

A Hybrid Architecture for Real-Time Mixed-Initiative Planning and Control

Steven W. Mitchell

Senior Scientist, Intelligent Systems Laboratory, 250/060
Lockheed Martin Federal Systems, Inc., 9500 Godwin Drive, Manassas, VA 20110-4157
<steve.mitchell@lmco.com>

Abstract

In many mission critical applications current technology is inadequate for fully automatic planning and control. In these applications society insists that planning and control be exercised by human minds. However, many such applications lie at the brittle edge of human capabilities. This has led to serious incidents such as those involving the USS Stark and the USS Vincennes in the Persian Gulf. In domains like these where full automation is unacceptable and purely human operation is inadequate, a promising approach is one which combines the strengths of humans and computers.

This paper describes one architecture for addressing this challenge. Interacting with human domain experts in a mixed initiative mode it combines elements of case-based and model-based reasoning in a hierarchical task network decomposition planner to generate plans, and uses multivariate utility theory to evaluate the plans. The architecture includes real-time monitoring of plan execution, and automatic replanning for plan failure or significant changes in the environment. The planner has been implemented in C and C++, and used as the Tactical Response Planner for the DARPA Ship Systems Automation (SSA) program.

Introduction: Characteristics of the Naval Command and Control Domain

Contemporary naval vessels are controlled in combat by a team of humans, assisted by a complex of computers and communications devices which act as their interface to a variety of sensors, soft- and hard-kill weapon systems, and other ship systems. This team is typically divided into sub-teams, and control is exercised by a hierarchy of human team leaders. Thus the sonar sensors are interfaced to computers which process their data and control the sensors. These computers function as an interface for the sonar operators, who are in turn controlled by the Sonar Officer. He may report in turn to another officer, who fuses the

sonar reports with reports from other sensor teams to provide a coherent tactical scene to the Captain, the Tactical Action Officer (TAO) or the Officer of the Deck (OOD). This officer sits at the top of the hierarchy and is responsible for overall control of the vessel in combat. Depending on the details of the combat system, the individual operators and team leaders may or may not have direct computer assistance. Typically, coordination within and between teams is primarily vocal, particularly with sub-teams such as ship propulsion and helm control who may not be collocated with the other ship system operators.

The increasing tempo and complexity of combat operations in regional and littoral scenarios, combined with the proliferation of ever more sophisticated threats among potential adversaries, creates a serious challenge to naval combat systems operating in this manner. The humans at the various levels in the control hierarchy are increasingly overwhelmed by the volume and complexity of the data they must absorb, analyze, and act on. As engagement times shorten, responses to this flood of data must be swifter than ever. This is particularly true in littoral environments where the tactical picture is more complex, and there is greater likelihood of multi-axis and multi-warfare attacks popping up at close range out of the clutter. Defending against such attacks requires precise coordination between hardkill and softkill systems across multiple warfare areas, which places a premium on rapid analysis and planning.

Further complicating this problem are the increasingly austere budgets for acquisition, training, and operations.

Answering this challenge requires system designers to build more capable combat systems in smaller hulls that can be operated by smaller crews. To facilitate this, future combat systems need to automate many tasks currently performed by operators, while providing additional functionality to maintain operational superiority. In particular, these systems must provide automated support for the conversion of the multi-sensor fused data into situation awareness, the real-time construction of a plan which balances the mission against the imperatives of the tactical situation, and the execution of that plan while continuously monitoring the situation and adapting or replanning as required.

The mission-critical nature of military operations prevents the full automation of this process, even if the

Copyright © 1997, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

technology was sufficiently mature. When the actions of a naval vessel can put human lives at risk or even initiate a state of war between nations, society demands a human being on the scene to assess the situation and assume responsibility for the consequences of any actions. Thus the most that automation is allowed to provide is support to the human or humans exercising their judgment.

The general problems of the naval command and control domain, along with some of the technologies being developed to deal with them, are discussed further in (Brander and Bennet 1991), (Mitchell and Anderson 1996), and (Vainshtein and Sublette 1992).

Planning Architecture Requirements in the Naval Command and Control Domain

Several requirements are apparent from the above domain description. What is needed is a real-time tactical planning and execution system (TPES) which continuously monitors its environment both for external changes and for the effects of its actions. Uncertainty abounds, both in the tactical picture and in the effects of actions. Planning and execution must take into account the current and projected status of own ship systems and consumables. The TPES must interact with a team of human operators, generating plans both under their interactive guidance and autonomously in response to exogenous events. Plan execution must be under the supervision of the human operators, but not absolutely dependent on their reflexes.

