

Multi-Agent Planning for Advanced Traffic Management

L. Tandy Herren, Ph.D. and Pamela K. Fink, Ph.D.

Southwest Research Institute
P.O. Drawer 28510
San Antonio, TX 78228-0510
(210) 522-3762
lherren@swri.edu

Abstract

This paper describes two multi-agent systems developed at Southwest Research Institute to address issues in automated planning. The Autonomous Underwater Vehicle (AUV) Planning system uses multiple cooperative agents to generate a path for the AUV. The Medical Test Planning System (MTPS) consists of multiple agents that incrementally build a schedule for a series of medical tests constrained by the resources available to perform them. The approaches used by these systems have considerable applicability to the problems faced in the area of advanced traffic management.

1. The Role of Planning and Scheduling in IVHS

One of the primary issues in the Intelligent Vehicle Highway Systems (IVHS) program is traffic synchronization. This issue encompasses requirements for automated check-in to the Automated Highway System (AHS) as well as requirements for lateral and longitudinal control of traffic flow. It also is related to incident detection and management. These requirements define a real-time planning and scheduling problem. At check-in, the automated system establishes the operating envelop of the vehicle and its driver and the desired destination. This information can be used by an automated system for planning routes and scheduling platoons of vehicles. Incident detection could provide additional data to the system, placing constraints on available routes and reasonable platoon schedules.

The planning and scheduling problem associated with traffic synchronization in an intelligent vehicle highway system is extremely complex, consisting of multiple conflicting goals and constraints. The overall goal is to construct a plan that moves vehicles to their destination safely and efficiently. This implies that both highway safety and throughput will be increased when, in the past, they have essentially been opposing goals. Planning for IVHS entails the resolution of multiple constraints under circumstances of resource limitations. Southwest Research Institute (SwRI) has investigated similar planning problems and determined that one viable method of automating complex planning and scheduling problems is the use of multi-agent architectures. This paper presents two planning systems that utilize advanced software technology in support of the types of problems that will be encountered in automated traffic management.

2. Characteristics of Plan Construction

A plan is a series of steps corresponding to possible actions in a domain that transform an initial state into a desired state. The need for good planning arises in many situations, from determining the path for an autonomous vehicle, to deciding what sequence of tests to run to diagnose a patient's illness. In each case, the planning technique must determine the sequence of intermediate steps required to move from the initial state to the goal state. Thus, planning can be characterized as a search of all possible combinations of steps, with the resulting plan being one that provides a way of getting from the initial state to the goal state. In most

real-world planning problems, the number of possible combinations of steps is prohibitively large. Therefore, artificial intelligence (AI)-based search techniques are used to limit the number of combinations that must be considered to find a plan. The key to automating intelligent planning is to minimize the search for a solution in such a way as to be sure that a "good" (possibly "best") plan is found without wasting too much time exploring impossible or inappropriate alternatives.

The planning problem is complicated when resources to carry out the plan are limited. Resource limitations exist under most real-world problem solving situations. However, in some circumstances, the difficulty of overcoming the limitations is so great that it prohibits many courses of action. For example, consider the typical set of errands that must be accomplished on a Saturday morning. One of the most limiting resources is time. To minimize the time required to complete the errands, optimal or near-optimal route planning is required. Other time constraints might include the schedules of various businesses, e.g., some might close early or open late on a Saturday. Resource restrictions complicate the planning process. The problem differs from the easier case in which resources can be considered unlimited and the goal is to design a plan, any plan, that will fulfill the defined needs. When resources are limited, the question becomes can a plan be developed to meet the defined needs with the resources available.

