

ONYX: An Architecture for Planning in Uncertain Environments

Curtis Langlotz, Lawrence Fagan,
Samson Tu, John Williams, and Branimir Sikic"

Medical Computer Science Group and *Department of Medicine
Knowledge Systems Laboratory (Incorporating the Heuristic Programming Project)
Stanford University

ABSTRACT

The ONYX program is designed to fill the need for planning in application areas where traditional planning methodology is difficult to apply. While the program being developed will assist with the planning of cancer therapy, its architecture is intended to be of use whenever goals are ill-specified, plan operators have uncertain effects, or trade-offs and unresolvable conflicts occur between goals. We describe a planning process which uses strategic information and a mechanistic model of the domain. The process consists of three steps: (1) generate a small set of plausible plans based on current data, (2) simulate those plans to predict their possible consequences, and (3) based on the results of those simulations, rank the plans according to how well each meets the goals for the situation.

I INTRODUCTION

Traditional planning programs attempt to satisfy a goal with a sequence of actions (for example, STRIPS (Fikes, 1971)). Each action is explicitly represented as an operator with preconditions and effects. The planner searches through a "state space" using operators to move between states. A successful plan is defined by a set of operators which generates a path of admissible states from the initial state to the goal state. The search through the state space is often aided by means-ends analysis. Frequently subplans must be built to satisfy the preconditions of the operators in a higher level plan.

More advanced planners have dealt with the problem of interacting subgoals. NOAH (Sacerdoti, 1977) uses a critic to repair incorrect or inconsistent plans. Stefik's MOLGEN (Stefik, 1981) uses constraints and a least commitment approach to handle subgoal interactions.

However, in medicine and many other application areas, the planning task cannot be represented in a form useful to a conventional planner. Often the goals are ill-specified and the operators have uncertain effects. Furthermore, incomplete and unresolvable interactions occur between parts of the goal, limiting the usefulness of least commitment and plan repair techniques. Consequently, medical therapy planning programs such as VM (Fagan, 1979), ONCOCIN (Shortliffe, 1981), and ATTENDING (Miller, 1983) have frequently relied on algorithms or step-by-step protocols to provide explicit guidelines in the construction of plans appropriate to a particular patient's condition.

Our work with ONCOCIN in the cancer therapy domain has revealed an important limitation of medical planning systems which use explicit criteria such as algorithms and protocols. The knowledge in these specifications is a "compiled" version of pathophysiological knowledge of the human body and of the strategic knowledge of the domain. In ONCOCIN,

plan elements are selected strictly according to the characteristics of the current treatment situation without considering the causal mechanisms of the domain or many of the strategies useful in prescribing therapy. Consequently, when a situation arises for which the algorithmic knowledge does not apply, the planning system often recognizes the problem but cannot propose alternative therapy. The ONYX system is designed to suggest expert quality therapy plans in such difficult cases.

II THE ONYX ARCHITECTURE

ONYX is faithful to the standard planning paradigm when possible, but it has been extended to provide for the uncertain and ill-specified nature of the problem and to accommodate the time-dependent nature of the planning problem.

The planning process used by ONYX consists of three steps:

1. *Plan generation.* Using current and past data about the patient, and exploiting the hierarchical nature of possible plan steps, generate a small set of "plausible plans" which are consistent with the patient's current state.
2. *Symbolic simulation.* Using knowledge about the structure and behavior of the human body, predict the future states of the patient after the execution of each plausible plan.
3. *Plan ranking.* Using knowledge about how patient data satisfy goals for the patient's progress, rank each of the plausible plans according to the extent that the simulation's predictions for each plan meet the therapy goals.

The following three sections will describe each of these phases in more detail.

A. Plan Generation

Cancer treatment strategies are often general statements which may be applied to a wide range of decisions in the plan generation process, from broad therapeutic choices (e.g., whether to give drug therapy or radiation therapy) to specific decisions about individual drug doses. One such strategy is: "If a problem is encountered with a treatment, try to eliminate the part of the treatment that might be causing the problem." In one context, this is interpreted as a suggestion to decrease or eliminate the previously administered drug that is the likely cause of toxicity. In another context, it may be used to help decide between continued drug therapy and alternative treatments. Currently, such a strategy must be represented in each context in which it applies, rather than as a single more general principle.

Strategies are represented as two kinds of production rules. Some strategies are *control rules* which guide the exploration of a therapy hierarchy, shown in Figure 1.

