

Coordinated Operations in Mixed Teams of Humans and Robots

Matthew J. Johnson, Koji Intlekofer, Jr., Hyuckchul Jung, Jeffrey M. Bradshaw, James Allen, Niranjan Suri, and Marco Carvalho

Florida Institute for Human and Machine Cognition, Pensacola, FL, U.S.A.

Abstract—The goal of this research was to facilitate the formation of mixed teams of humans and robots that could perform complex tasks in the real world. We developed a complete end-to-end system to meet this goal. We used a natural language based multi-modal interface to enable simple interaction for the people on the team. We used policy regulated behavior for the robots to ensure effective coordination from the robotic teammates. We also employed an advanced network infrastructure to enable robust and reliable communications among team members.

Index Terms—Robot and Human Cooperation, Team Formation, Teamwork policies, Agile Computing, Dialogue-Based Collaboration.

I. INTRODUCTION

The use of unmanned systems in the military is growing. In future military scenarios, large numbers of heterogeneous unmanned ground, air, underwater, and surface vehicles will work together, coordinated by a smaller number of human operators. A key requirement for such systems is for real-time cooperation with people and with other autonomous systems. While these heterogeneous cooperating platforms may operate at different levels of sophistication and with dynamically varying degrees of autonomy, they will require some common means of representing and appropriately participating in joint tasks. Just as important, developers of such systems will need tools and methodologies to assure that such systems will work together reliably and safely, even when they are designed independently.

Teamwork has become the most widely-accepted metaphor for describing the nature of multi-agent and multi-robot cooperation. The key concept usually involves some notion of shared knowledge, goals, and intentions that function as the glue that binds team members together [14; 35]. By virtue of a largely-reusable explicit formal model of shared intentions, team members attempt to manage general responsibilities and commitments to each other in a coherent fashion that facilitates recovery when unanticipated problems arise. For example, a common occurrence in joint action is when one team member fails and can no longer perform in its role [14].

A general teamwork model might entail that each team member be notified under appropriate conditions of the failure, thus reducing the requirement for special-purpose exception handling mechanisms for each possible failure mode.

Whereas early research on teamwork focused mainly on interaction among autonomous systems, there is a growing interest in various dimensions of human-agent-robot interaction [4; 7]. Unlike autonomous systems designed primarily to take humans out of the loop, many new efforts are specifically motivated by the need to support close and continuous human interaction [8; 30; 34].

Over the past several years, we have investigated how to facilitate such teamwork among humans, agents, and robots. To lay the groundwork for our research, we have studied how humans (and animals) succeed and fail in joint activity requiring a high degree of interdependence among the participants [13; 17; 18; 26].

There are several important challenges in making automation a team player [27]. These and other considerations played an important role in our design process (Section 2), and to implement a framework specifically designed to support joint activity among humans and robots (Section 3). Among the core components of this framework include an Agile Computing Infrastructure (ACI), the KAoS HART (Human-Agent-Robot Teamwork) services framework, and the TRIPS dialogue-based collaborative problem solving system. We describe the demonstration scenario and results (Section 4) and discuss related work (Section 5). We conclude with a summary and discussion of future research directions (Section 6).

II. DESIGN DESIDERATA

We developed an initial set of design desiderata to guide our work. These included (in no particular order):

- The approach must be suitable for real-time operations involving heterogeneous robots with widely-varying capabilities, and in both simulated and physical environments.
- In addition to teleoperation, the approach must support natural interaction through two-way spoken dialogue and gestures on a GUI.

- Multimodal interaction and the real-time use of maps and video must be supported.
- Robots must be capable of responding to direction from authorized humans whenever required and in whatever mode is most convenient: teleoperation, detailed instructions by voice or GUI interaction, or high-level guidance requiring significant robot autonomy. Transitions between modes of operation must be natural and efficient.
- The approach must work equally well in both local (i.e., with robots in sight of humans) and remote operation.
- In order to scale, the approach must support hierarchies of teams and subteams, which may be dynamic in composition.
- The approach must support the convenient packaging of permissions and obligations as part of team role descriptions. Teams should be able to coordinate behavior with or without team leaders, as the situation requires.
- The communication approach should rely on standard 802.11b protocols but provide a higher level of performance and reliability. We wanted the infrastructure to be able to handle multi-hop transmissions, manage bandwidth and deal with intermittent connectivity

III. DEMONSTRATION FRAMEWORK

Our demonstration framework included three major components:

- Agile Computing Infrastructure (ACI)
- KAoS HART (Human-Agent-Robot Teamwork) services framework
- TRIPS dialogue-based collaborative problem solving system.