During planning, two opposing forces must be carefully balanced to ensure both a good plan and an efficient method of finding it. The first is concerned with exploring as many options as is reasonable. This can help avoid having to, in some way, loosen the constraints on the problem and to undo and redo previous work towards finding a plan in a given situation. The second is concerned with constraining the problem as much as possible and as soon as possible to limit the amount of search required to find a solution. Thus, underconstraint can lead to extensive search, while overconstraint can result in expensive backtracking. Some

automated planning techniques lean toward underconstraint, such as special purpose subplanners, goal regression and reduction methods, and constraint satisfaction. Others, such as buggy planning and dependency-directed backtracking, lean toward overconstraint. The best approach is to combine the two extremes and to make them work together. This has been done by a number of researchers (Hayes-Roth et. al. [1979], Stefik [1981], Chang and Wee [1988], Fink, et. al. [1987], Herren, et. al. [1993]).

In general, the current approaches to automated planning can be divided into two categories: hierarchical planning and opportunistic planning. Hierarchical planning is characterized by a least commitment approach in which the general plan is fleshed out before precise details are filled in. Opportunistic methods are less rigidly structured, and plans are developed in separate pieces and assembled as opportunities arise. Multi-agent architectures for planning can be characterized generally as opportunistic. Each agent performs its component of the planning activity, then conflicts are resolved to assemble the final plan.

3. The Use of Multiple Agents in Planning and Scheduling

SwRI has investigated the use of multi-agent architectures in planning and scheduling in two software systems. The first system automated path planning for the Autonomous Underwater Vehicle (AUV) using multiple cooperative agents. Each agent had a particular goal to achieve that conflicted with the goals of the other agents. Thus, to find a solution, arbitration and compromise had to take place. The second system focused on scheduling under resource limited conditions. Planning was incorporated in this system primarily to allow for re-planning when an original plan could not be scheduled with the given resource constraints. Multiple independent agents implemented separate functions in schedule construction.

A number of factors must be considered when developing a multiple agent planning

system. These include the number and characteristics of the planning agents, how they communicate and cooperate, what types of conflicts can arise, and how conflicts are resolved. A multiple agent view of the planning problem involves one or more independent planning agents. These agents may or may not cooperate with each other during the planning process. Each agent has its own goals to be achieved and may represent the problem domain differently based on its function. The goals for each agent place constraints on the potential outcome of a plan. Such constraints can result in conflicts that must be resolved, but which may be known or unknown to each agent at the time it performs its function. Finally, what each agent knows about the others depends on the communication model available for requesting and sending information between agents. Each of these issues must be resolved in the design of a multiple agent planning system.

Issues in the number and characteristics of the planning agents involve how to break the planning process into units that can be implemented by a single problem solving approach. Generally, each agent has an independent goal, which leads to an independent view of the problem and possibly an independent problem solving method. The agents must cover all important functions of the overall planning problem and should not duplicate effort in generating the final plan. The problem solving method used by the agent should be selected based on the task it needs to perform in the planning process. Different problem solving methods often lead to different representations of the domain, complicating the communication process.

The communication model in a multi-agent system can be characterized by several factors. First, communication can be categorized as "direct" or "indirect," depending on whether the information can be sent via a communication mechanism or whether it must be acquired through observation or other inference mechanisms. Second, the quality of the communication can range from complete and error-free to potentially incomplete and/or

incorrect. The final attribute involves the difficulty of communication between agents, which varies depending on whether each agent represents the problem in the same manner, thus facilitating communication, or whether each represents the problem in a different manner, thus facilitating its own individual problem solving requirements.

Conflicts in a planning session can be the result of interaction between any number of constraints. The lowest level conflict is one between the goal of an agent and the requirements of physical law. For example, an agent must recognize that a single object can not occupy two distinct, disjoint locations in space at the same time. Other conflicts may arise between an agent's goal and the physical limitations of the problem. For example, an agent should be able to determine whether it is physically possible for an object to move from one location in space to another in a specified amount of time given the problem definition. Conflicts may arise between the goals of two agents. For example, one agent may have a goal requiring that the plan include passing through a particular location in space at a given time while another agent may have a goal requiring that the location be avoided at all times. Finally, a conflict could arise based on two different agents' preferences on how best to attain a particular goal. For example, one agent may prefer that an action occur within a given period of time, but another may prefer slower execution. Conflicts that arise during planning must be resolved to generate a final plan.

