

KNOWLEDGE MANAGEMENT FOR COMMAND AND CONTROL

Dr. Marion G. Ceruti,* Dwight R. Wilcox, and Brenda J. Powers*****

Space and Naval Warfare Systems Center, San Diego,
Code 24121,* *** Code 24122,** 53560 Hull Street, San Diego, CA 92152-5001
Tel. (619) 553-4068,* (619) 553-5467,** (619) 553-5538,***
marion.ceruti@navy.mil, dwight.wilcox@navy.mil, bpowers@spawar.navy.mil

Abstract

This paper highlights some major trends and developments in knowledge management with particular emphasis on knowledge capturing and authoring, and how this technology can be combined with intelligent agents to produce advanced capabilities for command and control systems. Past trends, present accomplishments and future work in knowledge management systems are covered. We also describe an in-house effort to accumulate, validate, and integrate the results of the Defense Advanced Research Projects Agency (DARPA) Rapid Knowledge Formation (RKF) program with a view toward the recommendation of specific technologies for transition to operational command centers. This technology will permit scientific, technical, and military experts, such as command-center and intelligence-center personnel to encode massive amounts of knowledge into reusable knowledge bases for application in many different tasks. This paper also presents current and future uses of adaptive and mobile intelligent agents to increase the dependability of FORCEnet by improving network-centric warfare in general, and the battlefield single integrated picture (SIP) in particular. Applications of mobile and adaptive intelligent agents are described. Further research is planned in the area of ontologies, expert systems, knowledge management, intelligent agents and user-interfaces for subject-matter experts under this program.

Keywords – command and control, FORCE-net, intelligent agents, knowledge base, knowledge-management system, network-centric warfare, subject-matter expert

I. Introduction to knowledge management

Knowledge has been defined as “an evolving mix of framed experience, values, contextual information and expert insight that provides a framework for evaluating and incorporating new ideas and information” [11, Davenport and Prusak, 1998]. Knowledge management is the process of creating value through the organizational integration of the knowledge in the organization [22, Wigg, 2003]. Knowledge management concerns KBs and much more. One type of a knowledge base is a state of information that consists of a collection of rules, axioms or assertions structured according to an ontology and a knowledge representation that allows knowledge to be stored explicitly, and from which conclusions can be drawn using an inference engine [9, Ceruti, 2004]. Another type of knowledge base is a structured acyclic graph, such as a Bayesian network that stores knowledge in its structure and in its associated conditional probability table. The discussion on knowledge bases in this paper is limited to the first type of

knowledge base. Intelligent agents are seen as another application of knowledge management. Because of its utility, this emerging technology gradually will be introduced into operational applications.

The remainder of the paper is organized as follows. Section I introduces some facts and observations about knowledge management and its benefit to command and control. Subsections include topics such as historical and current trends, clustering and partitioning, and measures of success in knowledge-management systems. Section II summarizes some successful knowledge-authoring technologies, including the DARPA Agent Integration Facility (DAIF). Section III describes intelligent-agent technology to manage knowledge in a command-and-control environment. Section IV presents directions for future research, including a discussion of how technologies can be used to produce future command-and-control systems with advanced capabilities.

I.A. Historical trends in information management

During the 1970s organizations that had computers had a distinct advantage over those that were not digitally automated. Today it is a necessity. During the 1980s, having a database management system was advantageous to many large organizations, the U.S. military included. Nowadays, it is a necessity. Today, the advantages of knowledge-based systems grow with each new application. In the future, knowledge-management systems (KMSs) and knowledge-based systems will become necessities. Many facts that apply to database systems also apply to knowledge-based systems. In terms of theoretical understanding and technological maturity, knowledge management is today where database management was 25 years ago. Given that this is the case, we can exploit some of the trends in knowledge management to enhance the single integrated picture (sometimes called the common-operating picture) in command and intelligence centers. The ultimate goal is to improve the performance of command- and intelligence-center personnel at multiple levels.

I.B. Current trends in knowledge management

Today, one important trend is to exploit existing knowledge and to capture new knowledge rapidly and with as much knowledge reuse as possible. As was demonstrated in the Knowledge Management for Command and Control Systems (KMCCS) project [20, Rubin et al., 2000], we also can generate new knowledge from existing knowledge. In this case, the new knowledge must be validated using a mixed-initiative process. Such mixed-initiative paradigms for knowledge generation have proven themselves to be very successful. (See, for example, [14, Harvitz, 1999]). In this paradigm, the system suggests candidates for new knowledge based on an inferential hierarchy of classes and instances of classes. The user validates the new knowledge and corrects it if necessary. Then the new knowledge is added to the KB.

Another trend in artificial intelligence today is to make expert systems and KBs more accessible to people who are experts in one or more areas, but who are not knowledge engineers or computer scientists. This was the goal of the DARPA RKF program. The objective of the RKF effort was to enable distributed teams of subject-matter experts (SMEs) to enter and modify knowledge directly, easily, and without the need for specialized training in knowledge representation, acquisition, or manipulation. The resulting knowledge bases and the tools that manage them are recommended for transition into operational command centers and can be made

available to provide specific answers to questions and could be applied in many different command- and intelligence-related problem-solving situations.