A. Agile Computing Infrastructure (ACI)

A challenging problem in dynamic tactical environments is the fact that unmanned vehicles are subject to communication constraints that limit bandwidth and increase latency. In addition, network disconnection is a concern, whether due to vehicles moving out of communications range, communications being obstructed by terrain, or a tactical need to minimize signal transmissions. Finally, communication may sometimes depend on peer-to-peer networks, where one robot communicates with another vehicle by using a third vehicle as a relay.

Agile computing may be defined as opportunistically discovering, manipulating, and exploiting available computing and communication resources in order to improve capability, performance, efficiency, fault-tolerance, and survivability [32]. The term *agile* is used to highlight the desire to both quickly react to changes in the environment as well as to take advantage of transient resources only available for short periods of time.

In our work on the ACI, we have demonstrated advances beyond the standard single hop network techniques to be able to proactively move resources in support a multi-hop ad hoc network. We also can provide filtering and transformation of data along the path, all governed by KAoS policies (see

below). Occasionally robots may move out of range and lose communication. If the robot were to simply move back into range communications would be restored, but our goal was to extend the range beyond single hop communications by recognizing the failure and attempting to proactively restore communications by trying to move an available resource [33].

ACI was also used in this exercise to optimize bandwidth to the varying requirements of different clients by dynamic transformation of video streams from the robots [12; 31].

B. KAoS HART Services Framework

The KAoS HART (Human-Agent-Robot Teamwork) services framework has been adapted to provide the means for dynamic regulation on a variety of agent, robotic, Web services, Grid Computing, and traditional distributed computing platforms [3; 23-25; 28; 30; 38]. It also provides the basic services for distributed computing, including message transport and directory services, as well as more advanced features like domain and policy services.

All team members, human and agent, register with the directory service and provide a description of their capabilities. This enables team members to query the directory service to find specific team members as well as match them based on capability. The domain and policy services manage the organizational structure among the agents, providing the specification of roles and allowing dynamic team formation and modification. A “KAoS Robot” extension [23] provides a generic wrapper for each type of robot and a consistent interface for client systems to access the robots. KAoS Robot enables detailed status monitoring in addition to policy checking and enforcement, providing essential ingredients for coordination.

Policies, implementing coordination constraints, are implemented in OWL (Web Ontology Language: <http://www.w3.org/2004/OWL/>), to which we have added optional extensions to increase expressiveness (e.g., role-value maps) [38]. A growing set of services for policy deconfliction and analysis are also provided [9; 38]. Policies are used to dynamically regulate the behavior of system components without changing code or requiring the cooperation of the components being governed. By changing policies, a system can be continuously adjusted to accommodate variations in externally imposed constraints and environmental conditions. There are two main types of policies; authorizations and obligations. The set of permitted actions is determined by *authorization policies* that specify which actions an actor or set of actors are permitted (*positive authorizations*) or not allowed (*negative authorizations*) to perform in a given context. *Obligation policies* specify actions that an actor or set of actors is required to perform (*positive obligations*) or for which such a requirement is waived (*negative obligations*). From these primitive policy types, we build more complex structures that form the basis for team coordination.

In addition to the considerations mentioned above, our research has been guided by three principles. First, we focus on situations where it is desirable for humans to remain “in-the-loop” and allow the degree and kind of control exercised

by the human to vary at the initiative of the human or, optionally, with automated assistance [5; 6]. Second, we assure that mechanisms for appropriate robot regulation, communication, and feedback in such situations are included from the start in the foundations of system design, rather than layered on top as an afterthought [23]. Third, working in the tradition of previous agent teamwork researchers (e.g., [14; 35]), we attempt to implement a reusable model of teamwork.