The short response timelines and extremely high information load imposed on those humans by the complexity and uncertainty of the littoral environment requires that the initiative shifts back and forth between the TPES and the human operators. The system must provide the operators with both alerts and plan recommendations when the situation changes in such a way as to require their intervention. Both to support routine activities, and in extremely urgent situations where the Captain or the OOD has pre-authorized a class of actions against a particular platform, the combat system may automatically plan and initiate actions, offering the operator the opportunity to command by negation rather than requiring him or her to review and authorize actions explicitly before they are taken.

In addition to responding to the tactical environment, the TPES must support operator queries for recommendations for hypothetical situations. Given the real time nature of tactical operations, exploring options this way may be interrupted at any time by additional alerts, actions or queries. A practical TPES must therefore be flexible enough to deal with multiple threads of initiative in multiple worlds, both real and hypothetical.

In addition to these operational requirements, several characteristics of the Naval procurement and operations culture constrain TPES design options. Procurement realities dictate that the architecture must accommodate integration of existing domain models and tactical decision aids. Naval tactical doctrine is formalized, and tactical planning consists largely of adapting doctrine to the

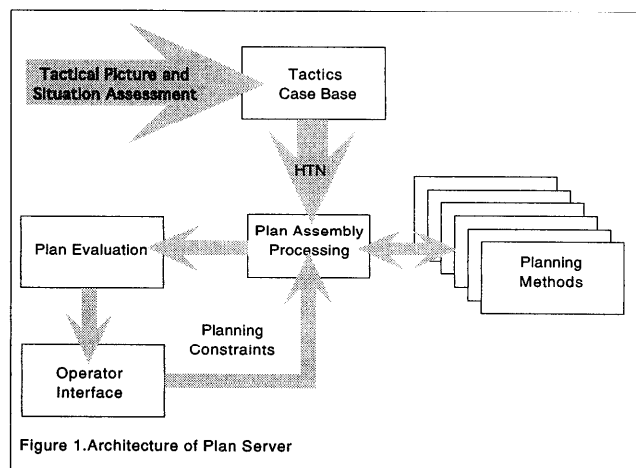
particulars of the immediate situation, constrained by rules of engagement imposed from higher echelons in the command structure and the expressed orders and guidelines of the commanding officer of the ship.

Architecture

Over the past seven years a series of prototype naval combat planners have been implemented. Starting with a relatively straight-forward rule-based submarine torpedo defense planner (Sublette and Vainshtein 1991), increasingly strict real-time constraints and complex multi-warfare test scenarios have forced these planners to evolve into the current mixed-initiative hybrid architecture for addressing the naval command and control problem.

The qualitative change from an automated planner to a mixed-initiative, multi-user planning and execution support system has been recent, and the implications of that phase transition are still being worked out. In what follows the process by which a plan is constructed will be detailed, followed by a description of how the plan server is used to support multiple planning and control users collaborating to manage a naval combatant during combat operations.

The Plan Server Architecture



The current plan server architecture is a blackboard system organized at the top level as a hierarchical task network (HTN) decomposition planner. The planner combines both reactive and deliberative planning elements. It uses case recognition and retrieval to select the tasks to be executed in the current tactical situation. These tasks are functionally equivalent to the goals of a classical planner, except they are subject to partial satisfaction by operators of uncertain effectiveness. The tasks are constrained by operator-specified Rules of Engagement and the Commanding Officer's Standing Orders/Orders of the Day. The tasks are organized into a partially-ordered HTN by the threat priorities assigned by the situation assessment team, and together with the associated tactical situation data (position, velocity and other characteristics of objects

being tracked by sensors, sea state, wind speed and direction, current course, speed, etc.) from the blackboard are passed to appropriate planning methods. The methods, which often compete with each other for ship resources, are invoked in the order specified by the HTN. This ordered invocation allows methods for higher priority tasks to preemptively tie up scarce resources, leaving lower priority methods to plan around the resulting resource constraints. The methods decompose the tasks into further method calls until the planner is dealing either with atomic actions or with pre-defined combinations of such actions, referred to as *tactics*. At this level domain models are called to compute the parameters of the actions. One such model might be used to calculate the minimum maneuver required to put the wind at a relative bearing of 20 degrees across the bow of the ship given a variety of constraints such as relative wind, locations of previously deployed chaff clouds, ship radar cross section as seen by various tracks, etc. Another model might be invoked to generate a schedule for a series of missile and gun engagements against a series of incoming anti-ship missiles which maximizes the probability of raid annihilation while the ship is engaged in a torpedo evasion maneuver. Generating that schedule requires calling other models to compute the launch time for an Evolved Sea Sparrow Missile to get the maximum probability of kill against an incoming anti-ship missile on a particular trajectory with the launch platform on a particular course, speed, and acceleration profile, given the resource tie-ups from previously scheduled actions.