In a discussion of current trends, one cannot ignore the issue of support from managers in the organization, which can make or break a technical knowledge-management effort. The most successful knowledge-management programs have been those that were supported by the managers of the organization. These managers must support knowledge sharing within the organization. This is a profound statement with multiple aspects, implications, and consequences of failure if this is not accomplished.

1. Managers must provide financial resources for the maintenance and upgrade of the knowledge-management systems in their organization.
2. A chief information officer or some other official in a position of both technical and political leadership must mandate some kind of standardization for knowledge exchange. A common ontology and a common knowledge representation are needed. A common query language also would be useful.
3. Senior management must select someone who can perform an authoritative knowledge-administration role to achieve tightly controlled knowledge federation, assuming that the leadership has considered the alternatives and selected this as policy.
4. Managers must become aware of the possible ways to manage information so that an informed and intelligent policy can be formulated and utilized.

For example, an integrated information system (for example, a KBS) can be tightly or loosely controlled and tightly or loosely coupled. These represent two orthogonal characteristics of KBSs. Managers must declare a policy on how knowledge will be entered and another policy for selecting the knowledge that will be shared. For example, in a loosely controlled KBS, the component systems decide what knowledge to share with the common integrated KB. In a tightly controlled KBS, the knowledge administrator decides what information the components will send to the common KB. The resulting KBs can be very different in the two cases. If an organization's management team wants a tightly controlled KMS, plans must be made to alter policy and to put the appropriate resources in place. This may involve some re-education of the members of the organization to assure they see the value of such a policy. Managers must support it or it will not happen.

Managers' support is critical to a tightly controlled KMS. An organization can have a loosely controlled KMS even with the support of its managers, if that is the official policy. Without any support from the organization's managers, particularly in the area of technical leadership and financial support, an organization with a KBS will have a loosely controlled system in which the component administrators determine the knowledge to share with the larger group. This is because without making provisions for tight control of information sharing, there will be no means to achieve it.

Loose and tight coupling of KMS each has advantages and disadvantages. For example, with tight coupling, a central knowledge administrator can exclude from the common operating picture the information that is not likely to be of any value to the group as a whole. The disadvantage with tight coupling is that a component knowledge administrator may not want to share some of the information that is considered at the component level to be sensitive for fear that the other components may gain some advantage in a competitive situation. Addressing the idea of loose vs. tight coupling, will need to be done to make the most of knowledge-based technologies in the context of an automated, single-integrated picture of the battle space.

1.C. Knowledge management by intelligent agents

The field of intelligent agents is a branch of knowledge management that is an emerging technology due to its usefulness in accomplishing a variety of tasks. Although a standard, unique and precise definition of “intelligent agent” is not available, Singh defines it as “a persistent computation that can perceive its environment, reason and act both alone and with other agents” [21, Singh, 1998]. Its utility in network-centric warfare in general, and in command and control in particular, is promising. In a dynamically changing environment, agents that monitor the battle space and facilitate knowledge management can be instrumental in providing relevant information, which enables the achievement of a greater degree of situational awareness for the entire network. This achieves dominant battle space knowledge over an adversary.

Agents permit data to be analyzed and processed locally so that the results can be sent over a network. This reduces bandwidth requirements because it obviates the necessity to send the entire raw data set. Intelligent agents can be used in command centers to issue alerts when events occur at remote sites that may be of interest to the commander. Intelligent agents also can be used to enhance force-protection and security because agents can monitor sensors, acquire and distribute sensor data, communicate among themselves and decide when an aggregate of events, variables or data constitute a threat to forces and friendly assets. Such analysis based on an aggregate of recently acquired data would be difficult for a human to do and even difficult to achieve with a single agent without the collaboration between agents.

This is one of many examples of how significant technological advantage is gained when technologies are combined, in this case, the technologies being intelligent agents and collaboration. This also occurs elsewhere in knowledge management. For example, agent technology can be combined with clustering and partitioning to reduce the size of the search space for each agent. Clustering and partitioning are covered in the next section.

1.D. Clustering and partitioning in knowledge bases

Large knowledge bases often have axioms that pertain to many different contexts and subject areas. For purposes of information comprehension, organization, integration, and maintenance, the knowledge base can be divided into sections that sometimes are called microtheories, partitions, models, or domains. Axioms that are true in one domain may not be true in another. For example, in the domain of Navy ships, nm is a unit that is used in the measurement of long distances that ships travel. Here nm means “nautical mile” whereas in the domain of molecular spectroscopy, nm means nanometer, a measurement of wavelength. To cite another example, in a microtheory about zoos, elephants do not fly, but in a microtheory about children’s stories and fantasies, they do. The apparent contradictions are resolved by the use of clusters and partitions. The axioms of one domain are not used together in the same line of reasoning with axioms from a totally disjoint domain. For example, it is very unlikely that both uses of nm would appear in the same application. Even more unlikely is the idea that zoo keepers should ever need to concern themselves with elephants that fly. Clustering of similar information in both databases [6, Ceruti, 2000] and knowledge bases [17, Mehrotra and Wild, 1995] also lends itself to better error detection and the resolution of semantic inconsistencies.

