

# Managing Collaboration in Hybrid-Agent Teams

John Allen<sup>\*</sup>, Abbe Mowshowitz<sup>#</sup>, Tim Norman<sup>†</sup>, Simon Parsons<sup>‡</sup>, and Alun Preece<sup>†</sup>

<sup>\*</sup>Honeywell, 3600 Technology Dr., Minneapolis, MN 55112, USA; <sup>#</sup>City College of NY (CUNY), Convent Ave. at 138th St., NY, NY 10031, USA; <sup>†</sup>King's College, University of Aberdeen, Aberdeen AB24 3UE, UK; <sup>‡</sup>Brooklyn College (CUNY), 2900 Bedford Avenue, Brooklyn, NY 11210, USA

**Abstract-**This paper proposes the use of the virtual organization framework in managing collaboration in a mixed team of software agents and humans aided by such agents. The paper argues that this framework facilitates an integrated management approach and sets the scene for experimental work to demonstrate the effectiveness of the approach.

## I. INTRODUCTION

This paper deals with managing collaboration in a team. In particular, we are interested in teams engaged in military missions, and teams in which members may come from different parts of an international coalition. In such situations, effective coordination can be problematic, with units unable to communicate easily, and handicapped by having been trained to operate under rather different doctrines. It is our contention that, with careful design, software agents can support effective collaboration in teams, and can overcome some of the problems with coalition forces. This paper reports on work towards such a design. In short, we are proposing that elements of the coordination are handled by software agents, agents which exchange and filter information for the units in the coalition force. The teams we are interested in are thus composed both of the human members of the units, software agents that support them, and possibly additional software team members (controlling, for example, autonomous vehicles). We use the term *hybrid-agent* to designate such teams.

We assume that a mission begins with the mission commander issuing the *commander's intent*, which we will further assume is broadly in line with the definition of this term provided by the U. S. Army: "...a clear, concise statement of what the force must do and the conditions the force must meet to succeed with respect to the enemy, terrain, and desired end state" [7,8]. In

our terms, the commander's intent is a high level statement of the goals of the mission, perhaps with some conditional aspects (which modify the goals if certain conditions apply). The high level problem of how to achieve the commander's intent then needs to be *decomposed* into lower level goals that capture the individual steps that must be taken to achieve the commander's intent. In a coalition context, the decomposition will then be followed by the assembly of units capable of achieving the goals identified by the decomposition, typically involving some form of coordinated allocation of resources, and these resources then have to be managed in a coordinated fashion through the execution of the mission.

This paper outlines an integrated approach to supporting this process. The approach derives from Mowshowitz's [16,17,18] switching model of virtual organization. A mission is interpreted as a goal-oriented task that is managed by dynamically assigning *satisfiers* to the *requirements* of the task. Within this scheme the problem of determining the commander's intent in a mission is equivalent to delineating the requirements of the given task, based on overall goals, that is by resolving complex or ambiguous goals into independent and consistent subgoals that specify what needs to be done. The resources available for realizing these subgoals (or requirements) are the satisfiers of the switching model.

A key function in hybrid-agent teams is determining the set of satisfiers at a given point in time, and establishing the way that they can be coordinated. This determination depends in part on the results of *negotiations* between agents and humans, and in part on the way that the forces will be *coordinated* through the execution of the operation. Thus, each of the principal elements

of collaboration — problem decomposition, negotiation strategies, and coordination mechanisms — in hybrid-agent teams finds a natural place in the switching model of virtual organization.

## II. REPRESENTATIONAL FRAMEWORK

Central to the switching model is the notion of a *virtually organized task*. Any goal-oriented activity can be expressed and managed as a virtually organized task. Such a task consists of a set of *abstract requirements*, a set of *concrete satisfiers*, a specification of the *satisficing criteria* to be achieved by assigning satisfiers to requirements, and a procedure for making the assignment. The strict separation of requirements from satisfiers allows resource allocation to be a dynamic process in which assignment of satisfiers to requirements may change over time as task objectives change. Managing such a dynamic process calls for dedicated, independent activities performed over time that aim to specify requirements and satisfiers, to examine the satisficing criteria in light of changing conditions, and to adjust the assignment algorithm or heuristic to insure compliance with the altered criteria.

The option of switching from one satisfier to another to meet a requirement allows for leveraging scarce resources and responding effectively to changing conditions. In military applications timely, effective and secure task performance may be critical success criteria. The satisficing criteria might specify choosing an available resource subject to constraints such as record of past performance, reliability, and so on.

