

Toward Practical Knowledge-Based Tools for Battle Planning and Scheduling

**Alexander Kott
Larry Ground
Ray Budd**

BBN Technologies
11 Stanwix St., Ste 1410
Pittsburgh, PA 15222
{akott,lground,rbudd}@bbn.com

**Lakshmi
Rebbapragada**

U.S. ArmyCECOM/RDEC/C2D
Ft. Monmouth, NJ 07703
Lakshmi.Rebbapragada@mail1.
monmouth.army.mil

John Langston

Austin Information Systems
Whispering Woods Cove
Parkville, MO 64152
Langstoj@swbell.net

Abstract

Use of knowledge-based decision aids can help alleviate the challenges of planning complex military operations. We describe the CADET system, a knowledge-based tool capable of producing automatically (or with human guidance) Army battle plans with realistic degree of detail and complexity. In ongoing experiments, it compared favorably with human planners. Tight interleaving of planning, adversary estimates, scheduling, routing, attrition and consumption processes comprise the computational approach of this tool. Although originally developed for Army large-unit operations, the technology is generic. In this paper, we focus particularly on the engineering tradeoffs in the design of the tool, and on the experimental comparative evaluation of the tool's performance.

The Problem and the Motivation

Influential voices in the US Army community argue for significant computerization of the military planning process (Wass de Czege and Biever 2001): "...the Army must create fast new planning processes that establish a new division of labor between man and machine. ... Decision aids will quickly offer suggestions and test alternative courses of actions."

The reasons for exploring potential benefits of such decision aids are multifaceted. The process of planning an Army operation remains relatively cumbersome, inflexible and slow. The planning process frequently involves significant disagreements on estimation of outcomes, attrition, consumption of supplies, and enemy reactions. There is a fundamental complexity of synchronization and effective utilization of multiple heterogeneous assets performing numerous, inter-dependent, heterogeneous tasks.

For the last several years our team was working on one such decision aid, called the Course of Action Development and Evaluation Tool (CADET), a tool for producing automatically (or with human guidance) Army battle plans. Our primary focus was on a particularly time-consuming phase of the Military Decision Making Process (MDMP), called the Course of Action (COA) analysis (Department of the Army 1997). More specifically, we focused on the COA analysis performed for relatively large and complex units of the US Army, such as a Division or a Brigade.

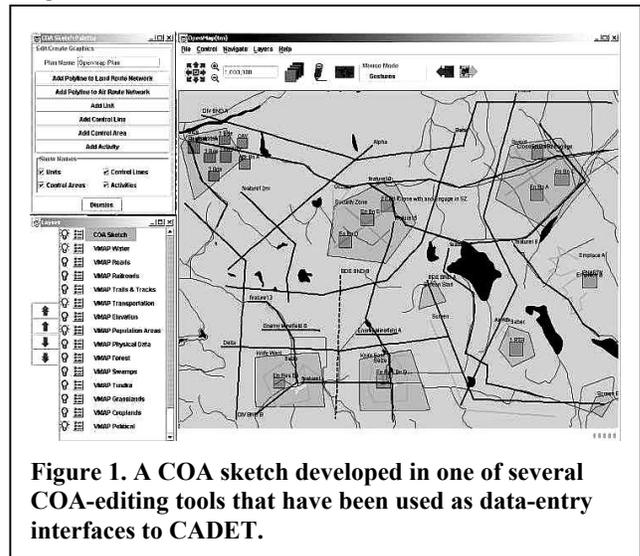


Figure 1. A COA sketch developed in one of several COA-editing tools that have been used as data-entry interfaces to CADET.

Done properly, in a setting such as an Army divisional or brigade planning cell, a detailed analysis of a tactical course of action involves a staff of 3-4 persons with in-depth knowledge of both friendly and enemy tactics.

