Toward an Intelligent Aviation Extranet

Jeffrey M. Bradshaw[1,2,4], Yuri Gawdiak[3], Alberto Cañas[1],Robert Carpenter[2], Jim Chen[3], Robert Cranfill[2], Jim Gibson[3], Kim Hubbard[3], Renia Jeffers[2], Nathalie Mathé[3], Luis Poblete[2], Tom Robinson[2], Amy Sun[2], Niranjan Suri[1], Shawn Wolfe[3], Isabelle Bichindaritz[4]

ABSTRACT[1]

This paper describes the Aviation Extranet project, a joint research and development collaboration between NASA Ames Research Center, The Boeing Company, and the Institute for Human and Machine Cognition at the University of West Florida. The intent of the collaboration is to develop a secure intelligent virtual private network framework based on Internet, software agent, and middleware technologies that can be commercially implemented and used by the aviation industry, the FAA, and by NASA. The architecture of the Aviation Extranet is described, with emphasis on the roles of software agents and mediating representations in defining a human-centered implementation approach.

Keywords

software agents, CORBA, extranet, virtual private networks, mediating representations, concept maps

INTRODUCTION

Commercial Aviation Maintenance Challenges

The aviation industry has developed a safe operating environment by continually refining the processes of operations and maintenance. However, it faces mounting pressures to be even more safe and efficient while incorporating significant technological changes. Moreover, despite major advances in equipment efficiency and reliability the industry is experiencing an increase in maintenance costs. For example, one airline recently found that no-fault-found part removals in maintenance situations alone cost them $12M over the course of a year. Maintenance personnel are challenged by increased systems complexity and the lack of ready-to-hand information and assistance.

Commercial Aviation Operations Challenges

The projected growth in air traffic over the coming decade will strain our already congested air traffic and ground management systems. A recent study by the Air Transport Association (ATA) highlights the major airline costs associated with aviation operations problems. According to the study, the total loss to ATA member airlines is $3.5B per year, with taxi delays ($1.6B) and indirect routes ($1.3B) being the major drivers. By broadening the picture to include other inefficiencies associated with the current air traffic control (ATC) system, one major airline alone estimates costs in excess of $2B annually ($700M directly attributable, plus $1,300M in lost productivity). One other major airline has reported similar internal estimates. The ATA study points to the centralized control and lack of distributed information in the current air traffic control (ATC) system as being the major source of these problems.

The Growing Information Pipeline

The most common response to these maintenance and operations problems is to say that aviation support personnel need more and better access to required information. In response to this perceived need, the aviation industry is spending hundreds of millions of dollars to produce and distribute digital maintenance data in a more usable form. On the operations side, the Federal Aviation Administration (FAA) has several initiatives underway that aim to improve real-time information access by upgrading the computing, network, and information infrastructure at airports nationwide. Add this to the burgeoning amount of information that can be made available through the public Internet and private intranets, and one gets the picture of a prodigious pipeline that might ultimately be able to provide just about any information one could want, anytime and anywhere.

Problems with Overabundant, Ubiquitous Information

Unfortunately, while it may be true that the universal information pipeline for aviation maintenance and operations is just around the corner, it must also be stated that overabundant, ubiquitous information can create its own set of problems. The amount of aviation-relevant data on private networks and the public Internet has grown so large that it is impossible to compass it by simple browsing methods. The fact that information is ubiquitous means that the hard problem has become not how to make data accessible, but rather how to control and secure it. Most important of all for the end user, we must find ways to better filter and present the wealth of information in a manner that is contextually relevant.

---

[1] 1. Institute for Human and Machine Cognition, University of West Florida, Bldg. 79/Rm. 196, Pensacola, FL 32514, jbradsha@ai.uwf.edu, (850) 474-3068; 2. Applied Research and Technology, Shared Services Group, The Boeing Company, P.O. Box 3707, M/S 7L-44, Seattle, WA 98124, jeffrey.m.bradshaw@boeing.com, (425) 865-6086; 3. NASA Ames Research Center, Code IP, Moffett Field, CA 94035; 4. Clinical Research Division; Long-Term Follow-Up, FB-600, Fred Hutchinson Cancer Research Center, 1124 Columbia Street, Seattle, WA 98104. Since requirements for core extranet technology are largely generic, we have been able to share results between projects in the aviation and health care fields. We acknowledge NASA's support through the Aviation Extranet Joint-Sponsored Research Agreement, and support from the Agency for Health Care Policy and Research through grant R01 HS09407 to the Fred Hutchinson Cancer Research Center.