The teamwork model for our coordinated operations exercise is implemented within various sets of KAoS policies. The intent of the policies is to provide information to establish and preserve common ground among both human and robotic team members, as well as helping to maintain organizational integrity. The policies are defined and enforced external to any specific robot API, so as new robots join, they automatically acquire all the teamwork intelligence possessed by the other robots. Issues addressed by our coordinated operations policies include the following:

- Providing feedback by acknowledging commands, except those where the relevant actions are directly observable.
- Only accepting commands from superiors in a chain of command. This helps prevent confusion and conflict among team members.
- Providing progress appraisal to the requestor of an action through notification when an action is finished, except when those actions are directly observable [16].
- Returning to base when the mission is complete or aborted and letting your teammates know (cmp. [14]).
- Notifying the team leader when there is a status change.

During the demonstration, we showed how teams, roles, and policies can be dynamically-created, deconflicted, and enforced at run-time.

KAoS HART has a modular architecture so that various reasoning components can be plugged in for new services that extend the existing services and augment the capabilities of team members. One example is a component called Kaa (KAoS Adjustable Autonomy) that is built on the foundation of the current policy mechanisms by performing adjustments of autonomy consistent with policy [5; 6]. The objective of Kaa is to be able to reason about and automatically or semi-automatically adjust autonomy along whichever dimensions (possibility, performability, authorization, obligation) are deemed to provide the most effective result. Kaa is supported by another component, Kab (KAoS Backup Planner), which computes backup plans for a failure situation. Kaa evaluates the suggested backup plans with other adjustment choices. For instance, when a robot is obliged to detect a motion but its camera (a main device for the role) fails, Kab could suggest using a different device (i.e., sonar) and/or using delegating the role to another robot within the current team or from a different team. Then, Kaa computes the utilities of those options as well as removing the obliged role, and selects the best choice based on decision theoretic reasoning.

KAoS Spatial Reasoning service is another component that augments the capabilities of team members. As described in the Section IV, it helps robots (without a sophisticated spatial reasoner) to process human commands with spatial references. By providing such a capability as a service, we can deploy less capable robots as competent team members.

C. TRIPS

The TRIPS system is a domain-independent dialogue-based collaborative problem solving system that has been tested in a range of different applications and domains [1; 19]. The core components of TRIPS include (i) a toolkit for rapid development of language models for speech recognition, (ii) a robust parsing system that uses a broad coverage grammar and lexicon of spoken language, (iii) an interpretation manager (IM) that provides contextual interpretation, (iv) an ontology manager (OM) that translates between representations, (v) a surface generator that generates system utterances to communicate with a user, and (vi) a GUI that allows humans to directly interact with a 2D map and a video display. The IM also draws from the Discourse Context module to help resolve ambiguities in the speech input, and coordinates the synchronization of the user's utterances and observed GUI actions. Based on the chosen hypothesis by IM, a domain-specific reasoner takes actions (e.g., sending commands to a specified robot) accordingly.

In this research, TRIPS provides the capability of dialogue-based interaction with robots, allowing the operators to speak naturally and incorporate GUI gestures to coordinate with the various team members (other humans as well as robots). Since TRIPS is a general-purpose system, a special interface is built to communicate with domain-specific robots. Given user requests, the interface translates the first order logic-based TRIPS internal representation (that corresponds to user commands) to robot-interpretable commands based on OWL. User queries about domain information such as the current team structure are translated into the KAoS Common Service Interface. Another main role of the interface is to update information from robots (e.g., robot action report, robot position, laser/sonar detection, video, etc.) by sending the information to related TRIPS components.



Figure 1. Team members in the demonstration: two humans and five heterogeneous robots

IV. DEMONSTRATION SCENARIO AND RESULTS



Figure 2. Overhead view of Bravo Pier with triangular icons representing robots.

The objective of the scenario we defined for the exercise was to isolate, find, and apprehend a human intruder on a pier. This hide-and-seek style task provided plenty of complexity for addressing teamwork issues. The location for the demonstration was Bravo Pier, an active military pier located at the US Naval Air Station in Pensacola, Florida. The intruder hid amongst the clutter on the pier which included buoys, palettes, blockades and some small structures.

As shown in Figure 1, the available team members consisted of two humans and five robots. The humans were given distinct roles. The “Commander” had the responsibility of establishing subteams and managing the overall search process. Relying on a combined speech and graphical interface the Commander operated remotely without direct sight of the area of operation. The second human played the role of “Lieutenant.” The Lieutenant was assigned to a team just like the robots and he worked in the field generally alongside and in sight of them. He wore a backpack that carried a laptop to provide a similar speech and visual interface as the Commander’s, through a head mounted display. The robot team members included four Pioneer 3AT robots variously equipped with different combinations of sonar, GPS, pan-tilt-zoom cameras, and SICK lasers. The fifth robot was an IHMC-designed and built robot called the tBot, a multi-role reconfigurable robot. All the robots had onboard computers and used wireless routers for communication.

