Use of Modeling and Simulation in Time-Critical Planning/Replanning

David R. Pratt, Ph.D.
SAIC
Orlando, Florida
prattda@saic.com

Robert W. Franceschini, Ph.D.
SAIC
Orlando, Florida
franceschinr@saic.com

ABSTRACT

The traditionally structured military planning process, and supporting toolsets, is largely based upon the traditional two-sided force-on-force actions that are paced by a timing cycle based upon the maneuver unit's echelon level. As we have seen in recent events, this paradigm has largely proven to be obsolete along several lines. First and foremost is the dynamic nature of current operations. Sensors, targets, and resources now present themselves in very narrow windows of availability. Second is the multidimensional nature of the applied force. Strategic forces can now extend down to an individual aircraft or squad, and the planning and replanning cycle is now measured in minutes versus days. Collectively, the increase in possible courses of action (COA) and the reduced planning/replanning time frames create a dramatic need for a tool to assist the mission planer. This is creating a new opportunity for modeling and simulation to become an integrated element of the overall operational command and control structure. This model is constantly updated using real-time information feeds to provide an ongoing assessment as to the viability of the current COA and possible alternatives. In this paper, we will address such a planning system; one that makes use of integrated modeling and simulation to gain insight into possible outcomes and to manage the possible futures. This paper will also contain the rational, structure, and implications of such a system.

ABOUT THE AUTHORS

DAVID R. PRATT (prattda@saic.com) is currently the chief scientist (Fellow) for SAIC's Analysis, Simulation, System Engineering, and Training Business unit. He also serves as the Forces Modeling and Simulation point of contact for DoD's High Performance Computing Modernization Program (HPCMP) User Productivity Enhancement and Technology Transfer Program (PET). A former captain in the United States Marine Corps and associate professor of computer science, Dr. Pratt was the technical director for the Department of Defense's Joint Simulation System (JSIMS) program. He received a Bachelor of Science in electrical engineering from Duke University and a Master of Science degree and a Ph.D. in computer science from the Naval Postgraduate School.

ROBERT W. FRANCESCHINI (franceschinr@saic.com) is a principal investigator with Science Applications International Corporation. He earned a Ph.D. in computer science from the University of Central Florida. He has directed research, leading to 50 scientific publications concentrated in the areas of computer-generated forces, distributed simulation, and multi-resolution simulation.

Use of Modeling and Simulation in Time-Critical Planning/Replanning

David R. Pratt, Ph.D.
SAIC
Orlando, Florida
prattda@saic.com

Robert W. Franceschini
SAIC
Orlando, Florida
franceschinr@saic.com

INTRODUCTION

The traditionally structured military planning process, and supporting toolsets, is largely based upon the traditional two-sided force-on-force actions that are paced by a timing cycle based upon the maneuver unit's echelon level. As we have seen in recent events, this paradigm has largely proven to be obsolete along several lines. First and foremost is the dynamic nature of current operations. Sensors, targets, and resources now present themselves in very narrow windows of availability. Second is the multidimensional nature of the applied force. Strategic forces can now extend down to an individual aircraft or squad and the planning and replanning cycle is now measured in minutes versus days. Collectively, the increase in possible courses of action (COA) and the reduced planning/replanning time frames create a dramatic need for a tool to assist the mission planer. This is creating a new opportunity for modeling and simulation to become an integrated element of the overall operational command and control structure. This model is constantly updated using the real-time information feeds to provide an ongoing assessment as to the viability of the current COA and possible alternatives. What is interesting to note is that this process is no different than what business leaders and program managers have been facing for years. In this paper, we will address such a planning system; one that makes use of integrated metrics and modeling and simulation to gain insight into possible outcomes and to manage the possible futures. This paper will also contain the rationale, structure, and implications of such a system.

DRIVING FACTORS IN C2 EVOLUTION

Perhaps the most significant change from the Cold War to our current operational environment is the change in the operational environment from a largely two-sided military operations-focused environment to the "three block war" paradigm (Krulak, 1999). While this transition has been widely reinforced by our experience in both Iraq and Afghanistan, many of our systems have not kept pace with the change in the environment.

We are now operating in a continuous operations mode where a unit always has ongoing tasking and engagements. The historical concept of "movement to contact" has been largely replaced by "in theater, in contact." This is a natural by-product of the shift from a purely military operation, where maneuver warfare was the operational paradigm, to that of stability and sustainment operations, where there is continuous involvement with friendlies, local populace, and hostiles. Since this environment is extremely dynamic, the cycle times for the commanders and planners are not measured in days or hours, but rather in minutes and seconds. Likewise, the range of things that must be considered include elements across the political, military, economic, social, infrastructure, and information (PMESII) spectrum. With all these systems interconnected, almost any action will trigger a set of ripple effects that might, and sometimes have, generated conditions much worse than the initial situation. Likewise, the advent of net-centric capabilities has increased both the amount and speed of data. This, in turn, has resulted in information overload, where vital information was lost in the sea of data.