Both requirements and satisfiers can be specified by sets of attributes, each of which may be represented by a variable with a given range of values. The set of attributes for a coalition operation might include ‘quality’ and ‘reliability’. The ‘quality’ attribute could be represented by a real variable taking values in the interval  $[0,1]$ ; ‘reliability’ could be measured by a finite, discrete variable with, say, the ten values 1 to 10, where 1 signifies least reliable and 10 most reliable. The attribute values may signify the degree to which given aspects of a requirement must be met by a satisfier, or possibly the strength of properties that contribute to the requirement’s importance or persistence. If requirement R has a reliability value of 8, then only a satisfier rated 8 or higher could be

assigned to R. Note that attribute sets for requirements and satisfiers will have a non-empty intersection but will not generally be identical.

Typical attributes of requirements are:

*Importance*: the importance of the given requirement among the list of requirements defining the virtually organized task – note that the importance of a requirement does not relate to any specific satisfier, but to the criteria for switching;

*Persistence*: the stability of the requirement or the likelihood that it will continue unchanged for a given period of time;

*Volume*: the quantity of the resource that needs to be processed in fulfilling the requirement;

*Tolerance level*: the specificity of the input to be processed in fulfilling the given requirement – the higher the specificity, the greater the difficulty in finding substitutes;

*Overall importance*: the relative significance of the virtually organized task of which the given requirement is a part. If the overall importance of a given requirement is low, the threshold for switching may be raised because there are always unpredictable consequences, and the risks posed might inhibit switching even if a small advantage could be gained.

Establishing the requirements is the first step of setting up a virtual organization. Once we have the requirements, the agents need to agree upon the specific capabilities to be assigned to the plan. In other words we need to identify satisfiers for these requirements.

Once the requirements and satisfiers have been specified, it is possible to formulate the resource allocation problem as time varying assignments of satisfiers to requirements. Since satisfiers (e.g., the particular aircraft available for sorties) and requirements (e.g., the numbers of reconnaissance sorties) may change over time, the allocation problem requires a dynamic solution. The switching model can be used to formulate a variety of dynamic solutions [2,24,26,27]. Figure 1 shows the overall structure of the switching model approach.

To build an algorithmic or heuristic allocation procedure, criteria for assigning satisfiers to requirements must be spelled out. In the case of coalition operations, the criteria might reflect such desiderata as accuracy, security, and possibly cost. If accuracy were of paramount

concern, requirements would be assigned satisfiers that maximize accuracy in targeting, subject to the constraints on requirement satisfaction.

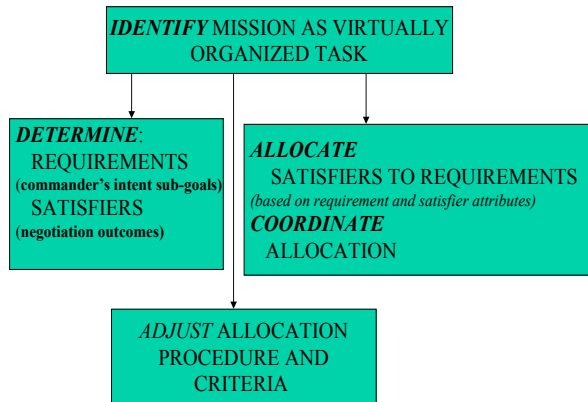


Figure 1. Switching model approach

Just as requirements and satisfiers are subject to change in the switching model, so are the assignment criteria. This feature of the model recognizes explicitly the need to modify the framework for resource allocation based on changing conditions in the field or alterations in strategy.

### III. PROBLEM DECOMPOSITION

As explained above (Section I), when planning an operation, the commander in charge will issue commander's intent. This intent and planning guidance statements are meant to provide focus and scope to the subsequent planning and mission execution phases of the operation, while allowing the sub-commanders the flexibility to achieve the goals in any way they see fit. This provides robustness in the operation by allowing for delegation of decision-making authority to the lower portions of the command hierarchy, thus removing potential communications/command choke-points and single points of failure endemic to centralized planning and control.

This approach is effective, and is embraced in the network centric visions of both the U. S. Army and the U.K. Ministry of Defence. However, it suffers from two weaknesses. First, as reported by Shattuck & Woods [29], sub-

commanders correctly interpret commander's intent statements only about 30% of the time, leaving considerable room for error. Secondly, in the context of mixed human/software agent teams, the commander's intent statements include information types that are not commonly represented in artificial intelligence planning or executions system.