Conflicts can be resolved by an a priori assignment of organization roles and precedence orderings among agents and the conflicts that can arise between them. If some method exists for applying an ordering to the agents that directs which one must attempt to resolve a conflict by altering its own plan, then conflict resolution is hierarchical. However, if a high-level governing agent dictates which agent will alter its goals, then conflict resolution is performed through arbitration. The following sections describe two multi-agent systems that develop plans for two very different domains.

In each case, the domain largely dictated the appropriate number and characteristics of the agents, the communication model between them, what types of conflict arise, and the way in which these conflicts are resolved.

Path Planning for the AUV. The AUV's top-level goal is to covertly maintain an array of electronic sensors under the polar ice cap. It is driven by three potentially conflicting goals: to accomplish a specified mission, to survive, and to remain undetected. The AUV is a battery-powered submarine with diesel recharge capabilities. Its range between battery recharges can be extended by lowering the speed of the vehicle. Cruising under battery power creates little noise, though the noise increases with the speed traveled. Much more noise is generated while running the diesel engine during recharge, thereby dramatically increasing the chance of detection.

The planning model developed for the AUV problem consists of a single planner (corresponding to the AUV) with three cooperating agents (corresponding to the three goals: mission, survival, and covertness). Each

agent has its own set of goals to be achieved and methods for accomplishing them (See Figure 1). The specific task that the AUV Mission Agent must perform is to determine the best path to take for placing a set of sensors at certain desirable locations, possibly performing maintenance on sensors already in place, and returning to port. The issues encountered by the Mission Agent involve planning a path around known and unknown obstacles, essentially the traveling salesman problem (see Fink, et. al. [1986]).

The goal of the AUV Survival Agent is to ensure that the AUV has enough fuel to complete the task and return to home port safely. At any point in a mission, the Survival Agent will wish to maintain a certain amount of battery charge as reserve in case of unforeseen trouble and to maintain a certain amount of diesel fuel. Based on these restrictions, it must determine travel speed at all points in the mission and the location where battery recharges should take place.

The Covertness Agent is responsible for ensuring that the plan will allow the AUV to