The task began with the Commander stating the task of finding an intruder, and querying for available resources and their capabilities. Given utterances such as “what resources are available” and “who has a laser”, TRIPS queried the KAoS directory service and informed the Commander of the robots and their positions/capabilities with speech and GUI. Then, using natural language, the Commander assigned two teams and their leaders. One team was fully robotic, two robots with one assigned as the leader. The other team was mixed, two robots with the Lieutenant assigned to lead. Acknowledgement policies provide useful feedback, since

there is no external indication from the robots that the team assignment has occurred.

After heterogeneous teams were dynamically created, the Commander defined an area of search interest on his display (shown in Figure 2) and used natural language to task each team to secure a particular side. The autonomous team used KAoS spatial reasoning to determine the location to secure, based on the area previously specified by the Commander. The robotic leader of the autonomous team was also aware that it must take its team members with it and tasked them accordingly. Here, when the commander issued the “secure area” command to the Lieutenant, the commander’s TRIPS system automatically notifies the Lieutenant’s TRIPS system of the area to secure so that the area appeared on the Lieutenant’s map display. After issuing the commands, the commander dynamically created a policy through speech to be notified when each team is in position. This is a normal synchronization tool employed by humans. We used natural language to dynamically create a KAoS coordination policy that enforced the communication requirement. Hence, once in position, the coordination policy took effect, and the robots reported.

With the boundary secure, the Commander directed each team to begin a search of the area. The Lieutenant used natural language to direct his team (shown in Figure 3) for the search. He had the option of using very specific teleoperation style commands like “turn left 45 degrees” or “move forward five meters,” or more qualitative commands like “move forward slowly” or “speed up.” Spatial reference from direct GUI interaction could be also used (e.g., “move here” by clicking a position on the map).



Figure 3 Lieutenant in the field with robots

During the mission, the Lieutenant could dynamically create objects based on sensor data and name them for future reference, for instance, by saying “this is a shed” and specifying the outline of an object on the GUI. Then, the Lieutenant could command a robot using object names as spatial references, like “SILVER, move to the right side the shed.” When computing the destination position (i.e., a GPS point) from “the right side of ...,” robots that were not equipped with a sophisticated spatial reasoner could rely on the KAoS Spatial reasoning services. The reasoner could handle any perspective, because the interface uses a subject and an object. In this case the subject was the robot and the object was the shed. The robot could have taken the Lieutenant’s perspective by simply defining the subject to be the Lieutenant instead.

In this scenario, the intruder was found by either a robot or a human and the other team members were notified. To apprehend the intruder who was hiding in a place with narrow passages covered by obstacles, the Lieutenant tried to use the tBot that was designed to operate in such an environment. However, the robot was not currently assigned to his team, which made the coordination services enforce the chain of command and prevent the action. This negative authorization was reported to the command issuer (i.e., the Lieutenant) with speech (e.g., “you are not authorized for that action”). The Lieutenant then proceeded through the correct chain of command for permission. The Commander dynamically reassigned the tBot to the Lieutenant’s team (e.g., by saying “TBOT, join team alpha”). The tBot seamlessly transitioned from voice teleoperation to joystick control to negotiate the tight obstacles and then apprehend the target successfully.

V. RELATED WORK

Previous work on multi-robot coordination has generally involved situations where the teams are fully autonomous or have a very minimal role for people. Such teams use a variety of approaches including low-level biologically-inspired reactive behaviors [2], algorithmic solutions [22], mission planning software [37], and market-based negotiation [15]. Although the pursuit of fully autonomous systems is a worthy goal, in many situations coordination with people can improve performance (e.g., [21]). In other situations, robots cannot yet be trusted to perform critical tasks on their own and must be teleoperated by one or more people [11]. There have been several examples of designing robots that will be involved in close and continuous interaction with people [10; 20]. Most of these have focused on interaction with one robot, although a few have attempted two or three [29]. Far rarer are studies involving multiple humans, except for the common situation of two or more people teleoperating a single robot [36].

