Integrated Design and Process Technology, IDPT-2006
Printed in the United States of America, June, 2006

APPLYING COMPLEX REASONING TOOLS TO MISSION EFFECTIVENESS

Andy Bevilacqua
Bevilacqua Research Corporation
P.O. Box 14207
Huntsville, Alabama 35815
256-882-6229 Extension 102
andyb@brc2.com

H. Barbara Sorensen Ph.D.
Warfighter Training Research Division
U.S. Air Force Research Laboratory
6030 South Kent Street
Mesa, Arizona 85212
480-988-6561 Extension 184
Barbara.Sorensen@mesa.afmc.af.mil

David J. Skipper Ph.D.
Bevilacqua Research Corporation
P.O. Box 14207
Huntsville, Alabama 35815
256-882-6229 Extension 105
davids@brc2.com

ABSTRACT

Survivability of military and civilian assets is a critical issue across the armed services, prompting the need for new technologies in automating the process of identifying, assessing, and targeting a threat to support combat operations. The overall effectiveness of combat operations in achieving mission objectives constitutes the Combat Assessment (CA). CA itself supports a Command and Control (C2) system that allows the Joint Services to assess the effectiveness of on-going target strikes. An integral part of CA is Battle Damage Assessment (BDA). Traditionally, analysts reviewed multiple sources of data and prepared written assessments of the best guess as to the damage and the need for re-strikes. Using new hybrid Artificial Intelligent (AI) techniques, an Integrated Battle Damage Assessment System (IBDAS) is a rapidly re-configurable and maintainable decision support architecture that can correlate large amounts of asynchronous data from disparate data sources in near real time. This paper describes key aspects of the reasoning used in the program that have applicability to other decision aid systems.

INTRODUCTION

Decision assistance tools have been used at least since early hunters stacked rocks to mark trails. The development of modern computers provided the decision maker with the ability to perform numerical calculations to provide additional insights for the decision maker. Similarly, the use of non-numeric or symbolic calculations provided more options for the decision maker. While relieving the decision maker of the requirement to perform tedious supporting computations, humans wanted more. They wanted decision aids that could organize and understand the data. They wanted decision aids that could suggest and explain alternatives as well as evaluate the options. Moreover, they wanted decision aids that could recognize and respond to changes in the world while grasping and absorbing new ideas and concepts. In short, they wanted an automated analyst capable of performing a job so well that the human can focus on those truly unusual and difficult problems that can make a job interesting, if not fun. They wanted a tireless associate to help them, which requires looking closely at problems and how these problems are represented and how "reasoning" is applied to solve them.

This paper considers the problem of "reasoning" in a general sense to support decision making. These techniques are being employed or are being developed for ongoing work at Bevilacqua Research Corporation. The following section describes the problem addressed by the system, with specific characteristics of interest to other decision support designers. The subsequent sections deal with our problem definition and the approach used to address this problem area.

PROBLEM OVERVIEW

Military forces perform combat assessment to understand how operations are affecting an opponent, which, in turn, leads to another decision, "What do we do next?" Consequently, combat assessment represents an important problem as well as one complex enough to benefit from decision support.

For this paper, combat assessment can be abstracted to three problem areas. These areas are situation awareness, situation understanding, and classification (or decision making). In the first area, the analyst or decision aid must sense the outside world and maintain a "mental" picture of every object and every activity in the combat area that could influence operations. This picture must be current and realistic at all times. In the second area, the decision maker must perceive what the situation means to the decision maker and the commander at the present time, and the implications for future times. Finally, the decision maker must decide on a course of action based on what the understanding is at that time. Sensing the world and maintaining the world situation picture require additional comments.

