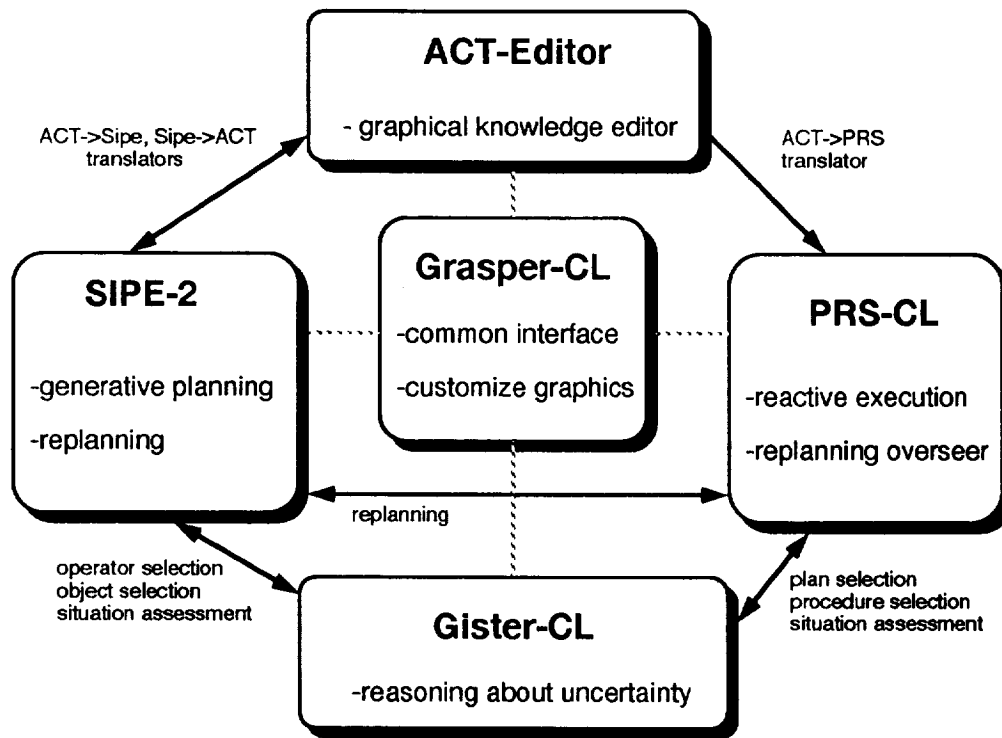


Figure 1: The Architecture of Cypress

systems of Cypress, but it is similar to a space mission. The most advantageous use of Cypress in space applications will most likely be in situations that do not directly involve humans. A planetary rover will certainly need the combination of plan-directed behavior with reactive response to the environment provided by Cypress, and can build directly on our use of Cypress modules to control an indoor mobile robot. Other appropriate space applications include control of a satellite or probe, controlling experiments on the shuttle or space station, and providing an assistant to astronauts to handle routine malfunctions and alert them of important events that affect the overall mission plan.

The military application domain knowledge includes approximately 100 plan operators, 500 objects with 15 to 20 properties per object, and 2200 initial predicate instances. Plans range in size from several dozen to 200 actions, including many that are to be executed in parallel [4].

The scenario begins with a goal request for deterring several military threats. SIPE-

2 uses a set of Acts previously input to the system to generate a plan with many threads of parallel activities. During the planning process, Gister-CL assists SIPE-2 in choosing appropriate military forces for particular missions, by analyzing uncertain information about the situation. Throughout the planning process, PRS-CL monitors the world for additional goals and events that might require immediate action. PRS-CL executes the plan by applying appropriate Acts to refine the plan to lower levels of abstraction, eventually bottoming out in actions that are executable in the world.

PRS-CL responds to many unexpected events by applying Acts representing SOPs. Sometimes an event causes an execution failure that cannot be repaired by any defined Acts (e.g., if transit approval is rescinded for air space that is being used). PRS-CL then invokes a second PRS-CL agent to issue a replanning request to SIPE-2. Meanwhile, the first agent continues execution of parallel threads of the plan not affected by the failure. The planner modifies the plan by eliminating

actions that use the air space in question and replacing them with an alternative mobilization. The actions in the new plan are selected so as not to interfere with the continuing execution of other actions in the original plan. The new plan is sent to the first agent, which integrates the new plan with its current activities and continues.

In a similar fashion, a Cypress agent controlling a planetary rover would have the executor handle unexpected obstacles in its path, and call the planner to modify the plan when progress can no longer be made in the desired direction. On a satellite, the executor could continue to monitor spacecraft systems while requesting the planner to modify the plan for transmitting pictures back to earth after a failure in one of the transmitters.

CONCLUSION

Cypress is a powerful framework in which to define agents that must accomplish complex goals in dynamic and unpredictable environments. The application of Cypress to the military domain and to the Space Shuttle's RCS (only the PRS-CL subsystem is used) attests to the system's usefulness.

The asynchronous replanning facility constitutes one important technological advance, providing flexible plan execution that can adapt to significant unexpected changes in the world. An interesting technical problem that had to be solved was the design of ACT as a common representation for both executors and planners. PRS-CL had to be extended in numerous ways to support the execution of plans employing constructs not found in the domain procedures defined for previous PRS-CL applications.

Several characteristics distinguish Cypress from other systems that provide both planning and reactive execution. Many systems do not use general-purpose planning and so cannot generate plans of sufficient complexity for interesting applications. Previous work in run-time replanning has either been limited to synchronous approaches [2] or focuses

on local, adaptive modifications to rule sets, rather than employing the full look-ahead reasoning of a planner [3, 1]. The ability to modify a complex, parallel plan at run time and adapt execution activity to the new plan is, to our knowledge, a new accomplishment.

ACKNOWLEDGMENTS

This research was supported by the ARPA and Rome Laboratory Planning Initiative (ARPI) under Contract F30602-90-C-0086. The military application domain is an extended version of the one used for the second IFD in ARPI. The RCS application was supported by NASA Ames Research Center and was done by a team led by Michael Georgeff. SIPE-2, PRS-CL, Gister-CL, Grasper-CL, and Cypress are trademarks of SRI International.

References

- [1] R. J. Firby. An investigation into reactive planning in complex domains. In *Proceedings of the 1987 AAAI*, pages 202-206, AAAI, Menlo Park, CA, 1987.
- [2] J.E. Laird. Integrating planning and execution in Soar. In *Proceedings of the AAAI Spring Symposium*, 1990.
- [3] D. M. Lyons and A. J. Hendricks. A practical approach to integrating reaction and deliberation. In *First International Conference on AI Planning Systems*, pages 153-162, College Park, Maryland, 1992.
- [4] David E. Wilkins and Roberto V. Desimone. Applying an AI planner to military operations planning. In M. Fox and M. Zweben, editors, *Intelligent Scheduling*. Morgan Kaufmann, San Mateo, CA, 1994.
- [5] David E. Wilkins, Karen L. Myers, John D. Lowrance, and Leonard P. Wesley. Planning and reacting in uncertain and dynamic environments. *Journal of Experimental and Theoretical AI*, to appear.