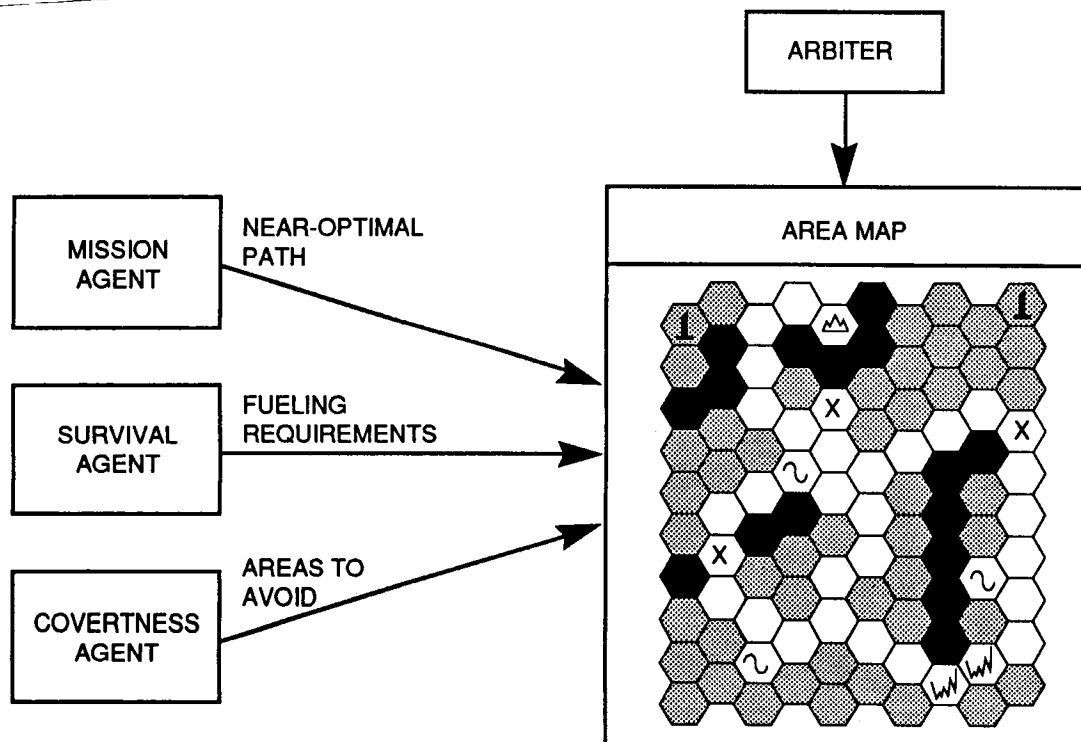


Figure 1. Overview of the AUV Planner

perform the mission and return safely to home port with a minimum risk of being detected or captured. It must analyze the plan to determine if a sufficient amount of distance exists between the AUV and any known risks along the mission path, taking into consideration not only the noise made by the AUV while it is moving, but also any noise produced while performing a specific task, such as recharging its battery or installing a sensor.

Each agent represents the problem domain differently in order to achieve efficient problem solving capabilities. The Mission Agent represents the world as a set of points that indicate free space, obstacles, and locations to be visited. The Survival Agent uses a set of mathematical equations that model fuel consumption related to the speed of the vehicle and frequency of battery recharging. The Coverttness agent models the signal-to-noise ratio of the AUV for performing the various operations in the Arctic Ocean environment and the passive and active detection capabilities of enemy submarines.

Communication takes place between the three agents via a common geographical representation. During planning, each agent annotates a stylized map representing the area of interest to indicate decisions based on its own internal problem-solving representation. However, conflicts between the goals of the agents do arise due to their independent requirements for a problem solution. For example, the Survival Agent may require a battery recharge at a given location while the Coverttness Agent may require a low-level of noise generation at the same location due to the potential proximity of an enemy submarine.

Based on the sources of conflict, a hierarchy of goals and constraints is generated. This hierarchy starts with the sources of the goals and their relative importance. During a planning session, each agent is assigned a level of importance for that particular mission, e.g., if the mission is not critical, then Coverttness and Survival may take precedence, so when conflicts arise, each constraint receives a priority ranking

based on the characteristics of the mission. Conflicts are resolved by arbitration. Each agent communicates its constraints and suggested plan to a higher level entity. Conflicts are resolved at this higher level through arbitration and selective relaxation of constraints. Initially, planning is based on the complete set of restrictions. In this way, if a plan exists that meets all the requirements of the agents then it is implemented. Otherwise, a less optimal plan is found by selectively relaxing constraints.

Scheduling Medical Tests. The Medical Test Planning System (MTPS) is a software system to plan and schedule medical tests for one or more patients. The user inputs disease hypotheses for one or more patients and specifies the availability of resources to perform medical tests. The system outputs a schedule of medical tests to test the disease hypotheses. The schedule depends on the medical tests that are selected and the resources that are available to perform the tests.

The MTPS consists of three major software modules or agents (See Figure 2). The first agent is the test selection facility (Herren, et. al. [1991]). When the user specifies one or more disease hypotheses, the test selection facility chooses a set of medical tests to cover the hypotheses. After the system selects the tests for a patient, the test planning facility reasons about how to order the selected tests for a given patient and whether or not the tests must be separated by wait periods. Finally, the test scheduling facility examines the test orders for all patients and arrives at a preliminary schedule based on resource and time constraints. Because the schedule generated by the test scheduling facility may be unable to conform to the specified constraints, an arbiter, the backtracking facility, evaluates problems in the schedule and determines which agent will most likely be able to modify the plan so that a solution is found.