SENSING THE WORLD

Combat assessment provides a commander with feedback on how well the commander's actions are affecting the opposing commander's actions. This feedback enters the decision loop through a variety of disparate sensor types and sensor locations, all acting at differing times. These sensors may range from on the spot human observation messages to live video feed from overhead platforms. Each sensing system has errors and constraints that must be considered. The result is that the sensors introduce "imperfections" into the information, which causes "noise" in the feedback loop. Further, each sensor must also face the possibility of deliberate deception on the part of the opponent. After the observations are collected, these sensor observations must be interpreted by an analyst to estimate changes and the amount of damage due to operations. A representation of these observations provides the current situation for the commander and the analyst.

MAINTAINING THE PICTURE

The temporal evolution of the current situation develops from what a commander imposes on an opponent and what the opponent imposes on the commander. The commander requires assessments of both in order to make informed decisions. This paper is concerned with the first assessment.

From the analysis of the sensor observations, the change in combat effectiveness of the opposing forces is estimated to form a Battle Damage Assessment (BDA) of operations. Both of theses estimates are strongly affected by the analyst's experience, training, and condition at the time the analysis is performed. In an effort to somewhat mitigate these variance factors, analysts work with standardized terminology. Unfortunately, there is still a considerable "art," as opposed to "science," in the overall process where good intelligence analysts are highly valued team members. As with many human endeavors, this art requires multiple understanding and processing modes in the human, and it is an extremely difficult computer problem to provide even an approximation to that art.

In order to prepare the BDA, the analyst must thoroughly understand the capabilities of the sensors used to obtain the pieces of information used in the BDA picture as well as the pieces themselves. These pieces must then be put together to form the overall picture. This welding together of the fragments of information, or fusion of the data, should provide better insights for the commander, based on the "whole is more than the sum of the parts" philosophy. Unfortunately, this also leads to editing based on the analysts' perceptions and shifts of focus in the commander's mind when viewing the composite. Automation of the fusion and understanding is a highly desirable goal.

RELATED ISSUES

Based on the previous discussion of the problem areas, we refine to certain key issues. Specifically, there are points buried in the previous areas that must be kept in mind while creating combat assessment decision support systems.

1. Time is limited. The time available to analyze and decide may vary from a few minutes to a few hours. Combat operations proceed, ready or not. Speed depends on algorithms, structures (representations), and computational environments.
2. Detailed knowledge and understanding of the opponents' systems may require large knowledge sets. The representation of the knowledge must be compact and easily exchanged between applications and systems. It must also be compatible with an array of reasoning approaches as we attempt to emulate human cognition.
3. Detailed knowledge of our own sensor systems and their constraints may also require large knowledge sets.
4. Fusion of multiple pieces requires understanding the data and the data imperfections as well as the ability to reason in the face of these imperfections. The representation of the pieces must accommodate complex aggregation of various data types and embody representations of the imperfections in the information.

5. Fusion of multiple pieces from multiple sensors also requires representations that have the ability to support reasoning with time to adjust observations to matching times and to project behaviors and functional components in time, so as to perform algorithmic, time based estimates.
6. Summarizing information for commanders requires the ability to abstract and to identify important issues; to defend the abstractions, issues, and decisions; and to explain the rationale. "Black box" decision processes are usually suspected as faulty where combat decisions must be made. Again, a clear easily understood representation is required.
7. The extent of the knowledge and the variety of knowledge types require multiple reasoning techniques to deal with different time constraints and quality of the answer issues. Reasoning in a system may well encompass more than basic deductions.

All of these are difficult issues to resolve, and the continuing reference to the importance of selecting a good representation offers a hope for making progress in combat assessment decision support systems.

GENERAL DESIGN APPROACH

While not claiming to have solved the previously listed issues, we did seek to alleviate some of the main difficulties by adopting a representation that supports multiple solution options through that representation, *i.e.* extensible hybrid representations. Our selection was to use the Conceptual Graph (CG) representation developed initially by Sowa (1984, 2000). This is a mature area of research, and it is also one of the three forms selected for incorporation into the draft ISO Common Logic standard (ISO FCD 24707 Common Logic). By providing a multidimensional representation for knowledge, the semantics, syntax, and topology of the knowledge are available for manipulation and comprehension. Pertinent portions of these capabilities are important to the discussions in the following sections.