An additional challenge is to develop effective methods for capturing and feeding back important aspects of real-time data flows for tracking purposes. The patchwork nature of airport systems and the emphasis on one-way information delivery rather than two-way information exchange currently inhibit such efforts. This fragmentation means, for example, that most airlines have difficulty in tracking the exact status of each aircraft, including such crucial information as surface and airborne operations/cycles per aircraft. The difficulty of accessing and integrating such information is not only a handicap to the airlines, but also to airframe manufacturers and the FAA who could use this data as feedback to enable better design of aircraft and of regulatory policy. To fully exploit the potential of abundant data and powerful technologies, we must begin to transform the dominant paradigm from that of information access to performance support; from data-centered computing to *human-centered computing.*

THE AVIATION EXTRANET COLLABORATION

Human-Centered Computing

NASA Ames Research Center (ARC) has recently identified human-centered computing as one of its three major information technology cornerstones.[2] As a complement to research aimed at developing technology for spacecraft and space vehicle autonomy, researchers have begun inventing and deploying sophisticated computational aids designed to amplify human cognitive and perceptual abilities and enable mixed-initiative human-machine interaction. Ken Ford, Associate Center Director for Information Technology at NASA Ames describes the new emphasis as follows:

> From an AI perspective, human-centered computing is aimed at building computational systems that amplify human intelligence, not substitute for it; the goal is not an *artificially* intelligent computer that would pass the Turing Test, but an *actually* intelligent system that wouldn't.
>
> Working in a human-centered framework the goal is to design and build "prosthetic devices" that are nothing like us— and that's the point— what would be more useless than a dry computer that was truly just like us wet ones? No advantage in that. This shift in perspective places human/machine interaction issues at the center of the subject and requires a "systems view" in which the system in question isn't "the computer" but instead includes the wet computer, the dry computer, and the larger system of which they are both part… The prostheses metaphor implies the importance of designing systems that fit the human and machine components together in ways that synergistically exploit their respective capacities [14].

The Role of NASA in Aviation Safety

What many in the public do *not* realize, is that the first "A" in the acronym "NASA" stands for aeronautics and that in addition to being a space agency, NASA is an *aeronautics* agency, and as such, has set ambitious goals for improving aircraft safety and design. Consistent with President Clinton's announcement of a major initiative to increase the safety and capacity of the aviation transportation system NASA, in collaboration with the FAA, is developing advanced performance support systems that will play a major role in realizing the twin goals of safer aircraft operation and higher throughput of the airport and ground control infrastructure.

One aspect of the effort is supported by a NASA Ames/FAA alliance to build a full-mission, real-time simulation facility that will be used to validate future airport surface and air traffic control tower technologies prior to field deployment. A second prong of this effort is the Aviation Extranet, described in the following section.

Overview of the Aviation Extranet Concept

The Aviation Extranet will provide integrated real-time aircraft and airport data to the airlines (for operations), manufacturers and vendors (for design and maintenance feedback), the FAA (for operations, safety analysis, and research and development), and NASA (for research and development).

Some of the core concepts of the Aviation Extranet approach were developed as part of a rapid prototype in the Surface Movement Advisor (SMA) project (a joint FAA/NASA project) at Atlanta Hartsfield International Airport. The SMA project provides a communication infrastructure that integrates planned, scheduled, and real-time aircraft data based on OAG, FIDS, HOST, ARTS, ACARS, and Airline Ramp and FAA Tower inputs. These data feeds are integrated and provided to all the users at the airport. The SMA prototype currently saves an estimated $50K per day at Atlanta and has the potential to provide commensurate savings at the other major hubs where it is now being deployed.