Along with information about the mission objectives, a commander's intent statement may include information about potential and acceptable risks, applicable rules of engagement, times and locations of coordination and areas for contingency planning. Traditional artificial intelligence planning systems [13,15] do not deal with this form of information. They are strictly concerned with the satisfaction of planning problems, and do not take into consideration anything outside of the domain language, e.g., the goal language and the operators are all defined in terms of ground literals that describe the domain, and nothing that describes the properties of the produced plan. Hierarchical Task Network planners [19,35] allow for a more elaborate representation language, which can include elements not in the ground instance language. These systems explicitly represent the decomposition of problems into collection of sub-problems that can be concerned with concepts other than those that may be represented in the ground literals. To rephrase this in terms of the switching, the representation language of HTN planners provides a way to capture the requirements.

Distributed agent systems are similar in representational power to HTN planners, but they typically do not represent explicitly the decomposition of problems in a single data structure such as a task network. However, this structure is implicitly represented in an agent's ability to generate *requirements* and provide *satisfiers*. Work in Design-To-Criteria scheduling [34] show methods of utilizing meta-information about the properties of plans or actions while developing those plans. Furthermore, agents utilizing *argumentation based negotiation* (discussed below) will have justifications for their *requirements*, allowing for greater creativity in finding resolution for problems such as the fact that Agent A can provide a aircraft for a surveillance operation, but cannot provide the communication media to affect coordination.

#### IV. NEGOTIATION

Following the decomposition of the commander's intent into a series of sub-goals, it is necessary to form a team to carry out the mission. The team will be formed from the various coalition units that are available, each of which, in the terminology of the switching model, is a satisfier that might be used, alone or in combination, to achieve a specific goal. Part of the formation of a team is picking a suitable set of units, but this is not all that is required.

Typically, the same pool of units will be required to carry out a number of missions simultaneously, and so there will be contention between missions for resources. We propose that this conflict be handled by negotiation between software agents representing the coalition units. These agents can increase effectiveness by automating otherwise complex and time-consuming negotiations — rather than take part in the back and forth necessary to reach agreement on, for example, which units will be supplied by which partners in a multi-national coalition. The mission commander can specify his preferences and delegate the negotiation to a software agent that will negotiate with agents representing the various multinational forces, and the result will be an allocation of resources that comes as close as possible to satisfying the requirements of each party. This process of negotiation will provide the means by which a specific set of satisfiers is selected, from the set that represents all the possible units that might be employed, for a specific operation.

Techniques for negotiation between software agents [9] proposed in the literature range from simpler approaches, such as auctions, through bilateral negotiations to more complex approaches, such as argumentation. In complex and uncertain environments, the need for the agents to interact efficiently, effectively and flexibly becomes paramount. Under such conditions, techniques from the more complex end of the spectrum have advantages. For example, while auctions are often efficient, they are typically inflexible and poorly suited to domains in which participants have incomplete knowledge [21].

Approaches in which more than just price of information is exchanged, such as that proposed in [5] provide greater flexibility, but the most powerful approaches are argumentation-based

negotiation (ABN) techniques [28]. In more detail [20,23,31], ABN allows agents to exchange information in addition to offers, including information such as justifications, critiques, and other forms of persuasive locutions within their interactions. These, in turn, allow agents to gain a wider understanding of the internal and social influences affecting their counterparts, thereby making it easier to resolve certain conflicts that arise due to incomplete knowledge [25]. At the same time, such negotiation models provide a means for the agents to achieve mutually acceptable agreements to their conflicts of interest. Kraus *et al.* [12], for example, have developed a model of ABN whereby agents may use promises of future reward, threats and various forms of appeals during a negotiation encounter.

More recently, research has focused on the use of formal dialogue games for tasks such as purchase negotiations [14] and multi-agent task planning and information exchange [10,22]. Furthermore, there is increased interest in the use of explicit models of the organizational context to inform negotiation [11]; in this way, agents may use references to command and authority relationships during negotiation to facilitate the generation of robust agreements and to inform the coordination mechanisms that will be employed during the operation of a virtual organization.

