From Tools to Teammates: Joint Activity in Human-Agent-Robot Teams

Jeffrey M. Bradshaw, Paul Feltovich, Matthew Johnson, Maggie Breedy,
Larry Bunch, Tom Eskridge, Hyuckchul Jung, James Lott, Andrzej Uszok,
and Jurriaan van Diggelen*

Florida Institute for Human and Machine Cognition (IHMC)
40 South Alcaniz Street, Pensacola, FL 32502
{jbradshaw,pfeltovich,mjohnson,mbreedy,lbunch,teskridge,
jlott,auszok}@ihmc.us, jurriaanvandiggelen@gmail.com

Abstract. Coordination is an essential ingredient of joint activity in human-agent-robot teams. In this paper, we discuss some of the challenges and requirements for successful coordination, and briefly how we have used KAoS HART services framework to support coordination in a multi-team human-robot field exercise.

1 Introduction

Over the past several years, we have been interested in learning how to facilitate teamwork among humans, agents, and robots. To lay the groundwork for our research, we have studied how humans succeed and fail in joint activity requiring a high degree of interdependence among the participants [9; 18]. Such interdependence requires that, in addition to what team members do to accomplish the work itself, they also invest time and attention in making sure that distributed or sequenced tasks are appropriately coordinated [21].

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 the help of adjustable autonomy mechanisms [4; 6; 17]. 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 [14]. Third, working in the tradition of previous agent teamwork researchers (e.g., [10; 24]), we attempt to implement a reusable model of teamwork involving a notion of shared knowledge, goals, and regulatory mechanisms that function as the glue that binds team members together. This teamwork model is to a large degree independent from and complementary to the set of domain-specific reasoners (e.g., task scheduling/optimization, spatial reasoning) that might be needed to accomplish a particular task objective.

---

* Currently at TNO, Soesterberg, The Netherlands.

M. Kurosu (Ed.): Human Centered Design, HCII 2009, LNCS 5619, pp. 935–944, 2009.

Although there are several important challenges in making automation a team player [19], in this paper we focus on only on the problem of coordination. Following a brief description of this aspect of joint activity, we describe the KAoS HART (Human-Agent-Robot Teamwork) services framework, which has been developed as a means of exploring our ideas about the role of regulatory constraints in joint activity [3; 5; 11; 14; 25; 26]. Finally, we discuss example policies from a field exercise that allowed us to implement and explore many of these capabilities. This exercise involved mixed human-robot teams whose objective was to find and apprehend an intruder hiding on a cluttered Navy pier [16].

2 Understanding Coordination

Malone and Crowston [21] defined coordination as "managing dependencies between activities." Teamwork, which by definition implies interdependence among the players, therefore requires some level of work for each party over and beyond the carrying out of task itself in order to manage its role in coordination. Part of that "extra" work involves each party doing its part to assure that relevant aspects of the agents and the situation are observable at an appropriate level of abstraction and using an effective style of interaction [1].

Although coordination is as much a requirement for agent-agent teamwork as it is for human-agent teamwork, the magnitude of the representational and reasoning gulfs separating humans from agents is much larger. Moreover, because the agent's ability to sense or infer information about the human environment and cognitive context is so limited, agent designers must find innovative ways to compensate for the fact that their agents are not situated in the human world. Brittleness of agent capabilities is difficult to avoid because only certain aspects of the human environment and cognitive context can be represented in the agent, and the representation that is made cannot be "general purpose" but must include specific representations and optimizations for the particular use scenarios the designer originally envisioned. Without sufficient basis for shared situation awareness and mutual feedback, coordination among team members simply cannot take place, and, of course, this need for shared understanding and feedback increases as the size of the team and the degree of autonomy increase.

Notwithstanding these challenges, adult humans and radically less-abled entities (e.g., small children, dogs, video game characters) are capable of working together effectively in a variety of situations where a subjective experience of collaborative teaming is often maintained despite the magnitude of their differences. Generally this is due to the ability of humans to rapidly size up and adapt to the limitations of their teammates in relatively short order, an ability we would like to exploit in the design of approaches for human-agent teamwork.