CONCEPTUAL GRAPH SUMMARY

Conceptual graphs (Figure 1) are constructed from concepts (types), represented by rectangles, and relationships, represented by ovals, between concepts. Edges within the graph are permitted from concepts to relations and from relations to concepts. There is a special form of the relation called an actor, represented by a diamond, that provides dynamic relationships, interfaces, and manipulations of the graph by the graph at the, sometimes, cost of the loss of decidability in a logic sense. However, approximate or analogous decisions are potentially still possible. Sowa shows that Conceptual Graphs can perform first order logic by the application of combinations of six graph manipulations (Figures 2 – 4):

1. Join
2. Detach
3. Restrict
4. Un-restrict
5. Copy
6. Simplify

Note that operations 1 and 2 are opposites, operations 3 and 4 are opposites, and operations 5 and 6 are opposites. Consequently, applying any of these operations to manipulate a Conceptual Graph represents some form of reasoning, albeit sometimes limited, with the Conceptual Graph. There are other useful graph manipulations, but the most prominent, non-reasoning operation is graph (sub-graph) isomorphism that provides, among other things, graph query operations.

The principle difficulties encountered when working with Conceptual Graphs are:

1. The requirement to think in multiple dimensions when representing knowledge,
2. Automation of the reasoning,
3. The computational complexity of some of these graph operations.

The first issue is one of education and training that is slowly being addressed by the Conceptual Graph community. The last is the subject of algorithm research that continues in the conceptual graph community, but it is partially alleviated by partitioning problems into working size chunks. The following sections revisit elements of the Problem Overview section.

SITUATION AWARENESS COMPONENT

The Situation Awareness component consists of two key parts. The first is the background knowledge about the systems and situations of interest. These are built as systems dictionaries of templates that contain CG models of the expected situation components, targets, and sensors. Combined with the actual observations of the sensors, Conceptual Graph fragments of the current situation are constructed by filling in graph templates based on message traffic (Figure 5), then applying one or more of the previous six operations to construct a graph of the situation (Figure 6), *i.e.* the "mental" picture. By working with small fragments containing key information, the computational load from using Conceptual Graphs become tractable.

A problem for Situation Awareness is the issue of imperfect information. Once again, dictionaries of imperfections based on CG models can be pre-built and filled in as information arrives (Figure 7). Conceptual Graphs can then incorporate the imperfection representations as concepts and relations that may be manipulated by subsequent reasoning processes. With this approach to uncertainty, the management does not require a specific uncertainty calculus to be built into Conceptual Graphs.

SITUATION UNDERSTANDING COMPONENT

Situation understanding occurs at both a perceptual and a reasoned level. Perception is the recognition of something without employing a chain of logic to arrive at the understanding. Humans sometimes refer to "internalizing" or a "flash of insight" when talking about perception. Deliberate reasoning, by contrast, follows a forward or reverse inference path of logic to deduce an understanding of the current situation at a slower pace. We employ the perceptual approach for reaching less complex decisions but reaching them quicker. Conceptual Graphs support perception through actors, graph matching, and graph complexity metrics.

Actors (Figures 8, 9 and 10) provide a fast algorithmic response to conditions described by concepts, so that when conditions or combinations of conditions are present in the graph, the actor can exceed a threshold and report that the condition exists or may exist in sections of the graph. Graph matching extends what is accomplished by an actor by also considering the topology of a graph. Thus, not just sets of values are important, but the relationships between the various elements. This is accomplished by performing various types of sub-graph isomorphism (complete/partial and exact/approximate) using known key situations against the current situation graph. Note that the matching can be flexible, depending on how important a fast response is for a given situation. The third complexity metrics provide a very crude similarity measure between graphs and standard situations. While potentially not as precise as sub-graph isomorphism, they can be substantially faster, providing a good initial look at what may be an emerging problem in a current situation.

