

Metaknowledge Supporting Battlespace Planning, Execution Monitoring and Replanning

Gerald M. Powell, Ph.D.

U.S. Army Communications-Electronics Command (CECOM)

Software Engineering Center

Fort Monmouth, NJ 07703

732/532-3154

gerald.powell@mail1.monmouth.army.mil

Abstract

In its broadest sense, battlespace planning is arguably the central element of military command and control. A major objective of this paper is to promote the development of a clearer understanding of the computational issues associated with battlespace planning and thereby more appropriately focus research on human expertise with respect to these types of problems and clarify the kind of computational studies of planning needed to design and build planning systems of this kind. This paper describes how the set of conditions assumed to hold in the standard AI planning paradigm are violated in battlespace planning. A set of characteristics about the nature of battlespace planning are identified, and a set of assumptions about capabilities required to adequately address these characteristics is provided. Suggestions are made regarding the kinds of knowledge that are needed to achieve these capabilities. Examples employing nominal battlespace scenarios are provided to illustrate the characteristics and capabilities and to highlight the importance of these knowledge types and show how they can be used. The likely sources of these knowledge types and possible ways to collect or develop them as well as how the utility of this knowledge can be evaluated is briefly described.

Introduction

Battlespace planning is a complex process. The complexity stems from a number of sources. First, a substantial body of knowledge regarding friendly and enemy forces must be brought to bear. Second, a diverse set of information about a given operational context must be gathered, derived and organized. This information includes facts and estimates about enemy disposition, command relationships, capabilities and intent and similar, but more detailed, information about ones own force; facts and projections about terrain and weather; and implicit tasks, constraints and other information associated with a given mission. Much of the information about a given operational context is characteristically incomplete and uncertain. It is probably the case that completeness and certainty can never be achieved prior to the initiation of planning and execution of some actions that implement the plan. Even in a case where completeness and certainty hold at the time of planning, once the battle is engaged the plan is likely to require significant revision due to unpredicted events resulting from either friendly or enemy force actions.

In spite of the complexity of battlespace planning, there are human experts required to carry out the tasks to create, monitor and revise plans. An objective is to aid these planners and current-operations personnel by developing a computational framework that will support cooperative planning activities between human and machine experts. These activities are becoming increasingly time critical and our ability to collect information supporting these activities in near real time is continuously improving. These factors will significantly strain the capacity of human experts to successfully perform these tasks; a goal is to develop appropriate computational environments to support them.

In order to develop such environments we must achieve a computational understanding of this type of planning. Approaching this problem with a strategy like that used to develop expert systems has not panned out. Identifying the knowledge used by experts to carry out battlespace planning has been a much more difficult task than that involving the simpler kind of diagnostic problems attacked most successfully in the expert systems work. There are several reasons why battlespace planning is more difficult. First, battlespace planning requires multiple experts cooperating to develop a plan as opposed to a single expert. Few individuals have extensive experience in all of these tasks. Consequently, it would be unusual to find an individual with a deep understanding of the manner in which these tasks interact with one another. Second, human experts, in general, typically are not aware of how they represent their knowledge or how they use it to guide their reasoning. Third, battlespace planning occurs under the unique circumstances of the situation (particular opposing force, terrain, weather, etc.) unlike many diagnostic tasks solved by expert systems. Last, the set of human experts required to cooperatively carry out battlespace planning bring their own unique experiences and philosophies to the task.

Attempts to identify the basis for human expertise, especially in problem domains as complex as this one, depend on the identification of at least the general outlines of the kind of computational architecture in which the knowledge can be represented. The outline can serve as a model to guide questioning the expert as well as interpreting the data collected. Planning protocols collected from experts while solving a problem provide a highly informative body of data. Without such a model to guide analysis, it is extremely difficult at best to identify general and clear characteristics that might be present.

