Deep Green: Commander's tool for COA's Concept

COL John R. Surdu, Ph.D.
Defense Advanced Research Projects Agency
mailto://john.surdu@us.army.mil

Kevin Kittka
Science and Technology Associates, Inc.
mailto://kkittka@stassociates.com

Keywords: Discrete event simulation, command and control, decision support systems, qualitative reasoning, planning, Human, Computer Interface

Abstract

Deep Green is composed of tools to help the commander rapidly generate courses of action (options) through multimodal sketch and speech recognition technologies. Deep Green will develop technologies to help the commander create courses of action (COA) (options), fill in details for the commander, evaluate the options, develop alternatives, and evaluate the impact of decisions on other parts of the plan. The permutations of these option sketches for all sides and forces are assembled and passed to a new kind of combat model which generates many qualitatively different possible futures. These possible futures are organized into a graph-like structure. The commander can explore the space of possible futures, conducting "what-if" drills and generating branch and sequel options. Deep Green will take information from the ongoing, current operation to estimate the likelihood that the various possible futures may occur. Using this information, Deep Green will prune futures that are becoming very improbable and ask the commander to generate options for futures that are becoming more likely. In this way, Deep Green will ensure that the commander rarely reaches a point in the operation at which he has no options. This will keep the enemy firmly inside our decision cycle. An overall vision for the Deep Green concept is an innovative approach to using simulation to support ongoing military operations while they are being conducted. By using information acquired from the ongoing operation, rather than assumptions made during the planning phase, commanders and staffs can make more informed choices and focus on building options for futures that are becoming more likely. This paper will describe an overview of the Deep Green concept with a focus on the Commander's tool for COA's in detail.

1. OVERVIEW OF BASIC BATTLE COMMAND CONCEPTS

In this section the authors sketch out the basic battle command doctrine and how Deep Green supports enhance battle command. This section will be a little heavy on doctrinal terms in order to set the context for later discussions. It refers to components of Deep Green that are described fully in the paper but have not yet been introduced. This was done because it is easier for readers to understand the context without the detailed knowledge of the components than to understand the roles of the components without the operational context.

Situational awareness is defined as "a Soldier knowing what is happening around him or her right now. Situational awareness occurs in the Soldiers' minds. It is not a display or the common operational picture; it is the interpretation of displays or the actual observation of a situation."[1] (FM 3-0) Deep Green does not seek to solve issues of sensor, data, or information fusion. Instead, it will extract information from the fused data store. The Commander's Associate component of Deep Green will develop new paradigms for presenting information to the commander to enhance his visualization of the battlespace. "*Commander's visualization* is the mental process of developing situational understanding, determining a desired end state, and envisioning the broad sequence of events by which the force will achieve that end state."[2] (FM 6-0) This visualization leads to the development of *commander's intent*, which is a concise statement of what the force must do and the conditions the force must establish with respect to the enemy, terrain, and civil considerations that represent the desired end state."[2] (FM 6-0). During execution, knowledge of the commander's intent helps subordinates exercise initiative, even if the plan goes awry. Visualization also leads the commander to a concept of the operation, which "directs the manner in which subordinate units cooperate to accomplish the mission and establishes the sequence of actions the force will use to achieve the end state." The commander (and the staff, if available) translates the commander's visualization into a specific

course of action for preparation and execution, focusing on expected results."[3] (FM 5-0) The Sketch to Decide subcomponent of Deep Green is designed to make it easy for the commander to translate his visualization into a course of action through an intuitive, natural, multi-modal (free-hand sketching and speech) interface. Staffs refine plans for branches and sequels to a plan during an operation. The interleaved operation of Blitzkrieg, Crystal Ball, and Commander's Associate will facilitate constant, ahead of real time enhancement and modifications of the plan.