The input for their effort comes usually from the unit Commander in the form of a sketch and a statement -- a

Views expressed in this paper are those of the authors and do not necessarily reflect those of the U. S. Army or any agency of the U.S. government.

high-level specification of the operation. In effect, such a sketch and statement comprise a set of high-level actions, goals, and sequencing, referring largely to movements and objectives of the friendly forces, e.g., “Task Force Arrow attacks along axis Bull to complete the destruction of the 2nd Red Battalion.”

With this input, working as a team for several hours (typically 2 to 8 hours), the members of the planning staff examine the elements of a friendly COA in minute detail. The process involves planning and scheduling of the detailed tasks required to accomplish the specified COA; allocation of tasks to the diverse forces comprising the Division or the Brigade; assignment of suitable locations and routes; estimates of friendly and enemy battle losses (attrition); predictions of enemy actions or reactions, etc.

The outcome of the process is usually recorded in a synchronization matrix format (Department of the Army 1997), a type of Gantt chart. Time periods constitute the columns. Functional classes of actions, such as the Battlefield Operating Systems (BOS), are the rows (see Fig. 3). Examples of BOS include Maneuver, Combat Service Support (e.g., logistics), Military Intelligence, etc. The content of this plan, recorded largely in the cells of the matrix, includes the tasks and actions of the multiple sub-units and assets of the friendly force; their objectives and manner of execution, expected timing, dependencies and synchronization; routes and locations; availability of supplies, combat losses, enemy situation and actions, etc.

How CADET is Used

It is in this complex, difficult and time-consuming COA analysis process that CADET assists military planners by rapidly translating an initial, high-level COA into a detailed battle plan, and wargaming the plan to determine if it is feasible.

In brief, the human planner defines the high-level COA via a user interface that enables him to enter the information comparable to the conventional COA sketch and statement (e.g., Fig. 1). As a collection of formal assertions and/or objects, including typically on the order of 2-20 high-level tasks, this definition of the COA is transferred to CADET, which proceeds to expand this high-level specification into a detailed plan/schedule of the operation.

Within this expansion process, CADET decomposes friendly tasks into more detailed actions; determines the necessary supporting relations, allocates / schedules tasks to friendly assets; takes into account dependencies between tasks and availability of assets; predicts enemy actions and reactions; devises friendly counter-actions; and estimates paths of movements, timing requirements, attrition and risk. The resulting detailed, scheduled and wargamed plan often consists of up to 500 detailed actions with a wealth of supporting detail.

Having completed this process (largely automatically, in about 20 seconds on a mid-level modern laptop PC), CADET displays the results to the user (e.g., Fig. 3) as a synchronization matrix and sometimes as animated

movements on the map-based interface. The user then reviews the results and may either change the original specification of the COA or directly edit the detailed plan.

Overview of CADET's Algorithm

Inputs: An initial set of activities (*Acts*). An initial battlefield state (*S*), including units, and geographic information.

Outputs: The initial set of activities (*Acts*) with addition of derived activities, timing, routes and allocated resources. The new battlefield state *S'* that includes effects of the added activities.

Procedure expand (*Acts*, *S*)

1. If an activity exists that is eligible for expansion then find the highest priority activity, using a set of scheduling heuristics.
2. Call the highest priority activity's expansion method to create supporting activities
 - Analyze the battlefield state using the domain specific knowledge within the task definition. Based on the analysis, generate the derived activities that are required for the successful completion of the expanding activity.
 - Calculate Resource candidates appropriate to each derived activity (including routes). Knowledge base rules specific to the activity provide the type and amount of resources needed.
 - Generate temporal constraints between the derived activities. Update *Acts*.
3. Propagate time constraints among the new activities.
4. If the newly expanded activity is ready for allocation, perform resource allocation and scheduling for this activity.
 - Use the knowledge base rules specific to the activity to compute the duration of the activity, depending on the resource and battlefield environments.
 - Use scheduling heuristics to assign resources and the time window to the activity.
 - Compute the effects of the activity, including battle losses, supply consumption, changes in geo-locations. Update *S* to reflect the effects.
5. Go to 1.