The RKF program supported technology that implemented clustering in knowledge bases [17, Mehrotra and Wild, 1995]. Clustering is also very useful in grouping knowledge according to the way it is used independent of the way in which it was entered. Clusters bring together

assertions that have similarities and allow comparisons that can reveal errors. The analysis of properly constructed clusters enables the resolution of discrepancies in the knowledge base and increases the efficiency for groups of experts to interact with the same knowledge base in a collaborative mode.

I.E. Monitoring success in knowledge management

An important precursor to selecting the right metrics for success is the question: “What do you want the KMS to accomplish for the organization that cannot be accomplished more efficiently any other way?” A definition of success is needed. It is very important to decide who will be the users of the system. Not all user groups have the same notion of success. To a security specialist, success is blocking all intrusive attacks to the network. To a scientist dealing with deadlines, success is measured by the speed and efficiency with which tasks can be accomplished. Clearly what constitutes success in one area appears as an impediment in another.

When managers have a clear idea of what the KMS is supposed to accomplish, the right metrics can be selected to assist in ascertaining when the goal has been reached. One metric that has been used in the past, yet that may not be appropriate for all KMSs, is to count the number of axioms or assertions entered into the KB per unit time. When monitored over time, this will show how fast the system can acquire new knowledge. This metric is particularly useful when comparing knowledge acquisition methods or knowledge-entry user interfaces. Another metric is axioms or assertions validated per unit time. Still a third metric is errors detected per knowledge cluster. If this error rate is high, it indicates the need to validate knowledge clusters in the KB. Knowledge clusters can be used to form nuclei for partitions, domains, microtheories and even models. (See, for example, [17, Mehrotra and Wild, 1995].) Any knowledge base of any appreciable and useful size will have errors. As discussed in section I.D, clustering can assist in the detection of at least some of these errors.

Variables that influence the success of a KMS are of two kinds, external and internal. External influences can include political, economic and organizational factors. Internal influences can be technical details, implementation as well as hardware and software selection. No single variable can guarantee the success of a KMS; however, some single variables, when taken separately can guarantee failure. For success, multiple variables must contribute in concert.

Ease of use is important but by its nature is difficult to measure. By definition, metrics must be measurable. User surveys can help to compare ease of KMS use to the ease of use of other software tools. Knowledge reuse is an important metric and criterion for success. If experts who share the same or related ontology and knowledge representation can reuse knowledge, others will think that this knowledge is more important.

With KMS, as with any other technology that affects the way large groups of technical professionals conduct business, management personnel need to become more aware of how their decisions affect technical progress and accomplishments, not to mention operational efficiency.

II. Description of Specific Knowledge-Authoring Technologies

Several innovative knowledge-authoring technologies that were developed or further enhanced in the DARPA RKF program include the following technologies:

- Enhancement of a knowledge-authoring system that includes integrated technologies,
- Novel knowledge-base access such as sketching and mapping techniques,

- Knowledge-base analysis and structure improvement using knowledge clustering techniques,
- Use of partitioning methods to improve query response time,
- Automation of logical explanations of query results,
- Use-case development for testing the technology,
- Distributed knowledge-base connectivity for group collaboration

The subsections below provide details about selected examples of RKF-developed technologies.

II.A. SHAKEN Knowledge-Authoring System

SHAKEN is a knowledge-based system that allows experts to input knowledge into the knowledge base directly without the need for knowledge engineers to encode the knowledge into a machine-readable language. The word, "SHAKEN" is not an abbreviation and does not require further definition. SHAKEN's novel integration technology has made a significant impact on the field of knowledge analysis. (See, for example, [5, Barker et al. 2003]). The SHAKEN legacy includes the following features:

- Multiple knowledge acquisition and reasoning paradigms
- Logic, simulation, pattern analogy, and spatial reasoning capabilities
- Concept maps for graphic entry of ontological information
- Combines theorem proving and simulation
- Multiple input modalities
- Natural language sketch and graphical knowledge entry
- Question answering
- Dialog, knowledge and refinement
- User-friendly tools due to the improved knowledge representation.
- KB diagnostics have been integrated: Stanford's trust analysis and Pragati's Multi - ViewPoint Clustering Analysis (MVP-CA) tool
- Technology to provide pedigree information and explanation of the reasoning process behind SHAKEN's answers and pattern features. (See, for example, [15, Kim and Blythe, 2003]).
- Software is well documented with manuals
- Software and licenses are available