We see the function of negotiation mechanisms as being to devise, formalize and evaluate novel models of interaction between agents that provide the necessary flexibility and expressiveness, but that are also efficient, robust and effective for the support of team formation and adaptation. These mechanisms will fit into the context of the virtual organization switching model by providing the means by which agents can offer, justify and critique services that may be used as satisfiers for virtual organization requirements. Furthermore, as circumstances change, the coordination between agents within the virtual organization may lead to further negotiation focused on the adaptation to such changes. Our vision is for both the negotiation and coordination models to have a common underlying formalism so that a single coherent framework for team formation, adaptation and operation may underpin the solution.

#### V. COORDINATION MECHANISMS

Once the team has been assembled through some negotiated resource allocation, the mission can proceed. However, during mission execution there is a need to coordinate team activities, to ensure, for example, that if Unit A is required to support Unit B in the capture of objective C, then A attacks the objective at the same time as B even if this means delaying the previously agreed time of attack.

In the scenario that we are imagining, in which units are supported by software agents that ease the team collaboration, we can imagine this coordination being managed by the software agents using some of the many techniques that have been developed for managing collaboration and synchronization between such agents. In this section we sketch a classification of interaction mechanisms, and describe how they fit into the approach we are describing. The main classes of interaction mechanism are the following.

*Tacit agreements* are the minimal form of interaction mechanism; because there is no explicit communication between agents, social norms dictate what agents do [30]. This minimal form of interaction mechanism does not necessarily lead to the simplest form of coordination. For example, consider the way that pedestrians on a busy street coordinate themselves into “lanes” in a way that is robust enough to cope with flows that cross each other and deal with people traveling at different speeds [1]. Clearly such techniques could be adapted to coordinate the physical motion of team members.

*Environmental cues* are mechanisms in which one agent will signal to another by modifying the environment in some way, for example one agent may lay a trail that another can then follow [3]. In this case, a message is usually very simple. The content of a message can be binary, for example indicating the recent presence or absence of one of the agents at a certain location. The content of a message can also be some scalar value. For example, one agent might signal to others the distance of a resource from a given location through the value it leaves at that location. As in the previous case, we can easily imagine how such techniques might be used to synchronize team movements. Note that communication for this kind of interaction is generally broadcast, in the sense that all other agents in the environment might detect it — it is hard to modify the environment in a way that can

only ever be detected by a select handful of agents.

*Signal broadcasting* is an extension of environmental modification in which the messages are actively sent from one agent to another rather than indirectly by “drawing on” the environment. Apart from this, the mechanism is much the same as in the previous case. An interesting difference is the temporal dimension. A typical broadcast mechanism (like wireless radio communication) propagates fast, so in most environments all agents within range will receive a broadcast at nearly the same time and so will make their decisions almost simultaneously. In contrast, when using environmental cues, the message is only readable when an agent is within sensor-range and so some agents may not read the message until some time after it has been left.

*Auctions* are market mechanisms — for selling and/or buying some commodity — in which messages that agents send are indications of how much they are willing to pay (or accept in the case of a seller) and priority is given to high-price offers to buy and low-priced offers to sell [6]. Messages (bids) consist of a number indicating an agent's desired trading price and possibly an indication of what good is required (for multi-commodity auctions) and how many units are requested (for multi-unit auctions). For agent interaction, the mechanisms are adopted with some scalar playing the role of money (currency), and the result typically determines not the allocation of commodities but what the bidding agents will do (“role allocation” and/or “task allocation”) or what resources they are allocating.

*Negotiation* of the kind described above is utilized when agents need a richer interaction mechanism than is provided by an auction. For more complex reorganization of resources, for example, and, in the most sensitive cases, we may need to make use of *argumentation* again as described above.

## VI. EXAMPLE

A vignette taken from a military scenario will serve to show how the switching model can be used as a framework for representing dynamic resource allocation as well as the tasks of problem decomposition and negotiation. For purposes of this paper we use Vignette 01 of the well-established Binni scenario [32], but we

could equally well use the Holistan scenario currently under development [33]. This Vignette describes a mission that calls for the establishment of a Total Exclusion Zone (TEZ) in three days time. The mission Commander intends to carry out the order by creating “a corridor between the [warring] Gao and Agadez forces by using a ‘firestorm’ between the opposing forces.”

The Commander’s intent has three parts:

- a) “An air reconnaissance of the area of conflict to determine the situation of the opposing forces and the likely course of the conflict as the Agadez element retreats.”
- b) “A campaign plan for tactical bombing using a new class of incendiary weapon (ADM-162c). This plan needs to take account of terrain and meteorological factors in order to clear a corridor at least 5km wide and 100 km long in order to effectively separate the opposing forces for long enough to initiate cessation of hostility discussions.”
- c) “The actual campaign will take place on day three of the Vignette and must be completed within 24 hours. It is to comprise a sustained 12 hour bombing campaign followed by 12 hours for the resulting firestorm to produce an effective ‘zone of separation’, which can be policed by continuous air surveillance using helicopters and strike aircraft.”