Once a satisfactory product is reached (typically within 5 to 30 minutes), the user utilizes it to present the analysis of the COA(s) to the Commander, and to produce operational orders.

Recently, there were several efforts to utilize the planning capability introduced by CADET. Battle Command Battle Lab-Leavenworth (BCBL-L) chose CADET as a key element for its Integrated COA Critiquing and Evaluation System (ICCES) program (Rasch, Kott, and Forbus, 2002), including the nuSketch system (Ferguson et al. 2000). DARPA used CADET for its Command Post of the Future (CPoF) program as a tool to provide a maneuver course of action. There, CADET was integrated with the FOX-GA system (Hayes and Schlabach 1998) to provide a more detailed planner to couple with FOX's COA generation capability. Battle Command Battle Lab-Huachuca (BCBL-H) integrated CADET with All Source Analysis System-Light (ASAS-L) to provide a planner for intelligence assets and to wargame enemy COAs against friendly COAs. The Agile Commander program of Army CECOM selected CADET-based Task Expansion Engine (TEE) as a technology for a key decision-support element within the larger framework of the program.

At this time, CADET is apparently the first and so far the only tool that was demonstrated to generate Army battle plans with realistic degree of detail and completeness, for multiple battle operating systems, and for the large scale and scope associated with such large, complex organizations as an Army Division or a Brigade. In the related domain of small-unit operations, (Tate et al. 2000) has described a very mature work.

Although originally developed for Army large-unit operations, the CADET technology is largely generic and can be applied to a broad range of tasks that require interleaving of planning, resource scheduling and spatial movements. Being a knowledge-based tool, CADET is adapted to a new application domain by changing its knowledge base. In particular, we have already built exploratory demonstrations for such tasks as intelligence collection using scouts and Unmanned Aerial Vehicles; tasks of Special Operations Forces; combat tasks of robotic forces such as the Army's forthcoming Future Combat System; and responses to terrorism incidence in an urban environment.

Key Requirements and Challenges of the Problem Domain

The needs of the problem domain clearly do not allow one to focus a useful decision aid on a narrow slice of the problem, e.g., only planning or only scheduling or only routing. Strong dependencies, for example, between scheduling of resources and hierarchical decomposition in planning as well as the route chosen for a task, indicate that a strong integration (perhaps the word unification might be even better) of all these processes is required (Wilkins and Desimone 1994, Kott and Saks 1998).

Further, as part of this tightly integrated process, the decision aid must perform elements of adversarial reasoning such as determination of enemy actions and reactions to friendly actions.

Also significant is the breadth of coverage in terms of the functional classes of tasks (BOS) that must be explored and planned by the decision aid. While maneuver tasks are central to the battle, other BOS, such as logistics or military intelligence are interdependent with the maneuver BOS and must be all analyzed in close integration.

In spite of the complexity implicit in these multiple interdependent problem aspects, speed is extremely important. It is most reasonable for a user in field conditions to expect an extremely fast response measured in seconds.

Because of rapidly changing elements of tactics, often evolving as operations unfold, and the differences in styles and procedures of different units and commanders, it is also imperative to provide a decision aid with the means to modify its Knowledge Base literally in field conditions, by end user, non-programmer.

Given that a decision aid of this type is most likely to be used in a framework of a larger deployed system, with its own style and implementation of the user interface, it is important to make the decision aid largely independent of the user interface assumptions.

The Technologies and the Engineering Tradeoffs

Perhaps the most fundamental engineering choice had to do with the basic functional focus and concept of user operation of the tool. We elect to focus CADET on the most time-consuming aspect of the MDMP, on its COA analysis phase. Other researchers (e.g., Kewley and Embrecht 2002, Atkin et al. 1999, Hayes and Schlabach 1998) are addressing a different (and preceding) phase of MDMP, the very interesting and challenging problem of generating the high-level maneuver COA. In addressing the style of interactions between the human and the decision aid, we prefer to de-emphasize the mixed-initiative, incremental style (even though CADET allows such a style) in favor of a rapid style of generating a complete plan from a high-level COA, followed by manual modifications.