It all comes down to two very important factors: the operational environment is more complex than any one person or staff can hope to comprehend in the allotted time, and the need to manage the future – not the current or past. To adequately deal with these two issues, a new approach and set of tools must be introduced. First, we will briefly address *the approach* and then the supporting tool set, with an emphasis on the predictive planning/replanning aspects as they augment mission success.

INTEGRATED PLANNING CYCLE ARCHITECTURE

Figure 1 shows the notional system architecture for a command and control (C2) system designed to support the time-critical planning/replanning process. At the most abstract level, the two key differences between this and the current systems are the integrated data/course of action (COA) repositories and the use of integrated modeling and simulation to aid in both the determination of when the current plan is no longer executable and in the selection of modifications. The following sections will discuss the various elements of the overall architecture touching on implementation details.

Inputs

The inputs in the system are shown on the far left side of the figure. In this case, we are using the term "input" as data elements that are originating from external to the system and/or are provided by humans directly into the system. The "Desired State" can be thought of as the commander's guidance or the preferred state of the environment. This state is compared against the world state and differences between the two are highlighted in an attempt to identify areas for monitoring/modifying the current plan.

"Intel" differs from "World Events" in that the data are provided by associated collection agencies and organizations and may, or may not, be biased, filtered and analyzed, but ideally, is of known quality. "World Events," on the other hand, come directly from the sources, such as the media, Web pages, third-party polls, etc., and is, by its very nature, biased and of questionable validity and completeness. As such, its validity, accuracy, and overall quality are suspect. That being said, since perception defines reality, it cannot be ignored and must be considered in the planning and executing process.

"Force Status" and "Operational Commitments" are critical elements for any commander in that they provide an indication as to the disposition and availability of the forces. In the era of "smart power," this takes on a whole different meaning and complexity than the data traditionally contained in the Joint Planning and Execution System (JOPES). It now includes such items as other federal agencies, coalition partners, non-governmental organizations (NGOs), and the media.