Once all of the methods have run, the planner assembles all of the plan fragments generated by the various methods into a coherent plan for execution. Where conflicts are detected despite the prioritization process described above, the planner alerts the operator and asks for assistance in deconflicting the plan, typically by operator adjustment of the tasking constraints and priorities. Once an executable (i.e. conflict-free) plan is constructed, a multivariate utility function is used to evaluate the plan.

Typically multiple plans are generated to deal with different hypotheses about the tactical situation. These hypotheses may be alternate scene interpretations generated by the situation assessment team, or potential evolutions of the scene postulated by the operators of the planning system. The planner recommends the plan which maximizes utility against the currently-designated default situation assessment, while minimizing the potential cost of an erroneous choice of assessments, weighted by the likelihood of such errors.

If the plan does not involve pre-approved actions, the preferred plan and its evaluation is presented to the operator for approval. Otherwise, the pre-approved elements of the plan are placed into execution (with the opportunity for operator command by negation), while the plan as a whole is presented for consideration. In either case, once approval is received the plan is placed into execution, replacing the previously executing plan. If the

plan is not acceptable, the operator modifies the constraints, and the planning engine tries again.

During plan execution the system continuously monitors both the tactical environment, and the effects of its actions on that environment. This is accomplished by monitoring the tactical picture presented by on- and off-board sensors, as processed by the data fusion and situation assessment systems. In this domain there are many agents changing the environment, some independently and some in collaboration. Furthermore, the actions of each of these agents (including the ship hosting the tactical planner) are of uncertain effect for a variety of reasons: a radar may not switch mode when so commanded, an inbound air track may or may not respond to a warning, a torpedo may not be decoyed by a countermeasure. Engines break down, and humans sometimes fail to obey orders. Any of these things can cause a plan to break, and force replanning.

Aside from the uncertain environment and the chances of war, the presence of humans within the ship's combat system and chain of command can cause abrupt changes in the operational or planning context. Command authorities distant from the field of operations may change the Rules of Engagement, partially or completely invalidating the plan. Or the Tactical Action Officer or one of his warfare area coordinators may change their assessment of the intentions of the enemy, changing the threat assessments and disrupting the partial ordering of tasks upon which the plan was constructed.

Even when a plan is completely invalidated by one or more of these factors and a new, radically different plan is generated, the laws of physics can impose constraints which effect the execution of a new plan. Missiles in the air continue to fly until they either hit something or are deliberately destroyed. Maneuvers must be faired smoothly one into the next, and power plant configurations take time to change. The execution of one plan imposes constraints on the construction of the next, even if the actions initiated under the former plan are terminated as part of executing the next plan.

All of these constraints must be taken into account by the planner and its domain models which parameterize the actions during plan construction. These models must also take into account the current state of consumables (such as missiles, shells, and fuel), any system faults or breakdowns, and the anticipated state of the ship systems at the time for which the actions would be scheduled. The models used by the system to control execution of the plan actions must similarly take into account the switchology involved in transitioning the various systems from one mode to another, and monitor the systems as they transition through the various intermediate states.

Use of the Plan Server to Support Multiple Users

Given the ability of the plan server to construct a plan under all of the uncertainty and time varying constraints described above, one is still left with an organizational problem. In current naval command and control systems there is a team of operators who manage different parts of

the tactical situation under the control of the officer in charge. Even assuming the plan server can construct plans to satisfy each of these specialists, the problem of coordinating the various Warfare area coordinators and fusing their warfare area plans into a coherent whole for evaluation and modification by the TAO remains. The problem here is that while the goals at the top of the various warfare areas are superficially independent, as they are decomposed into lower level subgoals it turns out they are neither independent nor serial, due to shared shipboard assets which are required to achieve the various goals.