Our work differs from the previous work in several ways. First we insist on allowing humans to remain “in-the-loop” at any level of control they desire. We also allow multiple humans to participate on the team. Finally, we provide a regulatory mechanism to help the robots coordinate in a manner that can enhance performance and safety.

VI. CONCLUSIONS AND FUTURE WORK

One limitation of the current system is the use of a centralized directory service. KAoS has recently added the capability of a distributed directory service, but it was not used for our demonstrations. Some teamwork issues that remain are how to form teams that maximize potential performance, identifying additional coordination policies, and expanding the natural language model to support these enhancements. There are several network issues to address as well including how to handle cross network communications, how to isolate sub teams to manage bandwidth, and how to prioritize message traffic to ensure effective team performance

The innovation in this research is not the hierarchical labeling that normally goes as teamwork, but the *regulatory infrastructure* that affects behavior in a manner that facilitates teamwork. It is these regulatory obligations that define the team much more so than a label or a colored jersey. In our work, the teams are not merely groupings, but provide the framework to support advanced coordination policies typical in human-human teams. When a leader is assigned, it means more than just being authorized to task other agents. It also defines the expected communication pattern among pertinent team members. For example, as a team member, you are obligated to ensure your leader knows you are working and to keep other members updated about pertinent information. These types of coordination, natural to humans, will enable robots to perform more like teammates.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of this research by the Office of Naval Research Coordinated Operations Program and the Army Research Laboratories, and the participation of the Naval Surface Warfare Center-Panama City and the University of South Florida. Thanks also to other members of the Coordinated Operations team, including Cho-Man Teng, Anil Raj, Marco Arguedas, Maggie Breedy, Tom Eskridge, Paul Feltovich, Lucian Galescu, and James Lott Jerry Pratt, Robin Murphy, Signe Redfield, and Richard Smith.

REFERENCES

1. Allen, J., N. Blaylock, and G. Ferguson. "A problem-solving model for collaborative agents." Presented at the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002) 2002.
2. Balch, T. and R.C. Arkin. "Communication in reactive multiagent robotic systems." *Autonomous Robots* 1, no. 1 (1995): 27-52.
3. Bradshaw, J.M., P. Beautement, M.R. Breedy, L. Bunch, S.V. Drakunov, P.J. Feltovich, R.R. Hoffman, R. Jeffers, M. Johnson, S. Kulkarni, J. Lott, A. Raj, N. Suri, and A. Uszok. "Making agents acceptable to people." In *Intelligent Technologies for Information Analysis: Advances in Agents, Data Mining, and Statistical Learning*, edited by N. Zhong and J. Liu, 361-400. Berlin: Springer Verlag, 2004.
4. Bradshaw, J.M., P. Feltovich, M. Johnson, L. Bunch, M.R. Breedy, H. Jung, J. Lott, and A. Uszok. "Coordination in Human-Agent Teamwork. Invited Paper and Presentation.", AAI Fall Symposium, 8-10 November, 2007.
5. Bradshaw, J.M., P. Feltovich, H. Jung, S. Kulkarni, W. Taysom, and A. Uszok. "Dimensions of adjustable autonomy and mixed-initiative interaction." In *Agents and Computational Autonomy: Potential, Risks, and Solutions. Lecture Notes in Computer Science, Vol. 2969*, edited by M.