The Aviation Extranet project is a joint research and development collaboration between NASA Ames Research Center, The Boeing Company, and the Institute for Human and Machine Cognition at the University of West Florida. Building on the success of SMA, the intent is to prototype a secure intelligent virtual private network framework based on Internet, software agent, and middleware technologies that will eventually be commercially implemented and used by the aviation industry, the FAA, and by NASA.[3]

Figure 1 is a conceptual view of the Aviation Extranet architecture. A customized Web client presents interactive aviation information. An open "presentation layer bus" enables the creation of rich client applications. It allows HTML/XML and client-side components (e.g., Java, JavaScript, plug-ins and ActiveX components, ORBs) to share a common object and messaging model, enabling

[2] The concept of human-centered computing is of course not a new one. It has grown out of research traditions that have been going on for decades [5; 23].

[3] A similar architecture is being developed to help support large-scale collaboration between medical staff at the Fred Hutchinson Cancer Research Center and primary-care physicians worldwide [3; 4; 7; 13]

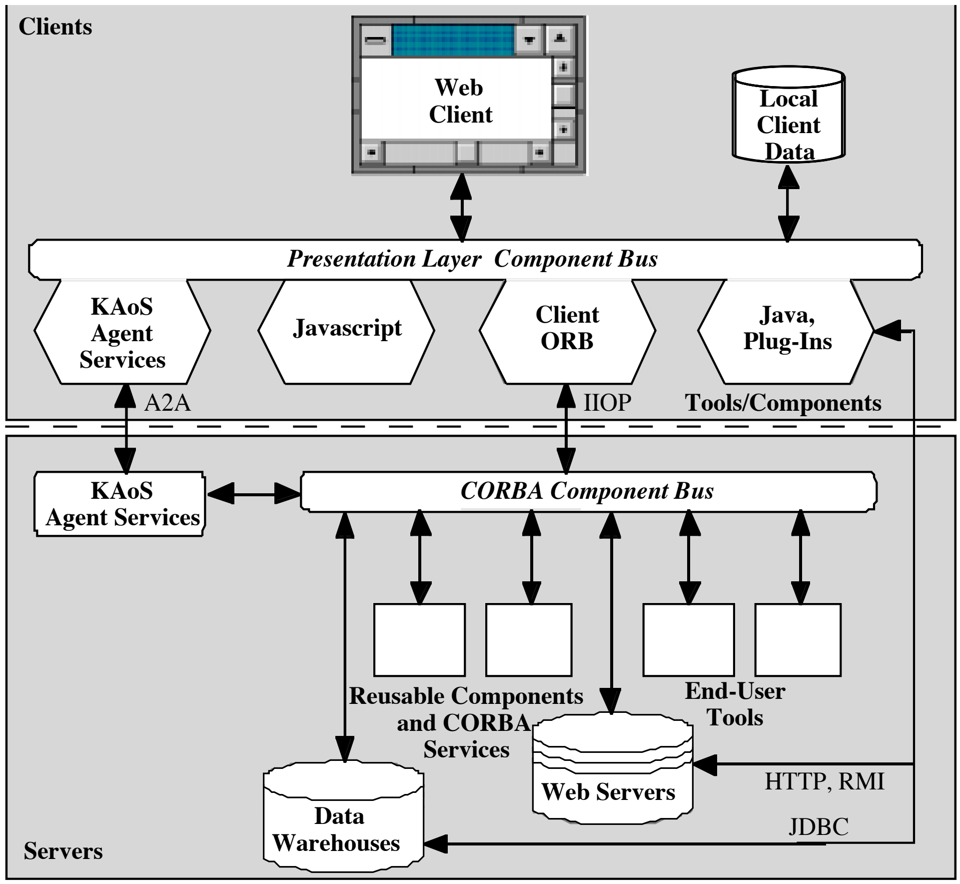

Figure 1. Aviation Extranet conceptual architecture

seamless integration of tools, services, and user-interface elements.

In addition to standard client-server connection protocols such as HTTP, RMI, and JDBC, IIOP connections to a server-side CORBA "component bus" is provided. Through products from various vendors, two-way interoperability between ORBs and Microsoft's COM (Component Object Model) can also be achieved.

End-user tools and reusable components may be implemented as any combination of Java agents or applets, plug-ins, and IIOP-enabled server-side components as desired. Component integration frameworks such as JavaBeans allow the incorporation of ActiveX components that can function as full peers to Java when necessary.