Each agent uses a different problem solving strategy and representation. Test selection is a structured selection problem in which tests are selected based on their relative

weighted value. Selecting medical tests for a patient involves five types of knowledge: knowledge of the disease hypothesis, knowledge of the patient's history and presenting symptoms, knowledge of test characteristics, knowledge of the current context (e.g., is the situation an emergency), and domain knowledge about the relationship between tests and diseases. This knowledge influences the weighted combination of test attributes by influencing the value and the weight of the attribute. The weighted combination of attribute values forms the basis of test selection.

The test planning facility determines the order of medical tests for a patient by reasoning about the pre-conditions of a test (i.e., the state the patient must be in at the start of the test) and the post-conditions of the test (i.e., the state the patient is in when the test is completed).

By assuming that a patient may exist in one of a set of defined physiological states, the intra-patient ordering process can be conceptualized implicitly as a form of model-based or functional reasoning, involving inputs,

outputs, and state changes. Once the patient is represented as a set of possible states, the system can reason about the order of medical tests by evaluating the tests' pre- and post-conditions. Interactions between tests exist if the post-conditions of one conflict with the pre-conditions of another. These tests should not be performed sequentially unless there is a sufficient wait period between them.

The test scheduling facility schedules the tests across all patients in the system. To schedule the tests, this facility must reason about resource and temporal constraints. Fundamentally, the task is to assign resources, in the form of time slots, equipment and personnel, to a given test. When a resource is exceeded, some assignments must be undone and reassigned differently. To approach this problem, the principles of most-constraint and least-impact were considered in the design of this facility. The most-constraint policy selects a task dynamically according to the criticality, which measures how a task is constrained by task interaction. The goal at each decision point is to select the task that is currently the most