Abstracting research in artificial intelligence (AI) on planning as a basis for such a model of human battlespace planning presents some problems because the characteristics of problem classes solvable by AI planners and those of battlespace planning are quite different in many respects. The remainder of this paper will begin by describing the set of conditions assumed to hold in the standard AI planning paradigm and show how these are violated in battlespace planning. This is followed by identifying a set of characteristics about the nature of battlespace planning, describing a set of assumptions about capabilities that seem to be required to adequately address these characteristics, and suggesting kinds of knowledge that are needed to achieve these capabilities. Examples employing nominal battlespace scenarios are provided to illustrate the importance of these knowledge types and to show how they can be used. Next, we briefly discuss the

likely sources of these knowledge types and possible ways to collect or develop this knowledge, and how the utility of this knowledge can be evaluated. Last, we take a look at these knowledge types from the perspective of how they relate to the concept of flexibility in a military context.

A major objective of this paper is to promote the development of a clearer understanding of the computational issues associated with battlespace planning and thereby more appropriately focus research on human expertise with respect to these types of problems and clarify the kind of computational studies of planning needed to design and build planning systems of this kind.

Standard AI Planning Paradigm

The standard AI planning paradigm has a set of strong simplifying assumptions. They are [2]:

- (1) There is a single agent and this agent is the planner;
- (2) The planner is given a well-defined goal which remains fixed over the course of planning;
- (3) The planner is assumed to have functionally complete and accurate knowledge of the starting situation;
- (4) The planner is assumed to possess the knowledge required to accurately model the world;
- (5) The planner is assumed to possess the resources (time and memory) required to use this model to reason about the possible worlds associated with different courses of action that might be pursued.

Assumption 1 does not hold in the battlespace context because there are multiple agents required to carry out a plan. The actions need to be appropriately allocated to each agent and coordinated across them. Each agent needs to be provided the information and resources required to execute its own actions and to accomplish the required coordination. Also, there are other causal agents (enemy forces, etc.) not under the control of the planner. In a multi-agent (and adversarial) world, it becomes questionable whether assumptions 3-5 hold. Violation of any one of these assumptions significantly affects the generalizability of standard AI planning strategies to battlespace planning. In addition, even assumption 2 does not typically hold. For example, the planning staff must identify tasks and constraints that are only implied by the mission statement and must make them explicit, and must further specify and develop more complete sets of other information elements. The mission statement must be analyzed, refined, elaborated and interpreted along with the commanders guidance in order to provide an adequate basis to begin formal planning.

In the standard AI planning paradigm, a complete plan is created prior to action execution. There is no attempt to gather information about the world (including the starting state) in which the plan will be executed. In planning paradigms that assume a complete plan will be generated prior to the initiation of any actions, the completed plan

is the basis for determining what aspects of the plan need to be monitored during execution. Plan failure is signaled by action failure.

Battlespace Planning: A Non-standard Planning Paradigm

A planner with an ability to react to plan failure may use the existing situation (a failed plan) as the new initial situation and completely replan from this new starting state. A battlespace planner would require a more complicated plan monitoring ability that could perform situation assessment, detect when a plan has failed and in addition provide information regarding the reason for failure. By knowing the reason for failure, it may be possible to revise the original plan thereby avoiding a total replanning effort. In order to know the reason for failure, the plan monitor would need to be given the reasons for an action's inclusion in the plan as well as a structure reflecting the dependencies between plan elements and assumptions on which the plan is based. Without such a structure, the point of failure detection may not be traceable to the state in which the cause was present. If given such a dependency structure, the plan monitor would need to translate it into a schedule of monitoring tasks. In fact, in a sophisticated monitor, the viability of a plan may be critiqued from the standpoint of whether, or how and to what extent, different monitoring tasks can be supported. The results of this type of analysis could be used as a basis for determining whether the information provided by the monitor about the plan or its execution would be adequate to recognize the need to revise a plan as well as to support plan revision. The specific nature of the communication (report types, electronic communications links to be used, etc.) between the monitor and the planner would need to be determined.