According to Army doctrine, "*information superiority* is the operational advantage derived from the ability to collect, process, and disseminate an uninterrupted flow of information while exploiting or denying an adversary's ability to do the same."[4,1] (JP 1-02 and FM 3-0) Information superiority leads to *situational understanding*, which "enables commanders to determine the implications of what is happening and forecast what may happen. Situational understanding enhances decision making by identifying opportunities, threats to the force or mission accomplishment, and information gaps. It helps the command identify enemy options and likely future actions, the probably consequences of proposed friendly actions, and the effects of the environment on both."[1] (FM 3-0) *Decision superiority*, then, is "the ability of the commander, based upon information superiority and situational understanding, to make effective decisions more rapidly than the adversary, thereby allowing one to dramatically increase the pace, coherence, and effectiveness of operations."[5] (JFCOM glossary)

Information superiority facilitates decision superiority, but the authors assert that we need to think about how to achieve decision superiority with information parity or information inferiority. The assumption that U.S. forces will always (or even most of the time) have information superiority, particularly in a counter-insurgency, is poor particularly against a patient, irregular force. In situations in which U.S. and coalition forces are engaged in peace making, peace keeping, or counter-insurgency, the enemy will often have better visibility of our tactics, locations, procedures, and patterns of operation than we will of them. The trick then is to achieve decision superiority in the absence of information superiority, to have good situational understanding even when information is scarce.

Decision superiority will require better predictive tools, and those tools must learn constantly. Decision support tools should learn enemy tactics and help the humans learn them as well. These same learning tools should be used to analyze our own operations to determine if we are becoming predictable -- like in Mogadishu.[6]

Deep Green seeks to achieve decision superiority in the potential absence of information superiority. The authors assert that the staff needs tools to help anticipate the information needs of the commander to find the relevant information. "Relevant Information is all information of importance to commanders and staffs in the exercise of command and control… Relevant information provides the answers commands and staffs need to successfully conduct operations."[1] (FM 3-0)

The plan induction subcomponent of Deep Green will induce the commander's intent for an operation from his sketching and speech. It will then go beyond simply monitoring commander-defined information requirements, known as commander's critical information requirements, or CCIR. Such a tool that understands the commander's intent and is also trying to predict which possible futures are becoming more likely could anticipate the commander's information needs, not merely monitor CCIR. The Crystal Ball component of Deep Green, described later, is designed to do just that. Crystal Ball tries to predict the most likely futures. With its understanding of the commander's intent, Deep Green will help determine what information is needed to narrow the set of likely futures. The authors assert that knowledge of which futures the operation is trending toward can greatly increase the quality of the commander's decisions. Commanders have a few finite points of influence in the course of a battle. Deep Green will seek the next points of influence, anticipating information requirements for the commander, based on learning his style and understanding his intent, and gather the information relevant to the next couple of decisions.

At the divisional staff level, a "red team" is used to critique the commander's and staff's choices, to provide a contrarian view, to avoid "group think." [7]. Deep Green also seeks to provide red team-like capabilities at lower echelons. The Blitzkrieg component of Deep Green generates more possible futures than humans can manage and then, using information about the current operation to predict the likelihood that those futures may occur, will provide a capability to do sensitivity analysis for the commander. "If this piece of information turns out to be false, what does that mean about my prediction of the likelihood, utility, and/or flexibility of some possible future?"

2. OVERALL VISION FOR DEEP GREEN

In a military operational environment the only invariant is constant change, particularly the situation and goals. Under uncertain and time-critical conditions, it is important for commanders to have the ability to rapidly understand the unfolding trajectory of the operation and generate options quickly. More importantly, however, in modern warfare, it is important for the commander to be able to proactively generate options well in advance of when those options are needed rather than generate options reactively as the situation forces him off the plan. In this situation, it is much more important for the commander to have options than to

have planned the optimum course of action in fine detail. Robust plans are those that provide not just good outcomes but maximum flexibility to adapt to unforeseen or unexpected situations.