DECISION MAKING COMPONENT

As seen in the previous sections, the actual decision making has already been partially distributed in the awareness and understanding of the graphs. At this point, the situation is seen and understood. Subsequently, the commander wants to know what to do next. In the simplest case, this is selecting courses of action that best fit our needs, that is, classification type of decision making. More detailed reasoning can accomplished by application of Sowa's six graph operations until the starting graph is transformed into a solution graph through the least common graph between the situation and the solution, or graph unification. Additionally, traditional rule based systems have been built using CG premises, CG conclusions with graph matching to accomplish a more traditional reasoning process. The advantage of this approach is that years of lessons learned about speeding up rule based inferences can be carried over into the CG world to obtain better real time response.

BRC BDA SYSTEM OVERVIEW

United States (US) forces rely on information superiority and information dominance on the battlefield to remain inside the operational tempo of their adversaries. Critical Target lists detailing platform, weapon, and navigation information are generated before a battle ever begins, and decisions are made beforehand about the courses of action as targets are destroyed. In the last decade, this combat assessment doctrine of "Decide-Detect-Deliver" was modified slightly by adding: "Assess," which means that during the Intelligence Preparation of the Battlespace (IPB) process, decisions are made concerning what to do if a target of opportunity is acquired. Therefore, when US assets detect a target of opportunity, the US forces can rapidly deliver their ordinance. Unless one can assess the damage, however, a re-strike recommendation cannot be made and critical opposing assets may escape. Thus, battle damage assessment (BDA) came to the forefront as a major factor in the overall combat assessment (CA) process. Since the Gulf War, BDA analysts have lacked a system allowing them to correlate and share the knowledge necessary to accomplish Phase I, II, and III assessments in a timely, accurate, and reliable way.

Bevilacqua Research Corporation is currently in the Phase II development process under the auspices of a Small Business Innovative Research (SBIR) program to develop a tool to answer this BDA gap for the U.S. Air Force Research Laboratory. This innovative software program, the Intelligent Battle Damage Assessment Station, utilizes conceptual graphs to provide assessments within and between distributed agents. This process follows the lines of the previous sections to correlate and assess large amounts of asynchronous inputs.

The Phase I SBIR program precursor to this effort concentrated on demonstrating a correlation/classification scheme based on conceptual graphs to accomplish rapid, accurate damage assessment for incoming ballistic missiles. Through the use of the hybrid cognitive processing techniques developed in Phase I, the Phase II program is building on Phase I research to develop a rapidly re-configurable and maintainable decision support architecture that can correlate large amounts of asynchronous data from disparate data sources in near real time for all phases of tactical and strategic battle damage assessments. Additionally, the cognitive reasoning engine can incorporate the knowledge of subject matter experts to provide the general order of battle analyst the capability to provide detailed, definitive assessments for equipment with which they may or may not be familiar.

The BDA tools (Figure 11) and the underlying correlation/classification engine developed in this program will be capable of correlating asynchronous data from disparate data sources in near real time, which will make it a valuable tool for real time damage assessment and operational assessment within the joint command and control structure. In an upgrade of the Phase I version, the current platform will operate in a collaborative environment, providing BDA analysts at separate locations easy access to the same informational database. In this way, the basic product developed under the program will prove more widely applicable throughout both the U.S. Air Force and the rest of the Department of Defense.

CONCLUSIONS

Intelligent decision aids appear to require a variety of cognitive skills to provide rapid, deep understanding of a subject. From our past developments, the key to providing these capabilities seems to lie with selecting a versatile knowledge representation. Conceptual Graphs provide the unique design options for developing complex systems. In a large part, this is due to the variety of ways in which "reasoning" can be implemented using this representation. That variety allows the developer to select the methods best fitting their needs.

REFERENCES

Sowa, J. F. 1984, *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley.

Sowa, J. F. 2000, *Knowledge Representation: Logical, Philosophical and Computational Foundations*, Addison-Wesley.