The integration of planning and scheduling is achieved via an algorithm for tightly interleaved incremental planning and scheduling (see side box). The HTN-like planning step produces an incremental group of tasks by applying domain-specific "expansion" rules to those activities in the current state of the plan that require hierarchical decomposition. This process is controlled by a mechanism that leads the algorithm to focus on most significant and most constrained tasks first, and to limit the decomposition to a limited incremental set of tasks. The scheduling step performs temporal constraint propagation (both lateral and vertical within the hierarchy, a fairly complex and partially domain-knowledge driven process) and schedules the newly added activities to the available resources and time periods. This interleaving approach descends conceptually from (Kott, Agin and Fawcett 1992) where similar interleaving applied to a design domain.

Although we originally planned to use a version of Constrained Heuristic Search (e.g., Kott, Saks and Mercer 1999; Kott and Saks 1998) for the scheduling step, we were led eventually to prefer computationally inexpensive scheduling heuristics. These combine domain-independent estimate of the degree to which an activity is constrained, the "earliest-first" rule, and the domain-specific ranking of activity priorities. This choice was driven partly by the rigorous performance requirements, and partly by the fact that the simpler approach tended to produce results more understandable to the users. No-backtracking approach (with a few minor exceptions) was chosen largely for the same reasons. More generally, we feel that given the compound complexity imposed by the need for tight interleaving of multiple, diverse problem-solving processes in CADET, it is prudent to avoid any unnecessary complexity within each of these individual processes.

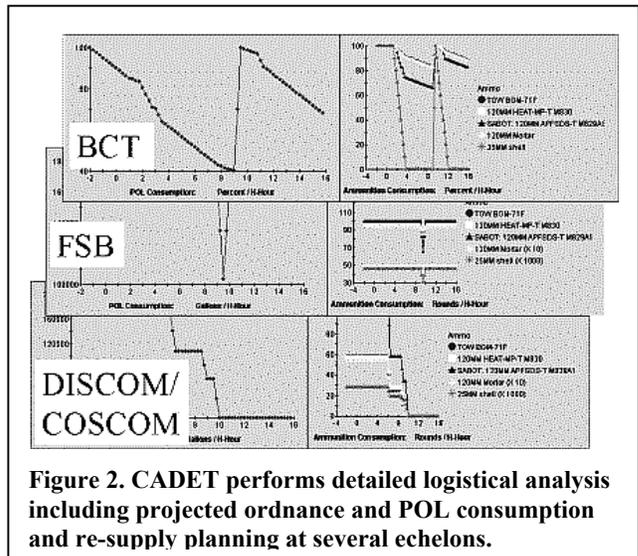
The same interleaving mechanism is also used to integrate incremental steps of routing, attrition and consumption estimate. A simplified, fast version of a Dijkstra routing mechanism is used to search for suitable routes over the terrain represented efficiently as a parameterized network of trafficable terrain. Optimization can be specified with respect to a number of factors, such as the overall speed of movement or cover and concealment, etc. For estimates of attrition, we developed a special version of the Dupuy algorithm (Dupuy 1990) that was calibrated with respect to estimates of military professionals, Army officers (Kott, Ground and Langston, 1999).

The adversarial aspects of planning-scheduling problems are addressed via the same incremental decomposition mechanism. CADET accounts for adversarial activity in several ways. First, it allows the commander and staff to specify the likely actions of the enemy. The automated planning then proceeds taking into account, in parallel, both the friendly and enemy actions. Further, the tool automatically infers (using its knowledge base and using the same expansion technique used for HTN planning) possible reactions and counteractions, and provides for resources and timing necessary to incorporate them into the overall plan. In effect, this follows the traditional MDMP's action/reaction/counter-action analysis (Department of the Army 1997).