The Defense Advanced Research Projects Agency (DARPA) has recently release a broad area announcement (BAA), 08-09 Solicitation[8] for a battle command technology program, called Deep Green. Going beyond IBM's "Deep Blue"[9] Supercomputer for Chess, Deep Green is meant to be a commander-driven technology, rather than on building technologies to remove the commander. The Deep Green program has the goal of providing tactical commanders a technology to:

- generate and analyze options quickly, including generating the many possible futures that may result from a combination of friendly, enemy, and other courses of action;
- use information from the current operation to assess which futures are becoming more likely in order to focus the development of more branches and sequels; and
- make decisions cognizant of the second- and third-order effects of those decisions.

Deep Green is composed of tools to help the commander rapidly generate courses of action (options) through multimodal sketch and speech recognition technologies. Deep Green will develop technologies to help the commander create courses of action (options), fill in details for the commander, evaluate the options, develop alternatives, and evaluate the impact of decisions on other parts of the plan. (See Figure 1.) The permutations of these option sketches for all sides and forces are assembled and passed to a new kind of combat model which generates many qualitatively different possible futures. These possible futures are organized into a graph-like structure. The commander can explore the space of possible futures, conducting "what-if" drills and generating branch and sequel options. Deep Green will take information from the ongoing, current operation to estimate the likelihood that the various possible futures may occur. Using this information, Deep Green will prune futures that are becoming very improbable and ask the commander to generate options for futures that are becoming more likely. In this way, Deep Green will ensure that the commander rarely reaches a point in the operation at which he has no options. This will keep the enemy firmly inside our decision cycle.

Figure 1: Operational Concept for Deep Green

The venerable Observe-Orient-Decide-Act (OODA) loops [10] no longer viable for an information-age military. Deep Green creates a new OODA loop paradigm. When something occurs that requires the commander's attention or a decision, options are immediately available. When the planning and execution monitoring components of Deep Green mature, the planning staff will be working with semi-automated tools to generate and analyze courses of action ahead of the operation while the command concentrates on the Decide phase. By focusing on creating options ahead of the real operation rather than repairing the plan, Deep Green will allow commanders to be proactive instead of reactive in dealing with the enemy.

Deep Green was inspired by two concepts: anticipatory planning and adaptive execution. **Anticipatory planning** can be described colloquially as "you know you're going to replay anyway, so why not re-plan ahead of time?" This drives the notion of generating options and futures before they are needed. To some extent Deep Green will trade depth for breadth. Today commanders plan a small number of options very deeply, i.e., all the way to the end of execution in great detail. Most of these deep plans are discarded once the plan goes awry. Sometime the commander and staff are unable to recognize that the plan is broken or is becoming broken. They are often unable to divorce themselves from the plan in order to seek new affordances based on the current state of the operation. By identifying the trajectory of the operation and focusing the commander and staff where to build (perhaps less deep) plans, the commander will have a broader set of options available at any time. This leads to the concept of **adaptive execution[11]**, which is similar to the AI planning concept of late binding. Adaptive execution intends to make decisions at the last moment in order to maintain flexibility to adapt to updated trajectories of the operation.