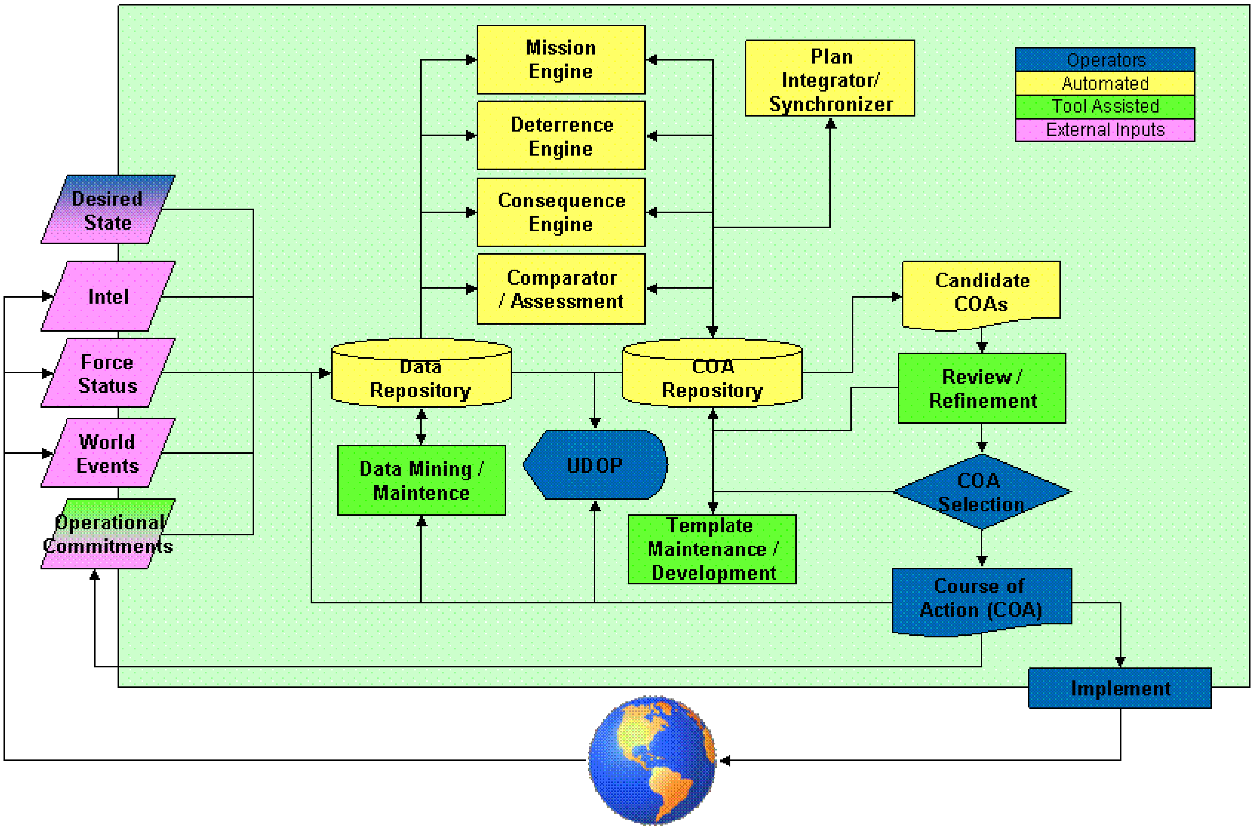

Figure 1. The Notional System Architecture

Data Repositories

In the center of the figure are the two logical data repositories. It is important to note that these are logical in the sense that they could be spread out over many physical databases. As long as the data elements exhibit location transparency, the actual implementation is an engineering exercise. We have shown the data in a separate repository from the COAs as a diagrammatic convention to highlight the difference between the COA (or plans) that make use of the data and the data itself. In reality, they are linked so that any changes in the data elements are reflected in the relevant plans. Also shown are the two data/plan template activities. These are both crucial activities to ensure the currency of the system. In the case of the data, the data-mining activity helps to discover the existing and emerging interrelationships between the data elements. The data maintenance activity provides a means of managing the pedigree and temporal validity of data. For example, a position taken with a Global Positioning System (GPS) device is going to be more accurate than that observed by a human. However, after several hours, both sightings are about equal as the possible location because movement dominates the initial positional error. Likewise, the confidence factor on a piece of data can go up as it is corroborated by other elements or go down if reports countering it come in to the system.

COA Development and Plan Issuance

The top center and far right of the figure contain the elements required for the generation, maintenance, and issuance of the plans. The far right is a modification to the traditional Military Decision Making Process in which the plan is developed, briefed to the decision maker, modified and issued. The top of the figure shows the three, integrated models that are used to do both forward projection of the current plan as well as ongoing excursions to provide options. While the models are shown as separate, in reality, they all interact. The "Mission Engine" is largely the traditional kinetic action model. The "Deterrence Engine" focuses on the PMESII/"smart power" elements. Finally, the "Consequence Engine" allows us to model the interconnections that exist in the environment to simulate the second, third, and fourth order effects that can prove to be quite detrimental in the current operational environment.

Finally, the "Plan Integrator/Synchronizer" performs two important functions. First, it takes the currently selected plan and compares it against the current world state, looking for areas where the operational reality has diverged from the plan. In doing so, it performs its second function of managing the futures graph (discussed below).

User-defined Operational Picture (UDOP)/Event Monitoring

Sandwiched between the two repositories is, arguably, the most important element of the entire system, the User-defined Operational Picture (UDOP). It is here where the warfighter monitors the world status, the unfolding of the operational picture, and then gains insights as to what elements need to be modified. The importance of the UDOP is that it allows the user to determine the element and structure of the data presented according to the user's preferences (Mulgund 2007). This aids in the cognitive process, and tasks focus by allowing the user to select the items that are most relevant to both them and the task at hand.

THE DATA TUPLE AND PROCESS

Like just about all modern computation systems, the system can be thought of as two main elements: the data and the actions performed on the data. While the architecture presented above provides the basic structure of the system, in this section, we will discuss the data and what processes are triggered as it is ingested and used.

The system will be receiving incremental and asynchronous data inputs during execution. The data will come into the system from a wide variety of sources, as shown in Figure 1, and in a wide range of formats. Since the purpose of the system is to reduce the cycle time needed to plan/replan, a degree of automation must be used. This can only be achieved if the data is in a structured format. For this reason, we have taken the leap of faith and assumed that unstructured-to-structured data converters will be available and reliable. The resulting data element is shown in Figure 2. The "source" field links the record to an external data base that contains the complete listing of the sources of the data. The use of indirection ameliorates some of the Intel classification issues when exposing sources of data. The "confidence," "aging," and "expiration" fields are all used to ensure the data is both current and valid. Since data is perishable, these fields, used in conjunction with the source, allow the decision makers to determine the underpinning of the plans and insights developed by the models and gain insight into the data's pedigree.

Record:
- DATE: Date / Time Stamp
- ENUM: Type (State, Status, Intel, Constraint, Goal)
- TEXT: Source (link to source database)
- FLOAT: Confidence (0..1)
- ENUM: Aging (None, Validity, Bounds)
- DATE: Expiration

Variable Record
  - Data content

Figure 2. The Structured Data Element

The variable part of the record contains data specific to the type of data. For example, unit status records would contain type, location, strength, etc., while records dealing with cultural modeling might contain the channel on which the message was disseminated, the speaker, the target audience, etc. Likewise, an infrastructure data record could contain the location, type, status, capacity, and estimated time to repair of a network node.