2.1 The Elements of Effective Coordination

Basic requirements. There are three basic requirements for effective coordination: interpredictability, common ground, and directability [18]:

- *Interpredictability:* In highly interdependent activities, it becomes possible to plan one's own actions (including coordination actions) only when what others will do can be accurately predicted. Skilled teams become interpredictable through shared knowledge and idiosyncratic coordination devices developed through extended experience in working together; bureaucracies with high turnover compensate for experience by substituting explicit, predesigned structured procedures and expectations.

- *Common ground:* Common ground refers to the pertinent mutual knowledge, beliefs, and assumptions that support interdependent actions in the context of a given joint activity [8]. This includes initial common ground prior to engaging in the joint activity, as well as mutual knowledge of shared history and current state that is obtained while the activity is underway. Unless I can make good assumptions about what you know and what you can do, we cannot effectively coordinate.
- *Directability:* Directability refers to the capacity for deliberately assessing and modifying the actions of the other parties in a joint activity as conditions and priorities change [7]. Effective coordination requires responsiveness of each participant to the influence of the others as the activity unfolds.

Coordination devices. People coordinate through signals and more complex messages of many sorts (e.g., face-to-face language, expressions, posture). Human signals are also mediated in many ways—for example, through third parties or through machines such as telephones or computers. Hence, direct and indirect party-to-party communication is one form of a coordination device, in this instance coordination by *agreement*. For example, a group of scientists working together on a grant proposal, may simply agree, through e-mail exchanges, to set up a subsequent conference call at a specific date and time. There are three other major types of coordination devices that people commonly employ: *convention*, *precedent*, and *situational salience* [9; 18].

Roles. Roles can be thought of as ways of packaging rights and obligations that go along with the necessary parts that participants play in joint activities. Knowing one's own role and the roles of others in a joint activity establishes expectations about how others are likely to interact with us, and how we think we should interact with them. Shoppers expect cashiers to do certain things for them (e.g., total up the items and handle payment) and to treat them in a certain way (e.g., with cheerful courtesy), and cashiers have certain expectations of shoppers. When roles are well understood and regulatory devices are performing their proper function, observers are likely to describe the activity as highly-coordinated. On the other hand, violations of the expectations associated with roles and regulatory structures can result in confusion, frustration, anger, and a breakdown in coordination.

Organizations. Collections of roles are often grouped to form organizations. In addition to regulatory considerations at the level of individual roles, organizations themselves may also add their own rules, standards, traditions, and so forth, in order to establish a common culture that will smooth interaction among parties.

Knowing how roles undergird organizations and how rights and obligations are packaged into roles helps us understand how organizations can be seen as functional or dysfunctional. Whether hierarchical or heterarchical, fluid or relatively static, organizations are functional only to the extent that their associated regulatory devices and roles generally assist them in facilitating their constituent responsibilities and their work in coordinating their actions with others when necessary.

The lesson here for mixed human-agent-robot teams is that the various roles that team members assume in their work must include more than simple names for the role and algorithmic behavior to perform their individual tasks. They must also, to be successful, include regulatory structures that define the additional work of coordination associated with that role.

3 KAoS HART Policy-Based Approach

The KAoS HART (Human-Agent-Robot Teamwork) services framework has been adapted to provide the means for *policy-based* dynamic regulation on a variety of agent, robotic, Web services, Grid services, and traditional distributed computing platforms [3; 13-15; 20; 23; 25]. In contrast to Asimov's laws of robotics—and similar in spirit to Grice's famous maxims [12]—the objective of KAoS teamwork policies is not principally to prevent harm but rather, in a positive vein, to facilitate helpful interaction among teammates. As a body, we call these policies the "Golden Rules of HART," recalling the biblical injunction: "Do unto others as you would have them do unto you" (Matthew 7:12). If robots (and people) were all sufficiently intelligent and benevolent, perhaps this abstract maxim would be the only rule needed for coordination. Since this is not the case, a number of more specific instantiations of the general principle are required as the basis of teamwork policy. Through a relatively small number of such policies, we can help assure that complex emergent activity can remain coordinated. Through policy learning mechanisms, additional constraints may emerge in an adaptive manner [22].

KAoS policies expressing these specific 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) [25]. A growing set of services for policy deconfliction and analysis are also provided [2; 25].

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 polices; 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.