3. BASIC SYSTEM ARCHITECTURE

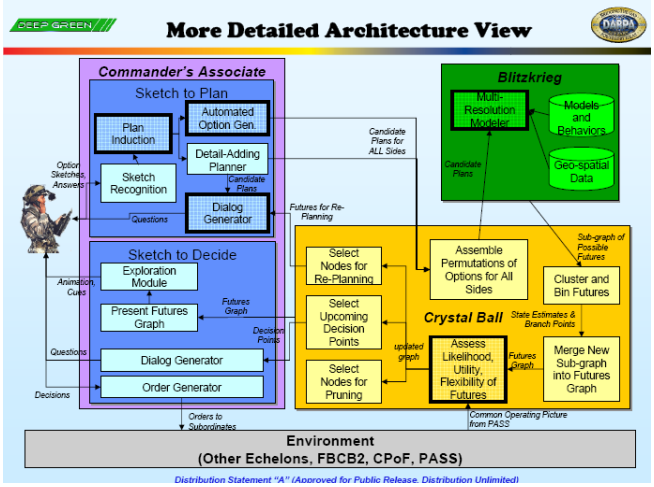

Figure 2: Architectural Overview of Deep Green

3.1. Blitzkrieg

Blitzkrieg is the simulation component of Deep Green. It is used to generate the possible futures that result from a set of plans (one plan for each side/force in the operation). Besides being very fast (the *blitz* in Blitzkrieg), it is designed to generate a broad set of possible futures. These futures should be feasible, even if not expected by human users. Over time, Blitzkrieg should learn to be a better predictor of possible futures, based on presented options. Blitzkrieg identifies branch points, predicts the range of possible outcomes, predicts the likelihood of each outcome, and then continues to simulate along each path/trajectory. Gilmer and Sullivan provide an example of a possible implementation of this idea [12] in which they determine branch points and continue to simulate along multiple paths. Blitzkrieg should reflect out-of-the-box thinking, rather than merely generating hundreds or thousands of "Monte Carlo" runs of a stochastic model and binning the outputs [13]. This will require an innovative hybrid of qualitative and quantitative technologies.

3.2. Crystal Ball

Crystal Ball serves several functions. First, it controls the operation of Blitzkrieg in generating futures. Second, it takes information from the ongoing operation and updates the likelihood metrics associated with possible futures. Third, it uses those updated likelihood metrics to prune parts of the futures graph and nominate futures at which the commander should generate additional options and invokes Sketch to Plan. Finally, it identifies upcoming decision points and invokes Sketch to Decide. While Crystal Ball has a moderate role prior to execution, it is the backbone of the system during execution.

3.3. Commander's Associate

The Commander's Associate has two major sub-components, Sketch to Plan and Sketch to Decide. (See Figure 2.) The two components are discussed separately because in an open, modular architecture, it is envisioned that one or the other must be able to be replaced with new technologies over time without disrupting the entire system. A goal of the Deep Green program is to develop and apply computer software technologies to develop a Commander's Associate that automatically converts the commander's hand-drawn sketch with accompanying speech of his intent into a Course of Action (COA) at the brigade level. The Commander's Associate must facilitate option generation, "what-if" drills, and rapid decision making.

Sketch to Plan

This component provides the commander the ability to generate quickly qualitative, coarse-grained COA sketches that the computer can interpret. Sketch to Plan will be multi-modal (both sketching and speech) and interactive. The computer will watch the sketch being drawn and listen for key words that indicate sequence, time, intent, etc. as the commander is creating the sketch. Sketch to Plan will induce both a plan and the commander's intent from the sketch and speech. Unlike other approaches that are optimized around machine interpretations[14] (i.e. constraining the sketching method to drag-and-drop modalities, forcing the human to learn the computer's 'language' to some extent), Sketch to Plan is optimized around the user free-hand sketching options over a map. In addition, the Sketch to Plan component must be imbued with enough domain knowledge that it knows what it doesn't know and can ask the user a small set of clarifying questions until it understands the sketch and can use it to initialize a combat model.

In many situations today, courses of action are sketched on butcher paper or white boards. The intent behind Sketch to Plan is to build an interface that is as easy to use as a piece of butcher paper. The sketch Recognizer converts a free-hand set of strokes, combined with speech, into a set of military objects, such as units and graphical control measures (MIL STD 2525b [15] and STANAG 2019 APP-6A [16]). This is a difficult problem, but one that is analogous to optical character recognition. Whereas alphabets have a small set of characters, there are thousands of military symbols, and in some cases their interpretation depends on other symbols around them. Sketch recognition results in the proper identification of a "bag of symbols."

The plan inducer, then, has the challenge of inducing the commander's plan, or scheme of maneuver, and intent for the recognized "bag of symbols." This induction of

higher levels of semantic abstraction and understanding is the truly "DARPA-hard" problem and will require a clever combination of domain knowledge, sophisticated reasoning technologies, and perhaps learning technologies.

We also envision a detail-adding planner within Sketch of Plan that adds details to the commander-generated option so that it can be modeled by Blitzkrieg. Finally, the dialog generator helps Sketch to Plan understand the commander's option by formulating clarifying questions when necessary.