II.B. Analogical processing and nuSketch Battlespace

1. Analogical processing: Developed at Cycorp, KRAKEN [18, Panton, et al. 2002] is another knowledge-authoring system like SHAKEN. The word, "KRAKEN" is not an abbreviation and does not require further definition. This software provides techniques for summarizing analogies, which were used in both SHAKEN and KRAKEN knowledge-authoring systems developed under the RKF program. These systems provided integration of the new technologies into a common user framework to enable experts to formulate, input, critique, correct and update their knowledge. The Analogical processor has a novel interface for knowledge entry via analogy called the Case Mapper. Other researchers outside Northwestern University (NWU) use the Case Mapper, including the Navy Research Laboratory. NWU also developed a next-generation federated reasoning system that uses analogical reasoning. This system's improved reasoning source architecture enables tight integration of analogical reasoning and multiple kinds of visual

and spatial reasoning with the more typical kinds of logical inference used in reasoning systems. The federated reasoning system is used in the Case Mapper, nuSketch Battlespace, and sKEA (another tool not quite ready for transition). Researchers at NWU demonstrated that a SME can be used successfully with sketched inputs, in nuSketch Battlespace and sKEA. (See, for example, [12, Forbus et al. 2001] and [13, Forbus et al. 2003]). For example, Prof. Ken Forbus developed a preliminary account of skolem resolution, which is an operation in analogical reasoning where new entities are identified by mining an analogy for information about them and then using other kinds of reasoning (including logical and spatial reasoning) to find or create such entities as necessary. In the case of Battlespace, for example, qualitative spatial constraints on the success of a task are mapped via analogy, and new corresponding entities in the current situation constructed by position- and path-finding visual operations. Also developed was a "Cyc-native" Analogy Server, as the next step toward tightly integrating analogical processing technology into the Cycorp, Inc. knowledge base. (N.B. Cyc is the knowledge-based system used in KRAKEN, mentioned above.)

2. nuSketch Battlespace: The nuSketch approach to multi-modal interfaces relies on rich visual and conceptual understanding of what is sketched, rather than the pattern-recognition technologies that most systems use. NWU researchers have demonstrated that this approach has several advantages:

- It can be used to create a sketching system that military personnel can learn to use quickly (in about two hours), without the onerous and repetitive training of statistical recognizers (e.g., speech and gesture recognizers) required by other approaches.
- The underlying conceptual representations that nuSketch uses enable it to serve as a front end to knowledge-based systems; this has been demonstrated by its integration into both the KRAKEN and SHAKEN knowledge-formation systems as a sketching front-end, as well as with Bolt Beranek and Newman (BBN)' s system in an experiment with active-duty military personnel.
- Visual computations and qualitative spatial representations can provide useful spatial reasoning abilities for course-of-action reasoning. This includes an initial demonstration of constructing enemy-intent hypotheses via mixed visual and conceptual analogies.

II.C. Multi-ViewPoint Clustering Analysis (MVP-CA) technology

To build reliable knowledge-based systems and utilize them maximally, the knowledge in the system must be abstracted, structured, and otherwise clustered in a suitable manner that facilitates its understanding, verification, validation, maintenance, management, testing and interoperability. The Multi-ViewPoint Clustering Analysis (MVP-CA) technology utilizes clustering analysis techniques to group rules of a knowledge base that share significant common properties. The MVP-CA tool is an analysis aid that exposes the semantics of the knowledge by clustering axioms in the knowledge base [17, Mehrotra and Wild, 1995]. The approach used here clusters the knowledge base into meaningful units to ease the task of software engineering. These clusters can form a basis for understanding the system both hierarchically (from detail to abstract) and orthogonally (from multiple viewpoints and perspectives).

The aim is to provide a semantic-mediation tool that enables knowledge engineers and SMEs, such as commanders, to perform the following actions:

- familiarize themselves rapidly with the terms and concepts in a knowledge base,

- exploit and reuse preexisting knowledge,
- merge and align concepts across different knowledge bases reliably and efficiently
- expose inconsistent and redundant axioms

The contribution of the MVP-CA technology falls into three broad categories related to the different stages: *development, maintenance and interoperation*, of a large-scale, knowledge-based development effort. During the *development phase* MVP-CA technology offers ontological engineering support for KBs by drawing attention to common functionalities as well as overlapping contexts across clusters of axioms in KBs authored by same or different teams. Important knowledge representational issues can be addressed. Higher-level knowledge entry aids can be formulated. The MVP-CA technology demonstrates the following features:

- contextual exposition of domain terms leading to formulation of intermediate concepts in the ontological hierarchy,
- legitimate alternatives on the placement of concepts in the ontological hierarchy,
- opportunities for component identification from clusters of axioms
- template formation from clusters of structurally similar axioms
- inverse representations on higher level concepts
- creation of higher-order predicates, such as macro predicates in Cyc and
- higher-level axioms in frame-based systems.

Such meta-annotations on existing concepts and axioms in the system can pave the way to meaningful knowledge management and higher-level knowledge entry by knowledge engineers during development of KBs.

During the *maintenance phase* clustering addresses the long-term quality assurance of KBs by providing a meaningful base for knowledge repair on the KBs. Juxtaposition of structurally similar axioms that focus on similar content helps a user identify problem areas very quickly. The problems may range from syntactic to semantic issues. They can be

- simple typographical errors that arise due to cut and paste operations,
- ontology hierarchy placement errors,
- inconsistencies across singly- and multiply-authored axioms,
- redundancies with multiple versions of same concept (both in single and multiple authoring modes)
- semantic gaps in the KB axioms

This approach provides inroads into meaningful verification and validation of KBs. MVP-CA's utility is not limited to KBs per se. It also could be used on XML-tagged documents and databases.