As noted above, the monitor would track more than the plan's actions. It would also need to track conditions reflecting facts or assumptions made about the world in which the plan was developed and will execute. For example, it may be assumed that weather conditions will be free of precipitation during a certain interval of time, or that it will be possible to maintain electronic communication among certain force elements at certain times and locations, or that the opposing force will be present in a particular region at a particular time. Monitoring particular battlespace conditions in addition to actions creates a situation where there is a potential increase in the need for interaction between monitoring and planning. It would be the monitor's responsibility to know how to evaluate these conditions and to return the appropriate values (or reports) to the planner. One way to address non-determinism in the plan would be to identify branches and sequels and to develop contingency plans for the appropriate points in the plan where that non-determinism is expected to be resolved. If the execution of a contingency is signaled, the planner would need to be able to refine it and communicate those modifications to the plan monitor.

In the battlespace, typically the initiation of some actions will commence prior to the completion of a plan due to constraints on the time available for planning, or due to lack of knowledge on the planner's part to be able to complete the plan. A planner may have the ability to create a complete plan but lacks certain information about the starting situation or mission to be able to generate it. Alternatively, the planner may not be able

to develop a complete plan because the planner lacks the knowledge needed to accurately model the world.

Recovering from failed actions in the real world can be quite different from recovery in a simulated world because actions may be irrevocable and/or resources involved in their execution get expended. Battlespacing planning requires a strategy using knowledge that gives it a means of choosing actions that are likely to be useful in achieving goals, or, minimally, avoids choosing actions that prevent or diminish the likelihood of goal attainment. Discovering and/or developing these types of knowledge and the control structures required to effectively use them will require a significant amount of research. For now, only some preliminary answers can be offered. In general, the answer is that these kinds of planning strategies will require knowledge of the problem domain (at the object level) and their own planning processes as well as properties of the object level knowledge; these are often termed metaknowledge.

Problem Reduction

Knowledge of global properties of the state space are needed to support the model-driven style of planning required for battlespace planning problems. Problem reduction is a model-driven, backward-reasoning method that can utilize its predictions to flexibly control search. The problem reduction approach utilizes operators that change problem descriptions into a set of subproblem descriptions; these successor problem descriptions are said to be reduced. The problem space in problem reduction search is an AND/OR graph (or tree) of state pairs. In a planning context, a problem state is represented by a collection of problem domain objects and relations between them. Each state pair is referred to as a subproblem. The first element of the pair represents the starting state of the subproblem; the second element represents the goal state of the subproblem. The reduction operators are of two types, namely, non-terminal and terminal rules. Non-terminal rules reduce, or decompose, a problem state-pair (s_0, g_0) into a set of AND nodes where the set consists of problem state-pairs $\{(s_j, g_j), (s_k, g_k), \dots\}$. The subproblems in the set are assumed to be simpler to solve than the predecessor node (i.e., the problem itself). The decomposition is such that the solution of all of the subproblems in this ANDed set implies a solution to the (parent) problem. A solution involves the expansion of the AND/OR tree such that a set of subproblems are generated that are said to be primitive. A primitive subproblem is one that cannot be decomposed and is reached via the application of a terminal rule. A terminal rule represents either an action or decomposition with a known solution. If the primitive subproblem is solved, a value of true is the interpretation given to the terminal rule. If no such set of primitive subproblems are found, the initial problem is unsolvable under the representation chosen.

The problem reduction search space is never explicitly represented at the time search begins. Instead, it must be grown by using some control strategy. The control strategy requires knowledge to determine if a non-terminal rule can be used to reduce some node in the AND/OR tree. This knowledge required to guide expansion of the tree is a type of metaknowledge.

Metaknowledge

Knowledge about subproblem dependence is a type of metaknowledge. This property of plan elements is relevant to plan generation, execution and revision. A typical way in which subproblem dependencies are addressed is to search for an ordering of subproblems that removes or minimizes the effects of such dependencies. Another type of metaknowledge pertains to the number of alternative decompositions that are possible with given a reduction rule. When planning with incomplete information such as in battlefield planning, the number of decompositions expressed by a rule provides information relevant to the range of choices available for extending the plan tree from that rule. This applies to the entire AND/OR tree as well. A (sub)problem that can be decomposed in multiple ways (OR nodes), and the number of ways a primitive subproblem can be generated, are an indication of the different types of solutions available to solve the (sub)problem. A subproblem having many alternative decompositions can be considered to be less critical to solving the problem than subproblems having only one, or a very few, decompositions. If there exists only one way to decompose a subproblem that is present in all alternative plans, then all plans depend on the achievement not only of this particular subproblem but of this particular solution to it. This constitutes a critical subproblem for all solutions (plans).