The approach taken is to use the plan server described above to mediate the multi-user planning process. Since the plan server can generate multiple plans in real time, each warfare area coordinator invokes the plan server with his or her own set of planning constraints to generate plans for their warfare area. These constraints include guidance from the warfare area coordinator and TAO-imposed constraints which limit the individual warfare area coordinator's use of ship resources. The warfare area coordinators individually interact with the plan server in the mixed-initiative manner described above. This interaction may include plan construction cycles both against the actual tactical scene, and against hypothetical extensions of the scene. When a coordinator is satisfied with their current warfare area plan and wishes to place it into execution, the system passes the associated planning constraint set to the TAO. He or she interacts with the computer to merge that set of planning constraints into the larger set used to generate an integrated plan for the whole ship. This merging process allows the TAO to balance the demands of the various warfare areas within the overall tactical context. The merged planning constraint set thus generated is the starting point for an interactive planning cycle involving the TAO and the plan server. If an acceptable integrated plan can be constructed, the TAO approves it and places it into execution. If not, he might levy additional constraints on some or all of the warfare area coordinators, and iterate the process at that level again.

This planning/executing/planning cycle continues throughout the mission. As discussed above, the cycle is subject to arbitrary interruptions as the tactical situation unfolds and as the initiative passes back and forth between the human operators and the automated planning and situation assessment systems. Plans, queries, explanations and justifications are constructed for time varying real and hypothetical tactical scenes, and plans are continually executed, monitored, and replanned in an intricate dance of human operators, computers, software agents, and objects and actors in the real world.

Experimental Results

As evaluated by US Navy domain experts during the US Navy's Advanced Ship Defense Combat System (ASDCS) Advanced Technology Demonstration program which ended in 1995, the planning engine used as a server in the

current mixed-initiative architecture functioned at the level of a competent Tactical Action Officer operating in real time through a graduated series of twelve test scenarios. This assessment was made on the same basis as similar assessments of human TAOs-in-training: the system was observed as it worked its way through the twelve scenarios operating in full automatic mode, and its performance was evaluated from the observed simulated ship behavior. Behavior was evaluated rather than outcomes, because the stochastic environment simulator used allowed the simulated ship to be hit sometimes even when the planner did everything right, and resource constraints did not accommodate sufficient simulation runs to allow statistically significant metrics to be computed on outcomes.

In the tests the planning system controlled an LPD-17 class surface vessel including helm control, three radar and one infrared surveillance sensors, an Identification-Friend-or-Foe radio interrogator, one missile targeting radar, two trainable missile launchers, three automated gun systems, six chaff launchers, an electronic-warfare suite, and a radio communications suite to send synthesized verbal messages. The Persian Gulf test scenarios ranged in complexity from managing the radar signature and sensors of the host platform against three hostile tracks while navigating the ship along a Path of Intended Maneuver (PIM), through fighting multiple dual-axis anti-ship missile engagements with both hard- and soft-kill weapons while conducting helicopter operations, managing ship sensors and signature, and navigating along the PIM. The test scenarios ranged from approximately five to fifteen minutes long. A final demonstration involved a thirty minute scenario with a total of approximately one hundred sensor tracks which combined amphibious operations with anti-ship missile engagements and anti-surface warfare. Under these conditions the planner was consistently able to generate a reasonable plan in less than 250 milliseconds. During these tests the planning engine was running on a Silicon Graphics *Indigo* workstation with a 150 MHz MIPS R-4400 processor. The planner and the associated display processes were implemented in approximately 70,000 lines of C and C++ code.

The performance of the planning engine was sufficient for the technology to be transitioned to the Demonstration and Validation phase of the Joint US/UK Surface Ship Torpedo Defense System program, and also to the Surface Ship Defense System program, which is developing a combat system for future non-AEGIS surface combatants such as the LPD-17 class of amphibious warfare ships.

The current prototype mixed-initiative planning system based on this planning engine performed in real time as part of the 1996 Advanced Combat System Demonstration of the DARPA Ship Systems Automation program. In this case, the planning system was running on a 180 MHz Silicon Graphic *Indigo II* workstation. This demonstration involved several intelligent software agents collaborating with each other and a varying number of human operators on a simulated DD-963 class surface combatant in a

challenging transition-to-war scenario set in the approaches to the Persian Gulf. Here the planning system controlled the ship's helm, three radar and one infrared surveillance sensors, an Identification-Friend-or-Foe radio interrogator, three passive sonars, two missile targeting radars, two trainable missile launchers, three automated gun systems, four chaff and acoustic countermeasure launchers, an on-board electronic-warfare suite, and a radio communications suite to send synthesized verbal messages. In addition the system was receiving inputs from two off-board electronic warfare receivers, and several off-board radar and sonar systems. The planning system interactively provided both "real-world" and hypothetical combat operations planning throughout the one hour scenario. In this configuration the plan server was able to generate an integrated multi-warfare plan in less than 500 milliseconds. This version of the planning engine and associated display processes was implemented with approximately 100,000 line of C and C++ code.