Additional information provided in the Vignette enlarges on the components of the plan and introduces constraints on execution. The resources available for the mission are given in Annex A of the scenario description.

The mission can be modeled as one or more virtually organized tasks. For simplicity, suppose it makes sense to define just one such task. As a first approximation to the requirements, one can make use of the three components of the Commander’s plan. However, these may need to be broken down further to be properly matched with any of the potential satisfiers. This is the job of “problem decomposition”. For example, the reconnaissance effort might be differentiated into high level and low level sorties. Specifying requirements of appropriate granularity makes it possible to consider assigning a specific satisfier (aircraft type in this case).

The satisfiers of the mission can be derived from the available resources given in the table. Some of the resources may in fact be considered

satisfiers, that is possible means for meeting the requirements of the mission. The status of other resources may have to be determined through a negotiation process. For example, Annex A lists 4 P3 Orion aircraft as reconnaissance resources. Negotiation between agents representing these resources would have to be carried out to determine their availability and condition for use in possible reconnaissance operations.

Once the resources are assembled, they need to be coordinated during the execution of the mission. In the case of this vignette, we could imagine employing some form of market mechanism to assign individual aircraft to specific flight-paths during the mission. For example, agents representing aircraft could “bid” for different flight-paths based on the specific resources that they have available. Thus, an initial allocation might give a shorter route to a plane with crew who have logged more hours recently (“bids” take the form of recent flight hours, and the auction awards the longer path to lower bids). However, this allocation could be reconsidered dynamically, so that if one aircraft becomes damaged, there is a new auction to reallocate routes in which effective range (modified by the damage sustained) becomes the new, more relevant, quantity that is bid.

## VII. CONCLUSION: MANAGEMENT STRUCTURES FOR HYBRID-AGENT TEAMS

This paper has argued for the use of a specific set of techniques for managing the use of hybrid-agent teams that support the operation of coalition forces. At a high level, we have described this set of techniques — the switching model, problem decomposition, negotiation, and coordination techniques — and sketched how they can be used in concert to provide this management. The great advantage of our approach is its flexibility. The switching model explicitly provides for switching dynamically between satisfiers, and the negotiation and coordination mechanisms, similarly, can respond to changing circumstances.

At the time of writing, we are refining the high-level description provided here, and working towards an implementation of the different parts of the model. For problem decomposition, we are combining and extending work from DTC scheduling and Honeywell’s *Policy* [4], for negotiation we are using an approach to argumentation-based negotiation, and for

coordination we are using a market-based approach.