Recall that battlespace planning is characterized by the need for planning to begin even though the knowledge required to develop a complete plan is lacking, and by the need to begin action execution prior to development of a complete plan. If a planner can be given knowledge of subproblem dependence or independence, and knowledge of subproblem or action criticality, it will have a basis for controlling its search in ways that tend to minimize the impact of incomplete knowledge on its ability to create and execute a successful plan.

One desirable characteristic of plans is to be able to make planning decisions as independently as possible. Subproblem independence is a property supporting this characteristic. It allows the choice of where to focus planning effort at any time to be governed by conditions such as the completeness of knowledge available to plan a subproblem or the requirement for one subproblem to be executed and accomplished prior to the execution of another. Also, if a subproblem fails, it does not necessarily indicate failure of the entire plan. The independence of the subproblem may simplify the problems that occur during plan revision. If a subproblem is critical to any potential solution, it allows planning activities to be focused upon this area of the plan so that failure might be detected early.

Subproblem dependence can be attributed to multiple factors. Before we look at those, let's consider what information types are likely to constitute a (sub)problem in a battlespace context. In the problem reduction paradigm, a (sub)problem is defined by a pair of states that represent the initial situation and the goal situation. In battlespace planning one might define the overall initial situation through the specification of: terrain information (maps and textual descriptions of terrain overview in the area of interest in the region); command and control measures (e.g., sectors and FEBA defined with operational overlays); information about opposing force (e.g., task organization, force

laydown overlay, and force intent given by the operational plan from higher headquarters); and information about own force (e.g., task organization for the larger force of which you are a part, and support priorities for own force).

The goal situation (or state) might be specified, for example, by the concept of the operation and guidance received from the echelon above. These typically will specify such factors as the type of defense to be employed; the terrain area for the main defensive effort; priorities for use of combat power against different assets; and constraints on certain resources during certain time intervals.

Note that we have found through the analysis of verbal protocols collecting during simulated planning exercises at the Army War College that the information characterizing the initial and goal states falls into three categories. Some of it is routinely given to planners, but other parts of it typically must be retrieved and/or derived in order to adequately formulate these states to be suitable for further decomposition [3].

Returning to the property of subproblem dependence, consider a scenario in which ones own force needs to perform a defense and counter-attack. Regarding the defense, decisions about the composition required for each unit; the relative combat power advantages of each unit against opposing force elements; the partitioning of the terrain into sub-areas; and the assignment of responsibility for different terrain areas to each unit could be made using a problem reduction approach. The need for a counter-attack may be identified as a constraint on resources available for the defense; more will be said about this when we discuss ordering of subproblems and execution of actions. Different solutions would be represented by decompositions that differ from each other in terms of their sets of subproblems. For this problem, the degree and nature of coordination required between multiple subproblems could be an indicator of subproblem dependence. The nature of the operational context may indicate that solutions that minimize coordination are to be sought. A high degree of coordination may require frequent acts of communication. Frequent communications would increase the risk of being detected and located. This line of reasoning would indicate that decompositions with schemes of maneuver maximizing independence of operations within a subproblem and minimizing coordination between subproblems should be selected over schemes that do not support these characteristics. An evaluation function to appropriately score or rank decompositions would need to be developed.

Consider another property of plans. Imagine a scenario in which the overall goal involves transporting supplies of different types (e.g., rations and ammunition) to a particular location. One decomposition may require that the actions for one subproblem, *i*, be executed and completed prior to the initiation of actions of other subproblems, *j* and *k*. A second decomposition to achieve the same overall goal may impose no order on the execution of actions for these three subproblems. Benefits can be derived from plans that impose little or no order on action execution. For example, action types where there is insufficient information to decide on a particular instantiation can be deferred while planning for the other actions can be completed and action execution can begin.