Sketch to Decide

When the commander is asked for a decision, Sketch to Decide will allow him/her to explore the future space to gain an appreciation for the ramifications of a choice. It is envisioned as similar to a comic strip with branch points that correspond to branch points in the futures graph. Scott McCloud [17] asserts that the idea of a comic in which the readers get to make a choice at the branch points is today "exotic" but may well become common in the future. Since the 1970s (and perhaps earlier), there have been novels and game books in which the reader is asked to make a decision and then is directed to a different page or paragraph, depending on the choice made, such as the 1980's children's Choose Your Own Adventure gamebook series or the DVD movie Clue based on the board game Clue as examples. Recently Forbus has explored the idea of a comic graph [18]. The idea here is the same: the user gets to choose which path to follow at a branch point. One can imagine the commander exploring the future space to understand how his courses of action may play out and identifying the critical branch (decision) points.

Sketch to Decide is designed to allow the user to "see the future," but this capability must be developed with care to prevent confusing the decision space. Humans are notoriously bad at thinking through probabilistic choices and even more so when there are competing outcome utilities. At each branch point, there are multiple decision dimensions/utilities that have to be considered, such as likelihood, risk, utility, resource usage, etc. In addition, the abstract nature of the state and the uncertainty of predictions, locations of units, etc. must be portrayed intuitively. Therefore, at any "frame" in the Sketch to Decide graph, the user can perform Sketch to Plan actions, allowing the commander to conduct "what-if" drills wherever he wants in the future space. The user is going to need a lot of help in evaluating these options, especially because they are probabilistically weighted. By presenting decisions early and allowing the commander to explore the future space, Sketch to Decide supports adaptive execution, allowing the commander to make decisions when they are needed, rather than committing too early.

As described earlier, visualization occurs not on the screen but in the commander's head. We will go beyond common paradigms for displaying information and develop technologies that tailor the presentation of information. This will be enabled by the system understanding the commander's scheme of maneuver and intent, as induced by Sketch to Plan. This knowledge of what the commander is trying to accomplish and what he is trying to accomplish should enable to system to not only present the right information at the right time, but to also present it in the best way to present it to aid in cognition. This will be further enhanced by imbuing Sketch to Decide with the ability to learn the commander's preferences and refine itself over time.

3.4. Automated Option Generation

The focus of Deep Green is on tools to help the commander (and staff) generate options quickly. Leaders from the field generally do not want machine-generated courses of action. Nevertheless, under Deep Green, we intend to sponsor a small set of modest efforts to generate options automatically. The long-term vision of Deep Green is for options to be generated by both the commander and the computer. Initially we expect the machine generation of options to be centered on making clever "mutations" of the human-generated options to increase the breadth of the futures generated. This highlights the need for Sketch to Plan to induce the commander's intent from the free-hand sketches. Any options generated by the computer should feasibly meet the commander's intent.

4. FUNDAMENTAL SHIFT AWAY FROM THE TRADITIONAL OODA PARADIGM

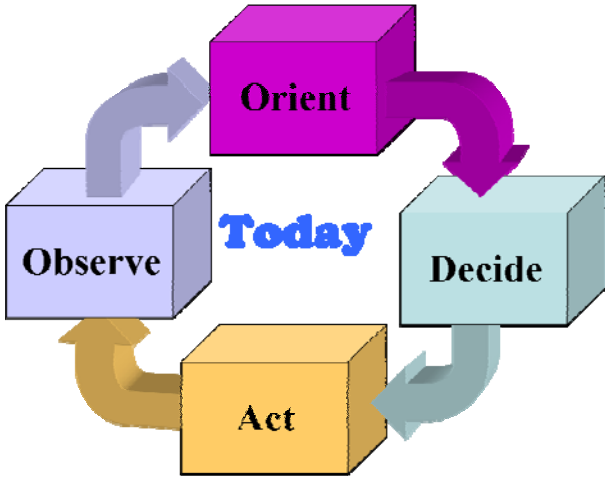

Figure 3: The OODA Loop

The OODA loop concept [19] was first introduced by Col John Boyd, U.S. Air Force fighter pilot ace, in 1986 in his presentation entitled "Patterns of Conflict" (POC). (See Figure 3) Since then there have been many variations of this