To enable *interoperability* of a command-and-control information system, the MVP-CA tool can be used in a discovery mode. Similarities across different ontologies can be extracted using the same technology as template formulation and redundancy detection. Clustering helps support mapping of terms across ontologies by highlighting terms that are used in similar contexts. Some of the possibilities for the terms are

- lexically distant from but semantically close to each other
- lexically close to but semantically distant from each other
- lexically reversed forms of each other

In the RKF project, the MVP-CA technology was used to analyze both SHAKEN and KRAKEN KBs. The results included findings from each of the categories mentioned above. The analysis provided valuable support for knowledge base quality evaluation by exposing important

ontological issues inherent in building large, reliable KBs quickly. The MVP-CA tool can be useful in augmenting the information-centric base ontologies with a knowledge-centric view by extracting clusters of rules having a functional basis or role in the system. This could form a semantic basis for better utilization and reuse of the system.

MVP-CA technology can pave the way for a cost-effective, high-quality solution towards building large, interoperable KBs rapidly for the purpose of threat awareness in the command-center environment. For example, MVP-CA technology can assist future efforts to research and develop fusion of human awareness and automatic processing of information – to detect changes in the body of modeled knowledge. Specifically, MCP-CA can be used to group knowledge and thus sort out similarities, conflicts, differences, and assess completeness of the knowledge. By grouping knowledge according to the way an analyst wants to use it as opposed to the way in which it was received into the KB, patterns can be discovered that would not otherwise be readily apparent.

II.D. DARPA Agent Integration Facility (DAIF) Status

The establishment of this facility involved not only accumulating hardware and software but also becoming familiar with the RKF-developed software, documentation, and installation procedures. Requirements were defined for the processor to run SHAKEN and nuSKETCH. The facility has a new PC running Windows XP, a flat-panel LCD, a scanner, and a Wacom graphics tablet. Internet downloads of software were performed from various RKF contractor internet sites, especially the SHAKEN products. A second PC obtained from in-house equipment will be part of the test bed to run Linux in support of the KRAKEN software. Linux version 7.1 and ARC MAP version 8.3 with Military Analysis and Spatial Analysis modules would be useful additions to the software, funds permitting.

A. nuSketch Battlespace is installed as well as the User's Manual and a large quantity of other documentation. The executable is straightforward to download and install on Windows from the NWU web site, which is protected by a password.

B. The federated reasoning system **Analogy Server**, the user's manual and other documentation also are available. Here again, the executable is straightforward to download and install on Windows from the NWU web site. Both SHAKEN and KRAKEN interface versions are available.

C. The **SHAKEN** software is easy to download from the SRI International web site, which is protected by password. Documentation is available. The Library version requires a Franz Common Lisp License. Franz Lisp has been installed. (N.B. SRI previously was an acronym that stood for something but due to a court decision, it is not allowed to stand for that any longer. Therefore, legally it cannot be used as an acronym and is considered simply to be the name of an independent company that no longer retains its former association.)

D. JTP (Java Theorem Prover from Stanford University) - To run JTP, it is necessary to download Java for Windows from the Sun web site. The latest version 04-18-03 is available but a new version expected. The executable was easy to download and install on Windows from

Stanford's web site. The download includes source code, MS-DOS, and Unix interface versions. No password is required to access the web site.

E. Multi-Viewpoint Clustering Analysis (MVP-CA) tool - The MVP-CA tool is not a stand-alone product and no simple download access exists. However, some component files and complex installation instructions for the SHAKEN version 2.57 are available. (At the time of this writing, Shaken 3.2 is the latest version.)

F. Concept Map SHAKEN Applet from the University of West Florida - We have the draft manual that was in the appendix to SRI March-May SHAKEN status report. The non-knowledge version was downloaded from the University of West Florida web site, but a separate load of concept map is not really necessary because the concept-map products already in SHAKEN are sufficient.

III. Intelligent agents for knowledge management in command and control to support FORCEnet.

III.A. Overview

This section describes how intelligent agents can support FORCEnet as it relates to Network-Centric Warfare (NCW). The focus of this effort was on the data available at the message level of granularity, with a view toward assessing the value added of process automation in the building of the Single-Integrated Picture (SIP), which sometimes is called the common-operating picture, of the battle space.

III.B. FORCEnet and network-centric warfare

Sea power 21 defines a navy with three fundamental concepts: sea strike, sea shield, sea basing, enabled by FORCEnet, which is the operational construct and architectural framework for naval warfare in the information age [10, Clark, 2002]. The goal of FORCEnet is to integrate warriors, sensors, command and control, platforms, and weapons into a networked, distributed combat force [10, Clark, 2002]. FORCEnet integrates warriors, sensors, weapons, and platforms to provide a common-operational picture and to facilitate integrated naval operations and forces that are fully interoperable with other joint forces [16, Mayo and Nathman, 2003].