*Support for this work was provided by the National Library of Medicine under Grant LM/CA-04136, the National Science Foundation under Grant IST83-12148, and Ford Aerospace under Grant SP882573CJ. Computing facilities were provided by the SUMEX-AIM resource under NIH Grant RR-00785 and by an equipment grant from the Xerox Corporation.

Others are *generation rules* which propose plan steps associated with the leaf nodes of the therapy hierarchy. At present, we primarily use the general strategy "keep close to the standard plan" in guiding the exploration of the hierarchy. Other strategies can be added as appropriate.



Figure 1: The therapy hierarchy used in the plan generation phase of the therapy planning process.

The plan generation process begins at the root of the hierarchy. As each node is explored, the control rules associated with the node are examined to determine whether that node's descendants should be evaluated. For example, control rules associated with the *Drug Treatment* node determine if drug therapies could plausibly be given to the patient. We have previously used a similar process for plan analysis in an adaptation of the ONCOCIN program (Langlotz, 1983).

When a leaf node in the hierarchy is reached, and the control rules suggest that the node corresponds to a reasonable class of therapies, generation rules are used to propose plausible plan steps. The conditions of generation rules test for patterns in past and present patient data. Their actions propose plan steps based on interpretations of the data. An example generation rule is shown in Figure 2.

```

    If *drug 1s a drug in *chemotherapy,
    *X 1s one of the current problems
    *drug 1s not one of the causes of *X, but
    *drug can contribute to *X
    then reduce the dosage of *drug
  
```

Figure 2: A generation rule in ONYX. Variable names are preceded by asterisks.

Complete therapy plans are subsequently formed by taking consistent collections from the proposed plan actions. The consistency is enforced by checking that the conditions of the rules which proposed each plan action are not contradicted by other actions in a collection.

B. Symbolic Simulation

Simulations are useful in predicting the consequences of carrying out proposed plans. But since knowledge of causal mechanisms in the medical domain is both uncertain and incomplete, it is difficult to find invariant quantitative mathematical relationships between parameters in our models. Consequently, deterministic mathematical models like those used in AI/MM (Kunz, 1984) are not acceptable.

The work of Kuipers (Kuipers, 1984) suggests useful alternatives to a completely quantitative approach. *Qualitative* simulations provide a way of expressing qualitative values for the states and trends of a set of interrelated variables in a system. Heuristics are used to help envision possible future states for the system. But these techniques cannot resolve the conflicting trends which occur so frequently in oncology. Furthermore, they do not represent the uncertainty in the processes they attempt to describe.

The need to address the problems of conflicting trends and uncertain relationships prompted us to develop a simulation architecture for ONYX in which both the structure and behavior of the model are represented symbolically. Simulation models are organized hierarchically according to part-of relationships. The behavior of each model is determined by the behavior and interconnections of its parts and by three knowledge bases, described below, which specify the model's behavior in response to stimuli. The state of each model is represented by a group of state variables, and by the state of its parts. Each model has *ports* through which it communicates with other models using simple message passing.

Hierarchical models are built interactively on a Xerox 1108 LISP workstation. A representation of three such models as they appear on the 1108 screen is shown in Figure 3. The box on the left shows the *body model*. One of the parts of the body model, the *bone marrow model*, is shown at the center. The *marrow space model*, which is a part of the bone marrow model, is shown at the right. The boxes containing X's are the ports through which the models communicate with one another.

The behavior of each model is described by three rule bases containing production rules. The first rule base




Figure 3: Each of the three large rectangular boxes signifies a model. Solid lines represent connections between models.

dictates how a model will respond to the stimuli it receives from other models through its ports. The second rule base contains knowledge about how to make further conclusions about the model's state based on any recent changes. The third rule base dictates how the new state of the model will be transmitted to neighboring models using a simple message passing scheme which acts along connections between models.

At present, a simulation of the models shown in Figure 3 can predict the behavior of the bone marrow in response to the intravenous administration of a hypothetical anti-cancer drug. This initial stimulus is propagated throughout the models using rules in each about the flow of the blood throughout the body. In the bone marrow, the drug moves from the blood space to the marrow space, where it is transmitted from the extracellular fluid to each cell population. Rules in the model of each cell population are used to predict the drug's effect on the size, growth, and maturation rates of these cells. These rates are currently represented by simple numeric constants, which are modified in response to the drug. As the simulation proceeds, visual representations of the simulation events are shown on the 1108 display. For example, a graphical image of the drug is shown moving along the connections of the model to represent its movement in the body.