Rigid policy constraints and roles cannot cope with unanticipated situations in a dynamically changing environment. This is particularly important in teamwork situations where multiple agents have to cooperate to achieve a common goal. In addition to an agent's own capabilities and constraints, we also need to take into account the capabilities and constraints of team members that might help or impede a joint task. IHMC's Jung and Teng have devised a methodology to adjust a team of agents' autonomy constraints and execution plans on the fly to avoid total task failures and performance degradation in teamwork situations [4; 6; 17].

4 Coordination Policy Examples

We developed a series of human-agent-robot coordination policies to support a scenario in which an intruder must be discovered and apprehended on a cluttered Navy pier with the assistance of two humans and five robots. Issues included robot capabilities, sensor limitations, and localization, however we focused on the coordination aspects of the task. We specifically included multiple humans and robots, and more robots than a single individual could easily handle by teleoperation.[1]

The teamwork model for our coordinated operations exercise was 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.

4.1 Cohen-Levesque Notification Obligation Policy

One of the most well known heuristics in team coordination was originally formulated by Cohen and Levesque as follows: “any team member who discovers privately that a goal is impossible (has been achieved, or is irrelevant) should be left with a goal to make this fact known to the team as a whole” [19, p. 9]. We have implemented our version of this heuristic in the form of an obligation policy that can be roughly described as follows:

> *A Robot is obligated to notify its Teammates when Action is Finished (whether Successfully Completed, Aborted, or Irrecoverably Failed)*

For example, in our field experiments, this policy ensured that, once an intruder has been apprehended, robot and human members of all teams, are notified [13]. This obligation would be triggered as soon as one robot became aware of this fact, and each robot would begin executing the appropriate task it was designed to perform following successful completion of the team goal (e.g., return to base, resume patrolling). If, on the other hand, the team commander were to abort the task due to a higher priority objective, or if any of the robots became aware that failure was inevitable, they would let their teammates know so that the appropriate behaviors for this situation would be triggered for the other members of the team. This single policy obviated the need to write a large number of special-purpose procedures for each possible success or failure mode.

---

[1] In addition to KAoS, which is our focus in this paper, the exercise involved the Agile Computing Infrastructure (ACI), the TRIPS dialogue-based collaborative problem solving system, Kaa/Kab adjustable autonomy and backup planning components, and an advanced multimodal display capability. For an overview of the entire system and scenario, see [16]. Robin Murphy’s group at the University of South Florida participated in a USV scenario that was loosely-coupled to our ground-based story.

4.2 Acknowledgements and Policy Deconfliction

We implemented a basic policy that requires robots to acknowledge requests. While this seemed a good general rule, there are important exceptions that need to be handled through KAoS policy deconfliction capabilities [7].

One reasonable exception to the acknowledgement policy is that people do not always verbally acknowledge requests, particularly when they are directly observable. Direct observability means that when a human requestor sends the communication to a robot receiver, the fact that the request was received, understood and being acted upon is observable by the requestor. For example, when a robot is told to move forward five meters, and then can be seen starting to move forward, there is normally no need for the robot to state "I have received your request to move forward and have begun." The same applies to queries. When somebody asks a robot "where are you," it is unnecessary for it to reply "I have heard your question and am about to reply", if it, alternatively, simply says "in the library." We implemented two additional policies to waive the obligation to acknowledge requests when the request is either a teleoperation command or a query.

1) *A Robot is obligated to acknowledge to the Requestor when the Robot Accepts an Action*
2) *A Robot is not obligated to acknowledge Teleoperation requests*
3) *A Robot is not obligated to acknowledge Query requests*

The two policies do indeed conflict with the original, but by assigning the more restrictive polices a higher priority (which can be done numerically or logically), it is possible to automatically deconflict these policies and achieve the desired behavior. Note an additional advantage in the use of ontologies of behaviors is the fact that we can define policies for abstract classes of actions (e.g., Action, Teleoperation requests, Query requests) that will be enforced on every specific action that falls in that class.