III.C. US DOD agent research

A number of research efforts are in progress to support FORCEnet applications in the field of intelligent-agent technology, relative to the development of agent architectures, to support the requirements needed for agents to operate in dynamic environments. Because uncertainty and non-determinism characterize command-and-control environments, agent architectures of this type are described below [19, Powers 2001]. This section presents an overview of a research effort that has developed agent architectures designed to provide agents the attributes they need to operate in dynamic environments. Based on these architectures, we discuss examples of how the agents are used in systems that support the human-computer collaboration so essential to today's command centers.

III.D Sensible Agent program

The Laboratory for Intelligent Processes and Systems at the University of Texas at Austin, sponsored by DARPA, is conducting research in distributed agent-based systems. The primary focus is on agent autonomy and multi-agent decision-making interactions in adaptive multi-agent systems. They also specialize in the area of agent-conflict resolution and communication. Of particular interest here is the Sensible Agent Program. The Sensible Agent architecture was designed specifically for use in complex and dynamic problem-solving environments [1, Barber and Martin]. Sensible Agents usually reside in multi-agent systems and can react, adjust and respond to the changing and unpredictable events of dynamic-application environments. Therefore, they are perfect candidates for implementation in command-and-control application domains. The level of autonomy discussed here refers to the degree of independence of an agent' s decisionmaking process in determining how to pursue its goals.

In their paper on agent autonomy [2, Barber and Martin, 1999], Barber and Martin state that "An agent' s degree of autonomy, with respect to some goal that it actively uses its capabilities to pursue, is the degree to which the decision-making process, used to determine how that goal should be pursued, is free from intervention by any other agent." When situations change in the battle space, goals will most certainly change. Thus, the ability of an agent to adjust its autonomy level based on the current situation is imperative. The research has coined the term Dynamic Adaptive Autonomy (DAA), which is the technology behind the ability of a sensible agent to modify dynamically its autonomy level in response to changes in circumstances or occurrences of unexpected events.

Sensible agents have been implemented in military application domains, such as Joint Force Air Component Commander, Naval Command and Control - Call for Fire, and Naval Radar Frequency Management. The naval radar-frequency-management domain consisted of a set of geographically distributed ships, each with one radar. The ships needed to maintain position and frequency relationships to minimize radar interference [3, Barber and Kim, 1999]. The experiment conducted in this domain showed that when situations changed during run time, the most feasible decision-making framework for agents in a multi-agent system varied [4, Barber et al. 1999]. The system featured agents that managed radar frequencies autonomously without human intervention. The experiment highlighted one of the most important features of DAA with respect to their potential use in command-and-control environments - the restructuring of collaborative decision making to exclude failed agents or failed communication links. In the case of a failed communication link, the locally autonomous/master agent assumed responsibility for it, and successfully completed the goal originally assigned to its command-driven agent.

This is significant, as it is highly likely that during a contingency, communication systems might fail resulting in a situation where connectivity between agents would be severed. Sensible agents using DAA can also be implemented in command-and-control application environments to perform monitoring and decision-aiding activities. For example, as assistants to decision-makers, they can evaluate the impact of decisions and dynamically determine and prioritize plan options based on changes in the environment.

In summary, sensible agent DAA provides a problem-solving framework that has been designed to operate in uncertain and unpredictable environments, which makes them ideal candidates for use in command-and-control application environments.

III.E. Intelligent agent in network-centric environments

Software agents can increase dependability in a networked environment in many ways. The emergence of intelligent, agent-based, adaptive software greatly improves Joint Task Force capabilities at the operational level of war by providing decision support for both planning and execution. Agents that continuously monitor the events in the operational environment can assist in providing information for staff planners to conduct threat analysis, terrain analysis, asset scheduling and tracking, route planning, logistical planning, fires coordination, communications planning, force protection planning, and coordination with allied troops and coalition forces.

A networked environment such as FORCEnet, allows information to be shared more easily and consequently enables more collaborative interactions to occur. Thus, creating a natural habitat for agent societies to exist and interact to achieve some common end goal has the potential to increase the dependability of FORCEnet. For example, agents in a multi-agent system could monitor and generate alerts when enemy units come within a specified radius of its position. Other agents could be responsible for determining the best weapon available, deliverable and acceptable in response to Call-for-Fire (CFF) messages. Other agents could monitor CFF messages, and issue alerts when the rules of engagement were violated and enemy or friendly units were directly or indirectly targeted. Still other agents assigned logistical duties could monitor the level of friendly supplies. Moreover, agents also can monitor environmental data in the battle space for indications of biological, chemical or nuclear weapons.

In addition to collaborating among themselves, agents can collaborate with commanders to provide continuous tailored decision support as the situation changes, and assist in rapid re-planning to deliver the right supplies and equipment to the correct place, in the right quantity and at the right time. Mobile agents can be deployed in parallel to retrieve the same information. This avoids a single point of failure. Agents can be used to perform tasks in real-time that are difficult for humans to do accurately in a timely manner. (See, for example, [7, Ceruti, 2001] and [8, Ceruti and Kaina, 2003]).

A networked environment also promotes and facilitates the acquisition, sharing, and application of knowledge. Intelligent agents offer the potential for large volumes of data to be collected, processed, and displayed without overloading users. Further, traversal, retrieval and processing of data in the environment, leads to information at a higher level of aggregation. This has the direct effect of reducing bandwidth requirements because agents can perform some of the required data processing locally, thus sending only finished data products over networks to command and intelligence centers. The result is more efficient situational awareness in both afloat and ashore command centers all over the battle space.