Related Work

Most automatic planning research has been done under the assumptions that actions are atomic and uninterruptable, the state of the world is completely known and effected only by the system executing the plans, and that the results of actions are deterministic (Hanks and Firby 1990). Clearly, these assumptions do not apply to the domain of naval combat and control. Recently, research in automatic planning has expanded to include domains where plan elements may consist of multiple actions which might be interrupted, both the state of the world and the results of actions are uncertain (Chrisman 1992), and where the planner under consideration is not the only active agent (Gmytrasiewicz and Durfee 1992).

Given the combinatorial nature of generative planning, considerable research has been focused on finding alternatives which are better behaved. One thread of research which is of interest is case-based planning (Blau, Bonissone and Ayub 1991). This approach attacks both the problems of knowledge acquisition and of bounding the planning process by structuring the plan knowledge base into situational similarity classes. The usage in the current planner of situational similarity classes to select the doctrinally correct tasks to be performed in a particular tactical situation is similar to the adaptation of case-based reasoning in (St. Amant and Cohen 1996).

Another important approach to bounding the computational requirements of planning is by hierarchical task network decomposition (Sacerdoti 1990) (Tate 1990) (Wilkins 1988). This decomposition can be done either top-down (Erol, Hendler and Nau 1994) as in the current planner, or bottom-up (Barrett and Weld 1994). In either case it provides a powerful method for organizing the generating of plans, or of controlling the adaptation of plan cases as in the current work.

The way the current planning engine supports multiple users cooperating on a single planning task resembles the

multi-agent planning work of (Ephrati and Rosenschein 1994). The use of constraint sets to communicate between the agents is similar to the work on distributed planning of (Tate 1996).

Decision theory is concerned with making rational decisions in an uncertain world (Raiffa 1968), and as uncertainty is one of the few constants in combat, decision theory is a potentially useful tool for planning actions in the tactical domain (Horvitz, Breese and Henrion 1988). This is particularly true since military operations occasionally require deliberately going in harm's way to accomplish a mission, and the goals involved in accomplishing that mission must be continually balanced against the goals of self-preservation and husbanding of resources for future needs. Both the priority of the various goals involved and the list of active goals (particularly mission goals) change with time. Similar balancing problems are common in many planning domains. While one can do such trade-offs using heuristic or purely priority-driven measures, neither approach accounts for the cost of information or the lack thereof, nor are they sufficiently principled to be broadly persuasive. For these reasons, decision theory is being increasingly used for planning both in the laboratory, and in real-world domains (Boddy and Dean 1989), (Wellman 1988).

Multivariate utility theory is an element of modern decision theory, and is used by the current planner for plan evaluation (Wellman and Doyle 1992).

Acknowledgments

The mixed initiative real-time automated planning and control technology described herein has been and continues to be developed under Lockheed Martin internal Independent Research and Development funding, as well as under contract funding. The support of USN Captain Robert L. Lowell, Jr. of the US Defense Advanced Research Projects Agency who supported early development and applications of this technology through Contract # DARPA/MDA 972-92-C-0068), is gratefully acknowledged. Additional thanks go to Mr. Michael Buckley and Mr. Lawrence Weeks of the US Naval Surface Warfare Center, Dahlgren Division, who supported further development and applications of this technology to surface multi-warfare under Contracts # N60921-93-C-A120 and N00178-96-C-2013, respectively. Notwithstanding the generous support of these gentlemen and organizations, the views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the United States Navy, or the US Government.