Unfortunately, space limitations require us to limit the discussion in this paper to selected technical aspects of the core extranet infrastructure, neglecting other important topics such as design, testing, and use scenarios for the system. In the next sections we provide brief descriptions of the application of software agents and mediating representations to support a context-sensitive approach to information navigation and delivery:

SOFTWARE AGENTS IN THE AVIATION EXTRANET

Motivation for Software Agents

Software agents have been proposed as one way to help resolve problems in digital information delivery [6]. Simple forms of such agents are springing up in great numbers as part of Internet search and resource discovery services. In the future, more sophisticated agents will become widespread; agents that can seamlessly integrate public and private information sources, personalize the content and presentation of the data for a given user and situation, negotiate and feed data back from the real world to improve future performance (Figure 2).

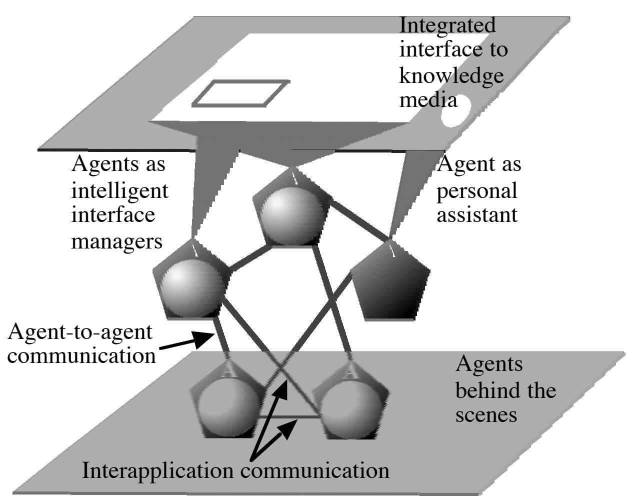

Figure 2. An agent-enabled system architecture.

The KAoS Open Agent Framework

Software agents in the Aviation Extranet are built using the KAoS (Knowledgeable Agent-oriented System) framework [8; 10]. KAoS is an embodiment of lessons learned in the endeavor to investigate principles on which industrial-strength open agent frameworks can be developed. It is unique in that it attempts to draw on both the latest research in intelligent systems and the continuing evolution of robust distributed object, middleware, and Internet technology and standards (e.g., CORBA, Java, COM).

The KAoS framework provides a common set of services to each agent, including registration and naming (as facilitated by a special agent called the *Domain Manager),* and location of other agents providing needed services (as facilitated by *Matchmakers* and *Broker agents).*

Each KAoS agent contains to components: the *generic agent,* which implements behavior that is shared and reused across all agents, and the *extension* which implements agent-specific capabilities. The generic agent includes the basic infrastructure for security, mobility, transport-level communication, and agent conversations.

KAoS differs from most other agent frameworks in that its support for agent conversations explicitly takes into account not only the individual message (e.g., *request, promise*), but also the various sequences of messages in which it may occur. In our current implementation, shared knowledge about message sequencing conventions *(conversation policies)* enables agents to coordinate frequently recurring interactions of a routine nature simply and predictably. *Suites* provide convenient groupings of conversation policies that support a set of related services (e.g., the *Matchmaker suite*). A starter set of suites is provided in the architecture but can be extended or replaced as required. With respect to security, we have begun work on explicit declarative *security policies* to enable communicating agents (and their hosts) to interact according to the safety and confidentiality requirements of their designers. Ongoing work on conversation and security policies and on a high-level agent design toolkit is described in [10; 25].

Within the Aviation Extranet framework, KAoS agents are being developed to fill the following roles:

1. Managing dynamic loosely-coupled services and information sources. A major challenge of building and maintaining dynamic loosely-coupled distributed systems is finding the required information sources and computing services on-the-fly. We have been using agent technology in two different ways to address this problem: 1. Matchmakers, whose major function is to help client agents find information about agents within the domain that have advertised their services, and to forward such requests to Matchmakers in other domains where appropriate; and 2. Agent-assisted external linking. Unlike the typical "launch-and-forget" interaction between applications or plug-ins supporting hyperlinks, each active component is assigned one or more agents to be aware both of what is going on in the application and what is going on in the rest of the agent world. Thus, once links are established and traversed, back-channels of communication can be used to keep all active applications and documents "in synch" with the current context. The new XML link model greatly facilitates this process.