IV. Directions for future research and transition applications

Future work in the area of analogical processing could include the following enhancements. Knowledge entry via analogy needs to be pursued seriously in both the SHAKEN and KRAKEN integrated knowledge-authoring system. The Case Mapper needs to be developed to the point of expert use (which would require restricted natural language and/or concept map interfaces for statement editing and predicate specification). Aside from similarity and/or difference summarization, the other ways that analogy could be used in knowledge formation need significantly more exploration.

The area of sketching tools could use the following improvements. More Battlespace reasoning needs to be built into nuSketch Battlespace. This includes improved pathfinding, combat-power calculations and integrating NWU's detailed trafficability domain theory from previous DARPA programs. This will provide the knowledge infrastructure for additional support tools that our military users have requested. Sketching tools could benefit from integration with sources of real-time data feeds, such as simulations, and the ability to communicate with wargames. The current enemy-intent hypothesis system must be improved in several ways. First, more Battlespace reasoning is needed to construct better explanations about why actions were successful. This would improve the accuracy of the hypothesized actions.

As part of a future research effort to upgrade command-center capabilities, an integration plan can be developed for combining the MVP-CA tool with the intelligent-agent technologies. Moreover, tighter integration of MVP-CA with SHAKEN can provide users more capabilities and options for knowledge analysis in a very user-friendly manner.

Pattern recognition, (e.g. threat awareness) depends on getting the right pieces of the puzzle together at the right time. Success in the battlespace depends on being able to predict events based on the state of the knowledge available. Future work in this could be to research ways to characterize the state of information systems and to correlate this state with observed events. Statistical as well as deterministic methods could be used to summarize, characterize and correlate state information.

Technologies that support data fusion through the use of knowledge-based systems and intelligent agents need to be investigated further. This paper describes the MVP-CA tool that is used to form clusters of axioms in knowledge bases [17, Mehrotra and Wild, 1995]. One direction of further inquiry could be to determine how clustering techniques can be applied to agent clustering of data into data sets for fusion algorithms, and the automatic retrieval of data using intelligent agents for fusion tasks via intelligent agents.

On the applications side, the following RKF technologies are recommended for transition into operational C2 systems:

- SHAKEN Knowledge-Authoring System, which has multiple knowledge acquisition and reasoning paradigms.
- Natural-language sketch and graphical knowledge entry with logic, simulation, pattern analogy, and spatial reasoning capabilities.
- Integrated knowledge-base diagnostics for trust analysis and knowledge-base improvement.
- Technology provides pedigree information and explanation of the reasoning process behind the system's answers and pattern features.

Acknowledgements

The authors thank the Defense Advanced Research Projects Agency, the Office of Naval Research and the Space and Naval Warfare Systems Center, San Diego, Science and Technology Initiative for their support of this work. This paper is the work of U.S. Government employees performed in the course of employment and no copyright subsists therein. It is approved for public release with an unlimited distribution.

References

- [1] K. S. Barber and C. E. Martin, "Autonomy Reasoning in Sensible Agent-based Systems," The University of Texas at Austin, Austin, TX.
- [2] K. S. Barber and C. E. Martin, 1999. "Agent Autonomy: Specification, Measurement, and Dynamic Adjustment," The University of Texas at Austin, TX, Technical Report TR99-UT-LIPS-AGENTS-09, May 1 1999.
- [3] K.S. Barber and J. Kim, 1999. "Constructing and Dynamically Maintaining Perspective-based Agent Models for Command and Control Applications in a Multi-Agent Environment." DARPA-JFACC Symposium on Advances in Enterprise Control, November 1999.
- [4] K. S. Barber, A. Goel, and C. E. Martin, 1999. "The Motivation for Dynamic Adaptive Autonomy in Agent-based Systems," in *Intelligent Agent Technology: Systems, Methodologies, and Tools. Proceedings of the 1st Asia-Pacific Conference on IAT*, Hong Kong, December 1999, J. Liu and N. Zhong, Eds. Singapore: World Scientific, pp. 131-140, 1999.
- [5] K. Barker, J. Blythe, G. Borchardt, V.K. Chaudhri, P.E. Clark, P. Cohen, J. Fitzgerald, K. Forbus, Y. Gil, B. Katz, J. Kim, G. King, S. Mishra, C. Morrison, K. Murray, C. Otstott, B. Porter, R.C. Schrag, T. Uribe, J. Usher, and P.Z. Yeh, 2003. "A Knowledge Acquisition Tool for Course of Action Analysis," *Proceedings of the American Association for Artificial Intelligence (AAAI)*, 2003.
- [6] M.G. Ceruti, 2000. "Restricting Search Domains to Refine Data Analysis in Semantic -Conflict Identification," *Proceedings of the 17th Armed Forces Communications and Electronics Association (AFCEA) Federal Database Colloquium and Exposition*, pp. 211-218, San Deigo, CA, September, 2000.
- [7] M.G. Ceruti, 2001. "Mobile Agents in Network -Centric Warfare," *IEICE Trans. Commun.* vol. E84-B, no. 10, pp. 2781-2785, October 2001.
- [8] M.G. Ceruti and J.L. Kaina, 2003. "Enhancing Dependability of the Battlefield Single Integrated Picture through Metrics for Modeling and Simulation of Time-Critical Scenarios," *Proceedings of the Institute of Electrical and Electronics Engineers (IEEE) Ninth International Workshop on Real-time Dependable Systems, (WORDS 2003F)*, October 2003.
- [9] M.G. Ceruti, 2004. "States of Matter, Information Organization and Dimensions of Expressiveness," to appear in the *Proceedings of the Association for Computing Machinery (ACM) International Conference on Computing Frontiers*, April 2004, (CF04).
- [10] V. Clark, ADM, USN, 2002. "Sea Power 21 Series – Part I: Projecting Decisive Joint Capabilities," *Naval Institute Proceedings*, vol. 128, no. 10, pp. 32-41, October 2002.
- [11] T.H. Davenport and L. Prusak, 1998. *Working Knowledge: How Organizations Manage What They Know*, Boston, MA, Harvard Business School Press, 1998.
- [12] K.D. Forbus, R.W. Ferguson and J.M. Usher, 2001. "Towards a Computational Model of Sketching," *Proceedings of IUI'01*, pp. 77-83, Santa Fe, NM, January, 2001.
- [13] K.D. Forbus, J. Usher and V. Chapman, 2003. "Qualitative Spatial Reasoning about Sketch Maps," *Proceedings of the AAAI IUI'03*, pp. 61-68, Miami, FL, 2003.
- [14] E. Harvitz, 1999. "Principles of Mixed-Initiative User Interfaces," *Proceedings of ACM CHI '99*, pp. 159-166, Pittsburgh, PA. USA, May 1999.