In the object-oriented fashion, the knowledge base of CADET is a hierarchy of classes of Activities. A class of activities contains a number of procedures (methods) responsible for: computing conditions of applicability of a decomposition method; generating sub-activities of an activity depending on such factors as the available assets, the terrain or the location and type of the enemy forces; adding temporal constraints; estimating timing and resources required for the activity; finding suitable routes and locations; etc.

In practice, the most expensive (in terms of development and maintenance costs) part of the KB is the rules responsible for expansion (decomposition) of activities. We find a great variability in the procedures used to

evaluate preconditions of decomposition and the decomposition itself. Some of them, for example, refer to qualitative geographic relations between units of force and features of the battlefield, similar to the type described in (Ferguson et al, 2000). Others, however, are unique to each activity, require significant computations using general-purpose programming operators, and do not appear particularly amenable to generalization and formalization. This was one of the reasons we elected to use a general-purpose programming language, Java, rather than a specialized representational framework such as those of (Wilkins and Desimone 1994, Tate et al. 2000). Other reasons had to do with the programmatic necessity to use a broadly popular language for which experienced programmers are readily available on the labor market.



On the other hand, CADET includes a module for KB maintenance that allows a non-programmer to add new units of knowledge or over-write the old ones, in a very simple point-and-click fashion. Although necessarily limited by our decision to eliminate any direct programming features, the KB maintenance tool does allow an end-user to enter potentially a majority of the required activity classes.

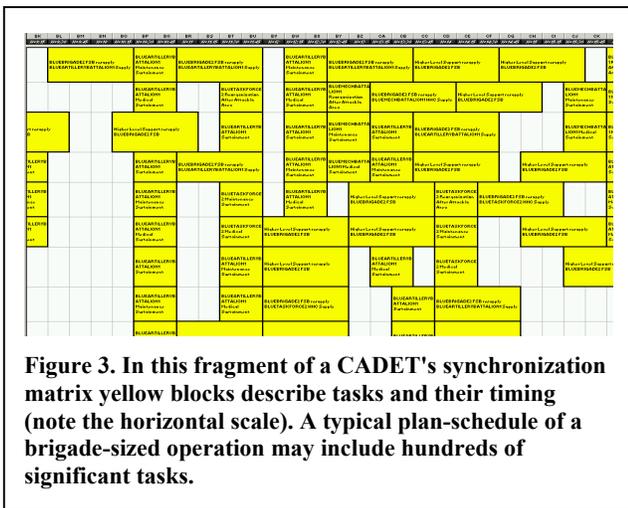
As a part of our developmental strategy, we elected to de-emphasize user interfaces and to develop no more than bare-bones, minimally necessary user interface features. These consisted mainly of an interface patterned after the synchronization matrix (Fig. 3), allowing the user to click on any cell (activity) and browse through the related network of domain-relevant objects (e.g., the units performing the activity, the location of the unit, etc.). This frugality with respect to user interfaces allowed us to devote a much greater fraction of the available funding to the primary functionality of CADET. Another reason is that we do not expect CADET to be deployed with a stand-alone user interface, but rather as a part of a larger framework with an existing interface.

Experimental Evaluation

Although several different experiments have explored applicability of CADET technology in the context of practical work of a brigade staff, e.g., (Rasch, Kott and Forbus 2002), here we discuss a particular experiment focused on evaluation of CADET-assisted planning process as compared to a conventional, manual one.

The experiment involved five different scenarios and nine judges (active duty officers of US military, mainly of colonel and lieutenant colonel ranks). The five scenarios were obtained from several exercises conducted by US Army, and were all brigade-sized and offensive, but still differed significantly in terrain, mix of friendly forces, nature of opposing forces, and scheme of maneuver. For each scenario/COA we were able to locate the COA sketches assigned to each planning staff, and the synchronization matrices produced by each planning staff. The participants, experienced observers of many planning exercises, estimated that these typically are performed by a team of 4-5 officers, over the period of 3-4 hours, amounting to a total of about 16 person-hours per planning product.