When data is received by the system, it triggers the process shown in Figure 3. We have already discussed the ingestion and formatting steps. In the Comparison step, the new data element is compared to the existing database. From this comparison, there are three possible alternatives shown below. A "conformation" verdict simply reinforces a current set of data, actions, and plans. As a result, this just triggers an update to the database. However, the other two options result in more processing.

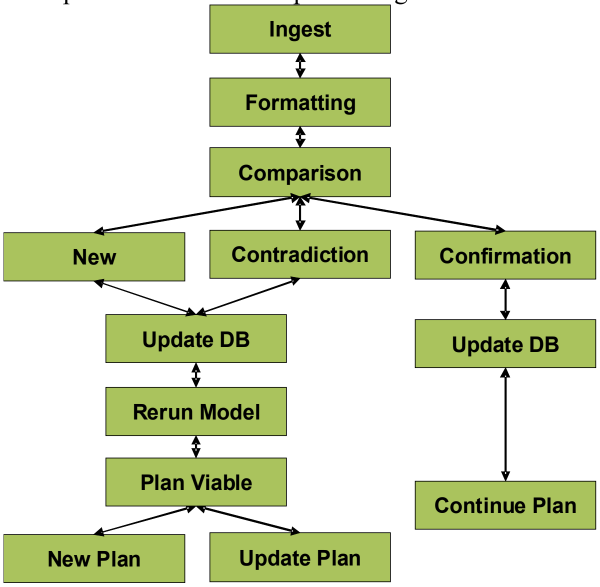

Figure 3. Actions taken upon receipt of data

The other two branches both require the database to be updated. Additionally, since they represent a shift in the knowledge base, the models need to be rerun to ensure their validity. Rerunning does not always require the complete futures graph (FG) to be recomputed. Rather, a comparison of the data elements against each of the segments' (a linked set of tasks and their outcomes) key data elements provides an indication if a given segment is dependent on the new/ revised data element. After the models have been rerun, a comparison of the current plan against the recomputed probabilities is performed to determine the continued viability of the current plan. Based upon this determination, a new plan could be selected or the current plan could be modified to reflect the change in the world state. The selection of a new plan, or even the modification of an existing plan, could be done automatically, but it is our contention that this is the point where a human-in-loop is required. As such, the system acts more in an advisory capacity than it does as an automated commander.

MODELING PARADIGMS

It is our expectation that no single model will ever be able to adequately represent the domain of operations we are likely to find ourselves in to any degree of accuracy. As such, we believe that an ensemble modeling approach is required (Henninger 2006). In current parlance, a "federation" is in many senses an ensemble, with the exception that an ensemble has multiple models modeling the same thing in different ways. While the control structure and flows are shown in Figure 4, the real power in ensemble modeling is the interplay between the models to provide both collaborative and competing insight in the actions. It is the function of the Plan Integrator/Synchronizer operating as the control layer to combine and control the three types of models into a coherent simulation environment. In our case, the models are linked via the Repository Layer as a shared data store. Logically, we have chosen the integration point across the models to by the physical "things." This flows logically from our ability to understand and deal with physical systems easier than we deal with and organize along abstract concepts.

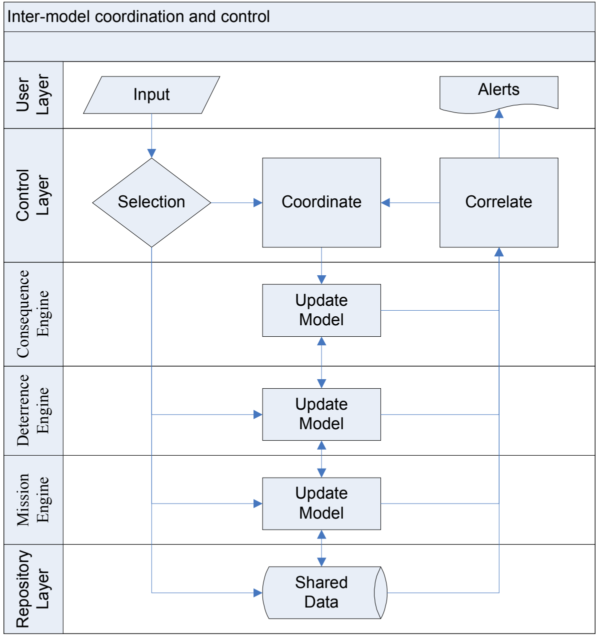

Figure 4. The use of ensemble modeling

The three gross characterizations of modeling environments or engines are based upon their type of modeling and their application to operational environments. While there are some models that cut across the lines, we have found that most models tend to focus in one of the areas. The first of these categories, and the simplest to understand, is the Mission Engine. Also known as "world models," these are the models of traditional military action that focus on the traditional kinetic actions and representations of physical things moving around and interacting. As such, they provide the context for most of the operations and visualization of things we do and see.