process. The venerable Observe-Orient- Decide-Act (OODA) loop is no longer viable for an information-age military. Previous work has centered on speeding up the overall loop or developing technologies that work within a single phase of that loop. Today, when the plan goes awry, we go into a reactive mode, in which we create courses of action, analyze them, and then choose.

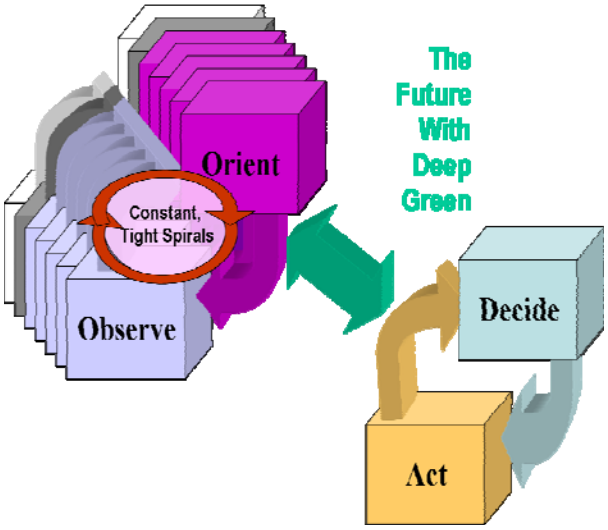

Figure 4: Multiple OO's, One DA Loop Processes

Deep Green creates a new OODA loop paradigm. (See Figure 4) Observe (execution monitoring) and Orient (options generation and analysis) phases run continuously and are constantly building options based on the current operation and making predictions as to the direction the operation is taking. When something occurs that requires the commander's attention or a decision, proactive options are immediately available. Ideally, the OO part of OODA is done many times prior to the time when the commander must decide. When the planning and execution monitoring components of Deep Green mature, the planning staff will be working with semi-automated tools to generate and analyze courses of action ahead of the operation while the command concentrates on the Decide phase. By focusing on creating options ahead of the real operation rather than repairing the plan, Deep Green will allow commanders to be proactive instead of reactive in dealing with the enemy.

5. SUMMARY

We are just getting started! Deep Green will provide technology to break the OODA paradigm. Deep Green enables the rapid construction of sophisticated planning and execution systems using existing technologies. The overall objective will be an open and scalable battle command decision support architecture that interleaves anticipatory planning and adaptive execution to stay inside the enemy's decision cycle. Deep Green will provide an implementation framework to enable rapid technology insertion into battle command systems today and in the future. When successful, we will build a revolutionary decision support system that will allow us to defeat peer competitors in the future. So long and thanks for all the fish.

6. REFERENCES

[1] Unknown. (2008). FM 3-0; OPERATIONS. In H. D. o. t. Army (Ed.) (pp. 220): U.S. Army Training and Doctrine Command..

[2] Unknown. (2003). FM 6-0; Mission Command: Command and Control of Army Forces. In D. o. t. Army (Ed.) (pp. 333): Dept of the Army.

[3] Unknown. (2005). FM 5-0: Army Planning and Orders Production. In D. o. t. Army (Ed.) (pp. 328): Department of the Army.

[4] Unknown. (2007). JP 1-0, Joint Warfare of the Armed Forces of the United States. In D. F. O. P. A. J. F. D. (J7) (Ed.) (pp. 155): Joint Chiefs Of Staff Publication.

[5] Unknown. (2008). USJFCOM Joint Forces Command Glossary. Retrieved 06 MAR 2008, from http://www.jfcom.mil/about/glossary.htm

[6] Major Timothy Karcher, US. Army. (1967). Understanding the "Victory Disease," From the Little Bighorn to Mogadishu and Beyond (pp. 64): Combat Studies Institute Press.