Using the same scenarios and COAs, we used the



CADET tool to generate a detailed plan and to express it in the form of a synchronization matrices. The matrices were then reviewed and edited by a surrogate user, a retired US Army officer. This reflected the fact that CADET is to be used in collaboration with a human decision-maker. The editing was rather light – in all cases it involved changing or deleting no more than 2-3% of entries on the matrix. The time to generate these products involved less than 2 minutes of CADET execution, and about 20 minutes of review and post-editing, for a total of about 0.4 person-hours per product. The resulting matrices were transferred to the Excel spreadsheet and "disguised," i.e., given the same visual style as that of human-generated sets.

The products of both the CADET system and of human staff were organized into a total of 20 packages and submitted to the nine judges, four packages to a judge.

Each package consisted of a sketch, statement, synchronization matrix and a questionnaire with grading instructions. The judges were not told whether any of the planning products were produced by the traditional manual process or with the use of any computerized aids. To avoid evaluation biases, assignments of packages to judges were fully randomized. Each judge was asked to review a package and grade the products contained in the package.

Not unexpectedly, data showed a significant scatter. While mean values for several experimental series ranged from 3.9-5.0, standard deviations ranged from 1.6-2.4. Judges comments also demonstrated significant differences of opinion regarding the same product.

Overall, however, the results demonstrated that CADET performed on par with the human staff - the difference between CADET's and human performance was statistically insignificant. Thus, based on the mean of grades, CADET lost in two of the five scenarios, won in two, and one was an exact draw. Taking the mean of grades for all five scenarios, CADET earned 4.2, and humans earned 4.4, with standard deviation of about 2.0, a very insignificant difference. Finally, comparing the "undisguised" series, we see that CADET earned the mean grade of 4.4 and humans earned 3.9, although the difference is still rather small.

The conclusion: CADET helps produce complex planning products dramatically (almost two orders of magnitude) faster yet without loss of quality, as compared to the conventional, manual process.

Strengths, Limitations and Future Directions

CADET shows a promise of reaching the state where a military decision-maker, a commander or a staff planner, uses it in field conditions, to perform planning of tactical operations, to issue operational plans and orders, and to monitor and modify the plans as the operation is executed and the situation evolves.

CADET generates Army battle plans with realistic degree of detail and completeness, for multiple battle operating systems, for the large scale and scope associated with such complex organizations as an Army Division or a Brigade, performing dramatically faster than a conventional human-only planning staff, with comparable quality of planning products.

Although originally developed for Army large-unit operations, the CADET technology is largely generic and can be applied to a broad range of domains that involve planning, resource scheduling and spatial movements.

CADET achieves its capabilities via a combination of approaches:

- Adopting a simple and transparent concept of user operation, which assumes literally no training;
- Using tightly interleaved incremental planning and scheduling, also integrated with route and attrition and consumption calculations;
- Adhering to computationally inexpensive algorithms that often trade optimality for speed and thereby assure an almost instantaneous response to the user;

- Integrating adversarial considerations into the solution process via the action-reaction-counteraction paradigm;

CADET's current state of capabilities also points toward several key gaps that must be overcome to realize the full potential of such tools:

Military decision-making commonly requires collaboration of multiple officers with distinct functions, responsibilities and expertise. In the near-future warfare, these officers will often collaborate while dispersed over the battlefield, communicating over the tactical Internet, possibly in asynchronous mode. Tools like CADET must support such forms of collaboration.

Presentation of CADET's products requires qualitatively different user interfaces and visualization mechanisms. Our experiments suggest that the users had difficulties comprehending a synchronization matrix generated by the computer tool, even though it was presented in a very conventional, familiar manner.

It is often said "no tactical plan survives first contact with the enemy intact." Combat planners must be able to plan rapidly, communicate orders to subordinates clearly and react without delay to changes in the situation. Planning tools like CADET should give the commander the ability to accelerate this cycle of recognition, re-planning and reaction, i.e., the capability of continuous re-planning during execution.