2. Providing a secure unifying framework for distributed heterogeneous components. We think of the agents as being sorted into functional layers: presentation services, application services, generic agent services (e.g., naming, security), and data services. The *Domain Manager* controls the entry and exit of agents in a domain according to policies set by the domain administrator. We are in the process of developing certificate-based authentication and encryption capabilities for agents. We are also looking at the possibility of allowing developers to optionally layer KAoS services on top of emerging standards such as the OMG Naming, Trader, and Security Services, and the Mobile Agent Facility.

3. Coordinating interaction at the knowledge-level. Though the lowest-common-denominator methods provided through the component bus will be adequate to enable interoperability between most system components, some specialized intelligent software modules require a higher level of communication semantics than can be directly supported by IDL alone. As has been frequently argued [8; 18; 19; 20; 21], there is still much work to be done on "knowledge-level" methodologies and standards that can ensure that the operational semantics (and pragmatics) of these components are explicitly represented.

KAoS agents on the client and on the server can communicate using an agent-to-agent (A2A) protocol that runs on top of standard transport-level protocols such as IIOP, RMI, COM+, and sockets. Generic agent services built on the foundation of existing distributed object services allow software components the option of using a common agent-to-agent interlingua (an "agent bus") to communicate and coordinate their actions at the "knowledge-level" rather than relying on more primitive program-to-program protocols. Shared ontologies provide a common vocabulary for collaboration on problems that span different toolsets and information sources.

In addition, KAoS agents are currently being developed to support the notion of a dynamically scalable client. Because agents can transparently migrate from host to host of their own volition, pieces of the application can incrementally be offloaded to the client or the server depending on the characteristics such as processor speeds and network bandwidths. Additionally, user interface agents can begin to adapt the presentation of the client application according to parameters such as screen size and I/O modalities (e.g., pen vs. keyboard).

CONCEPT MAPS AS A MEDIATING REPRESENTATION

One important aspect of a human-centered approach to aviation information delivery systems is the design of effective modes of presentation and navigation. Within the Aviation Extranet, *mediating representations* (e.g., repertory grids, concept maps, activity graphs) are designed to reduce the problem of representation mismatch, the disparity between a person's natural description of the problem and its representation in some computable medium [9; 15]. Mediating representations provide a bridge between informal diagrams or verbal reports and typical knowledge representation schemes such as production rules and semantic networks. The design of a mediating representation is optimized for human understanding rather than mere machine efficiency. Effective mediating representations make important things explicit while hiding unnecessary detail. They expose natural constraints, facilitate computation and analysis, and are rich enough for comfortable expression of knowledge in a given domain [26]. The choice of representation can have an enormous effect on human problem solving performance [22].

One important type of mediating representation that has received much attention in the last few years is the concept map. Concept maps were developed by Novak [24] on the foundation of Ausubel's assimilation theory [1; 2]. A concept map is an ideal mediating representation for the human-centered presentation of contextualized knowledge. It is a two-dimensional representation of a set of concepts that is constructed so that the interrelationships among them are evident. Their function is to help people clarify, organize, share, and explain their own conceptual framework, and to use it to find relevant information in context and collaborate with others.

Concept maps are being extended beyond knowledge representation, to serve as the navigation interface to a domain of knowledge. Based on many years of innovative research on concept map tools and applications [11; 12; 16; 17], the IHMC (Institute for Human and Machine Cognition) at the University of West Florida has been awarded a patent on enabling technology for next generation multimedia browsers. The approach permits the user to "browse" a multimedia (e.g., digital video, text, audio, pictures) model of expert domain knowledge. The research is aimed at avoiding several key problems associated with traditional browsers, including the navigation problem in which users find it difficult to stay oriented in a large scale, dynamic context-free information environment. The navigation problem is alleviated by the use of concept maps as a mediating representation facilitating the traversal of logical linkages among clusters of related objects.