- [15] J. Kim and J. Blythe, 2003. "Supporting Plan Authoring and Analysis," *Proceedings of the ACM Intelligent User Interface Conference*, Miami, FL. Jan. 2003.
- [16] R.W. Mayo, Vice Admiral (VADM), United States Navy (USN) and J. Nathman, VADM, USN, 2003. "Sea Power 21 Series – Part V: FORCEnet: Turning Information into Power," *Naval Institute Proceedings*, vol. 129, no. 2, pp. 42-46, February 2003.
- [17] M. Mehrotra and C. Wild, 1995. "Analyzing Knowledge-Based Systems Using Multi-ViewPoint Clustering Analysis," *J. of Systems and Software*, vol. 29. no. 3, pp. 235-249, 1995.
- [18] K. Pantou, P. Miraglia, N. Salay, R.C.Kahlert, D. Baxter, R. Reagan, 2002. "Knowledge Formation and Dialogue Using the KRAKEN Toolset," *Proceedings of the AAAI*.
- [19] B. Powers, 2001. "Adaptive Intelligent Agents in Informative Environments: Human-Computer Collaboration in Command and Control Application Environments," Space and Naval Warfare Systems Center, San Diego Biennial Review, TD 3117, pp. 203-207, San Diego, 2001.
- [20] S.H. Rubin, M.G. Ceruti and R.J. Rush, 2000. "Knowledge Mining for Decision Support in Command and Control Systems," *Proceedings of the 17th AFCEA Federal Database Colloquium and Exposition*, pp. 127-133, San Deigo, CA, September, 2000.
- [21] M.P. Singh, 1998. "Agent communication languages: Rethinking the principles," *IEEE Computer*, vol. 31, no. 7, pp. 26-47, Dec. 1998.
- [22] K. Wigg, 2003. *Knowledge Management Foundations*, ed. S. Press, Arlington, 2003.

Author biographies

Dr. Marion G. Ceruti is a scientist in the Advanced Concepts and Engineering Division of the Command and Control Department at the Space and Naval Warfare Systems Center, San Diego. She received the Ph.D. in 1979 from the University of California at Los Angeles. Dr. Ceruti's professional activities include information systems research and analysis for command and control decision-support systems, sensor fusion, and research management. She is the author of over 70 journal articles, conference proceedings, monographs and book chapters on various topics in science and engineering. She is a senior member of the IEEE and a member of the IEEE Computer Society, the ACM, the AFCEA, the International Society for Computers and Their Applications, and the New York Academy of Sciences.

Mr. Dwight R. Wilcox is a computer scientist in the Engineering Division of the Command and Control Department at the Space and Naval Warfare Systems Center, San Diego. His professional interests include many aspects of software and systems development.

Ms. Brenda J. Powers is a scientist in the Command and Control Department at the Space and Naval Warfare Systems Center, San Diego. She received the BS degree from the Point Loma University in 1984. She is currently working on her MSSE thesis at the Naval Postgraduate School in Monterey, CA in the area of collaborative interoperability for decision support. Her professional interests include intelligent agent research in the area of agent-oriented methodologies and ontologies, and object-oriented analysis and design. She is a member of the ACM and Artificial Intelligence Special Interest Group.