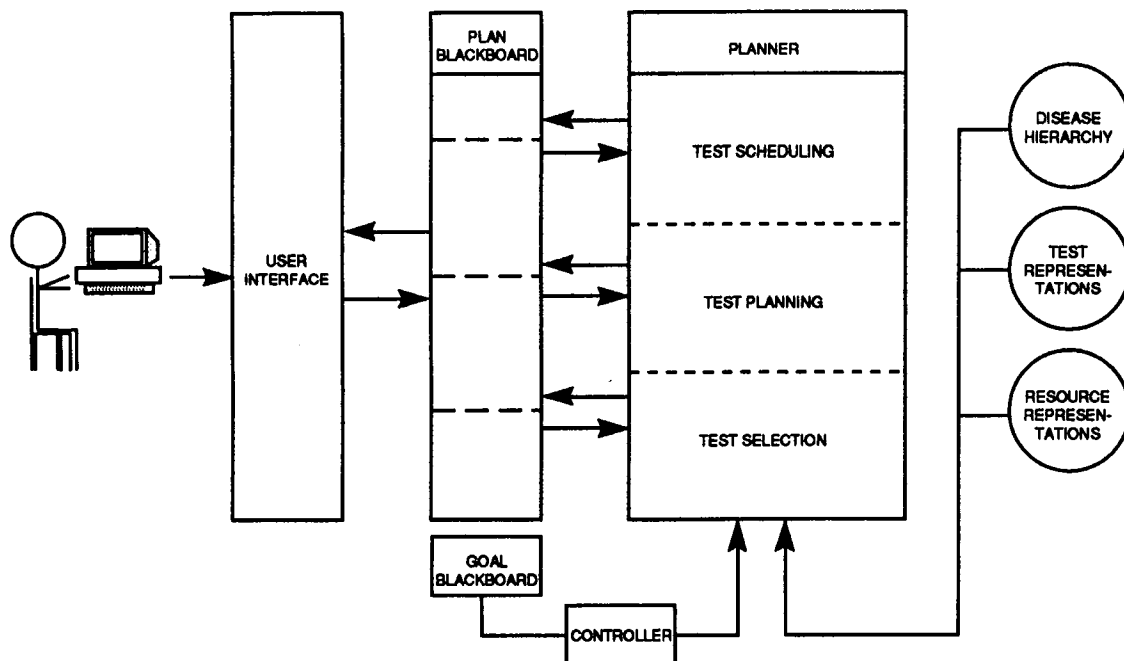


Figure 2. Overview of the Medical Test Planning System

constrained. The least-impact policy dynamically chooses a solution for the selected task according to the criticality of each possible solution, which evaluates the impact of a choice or change on the rest of the unachieved tasks. Thus, the goal is to select the approach to the task that generates the fewest additional constraints on the decisions left to be made. In addition to these two basic principles, a number of additional heuristics help assure that assignments are made that lower the probability of a conflict that requires undoing and redoing decisions.

The three major agents of the MTPS communicate via a global blackboard structure called the plan blackboard. Each agent posts its results on the blackboard and the other agents can evaluate and respond to these results. This shared data structure streamlines the communication between the separate agents. A global control structure selects the order in which the agents respond to information on the plan blackboard by examining a specialized goal blackboard. Thus, control is handled by posting data on the goal blackboard and evaluating it with rules for activating one of the major agents.

Conflicts arise when the planned tests can not be scheduled within the resource constraints. If resource contention arises, the backtracking facility determines how best to alter the schedule to resolve the conflicts identified by the test scheduling facility. If a resource is completely overloaded, i.e., overloaded in all time slots of the schedule, then one or more patients must relinquish their use of that resource. In that case, the test selection facility selects alternative tests that do not require that resource. The backtracking facility resets a patient's schedule and forces the test selection facility to select an alternative for the test in contention for the overloaded resource.

The backtracking facility has two alternative actions when a resource is overloaded in a single time slot. First, if the allotted duration of the schedule has not been exceeded, a patient's tests can be shifted to different time slots, preserving the original ordering of the

tests as well as the tests selected. After the shift, the tests are re-scheduled and evaluated by the test scheduling facility. If shifting the tests exceeds the allotted duration of the schedule, a patient's tests can be re-ordered by the test planning facility. This facility attempts to re-order a patient's tests minimally because changes in the order are likely to introduce additional wait periods. Therefore, this facility examines the set of tests for a patient that are in contention for scarce resources. Tests between the time period of the first contention and the time period of the last contention are cycled. That is, the last test in a contention period is moved in front of the first test and the test planning facility is invoked to evaluate the new ordering.

4. Comparison of the Systems

Both the AUV Planner and the MTPS employ multiple agents to develop a final plan. The multiple agents use separate representations based on the goal they attempted to fulfill. In the AUV Planner, the goals of the agents are potentially conflicting. Each of the three agents attempts to solve the whole problem from its perspective, leading to conflicts in the partial path plans that must be resolved. The goals of the agents in the MTPS do not conflict, instead they constitute building blocks of a final solution. Each agent attempts to solve just one portion of the problem that must be combined with the results of the other agents to reach a final solution. Each agent, however, can make certain choices that might lead to a resource conflict that must be resolved.

Although the agents use an individualized representation of the problem corresponding to the problem solving strategy that is appropriate for their function, both the AUV Planner and MTPS maintain a representation that is external to the agents and is used for communication between the agents. The agents in the AUV Planner communicate via a stylized map, and the agents in the MTPS communicate via a plan blackboard. These common external representations of the problem ensure that communication between the agents is direct and complete, thus simplifying the

communication process considerably.