Returning to the scenario involving a defense followed by a counter-attack, knowledge about counter-attacks may dictate that, although they are executed sometime after a defensive operation begins, counter-attacks need to be planned first to ensure an appropriate force is held in reserve. Here is an example where there is a particular order imposed on the execution of these subproblems, but it is the reverse of that used to plan them.

Another property of plans that may be desirable is robustness. In addition to minimizing subproblem dependence, minimizing the incorporation of critical subproblems or critical actions in plans is a way of increasing their robustness. Imagine a problem in which half of the decompositions involve a subproblem that requires crossing a particular body of water and, in addition, that the only solution to this crossing in all of these subproblems is by seizing and controlling a particular bridge. The remaining half of the decompositions do not involve crossing this particular body of water, and, in addition, they have more than one solution for every subproblem that appears in the decompositions. First, it is important to note that robustness stems from being able to choose from multiple alternative solutions. If a plan fails based on one of these choices, a failed subproblem may be repairable by using another potential solution which appears in one of the alternative decompositions. Second, the ability to identify a critical subproblem or critical action alerts the planner to a potential need for allocating more time to that subproblem or action for planning purposes as well as to a potential need to develop contingency plans for world states that would be anticipated should that subproblem fail.

The key point here is that subproblem independence, order of executing actions, and critical subproblems may be properties of plans that should be used to guide the search for successful solutions (plans) especially in contexts where planning decisions must be made without having complete knowledge available to make them. Their merit should be judged through human expert evaluation of their use, through simulations, through scientific experimentation, and through various types of warfighting exercises. Examples of the kinds of knowledge needed to implement them have been described above, but a body of knowledge adequate to use these properties to guide search for successful plans in real-world planning contexts will need to be identified. Some of this knowledge may exist in doctrinal sources such as field manuals. Some of it may exist in the form of expertise residing with human planners and may require eliciting it through techniques such as verbal protocols. And some of it may need to be developed through a process of discovery using experimentation, simulation and warfighting exercises. Clearly, with changes in missions and concepts of operations, changes in technology, and changes in threats, this knowledge is likely to be dynamic.

Concluding Remarks

Plans having a high degree of subproblem independence, that impose little ordering on (sub)problem planning and on action execution, and that minimize the presence of critical subproblems potentially afford a significant amount of flexibility. The flexibility is: present in being able to appropriately focus the planning process itself; in the autonomy given to the execution of different parts of a plan; derived from identifying multiple alternative solutions to choose from in developing the initial plan, and in creating plans

that potentially can be simpler to repair or replan both prior to action execution and while plan execution is underway.

These properties support the principle of Maneuver which is one of the American Principles of War [5] and Dominant Maneuver which is one of the “four emerging operational concepts” of JV2010 [4] and network-centric warfare. A recent paper published by the Strategic Studies Institute of the U.S. Army War College is entitled *The Growing Imperative to Adopt “Flexibility” as an American Principle of War* [1]. This paper underscores the importance of flexible planning, or if planning fails, flexible adaptation. It indicates “While flexibility’s strengths are most easily conceptualized in the “post-contact” phase of conflict, i.e., as a *response*, flexible planning can be the vital force that shapes the outcome. A flexible planning disposition is one that is not necessarily satisfied with the first answer nor the popular one. It represents a critical search for blind spots and overlooked branches and sequels within the time available to plan.”

References

1. Frost, Robert. The Growing Imperative to Adopt “Flexibility” as an American Principle of War. Strategic Studies Institute, U.S. Army War College, October 15, 1999.
2. http://www-rci.rutgers.edu/~cfs/472_html/Planning/Intro_Planning_472.html
3. Powell, Gerald. A Cognitive Task Analysis of Army Corps-level Planning, Computer Generated Forces Conference, Institute for Simulation and Training, University of Central Florida, 1998
4. Shalikashvili, John. Joint Vision 2010, Washington, D.C., U.S. Government Printing Office, 1996.
5. U.S. Joint Staff, The American Principles of War. Doctrine for Joint Operations, Joint Publication 3-0, Washington, D.C.: U.S. Government Printing Office, February 1, 1995, pp. A-1 through A-3.