FIGURES

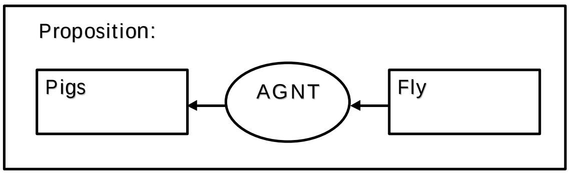

Figure 1 Conceptual Graph

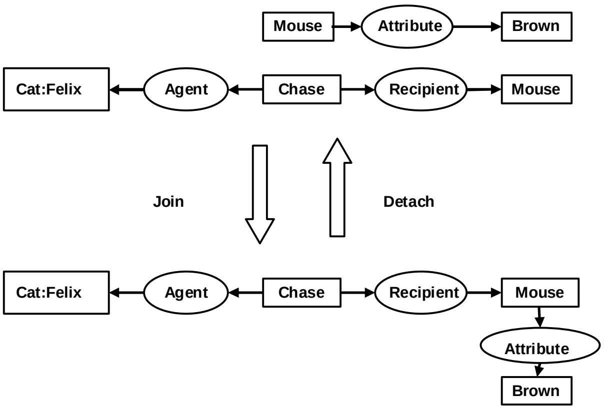

Figure 2 Join Detach

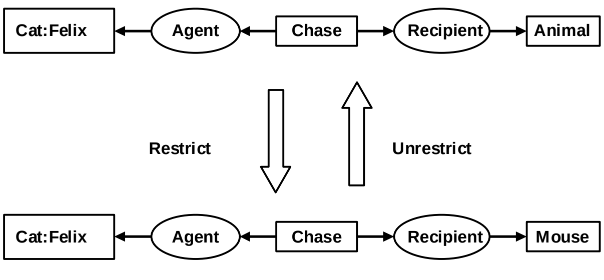

Figure 3 Restrict Un-Restrict

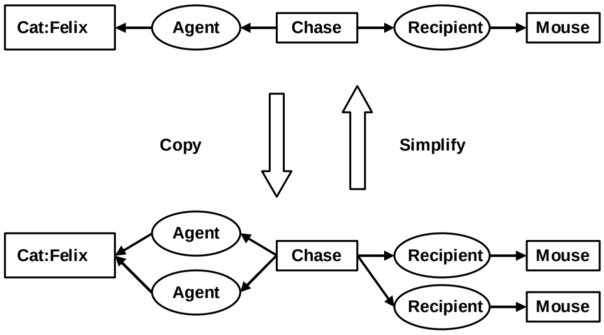

Figure 4 Copy Simplify

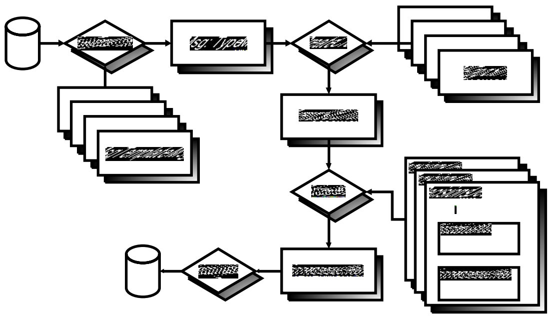

Figure 5 Template Usage

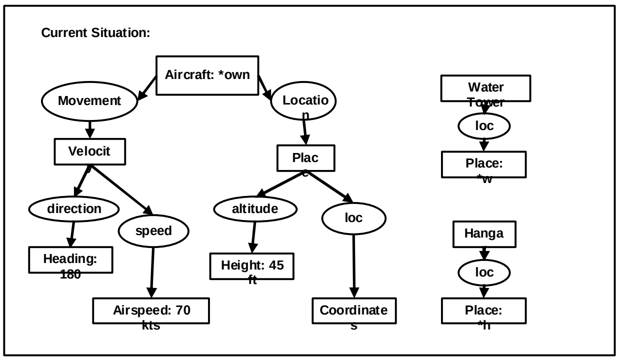