References

- Barrett, A., and Weld, D. S., 1994. Task-Decomposition via Plan Parsing. In Proceedings of the Twelfth National Conference on Artificial Intelligence, Vol. 2, 1117-1122, Seattle, WA, USA: AAAI.
- Blau, L., Bonissone, P.P., and Ayub, S. 1991. Planning with Dynamic Cases, In the Proceedings of the Case-Based Reasoning Workshop, 295 - 306, Washington, DC, May 1991: DARPA.
- Boddy, M. and Dean, T. 1989. Solving time-dependent planning problems. In Proceedings of the Eleventh International Joint Conference on Artificial intelligence.: International Joint Conferences on Artificial Intelligence, Inc.
- Brander, G. N. and Bennet, E. J. 1991. Real-Time Rule Based Resource Allocation in Naval Command and Control. Presented at the IEE Colloquium on Rule-Based Systems for Real-Time Planning and Control: IEE.
- Chrisman, L. 1992. Probabilistic Modeling of Action. In Proceedings of the 1st International Conference on Artificial Intelligence Planning Systems. 28-36, College Park, MD, USA: DARPA.
- Ephrati, E., and Rosenschein, J. S. 1994. Divide and Conquer in Multi-agent Planning. In Proceedings of the Twelfth National Conference on Artificial Intelligence, Vol. 1, 375-380, Seattle, WA, USA: AAAI.
- Erol, K., Hendler, J., and Nau, D. S. 1994. HTN Planning: Complexity and Expressivity. In Proceedings of the Twelfth National Conference on Artificial Intelligence, Vol. 2, 1123-1128, Seattle, WA, USA: AAAI.
- Gmytrasiewicz, P. J. and Durfee, E. H. 1992. Decision-theoretic recursive modeling and the coordinated attack problem. In Proceedings of the 1st International Conference on Artificial Intelligence Planning Systems, 88-95, College Park, MD, USA: DARPA.
- Hanks, S. and Firby, R. J. 1990. Issues and Architectures for Planning and Execution. In Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control, 59-70, San Diego, CA, USA: DARPA.
- Horvitz, E. , Breese, J., and Henrion, M. 1988. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, Vol. 2.
- Mitchell, S. W. and Corcoran, G. 1995. Developing advanced submarine combat systems utilizing real-time automated planning technologies. In Proceedings of the 1995 International Maritime Defense Exhibition & Conference, Vol. 3, 165-185, London, England.
- Mitchell, S. W. and Anderson, J. L. 1996. Submarine Combat Operations Planning and Execution: Integrating an advanced submarine combat system using real-time automated planning and control. In Proceedings of the Submarine Technology Symposium, Columbia, MD, USA: The Naval Submarine League.
- Raiffa, H. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading, MA: Addison-Wesley.
- Sacerdoti, E. D. 1990. The nonlinear Nature of Plans. In Allen, J., Hendler, J. and Tate, A. ed., 1990, *Readings in Planning*. Los Altos, CA, USA: Morgan Kaufman.
- St. Amant, R., and Cohen, P. R. 1996. A Planner for Exploratory Data Analysis. In Proceedings of the Third International Conference on Artificial Intelligence Planning Systems, 205-212, Edinburgh, Scotland: AAAI Press.
- Sublette, C. P. and Vainshtein, I. 1991. ARMOR: An Automated Decision Aid for Submarine Self-Defense Planning. In Proceedings of the 8th Annual Conference on Command and Control Decision Aids, Washington, DC, USA: Joint Services Decision Aids Working Group.
- Tate, A. 1990, Generating Project Networks. In Allen, J., Hendler, J. and Tate, A. ed., *Readings in Planning*. Los Altos, CA, USA: Morgan Kaufman.
- Tate, A. 1996. Representing Plans as a Set of Constraints -- the <I-N-OVA> Model. In Proceedings of the Third International Conference on Artificial Intelligence Planning Systems, 221-228, Edinburgh, Scotland: AAAI Press.
- Vainshtein, I. and Sublette, C. 1992. Automated Real-time and Contingency Planning in Submarine Combat Systems. In Proceedings of the Submarine Technology Symposium, Columbia, MD, USA: The Naval Submarine League.
- Wellman, M. P. 1988. Formulation of tradeoffs in planning under uncertainty. Technical Report MIT/LCS/TR-427, Aug. 1988: MIT.
- Wellman, M. P. and Doyle, J. 1992. Modular Utility Representation for Decision-Theoretic Planning. In Proceedings of the 1st International Conference on Artificial Intelligence Planning Systems. 236-242, College Park, MD, USA: DARPA.
- Wilkins, D. 1988. *Practical Planning: Extending the classical AI planning paradigm*. Los Altos, CA, USA: Morgan Kaufman.