The Deterrence Engine is focused on the cultural and influence operations that operate more in the cognitive space and are often not visible in the traditional sense. As such, these models tend to represent the social, cultural, and economic elements that impact the behaviors that are manifested in the mission family of models. It is interesting to note that since these models are often based upon human cognition and social theory there is little or no traditional validation or experimental basis for these models. As such, they are often discounted to the determent of the operation.

Finally, the Consequence Engine operates in a much longer time frame, as it models the ripple effects of actions as they cause indirect follow-on effects. These models are often implemented as a series of interconnected network models. A classic case is a hydroelectric dam that is modeled as a node on both the hydrological network and the electrical network. In this case, disabling the dam's capability to generate electricity will cause the factories that are consumers to go dark and stop production of their items. Physically destroying the dam will not only cause that effect, but it will also case the water to rise downstream, potentially flooding the low laying areas.

PLAN REPRESENTATION

Perhaps the most critical aspect of doing predictive planning is the representation of the plan. This forms the basis of both the content of the plan and the processing model. A more detailed description of the plan representation can be found in (Pratt, et. al. 2008). For the sake of simplicity, we are going to assume a single-level plan representation. In reality, each of the plan nodes can be decomposed down to multiple subordinate plans. While this fractal nature is a key element of the C2 process, and a topic of follow-on work, we can discuss single-level plans without loss of generality. Likewise, other than synchronization points, we are not considering adjacent and supporting units plans.

At the lowest level, a plan is constructed of a series of task nodes, shown in Figure 5. They are similar in concept to the fluents described in (Cohen 2001). At the level we are planning, these represent atomic actions. The contents of the nodes can be divided up into roughly three categories. The static information contains the task name, activity, resources, preconditions, estimated duration, and stopping condition. These are planned during the initial planning process and could be updated during the replan process. Of key interest is the Stopping condition that signifies a planned transition out of the current task. This could be time-driven (Attack at 1000), geographic-based (When reaching a phase line), coordination-driven (Link up with 2nd squad), event-driven (When you take enemy fire), etc. Obviously, a time or distance transition will provide a better Estimated Duration value.

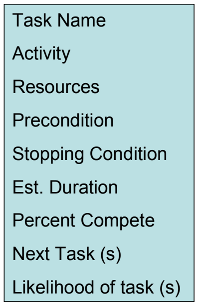

Figure 5: Task Node Representation

The dynamic portion of the plan is the Percent Complete slot. As the node is being executed, the value is updated to project progress to the planned stopping condition. This value provides a metric into the status of the plan and allows the commander and staff to monitor the execution progress.

The remaining part of the plan node is of variable length. Initially, as shown in Figure 6, there are follow-on single tasks to each task. This is consistent with the military planning process, where that plan is a linear sequence of tasks. However, as the mission unfolds, possible alternative futures are generated through plan perturbation (initially, this is largely a geographic consideration such as taking a left or right at a road junction) or interaction with the enemy (spotting, attack/defend, etc.). This results in a possibility of having more than one possible follow-on task. Thus, the link to the next task and the probability of taking that link are stored with the task node.

As shown in the figure above, there are three categories of task nodes. Previous nodes are nodes that have been executed and are preserved for historical and after-action review (AAR) purposes. The Current node is highlighted and is the one that the commander and staff believe we are currently in. This task is the one currently being executed and updated. This node serves as the root of the futures tree. The Planned future nodes, shown in blue are, as the name implies, the planned future tasks and have not been executed.

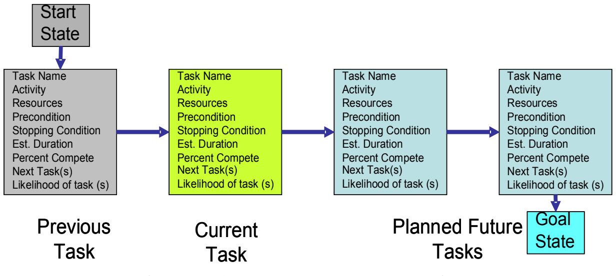

Figure 6: Plan Representation

The Start and Goal States are just tags to indicate where the plan begins and the expected outcome if the plan executes as anticipated. Since we will assume, for the remainder of the paper, that the monitor/modify/execute loop is being used during actual events, we can assume that we are in mid plan. This allows us to assume that the initial plan has been developed and the execution of it has started. Likewise, we have not reached the goal state, so successful termination is not an issue.