4.3 Role Management and Progress Appraisal

Groups often use roles to perform task division and allocation. Roles provide a membership-based construct with which to associate sets of privileges (authorizations) and expected behaviors (obligations). When an actor is assigned to a role, the regulations associated with the role automatically apply to the actor and, likewise, are no longer applicable when the actor relinquishes the role. These privileges and expectations that comprise a role may be highly domain dependent. For example the role "Team Leader" in a military domain is significantly different from "Team Leader" in sports. Roles may also specify expected behaviors. For example, if your role is a "Sentry," then you are obligated to remain at your post, and other actors will expect you to fulfill that obligation. Roles can also affect other behaviors such as expected communications. If you are assigned to be a "Sentry", you are obligated to announce any violations of your boundary and report these to your immediate superior.

Taking advantage of the extensibility and inheritance properties of OWL ontologies, we defined roles at various levels of abstraction with sub-roles refining the regulations pertinent to more generic super-roles. In this way, some high-level roles need not be domain specific or involve specific tasking, but they are still defined by their associated regulations. "Teammate" can be considered a generic role that has some of its regulations already noted. We view this level of abstraction as appropriate for expectations that facilitate coordination such as acknowledgements and progress appraisals. The obligation to acknowledge requests can be thought of as a policy associated with being a teammate. We have developed two policy sets that we feel apply generally to robots assigned to the role of "Teammate." The first is the acknowledgement policy set discussed above. The second involves progress appraisal:

1) *A Robot is obligated to notify the Requestor when requested Action is Finished (includes Completed, Aborted, and Failure).*
2) *A Robot is not obligated to notify the Requestor when a requested Teleoperation Action is Completed.*
3) *A Robot is not obligated to notify the Requestor when a requested Query Action is Completed.*

The first policy ensures that the requestor of a task is notified when the tasked robot encounters problems or successfully completes the task since the action status of Finished is ontologically defined as a super-class of the statuses Completed, Failed, and Aborted. The second two policies in this set are exceptions similar to those in the acknowledgement set. With knowledge that these policies are in place, human and robotic team members have the mutual expectation that these progress appraisals will be performed. This interpredictability removes the need to explicitly ask for such communication and, perhaps just as importantly, the absence of these obligatory communications becomes an indicator that additional coordination may be necessary. For example, a robot is commanded to autonomously navigate to a distant location. Since it is known that the robot would notify team members if it had arrived, or it was stuck, or had otherwise failed, the others can assume that it is still moving toward the goal. If team members were concerned with an approaching deadline or that the task was taking too long, they would query for the robot's position and create a new estimate of when it should reach the goal.

The policies outlined here are just two of several sets that we have explored, informed by previous theoretical work, simulations, and field experiments performed by ourselves and by others [1, 3, 15-17, 19, 26-28]. As we encounter new challenges in future work, we will continue to revise and expand such policy sets.

4.4 Policies Relating to Team Leaders

In contrast to our previous work on human-robot teams, where all team members were "equal," we decided to explore the role of team "leaders." Leaders not only must adhere to their own regulations, but they also impact the regulatory structure of all the other roles in the group. Peer interaction may be undirected, but Leaders tend to alter the pattern of activity, with themselves becoming the focal point. In particular we have identified several policy sets particular to leaders. The first set is about the chain of command:

1) *A Robot is authorized to perform Actions requested by its Team Leader*
2) *A Robot is authorized to Accept Actions requested by a higher authority*
3) *A Robot is not authorized to perform Action requests from just any Requestor*
4) *A Robot is authorized to Accept Actions that are self-initiated*

The first policy gives team leaders the authority to command their team. The second gives the same authority to anyone directly higher in the chain of command. The third policy explicitly restricts access to the robots from those outside of the chain of command. The fourth policy makes self initiated actions an exception to the third policy.

Another policy set was used to explore notification to help maintain common ground between the team leader and each of the team members:

1) *A Robot is obligated to notify its Team Leader when an Action is requested by a higher authority*
2) *A Robot is obligated to notify Its Team Leader when starting a self-initiated Action*
3) *A Robot is obligated to notify its Team Leader when a self-initiated Action is Finished (includes statuses of Completed, Aborted, and Failure).*

4.5 Team Creation and Management

The KAoS Directory Service manages organizational structure, allowing dynamic team formation and modification. Teams and subteams can be created dynamically, allowing for the creation of complex organizational structures. Agents can join and leave teams as necessary to support the desired structure. Actors can be assigned roles including Team Leader, affecting the dynamics of coordination as discussed in the previous section. Queries can be made to identify current team structure, who is on a certain team currently, or who is team leader.