Software agent technology will be used to complement the use of mediating representations in the Aviation Extranet. We are applying and extending the results of previous research to increase the effectiveness of tools for the design, validation, and deployment of performance support systems. Concept maps will be employed as shared cognitive artifacts on which both humans and agents can operate, facilitating collaboration at the knowledge level. Development tools for the design of plans, and conversation and security policies are being investigated as a way to help people without specialized training create domain-specific agents [10]. The web-based performance-support-oriented browser will make both the development and deployment tools available in a universal client interface.

STATUS AND FUTURE DIRECTIONS

We have recently completed the initial phase of the prototype, demonstrating KAoS agent-mediated access of Boeing maintenance databases and live operations data from a Java-based web application. The interfaces to maintenance and operations data are exposed to the agents through CORBA. The agents use the registration, advertising, and recruitment capabilities of the Domain Manager and Matchmaker to establish and locate needed resources.

Developing appropriate security mechanisms to seamlessly connect Internet clients to confidential data behind one or more firewalls has been a major focus of the initial effort. User authentication and encryption is based on public-key certificate (digital ID) technology and Secure Sockets Layer (SSL) protocols. At login time, a "personal certificate" establishes the identity of the user to a particular server. Depending on the administrative policy, this "single sign-on" may grant access to additional servers provided that they also recognize the authority who signed the certificate. Access rights to particular information and services can be selectively granted based on properties of personal certificates. Other types of certificates are used to verify the identity of servers (site certificates), of message-senders (e.g., S/MIME encrypted email), and of running code (signed code).

We are enhancing KAoS for use in secure web applications. In addition to requiring agent code downloaded to the client to be digitally signed, any use of agents in Aviation Extranet applications will also require that communicating agents mutually authenticate each other before being allowed to engage in conversations. More sophisticated security mechanisms to allow fine-grained resource management and "anytime" mobility are planned for the future [10].

We have recently begun investigation of specific mechanisms for combining agents and concept maps in the development of performance-support-oriented browsing tools. This enhanced technology foundation along with additional data sources and services that will be integrated into the testbed will allow us to tackle successively more ambitious scenarios, such as the use of aviation extranet clients as part of a simulation of portions of the national airspace.

REFERENCES

[1] Ausubel, D. P. (1963). *The Psychology of Meaningful Verbal Learning*., New York: Grune & Stratton.

[2] Ausubel, D. P., Novak, J. D., & Hanesian, H. (1978). *Educational Psychology: A Cognitive View*. (Second ed.).

[3] Bichindaritz, I., Bradshaw, J. M., & Sullivan, K. (1997). Distributed reuse of knowledge in a computerized decision support system for bone-marrow post-transplant care over the World Wide Web. *Proceedings of the 1997 American Medical Informatics Association (AMIA) Annual Fall Symposium,* (pp. 844). Nashville, TN, 25-29 October.

[4] Bichindaritz, I., & Sullivan, K. (1998). Reasoning from knowledge supported by more or less evidence in a computerized decision support system for bone-marrow post-transplant care. *Proceedings of the AAAI 1998 Spring Symposium on Multimodal Reasoning,* . Stanford, CA: Stanford University, AAAI Press, March.

[5] Boy, G. A. (1991). *Intelligent Assistant Systems.*, London: Academic Press.

[6] Bradshaw, J. M. (Ed.). (1997). *Software Agents.* Cambridge, MA: The AAAI Press/The MIT Press.

[7] Bradshaw, J. M., Chapman, C. R., Sullivan, K. M., Almond, R. G., Madigan, D., Zarley, D., Gavrin, J., Nims, J., & Bush, N. (1993). KS-3000: An application of DDUCKS to bone-marrow transplant patient support. *Proceedings of the Sixth Annual Florida AI Research Symposium (FLAIRS '93),* (pp. 89-95). Ft. Lauderdale, FL.

[8] Bradshaw, J. M., Dutfield, S., Benoit, P., & Woolley, J. D. (1997). KAoS: Toward an industrial-strength generic agent architecture. In J. M. Bradshaw (Ed.), *Software Agents.* (pp. 375-418). Cambridge, MA: AAAI Press/The MIT Press.