Futures Representation

As discussed above, the initial plan is a sequential organization of the task. One of the old adages is "no plan survives first contact with the enemy." Well, this is not quite true; many plans do not even make it that far. So experience has taught us that the future is much more like what is shown in Figure 7. The Futures Graph (FG) can best be thought of as a directed cyclic graph rooted at the current node.

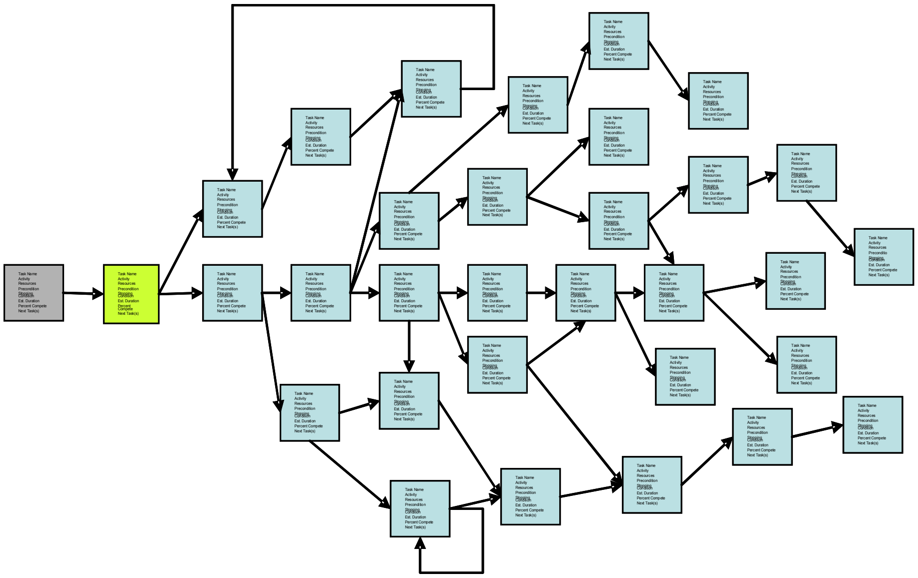

Figure 7: The Futures Graph (FG) showing the different paths

There are several key things to note about the FG. First, it is cyclic. There are situations where a node loops back on itself (bottom center of the picture). An example would be "Conduct patrolling operations until enemy contact is made." This is different from a holding task such as "establish a defensive posture," in the sense that the unit subordinates are actively doing things that are repetitive. Larger loops, like the one shown in the center, are when a series of tasks link back on themselves. An example of this would be to clear out a building and then have to come back and do it again after the adversary had reoccupied it. These loop constructs are quite common for asymmetric and peacekeeping operations, where it seems that some activities repeat.

Also shown in the figure is the dead end. As shown in the center top of the figure, these are one or more tasks that lead to a terminator that is not the goal state, and there are no more viable tasks that can be performed to reach the goal state. Obviously, these are options that, if known about in advance, should be avoided.

The final construct to note is the placement of the arrows. The majority of the tasks have the arrow leaving the right side and entering the left. This is to indicate normal commencement and completion of the task. As shown in a few of them, the arrows enter or leave the top or bottom. The leaving indicates a task that was interrupted prior to completion and another action commenced. An example of this is the enemy setting up an ambush along a trail. Likewise, when the task is entered through the top or bottom, it is joined in progress without preparation.

The two hardest issues to deal with are accurate predictions and timely execution. As the commander's plan progresses, variability and decisions occur that may permit alternative future paths. Accurate predictions of the variability and potential decisions are modeled. Each point of variability assumes particular conditions and can be assigned initial assessments based on analysis of the world state against those conditions. These assessments are used to determine if there needs to be a parallel split to the world state.

As part of the construction of the FG, we make use of Multi-Trajectory Control to help guide the generation of new trajectory and tasks. Multi-trajectory simulation futures heuristics are embedded in the branch-chooser capability. During simulation of a trajectory, a branch-chooser is invoked at the conclusion of a task or at the arrival of an event (subsumption) and will choose between several alternatives.

The first alternative is that branching may be "turned off" for that type of branch. In this case, the branch-chooser does not generate a node for the FG. It invokes a deterministic or stochastic response in the execution of the simulation, according to the features of the underlying simulation model.

The second alternative is that the branch-chooser may determine that a trajectory should be "pruned." In this case, the simulation will stop the execution of that trajectory and add a node to the Futures Graph, with information about the applicable state and an explanation of the termination criteria.

The third alternative that the branch-chooser would make would be to determine the level of fidelity of the multi-resolution model used to estimate future states emanating from the node in question.