5 Conclusions

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, this means more then just being authorized to task other agents. For instance, it also defines the expected communication pattern among pertinent team members. As a team member, you are obligated to ensure that 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 less like tools and more like teammates.

References

1. Bradshaw, J.M., Sierhuis, M., Acquisti, A., Feltovich, P., Hoffman, R., Jeffers, R., Prescott, D., Suri, N., Uszok, A., Van Hoof, R.: Adjustable autonomy and human-agent teamwork in practice: An interim report on space applications. In: Hexmoor, H., Falcone, R., Castelfranchi, C. (eds.) Agent Autonomy, pp. 243–280. Kluwer, Dordrecht (2003)
2. Bradshaw, J.M., Uszok, A., Jeffers, R., Suri, N., Hayes, P., Burstein, M.H., Acquisti, A., Benyo, B., Breedy, M.R., Carvalho, M., Diller, D., Johnson, M., Kulkarni, S., Lott, J., Sierhuis, M., Van Hoof, R.: Representation and reasoning for DAML-based policy and domain services in KAoS and Nomads. In: Presented at the Proceedings of the Autonomous Agents and Multi-Agent Systems Conference (AAMAS 2003), Melbourne, Australia, July 14-18, pp. 835–842 (2003)

3. Bradshaw, J.M., Beautement, P., Breedy, M.R., Bunch, L., Drakunov, S.V., Feltovich, P.J., Hoffman, R.R., Jeffers, R., Johnson, M., Kulkarni, S., Lott, J., Raj, A., Suri, N., Uszok, A.: Making agents acceptable to people. In: Zhong, N., Liu, J. (eds.) Intelligent Technologies for Information Analysis: Advances in Agents, Data Mining, and Statistical Learning, pp. 361–400. Springer, Berlin (2004)
4. Bradshaw, J.M., Feltovich, P., Jung, H., Kulkarni, S., Taysom, W., Uszok, A.: Dimensions of adjustable autonomy and mixed-initiative interaction. In: Nickles, M., Rovatsos, M., Weiss, G. (eds.) AUTONOMY 2003. LNCS, vol. 2969, pp. 17–39. Springer, Heidelberg (2004)
5. Bradshaw, J.M., Feltovich, P.J., Jung, H., Kulkarni, S., Allen, J., Bunch, L., Chambers, N., Galescu, L., Jeffers, R., Johnson, M., Sierhuis, M., Taysom, W., Uszok, A., Van Hoof, R.: Policy-based coordination in joint human-agent activity. In: Presented at the Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, The Hague, Netherlands, October 10-13 (2004)
6. Bradshaw, J.M., Jung, H., Kulkarni, S., Johnson, M., Feltovich, P., Allen, J., Bunch, L., Chambers, N., Galescu, L., Jeffers, R., Suri, N., Taysom, W., Uszok, A.: Toward trustworthy adjustable autonomy in KAoS. In: Falcone, R. (ed.) Trusting Agents for Trustworthy Electronic Societies. Springer, Berlin (2005)
7. Christofferson, K., Woods, D.D.: How to make automated systems team players. In: Salas, E. (ed.) Advances in Human Performance and Cognitive Engineering Research, vol. 2. JAI Press/ Elsevier (2002)
8. Clark, H.H., Brennan, S.E.: Grounding in communication. In: Resnick, L.B., Levine, J.M., Teasley, S.D. (eds.) Perspectives on Socially Shared Cognition. American Psychological Association, Washington (1991)
9. Clark, H.H.: Using Language. Cambridge University Press, Cambridge (1996)
10. Cohen, P.R., Levesque, H.J.: Teamwork. SRI International, Menlo Park (1991)
11. Feltovich, P., Bradshaw, J.M., Jeffers, R., Suri, N., Uszok, A.: Social order and adaptability in animal and human cultures as an analogue for agent communities: Toward a policy-based approach. In: Omicini, A., Petta, P., Pitt, J. (eds.) ESAW 2003. LNCS (LNAI), vol. 3071, pp. 21–48. Springer, Heidelberg (2004)
12. Grice, H.P.: Logic and Conversation. In: Cole, P., Morgan, J.L. (eds.) Syntax and Semantics 3: Speech Acts. Academic Press, New York (1975)
13. Johnson, M., Chang, P., Jeffers, R., Bradshaw, J.M., Soo, V.-W., Breedy, M.R., Bunch, L., Kulkarni, S., Lott, J., Suri, N., Uszok, A.: KAoS semantic policy and domain services: An application of DAML to Web services-based grid architectures. In: Presented at the Proceedings of the AAMAS 2003 Workshop on Web Services and Agent-Based Engineering, Melbourne, Australia (July 2003)
14. Johnson, M., Bradshaw, J.M., Feltovich, P., Jeffers, R., Uszok, A.: 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)
15. Johnson, M., Feltovich, P.J., Bradshaw, J.M., Bunch, L.: Human-robot coordination through dynamic regulation. In: Proceedings of the International Conference on Robotics and Automation (ICRA), Pasadena, CA (2008) (in press)
16. Johnson, M., Intlekofer Jr., K., Jung, H., Bradshaw, J.M., Allen, J., Suri, N., Carvalho, M.: Coordinated operations in mixed teams of humans and robots. In: Marik, V., Bradshaw, J.M., Meyer, J. (eds.) Proceedings of the First IEEE Conference on Distributed Human-Machine Systems (DHMS 2008), Athens, Greece (2008) (in press)