- Nickles, M. Rovatsos, and G. Weiss, 17-39. Berlin, Germany: Springer-Verlag, 2004.
6. Bradshaw, J.M., H. Jung, S. Kulkarni, M. Johnson, P. Feltovich, J. Allen, L. Bunch, N. Chambers, L. Galescu, R. Jeffers, N. Suri, W. Taysom, and A. Uszok. "Toward trustworthy adjustable autonomy in KAoS." In *Trusting Agents for Trustworthy Electronic Societies*, edited by R. Falcone. Berlin: Springer, 2005.
 7. Bradshaw, J.M., M. Sierhuis, A. Acquisti, P. Feltovich, R. Hoffman, R. Jeffers, D. Prescott, N. Suri, A. Uszok, and R. Van Hoof. "Adjustable autonomy and human-agent teamwork in practice: An interim report on space applications." In *Agent Autonomy*, edited by H. Hexmoor, R. Falcone, and C. Castelfranchi, 243-280. Kluwer, 2003.
 8. Bradshaw, J.M., M. Sierhuis, Y. Gawdiak, R. Jeffers, N. Suri, and M. Greaves. "Teamwork and adjustable autonomy for the Personal Satellite Assistant." Presented at the Workshop on Autonomy, Delegation and Control: Interacting with Autonomous Agents. Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-2001), Seattle, WA, 6 August, 2001.
 9. Bradshaw, J.M., A. Uszok, R. Jeffers, N. Suri, P. Hayes, M.H. Burstein, A. Acquisti, B. Benyo, M.R. Breedy, M. Carvalho, D. Diller, M. Johnson, S. Kulkarni, J. Lott, M. Sierhuis, and R. Van Hoof. "Representation and reasoning for DAML-based policy and domain services in KAoS and Nomads." Presented at the Proceedings of the Autonomous Agents and Multi-Agent Systems Conference (AAMAS 2003), Melbourne, Australia, 14-18 July, 2003, 835-842.
 10. Breazeal, C., A. Brooks, J. Gray, G. Hoffman, C. Kidd, H. Lee, J. Lieberman, A. Locker, and D. Mulanda. "Humanoid robots as cooperative partners for people." *IJHR* (2004): submitted for publication.
 11. Burke, J.J., R.R. Murphy, M. Covert, and D.L. Riddle. "Moonlight in Miami: A field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise." *Human-Computer Interaction* 19, no. 1-2 (2004): 85-116.
 12. Carvalho, M. "FlexFeed: A Mobile Agent-Based Framework for Dynamic Sensor Networks." Masters, University of West Florida, 2004.
 13. Clark, H.H. *Using Language*. Cambridge, UK: Cambridge University Press, 1996.
 14. Cohen, P.R. and H.J. Levesque. "Teamwork." Menlo Park, CA: SRI International, 1991.
 15. Dias, M.B. "Dynamically formed human-robot teams performing coordinated tasks." Presented at the AAAI Spring Symposium on Human-Robot Teams, Stanford, CA, March, 2006.
 16. Feltovich, P., J.M. Bradshaw, W.J. Clancey, M. Johnson, and L. Bunch. "Progress appraisal as a challenging element of coordination in human and machine joint activity." Presented at the Engineering Societies for the Agents World VIII, Athens Greece, October, 2007.
 17. Feltovich, P., J.M. Bradshaw, R. Jeffers, N. Suri, and A. Uszok. "Social order and adaptability in animal and human cultures as an analogue for agent communities: Toward a policy-based approach." In *Engineering Societies in the Agents World IV. LNAI 3071*. Lecture Notes in Computer Science, 21-48. Berlin, Germany: Springer-Verlag, 2004.
 18. Feltovich, P.J., J.M. Bradshaw, W.J. Clancey, and M. Johnson. "We regulate to coordinate: Limits to human and machine joint activity." Presented at the Proceedings of ESAW 2006, Dublin, Ireland, 6-8 September, 2006.
 19. Ferguson, G. and J. Allen. "TRIPS: An integrated intelligent problem-solving assistant." Presented at the Proceedings of the National Conference on Artificial Intelligence (AAAI 98), Madison, WI 1998.
 20. Fong, T.W. "Collaborative Control: A Robot-Centric Model for Vehicle Teleoperation. Technical Report CMU-R1-TR-01-34." Carnegie Mellon University, 2001.
 21. Fong, T.W., C. Thorpe, and C. Baur. "Robot as partner: Vehicle teleoperation with collaborative control." Presented at the Workshop on Multi-Robot Systems, Naval Research Laboratory, Washington, DC, March, 2002.
 22. Howard, A., L.E. Parker, and G.S. Sukhatme. "Experiment with a large heterogeneous mobile robot team: Exploration, mapping, deployment, and detection." *International Journal of Robotics Research* 2, no. 5-6 (May-June 2006): 431-447.