From this discussion, it is easy to see that the definition of a task and the linking of the task are central to the execution and management of the FG and, in turn, the success of the mission. To that end, we have developed a set of rules for the construction and linking of tasks:

- Tasks should be atomic for the level of unit being planned.
- No decision that could result in different follow-on tasks should be made in task.
- If a branching point is established within a task, the task is broken into two tasks.
- Linear tasks are not combined into a compound task.

From these rules, it is easy to see that even for simple military operations, there will be a large number of tasks. We have made an overt trade-off in the number of tasks for simplicity of the task itself. This follows the agent paradigm: It is easier to manage a lot of simple, well-behaved things than it is to manage a few complex things with lots of internal exceptions and branching. Thus, as the world state evolves, new states are added and old, non-reachable states are pruned from the FG without having to break apart current states.

The Augmented FG

The FG provides a good representation as to what might happen, but is actually quite sterile in its meaning. For that reason, we have augmented it with some additional plan-level information. As shown in Figure 8, the coloring of the futures boxes allows us to rapidly see which of the states fall into the three desirability categories. A simplistic taxonomy of Desirable, Borderline, and Undesirable provides the commander and staff an easy way of annotating each task as to how they feel about executing each of them. In some cases, there is a natural deterioration of the tasks. In others, things go bad quickly.

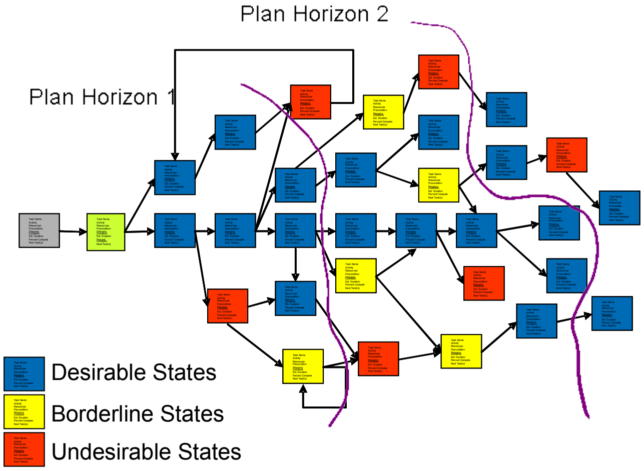

Figure 8: The Annotated FG

The final annotations are the two Plan Horizons. These are how far out we are planning with a degree of certainty. They can be expressed temporally (the next 72 hours) or based upon some task attainment (when we get to Phase Line Delta). Either way, the further out we look, the less certain we are as to the correctness and viability of the plan. So, while we should consider these tasks, the amount of resources we should spend evaluating them is reduced. In our design, we have implemented two Plan Horizons, since this matches with our evaluation process. Note that the Plan Horizon lines do not break up tasks; rather, they assume the completion of the task. As the system matures, we expect to be able to push the plan horizons to the left and provide increasingly more valuable and insightful plan evaluations to the commander and staff in a timely manner.

MANAGING THE COMPLEXITY OF THE FUTURE

Given the complexity of operations, it is only to be expected that the FG exhibits a tendency to become computationally unmanageable due to the branching factors associated with the number of options and possible outcomes. For this reason, was have made several simplifying assumptions that enable us to manage this complexity and produce results within manageable timelines. Key among these simplifications are the allocation of computational resources to most likely and most relevant actions and the use of parallel process, simulation cloning, and a consistent simulation platform to rapidly evaluate futures.

The allocation of the computational resources is driven by the selection of the type of modeling used to evaluate the particular a segment of the FG under consideration. As shown in Table 1, we propose using three different, but related, modeling paradigms that share a common underlying data store to ensure consistency across the levels. The most computational demanding, Entity modeling, the explicit modeling of the battlespace, is reserved for segments that are both close in and likely to happen. For things both further out and less likely, the use of aggregate modeling, modeling of the units/groups and their abstract actions, provides insights into what might happen. Finally, for the lease likely and most distance segments, qualitative modeling, implicit modeling of organization, actions, and environment, can be used to ensure that outcomes are not arbitrarily cut off from consideration.

Table 1. Allocation of Types of Models

	Prior to Plan Horizon 1	Prior to Plan Horizon 2	Post Plan Horizon 2
Most Probable Future (>50%)	Entity Modeling	Aggregate Modeling	Qualitative Modeling
Probable Future (50% > X > 25%)	Aggregate Modeling	Qualitative Modeling	Qualitative Modeling
Least Probable Future (<25%)	Qualitative Modeling	Qualitative Modeling	Qualitative Modeling

Entity Modeling, the modeling of discrete entities in the battlespace, is the traditional OneSAF® (Department of the Army) space. In this realm, the physical constructs (tanks, soldiers, airplanes, etc.) and their interactions are explicitly modeled. Intuitively, it is the most realistic and understandable of the simulation classes. It is also, by far, the most computationally complex.