Ongoing work on CADET technology focuses on closing these gaps.

Acknowledgements

The work described in this paper was supported by funding from US Army CECOM (DAAB 07-96-C-D603 - CADET SBIR Phase I, DAAB 07-97-C-D313 - CADET SBIR Phase II and DAAB 07-99-C-K510 - CADET Enhancements); DARPA (DAAB07-99-C-K508 Command Post of the Future); and TRADOC BCBL-H (GS-35-0559J/DABT63-00-F-1247). TRADOC BCBL-L provided additional funding under DAAB 07-99-C-K510. MAJ R. Rasch of the US Army BCBL-L provided important advice.

References

Atkin, M. S.; Westbrook, D. L.; and Cohen, P. R. 1999. Capture the Flag: Military simulation meets computer games. In Proceedings of AAAI Spring Symposium Series on AI and Computer Games, 1-5.

Bohman, W. E. 1999. STAFFSIM, An Interactive Simulation for Rapid, Real Time Course of Action Analysis by U.S. Army Brigade Staffs. Thesis, Naval Postgraduate School, Monterey, CA.

Department of the Army 1997. *Field Manual 101-5, Staff Organization and Operations*, Washington, D.C.

Dupuy, Trevor N. 1990. *Attrition: Forecasting Battle Casualties and Equipment Losses in Modern War*, Hero Books, Fairfax, Va.

Ferguson, R.W., Rasch, R.A., Turmel, W., and Forbus, K.D. June, 2000. Qualitative Spatial Interpretation of Course-of-Action Diagrams. In Proceedings of the 14th International Workshop on Qualitative Reasoning, Morelia, Mexico.

Hayes, C. C., and Schlabach, J. 1998. FOX-GA: A Planning Support Tool for assisting Military Planners in a Dynamic and Uncertain Environment. In Integrating Planning, Scheduling and Execution in Dynamic and Uncertain Environments, R. Bergmann and A. Kott, eds. AAAI Press, Madison, Wisconsin.

Kewley, R. and Embrecht, M. 1998. Fuzzy-Genetic Decision Optimization for Positioning of Military Combat Units. In Proceedings SMC'98, pp. 3658 - 3664, 1998 IEEE International Conference on Systems, Man, and Cybernetics, October 11-14, La Jolla, California

Kott, A., Ground, L., and Langston, J. 1999. Estimation of Battlefield Attrition in a Course Of Action Analysis Decision Support System. 67th Military Operations Research Society Symposium, West Point Military Academy.

Kott, A., Saks, V. and Mercer, A. 1999. A New Technique Enables Dynamic Replanning of Aeromedical Evacuation, *AI Magazine*, v.2, n.1, pp.43-54.

Kott, A. and Saks, V. 1998. A Multi-Decompositional Approach to Integration of Planning and Scheduling - an Applied Perspective. In Integrating Planning, Scheduling and Execution in Dynamic and Uncertain Environments, R. Bergmann and A. Kott, eds. AAAI Press, Madison, Wisconsin.

Kott, A., Agin, G., and Fawcett, D. 1992. Configuration Tree Solver. A Technology for Automated Design and Configuration, *ASME Journal of Mechanical Design* 114(1): 187- 195.

Rasch, R., Kott, A., and Forbus, K. 2002. AI on the Battlefield: an Experimental Exploration. In Proceedings of the IAAI 2002 Conference, Edmonton, Canada.

Tate, A., Levine, J., Jarvis, P., and Dalton, J. 2000. Using AI Planning Technology for Army Small Unit Operations. In Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems (AIPS_2000).

Wass de Czege, H. and Biever, J. D. 2001. Six Compelling Ideas On the Road to a Future Army, *Army Magazine*, Vol.51, No.2, , pp. 43-48.

Wilkins, D. E. and Desimone, R. V. 1994. Applying an AI planner to military operations planning. In Fox, M. and Zweben, M., eds. *Intelligent Scheduling*, 685--709. Calif.: Morgan Kaufmann.