Figure 6 Situation Graph

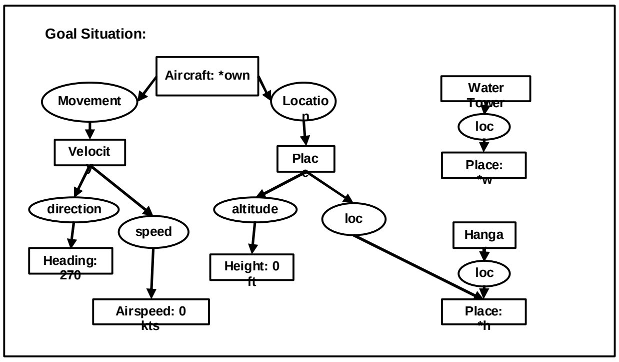

Figure 7 Dictionary Usage

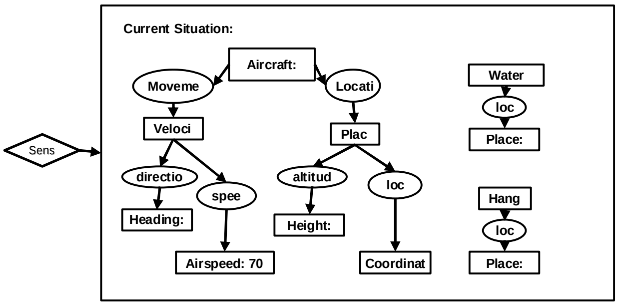

Figure 8 Interface Actors

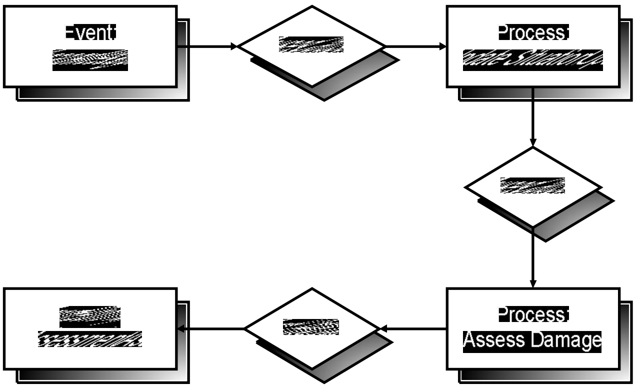

Figure 9 Active Relation

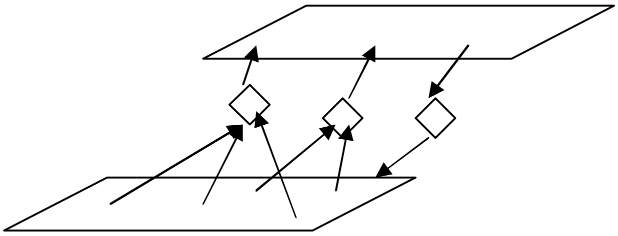

Figure 10 Perception Actor

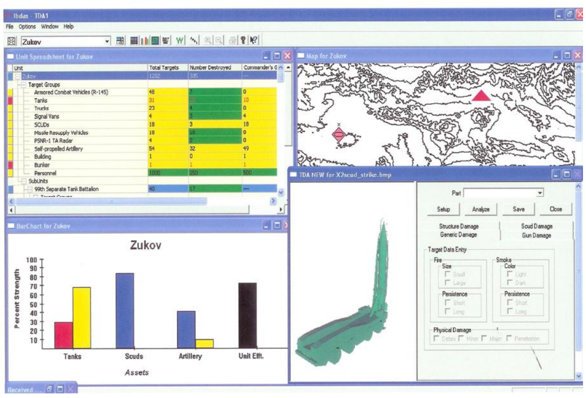

Figure 11 BDA Tools