[9] Bradshaw, J. M., Ford, K. M., Adams-Webber, J. R., & Boose, J. H. (1993). Beyond the repertory grid: New approaches to constructivist knowledge acquisition tool development. In K. M. Ford & J. M. Bradshaw (Ed.), *Knowledge Acquisition as Modeling.* (pp. 287-333). New York: John Wiley.

[10] Bradshaw, J. M., Greaves, M., Holmback, H., Carpenter, R., Cranfill, R., Jeffers, R., Poblete, L., Robinson, T., Sun, A., Gawdiak, Y., Cañas, A., Suri, N., Silverman, B., Brooks, M., Wong, A., Bichindaritz, I., & Sullivan, K. (1998). Mediating representations for an agent design toolkit. *Proceedings of the Eleventh Annual Knowledge Acquisition Workshop (KAW98).* Banff, Alberta, Canada, 18-23 April.

[11] Cañas, A. J., & Ford, K. M. (1992). An environment for collaborative knowledge building. *Presentation at the Workshop on the Technology and Pedagogy for Collaborative Problem Solving as a Context For Learning,* . Toronto, Ontario, Canada.

[12] Cañas, A. J., Ford, K. M., Hayes, P. J., Brennan, J., & Reichherzer, T. (1995). Knowledge construction and sharing in Quorum. *Proceedings of the AI in Education Conference,* (pp. 218-225).

[13] Chin, T. (1997). Federal grants will help develop web-based decision support tools. *Health Data Network News*, 20, January 20, 1, 8.

[14] Ford, K. M. (1997). Information technology at NASA: Accepting the challenge of excellence. Invited presentation. *Conference on Information and Knowledge Management (CIKM-97).*

[15] Ford, K. M., Bradshaw, J. M., Adams-Webber, J. R., & Agnew, N. M. (1993). Knowledge acquisition as a constructive modeling activity. In K. M. Ford & J. M. Bradshaw (Ed.), *Knowledge Acquisition as Modeling.* (pp. 9-32). New York: John Wiley.

[16] Ford, K. M., Cañas, A. J., Andrews, E. J., Schad, N., Bezdek, J. C., Brennan, J., & Coffey, J. (1991). NUCES: Nuclear Cardiology Expert System: A preliminary report. *Fourth International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, Kauai, Hawaii, June.

[17] Ford, K. M., Stahl, H., Adams-Webber, J. R., Cañas, A. J., Novak, J., & Jones, J. C. (1991). ICONKAT: An integrated constructivist knowledge acquisition tool. *Knowledge Acquisition*, 3(2), 215-236.

[18] Gaines, B. R., & Shaw, M. L. G. (1996). A networked, open architecture knowledge management system. B. R. Gaines & M. Musen (Ed.), *Proceedings of the Tenth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop,* 2 (pp. 45:1-45:22). Banff, Alberta, Canada.

[19] Genesereth, M. R. (1997). An agent-based framework for interoperability. In J. M. Bradshaw (Ed.), *Software Agents.* (pp. 317-345). Cambridge, MA: The AAAI Press/The MIT Press.

[20] Gennari, J. H., Stein, A. R., & Musen, M. A. (1996). Reuse for knowledge-based systems and CORBA components. B. R. Gaines & M. Musen (Ed.), *Proceedings of the Tenth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop,* 2 (pp. 46:1-46:16). Banff, Alberta, Canada.

[21] Kremer, R. (1996). Toward a multi-user, programmable web concept mapping "shell" to handle multiple formalisms. B. R. Gaines & M. Musen (Ed.), *Proceedings of the Tenth Banff Knowledge Acquisition for Knowledge-Based Systems Workshop,* 2 (pp. 48:1-48:20). Banff, Alberta, Canada.

[22] Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.

[23] Norman, D. A., & Draper, S. W. (Ed.). (1986). *User-Centered System Design.* Hillside, NJ: Lawrence Erlbaum.

[24] Novak, J. D., & Gowin, D. B. (1984). *Learning How to Learn.*, New York: Cambridge University Press.

[25] Smith, I. A., Cohen, P. R., Bradshaw, J. M., Greaves, M., & Holmback, H. (1998). Designing conversation policies using joint intention theory. *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98).* Paris, France, 2-8 July.

[26] Winston, P. H. (1984). *Artificial Intelligence.* (Second ed.), Reading, MA: Addison-Wesley.