When the simulation is complete, the history of state variables can be plotted against time. The response curve shown in Figure 4 shows the simulation's prediction of the amount of cells in a particular marrow population, as plotted on the screen. Because these cells are key indicators of toxicity caused by drug therapy, information of this kind can ultimately be used to help estimate the extent to which a plan has met the treatment goals.




Figure 4: A plot generated by the simulation which shows the size of a population of bone marrow cells over 27 days.

C. Plan Ranking

While the plan ranking phase of ONYX remains substantially unimplemented, we are experimenting with decision analytic techniques which rank alternative plans. These techniques require knowledge of the alternatives, probabilities, and preferences inherent in the decision situation (Howard, 1984). The current ONYX architecture can be readily augmented to provide this information.

Treatment alternatives are already provided by the plan generation mechanism. We intend to augment the symbolic simulation so that it can provide the probabilities of possible outcomes, rather than simple deterministic predictions. Preferences for the patient will be explicitly represented as at least three general goals: (1) to improve the patient's prognosis, (2) to reduce the treatment risk for the patient, (3) to remain close to the protocol guidelines for treatment when possible. Knowledge about how significant patient outcomes affect each of these goals, together with the relative importance of these goals, will allow computation of the relative utility of alternative plans.

III CONCLUSION

Traditional techniques cannot solve planning problems in many application areas. ONYX's three-step planning

process has been developed as an alternative. The first step in the planning process constrains the generation of possible plans. The second step attempts to predict how each plan might affect the patient. The third step will rank how well the possible future consequences of plans satisfy the goals for the patient. These three components of the ONYX system are not yet sufficiently complete to cooperate in generating a ranked set of plans. Our next task is to combine them to form a coherent planning system. We then intend to integrate the strategic and mechanistic knowledge in ONYX with the protocol-based knowledge in ONCOCIN, thereby augmenting its decision-making and explanatory capabilities.

ACKNOWLEDGEMENTS

Special thanks are due to Mark Stefik, Danny Bobrow, and Sanjay Mittal for their help with the initial design of ONYX and with the LOOPS programming environment. Mark Musen, Glenn Rennels, Ted Shortliffe, and Cliff Wulfman provided helpful comments on earlier drafts of this paper. We are also grateful for the collaboration of project members Michael Kahn and Peter Tarczy-Hornoch, and for the helpful support of the members of the ONCOCIN project.

REFERENCES

- Fagan, L. M., Kunz, J. C., Feigenbaum, E. A., and Osborn, J. J. *Representation of dynamic clinical knowledge: Measurement Interpretation in the intensive care unit, in Proceedings of the 6th International Joint Conferences on Artificial Intelligence, Tokyo*, pages 260-262, 1979.
- Fikes, R. E., and Nilsson, N. J. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 1971, 2, 189-208.
- Howard, R. A. *The Foundations of Decision Analysis. In Readings on the Principles and Applications of Decision Analysis*, chapter 31. Strategic Decisions Group, Palo Alto, California, 1984.
- Kuipers, B. Commonsense Reasoning About Causality: Deriving Behavior from Structure. *Artificial Intelligence*, 1984, 24(1-3), 169-203.
- Kunz, J. *Use of Artificial Intelligence and Simple Mathematics to Analyze a Physiological Model*. PhD thesis, Stanford University, 1984. Rep. No. STAN-CS-84-1021.
- Langlotz C. P. and Shortliffe, E. H. Adapting a Consultation System to Critique User Plans. *International Journal of Man-Machine Studies*, 1983, 19(5), 479-496.
- Miller, P. L. ATTENDING: Critiquing a Physician's Management Plan. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1983, 5, 449-461.
- Sacerdoti, E. D. *A Structure for Plans and Behavior*. New York: Elsevier 1977.
- Shortliffe, E. H., Scott, A. C., Bischoff, M., Campbell A. B., van Melle, W., Jacobs, C. *ONCOCIN: An Expert System for Oncology Protocol Management, In Proceedings of the Seventh International Joint Conference on Artificial intelligence*, pages 876-881, 1981.
- Stefik, M. J. Planning With Constraints (MOLGEN: Part 1). *Artificial Intelligence*, 1981, 16(2), 111-139.