Aggregate modeling takes a slightly different approach in that it models the logical entities in the battlespace. Most commonly, these are the units (squads, platoons, companies, flights, cultural groups, etc.) that are involved in the action. We have implemented aggregate models in the OneSAF proof of principle, making use of the entity composer and developing simplistic Lanchester equation-based attrition for the models. (Pratt 2006). In that experimental prototype, we made use of the OneSAF terrain without modification, allowing for a consistent representations and results in both the entity and aggregate spaces.

One of the best analogies for qualitative modeling is the Risk® (Hasbro, Inc.) board game (Wikipedia 2008). It is a fairly simplistic model that uses a very simple terrain representation and arbitrates the interactions via a comparison of the relative strengths and a roll of five dice. Using simplified pixel-based terrain for movement and intervisibility and table-based interactions of the effects, the resulting modeling, when executed over multiple times, proves a reasonable, low-fidelity approximation of the possible futures. For both the deterrence and consequences engines, a combination of simplified representation and gross level effects modeling for qualitative models complements the mission engine.

This schema allows us to allocate our computational resources in a manner where we have the greatest return on investment. What we are doing is putting the cycles where we have the most certainty of having the task take place. Since we have vague notions of what will happen the further out we look, it makes less sense to apply precision to analyzing these tasks. However, there has to be a third plan horizon that is the cutoff point for the simulation. The simple cutoff point is anything after the goal state or below a minimal probability of occurrence, which does not get simulated. As time elapses and the plan is executed, the probability of selected nodes will increase, and they can be expanded at that time.

Taking advantage of the independence of the segments allows us to make use of parallel processing to execute both multiple segments and multiple runs of the same segment concurrently. The multiple runs of the same segments are required to generate the probability distributions needed to determine the likelihood of the following actions. The segments are chosen for evaluation when they cross one of the Plan Horizon boundaries or when a critical data element is modified.

PARTING THOUGHTS

In this paper, we have presented the rationale, architecture, and approaches for integrating modeling and simulation as an integrated element in the planning and replanning cycle. The need for such a system is driven by the fact that, on a daily basis, decision makers are now facing challenges far beyond the capabilities and designs of their support systems. Given the complexity and visibility of modern day operations, the range of possible outcomes far exceeds what any person or staff reasonably can be expected to comprehend and manage effectively. The use of automated tools, such as simulation, is the only way to gain insight into potential futures without being inundated by the data and complexity of the potential futures.

ACKNOWLEDGEMENTS

The authors would like to thank Ted Hardebeck, Ph.D., Rob Alexander, Ph.D., and Bob Burch for their discussions and insights that led to the concepts and approaches described in this paper.

REFERENCES

Cohen, P. R. 2001 Fluent Learning: Elucidating the Structure of Episodes, Proceedings of the Fourth Symposium on Intelligent Data Analysis, Lecture Notes in Computer Science. Springer.

Henninger, A., Pratt, D., and Roske, V. (2006). “A Case for Using Ensembles in Analyses for Military Decision Making”. Proceedings of the North American Association for Computational Social and Organizational Sciences Conference 2006 (NAACSOS Conference 2006). South Bend, IN.

Hybinette, M and R. Fujimoto. 1997. Cloning: A Novel Method for Interactive Parallel Simulation. In Proceedings of the 1997 Winter Simulation Conference, ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Krulak, Charles C. (1999). "The Strategic Corporal: Leadership in the Three Block War". Retrieved on March 12, 2009. Marines Magazine, on Air University. http://www.au.af.mil/au/awc/awcgate/usmc/strategic_corporal.htm.

Mulgund, Sandeep and Seth Landsman, “User Defined Operational Pictures for Tailored Situation Awareness”, The MITRE Corporation, 2007, from http://www.mitre.org/work/tech_papers/tech_papers_07/07_0093/

PEO STRI, OneSAF Home Page, http://www.onesaf.net

Pratt, D.R., E. Towers, D.R. Shires, and K.T. Kick. 2006. l. Progressive, Multi-Resolution Course of Action Analysis. In Proceedings of the 2006 Interservice/Industry Training, Simulation and Education Conference.

Pratt, David R., et al, “A Multi Threaded and Resolution Approach to Simulated Futures Evaluation”, Winter Simulation Conference, Miami, Florida, December 2008.

Pratt, David R. and Robert W. Franceschini, “Use of Predictive Planning/Replanning to Augment Mission Success”, Proceedings of the Special Operations Forces Industry Conference (SOFIC) 2009, Tampa, Florida, June 2009.

Wikipedia entry on Risk (game), from http://en.wikipedia.org/wiki/Risk_(game) [accessed July 23, 2008]