17. Jung, H., Bradshaw, J.M., Kulkarni, S., Breedy, M.R., Bunch, L., Feltovich, P., Jeffers, R., Johnson, M., Lott, J., Suri, N., Taysom, W., Tonti, G., Uszok, A.: An ontology-based representation for policy-goverened adjustable autonomy. In: Proceedings of the AAAI Spring Symposium. AAAI Press, Stanford (2004)
18. Klein, G., Feltovich, P.J., Bradshaw, J.M., Woods, D.D.: Common ground and coordination in joint activity. In: Rouse, W.B., Boff, K.R. (eds.) Organizational Simulation, pp. 139–184. John Wiley, New York City (2004)
19. Klein, G., Woods, D.D., Bradshaw, J.M., Hoffman, R., Feltovich, P.: Ten challenges for making automation a "team player" in joint human-agent activity. IEEE Intelligent Systems 19(6), 91–95 (2004)
20. Lott, J., Bradshaw, J.M., Uszok, A., Jeffers, R.: Using KAoS policy and domain services within Cougaar. In: Presented at the Proceedings of the Open Cougaar Conference 2004, New York City, July 20, pp. 89–95 (2004)
21. Malone, T.W., Crowston, K.: What is coordination theory and how can it help design cooperative work systems? In: Presented at the Conference on Computer-Supported Cooperative Work (CSCW 1990), Los Angeles, CA, October 7-10, pp. 357–370 (1990)
22. Rehak, M., Pechoucek, M., Bradshaw, J.M.: Representing context for multiagent trust modeling. In: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Hong Kong (2006)
23. Sierhuis, M., Bradshaw, J.M., Acquisti, A., Van Hoof, R., Jeffers, R., Uszok, A.: Human-agent teamwork and adjustable autonomy in practice. In: Presented at the Proceedings of the Seventh International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS), Nara, Japan, May 19-23 (2003)
24. Tambe, M., Shen, W., Mataric, M., Pynadath, D.V., Goldberg, D., Modi, P.J., Qiu, Z., Salemi, B.: Teamwork in cyberspace: Using TEAMCORE to make agents team-ready. In: Presented at the Proceedings of the AAAI Spring Symposium on Agents in Cyberspace, Menlo Park, CA (1999)
25. Uszok, A., Bradshaw, J.M., Johnson, M., Jeffers, R., Tate, A., Dalton, J., Aitken, S.: KAoS policy management for semantic web services. IEEE Intelligent Systems 19(4), 32–41 (2004)
26. Uszok, A., Bradshaw, J.M., Breedy, M.R., Bunch, L., Feltovich, P., Johnson, M., Jung, H.: New developments in ontology-based policy management: Increasing the practicality and comprehensiveness of KAoS. In: Proceedings of the 2008 IEEE Conference on Policy, Palisades, NY (2008)