[7] Ripps, T. S. G. (2008). Team to Provide Different Perspectives to Division Commander in Afghanistan. The Dispatch: The Monthly Magazine of the Texas Military Forces. Vol 2: 5.

[8] Surdu, John R., COL. (2007). Deep Green Broad Agency Announcement No. 08-09. Defense Advanced Research Projects Agency (DARPA), Information Processing Techniques Office (IPTO), Website: http://www.darpa.mil/ipto/solicitations

[9] Unknown. (1997). Kasparov Vs. DEEPBLUE: The Rematch – Overview. IBM. Website http://www.research.ibm.com/deepblue/

[10] John R. Boyd, C. (1986). "Patterns of Conflict." 196. http://www.d-n-i.net/boyd/pdf/poc.pdf

[11] Brigadier General Huba Wass de Czege, U. A., Retired and U. A. Major Jacob Biever (1998). "Optimizing Future Battle Command Technologies." Military Review(March/April 1998): 6.

[12] Gilmer, J. B., Jr. and F. J. Sullivan (2005). "Issues in Event Analysis for Recursive Simulation." Proceedings of the 2005 Winter Simulation Conference: 8.

[13] Law, A. M. and W. D. Kelton (2000). Simulation modeling and analysis. United States of America, The McGraw-Hill Companies, Inc.

[14] Forbus, K. D., J. Usher, et al. (2003). Sketching for Military Courses of Action Diagrams. IUI'03. Miami, Florida, USA, ACM.

[15] Unknown (1999). "Department of Defense Interface Standard: Common Warfighting Symbology - MIL-STD-2525B ": 556. http://assist.daps.dla.mil/docimages/0001/52/27/2525B.PD5

[16] Unknown (2000). "NATO Standardization Agreement 2019: APP6A Military Symbols for Land Based Systems." http://www.mapsymbs.com/app-6ahandbook.zip

[17] McCloud, S. (1994). Understanding Comics: The Invisible Art. New York, NY, HarperCollins Publishers, Inc.

[18] Forbus, K. D., J. Usher, et al. (2003). Qualitative Spatial Reasoning about Sketch Maps. Fifteenth Annual Conference on Innovative Applications of Artificial Intelligence, Acapulco, Mexico, American Association for Artificial Intelligence Press.

[16] Boyd, Col. J. R. (1986). "Patterns of Conflict." (Pg 132) 196 pages from http://www.d-n-i.net/boyd/pdf/poc.pdf

Biography

COL John R. Surdu, Ph.D., U.S. ARMY, Program Manager, DARPA Information Processing Techniques Office. COL “Buck” Surdu was commissioned a second lieutenant of infantry after graduating from the United States Military Academy in (1985).

COL Surdu has worked as a research scientist and team leader at the Army Research Laboratory, focusing on unique uses for virtual reality technologies for command and control applications. As a senior research scientist at the Information Technology and Operations Center, he directed several applied research efforts. From 2003-2006 COL Surdu has been the Product Manager for the One SAF program office.

In addition to a Bachelor of Science degree in computer science from the United States Military Academy, in (1985) COL Surdu earned a Master of Business Administration degree from Columbus State University. In (1990), COL Surdu earned Master of Science degree in computer science from Florida State University (focusing on artificial intelligence in 1995). He finalized his formal education with a doctoral degree in computer science from Texas A&M University in (2000) (focusing on simulation technology and its applications to command and control).

Kevin Kittka, Program Manager, Science and Technology Associates, Inc., has a B.S. from The Pennsylvania State University and a M.S. from The University of Maryland University College. Mr. Kittka has provided consulting support to NAVSEA for 8 years and Scientific, Engineering, Technical, and Administrative Support to the Defense Advanced Research Projects Agency for the past 20 years. His research focus is on Information Systems. Currently a member of AAAI and an associate member of IEEE.

Note: This document is Approved for Public Release, Distribution Unlimited (Case #11260)