The conflicts that can arise between agents differ between the systems. Conflicts in the AUV Planner can occur between the goals of the agents and between the agents' preferences. Alternatively, in the MTPS conflicts arise solely because the solution generated in concert between the agents violates the physical limitations of the problem, i.e., the resource constraints. This difference did not influence the method of conflict resolution. Both systems use a higher-level agent to arbitrate conflicts that occur. This approach provides more flexibility than assigning an a priori ordering between the agents to resolve conflict.

The approach used in the development of these two automated planning systems is generalizable to a wide variety of planning and scheduling problems. The AUV Planner is appropriate for problems that cannot be decomposed, such as path planning. The agents function to enforce separate goals for the final path (i.e., it must fulfill the mission, allow the AUV to survive, and keep the AUV from being detected). The approach implemented in the MTPS reflects the fact that scheduling is composed of three parts: deciding what to do, deciding how to do it, and assigning the resources to do it. Most planning and scheduling problems fall into one of these two categories.

5. Multiple Agents for Planning and Scheduling in the IVHS

A possible architecture for planning and scheduling for traffic synchronization could consist of numerous independent agents acting in concert to dynamically plan traffic flow. These agents may be of two types, one type being those that represent the interests of the overall highway system, such as safety and throughput, and a second type being those that represent the issues of individual vehicles, such as minimal distance to travel and minimal time required to travel to the destination. The goals of some of these agents will conflict. For example, getting vehicles on and off the highway safely may

result in a lowered throughput, and maximizing the highway's overall throughput may conflict with an individual vehicle's goal of minimizing distance and/or time. Information collected at check-in, such as the vehicle characteristics and destination, would guide factors such as lane selection and platoon formation. Thus, a model similar to that used in the AUV research would be appropriate for supporting the general problem of route planning, lane selection, and platoon formation.

The issue of limited resources in an intelligent highway vehicle system complicates the planning and scheduling task. For example, there will only be so many lanes and so many positions in those lanes on a given highway between a start location and a given destination. If more vehicles arrive than there are places for them at a given time and location, a resource conflict arises. Or, an accident could occur that temporarily removes some of the available resources. A planning and scheduling system for an IVHS will have to be capable of evaluating and appropriately responding to such situations, potentially modifying an existing plan already in progress based on the new resource constraints. Techniques studied and developed in MTPS are appropriate to addressing some of these issues.

References

- (1) K.H. Chang, and W.G. Wee, "A Planning Model with Problem Analysis and Operator Hierarchy", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Sept. 1988, Vol 10, No. 5, pp. 672-675.
- (2) P.K. Fink, J.C. Lusth, & H.W. Mullaney, "Automated Mission Planning in an Autonomous Underwater Vehicle," *Workshop on Future Directions in Artificial Intelligence*, Washington D.C., June 17-19, 1986.
- (3) P.K. Fink, H.W. Mullaney, and K.K. Palmer, "An Examination of the

Automated Planning Problem", Working Paper, 1987b.

- (4) B. Hayes-Roth, F. Hayes-Roth, S. Rosenschein, and S. Cammarata, "Modeling Planning as an Incremental, Opportunistic Process," *Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, Japan, August 1979*, pp. 375-383.
- (5) L.T. Herren, P.K. Fink, & B.L. Robey, "Medical Test Selection in a Knowledge-Based Scheduling System," *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Orlando, FL, October 31 - November 3, 1991*.
- (6) L.T. Herren, B.L. Robey, & P.K. Fink, "A Multi-Agent, Knowledge-Based System for Planning and Scheduling Medical Test," *Proceeding of the Sixth IEEE Symposium on Computer-Based Medical Systems, Ann Arbor, MI, June 13-16, 1993*.
- (7) M.J. Stefik, "Planning with Constraints (MOLGEN: Part I)," *Artificial Intelligence*, Vol. 16, No. 2, May 1981a, pp. 112-139.
- (8) M.J. Stefik, "Planning and Meta-Planning (MOLGEN: Part II)," *Artificial Intelligence*, Vol. 16, No. 2, May 1981b, pp. 141-169.