 23. Johnson, M., J.M. Bradshaw, P. Feltovich, R. Jeffers, and A. Uszok. "A semantically-rich policy-based approach to robot control." In *Proceedings of the International Conference on Informatics in Control, Automation, and Robotics*. Lisbon, Portugal, 2006.
 24. Johnson, M., P. Chang, R. Jeffers, J.M. Bradshaw, V.-W. Soo, M.R. Breedy, L. Bunch, S. Kulkarni, J. Lott, N. Suri, and A. Uszok. "KAoS semantic policy and domain services: An application of DAML to Web services-based grid architectures." Presented at the Proceedings of the AAMAS 03 Workshop on Web Services and Agent-Based Engineering, Melbourne, Australia, July, 2003.
 25. Johnson, M., P. Feltovich, J.M. Bradshaw, and L. Bunch. *Coordinating Human-Robot Teams through Dynamic Regulation*, 2007.
 26. Klein, G., P.J. Feltovich, J.M. Bradshaw, and D.D. Woods. "Common ground and coordination in joint activity." In *Organizational Simulation*, edited by W.B. Rouse and K.R. Boff, 139-184. New York City, NY: John Wiley, 2004.
 27. Klein, G., D.D. Woods, J.M. Bradshaw, R. Hoffman, and P. Feltovich. "Ten challenges for making automation a "team player" in joint human-agent activity." *IEEE Intelligent Systems* 19, no. 6 (November-December 2004): 91-95.
 28. Lott, J., J.M. Bradshaw, A. Uszok, and R. Jeffers. "Using KAoS policy and domain services within Cougaar." Presented at the Proceedings of the Open Cougaar Conference 2004, New York City, NY, 20 July, 2004, 89-95.
 29. Sellner, B.P., F.W. Heger, L.M. Hiatt, R. Simmons, and S. Singh. "Coordinated multi-agent teams and sliding autonomy for large-scale assembly." *IEEE ** 94*, no. 7 (July 2006).
 30. Sierhuis, M., J.M. Bradshaw, A. Acquisti, R. Van Hoof, R. Jeffers, and A. Uszok. "Human-agent teamwork and adjustable autonomy in practice." Presented at the Proceedings of the Seventh International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS), Nara, Japan, 19-23 May, 2003.
 31. Suri, N., J.M. Bradshaw, M.H. Burstein, A. Uszok, B. Benyo, M.R. Breedy, M. Carvalho, D. Diller, P.T. Groth, R. Jeffers, M. Johnson, S. Kulkarni, and J. Lott. "DAML-based policy enforcement for semantic data transformation and filtering in multi-agent systems." Presented at the Proceedings of the Autonomous Agents and Multi-Agent Systems Conference (AAMAS 2003), Melbourne, Australia, 14-18 July, 2003.
 32. Suri, N., J.M. Bradshaw, M. Carvalho, T.B. Cowin, M.R. Breedy, P.T. Groth, and R. Saavendra. "Agile computing: Bridging the gap between grid computing and ad-hoc peer-to-peer resource sharing." Presented at the Proceedings of the Third International Workshop of Agent-Based Cluster and Grid Computing. In conjunction with the Third IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2003), Tokyo, Japan, 11-16 May, 2003, 618-625.
 33. Suri, N., M. Carvalho, and J.M. Bradshaw. "Proactive resource management for agile computing." Presented at the Proceedings of the Tenth Annual ECOOP Workshop on Mobile Object Systems and Resource-Aware Computing, Oslo, Norway, 14 June, 2004.
 34. Tambe, M., D. Pynadath, C. Chauvat, A. Das, and G. Kaminka. "Adaptive agent architectures for heterogeneous team members." Presented at the Proceedings of the International Conference on Multi-Agent Systems (ICMAS 2000) 2000.
 35. Tambe, M., W. Shen, M. Mataric, D.V. Pynadath, D. Goldberg, P.J. Modi, Z. Qiu, and B. Salemi. "Teamwork in cyberspace: Using TEAMCORE to make agents team-ready." Presented at the Proceedings of the AAAI Spring Symposium on Agents in Cyberspace, Menlo Park, CA 1999.
 36. Tews, A., M.J. Mataric, and G.S. Sukhatme. "A scalable approach to human-robotic interaction." Presented at the ICRA 2003, 1665-1670.
 37. Ulam, P., Y. Endo, A. Wagner, and R.C. Arkin. "Integrated mission specifications and task allocation for robot teams — design and implementation." Presented at the IEEE International Conference on Robotics and Automation (ICRA 2007) 2007, 4428-4435.
 38. Uszok, A., J.M. Bradshaw, M. Johnson, R. Jeffers, A. Tate, J. Dalton, and S. Aitken. "KAoS policy management for semantic web services." *IEEE Intelligent Systems* 19, no. 4 (July/August 2004): 32-41.