#### REFERENCES

- [1] Borgers, A. and Timmermans, H.A model of pedestrian route choice and demand for retail facilities within inner-city shopping areas, *Geographical Analysis*, 18:115–128, 1986.
- [2] Dantzig, G. *Linear Programming and Extensions*. Princeton, NJ: Princeton University Press, 1963.
- [3] Dorigo, M., Maniezzo, V., and Coloni A. The Ant System: Optimization by a colony of cooperating agents, *IEEE Transactions on Systems, Man, and Cybernetics Part B, Cybernetics*, 26(1):29–41, 1996.
- [4] Dorneich, M. C., Whitlow, S. D., Miller C. A., Allen J. A. A Superior Tool for Airline Operations. In *Ergonomics in Design*, 12(2):18–23, 2004.
- [5] Faratin, P., Sierra, C., and Jennings, N. R. Using similarity criteria to make negotiation trade-offs, In *Proceedings of the 4th International Conference on Multiagent Systems*, pages 119–126, Boston, MA, 2000.
- [6] Friedman, D. The double auction institution: A survey, In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 1, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.
- [7] FM 5-0. *Army Planning and Orders Production*. U.S. Department of the Army. January 2005.
- [8] FM 6-0. *Mission Command: Command and Control of Army Forces*. U.S. Department of the Army, August, 2003.
- [9] Jennings N. R., Faratin, P., Lomuscio A. R., Parsons. S., Sierra C., and Wooldridge M. J. Automated negotiation: Prospects, methods and challenges. *Group Decision and Negotiation*, 10:199–215, 2001.
- [10] Kalofonos, D., Karunatillake, N. C., Jennings, N. R., Norman, T. J., Reed, C., and Wells, S. Building Agents that Plan and Argue in a Social Context. In *Proceedings of the First International Conference on Computational Models of Argument*, Liverpool, UK, 2006.
- [11] Karunatillake, N. C., Jennings, N. R., Rahwan, I., and Norman, T. J. Argument-based negotiation in a social context. In *Proceedings of the Second International Workshop on Argumentation in Multi-Agent Systems*, pages 74–88, 2005
- [12] Kraus, S., Sycara, K., and Evenchik, A. Reaching agreements through argumentation. *Artificial Intelligence* 104:1–69, 1998.
- [13] Lifschitz, V. On the semantics of STRIPS. In *Proceedings of the Workshop on Reasoning about Actions and Plans*, pages 1–9, Timberline, Oregon: Morgan Kaufmann, 1986.
- [14] McBurney, P., van Eijk, R., Parsons, S., and Amgoud, L. A dialogue-game protocol for agent purchase negotiations. *Autonomous Agents and Multi-Agent Systems* 7:235–273, 2003.
- [15] McAllester, D., and Rosenblitt, D. Systematic Nonlinear Planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 634–639, Menlo Park, CA: AAAI, 1991.
- [16] Mowshowitz, A. Virtual organization, *Communications of the ACM*, 40(9):30–37, 1997.
- [17] Mowshowitz, A. *Virtual Organization: Toward a Theory of Societal Transformation Stimulated by Information Technology*. Westport, CT: Quorum Books (Greenwood Publishing Group), 2002.
- [18] Mowshowitz, A. and Kawaguchi, A. Quantifying the switching model of virtual organization. *Journal of Information Technology Theory and Application*, 6(4):53–74, 2005.
- [19] Nau, D., Au, T., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., and Yaman, F. SHOP2: An HTN Planning System. In *Journal of Artificial Intelligence Research*. 20:379–404, 2003.
- [20] Norman, T. J., Carbogim, D. V., Krabbe, E. C. W., and Walton, D. N. Argument and Multi-Agent Systems. In C. A. Reed and T. J. Norman (editors) *Argumentation Machines: New Frontiers in Argument and Computation*, volume 9 of *Argumentation Library*, pages 15–54, Kluwer Academic Publishers, Dordrecht, 2004.
- [21] Norman, T. J., Preece, A. D., Chalmers, S., Jennings, N. R., Luck, M., Dang, V. D., Nguyen, T. D., Deora, V., Shao, J., Gray, W. A., and Fiddian, N. J. Agent-based formation of virtual organizations. *Knowledge Based Systems*. 17:103–111, 2004.
- [22] Oren, N., Norman, T. J., and Preece, A. D. Arguing with confidential information. In *Proceedings of the Seventeenth European Conference on Artificial Intelligence*, 2006.
- [23] Parsons, S., Sierra, C., and Jennings, N. R. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
- [24] Pinedo, M. *Scheduling: Theory, Algorithms, and Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [25] Pollock, J. L. Perceiving and reasoning about a changing world. *Computational Intelligence*, 14:498–562, 1998
- [26] Powel, W., Snow, W. and Cheung, R. Adaptive labeling algorithms for the dynamic assignment problem, *Transportation Science*, 34:67–85, 2000
- [27] Psaraftis, H. Dynamic vehicle routing: status and prospects, *Annals of Operations Research*, 61:143–164, 1995.
- [28] Rahwan, I., Ramchurn, S. D., Jennings, N. R., McBurney, P., Parsons, S., and Sonenberg, L. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(4):343–375, 2003.
- [29] Shattuck, L. G., and Woods, D. D. Communication of Intent in Military Command and Control Systems. In *The Human in Command: Exploring the Modern Military Experience*, McCann & Pigeau, (eds.) Kluwer Academic/Plenum Publishers, New York, 2000.
- [30] Shoham, Y., and Tennenholtz, M. Emergent conventions in multi-agent systems, In *Proceedings of the Conference on Knowledge Representation and Reasoning*, 1992.
- [31] Sycara, K. Argumentation: Planning other agents' plans, In *Proceedings of the Eleventh Joint Conference on Artificial Intelligence*, pages 517–523, 1989.
- [32] TTCP C3I. A Coalition Command and Control Mission Scenario and Vignettes to Support Coalition Situation Awareness Research and Experimentation. CSA White Paper No. 4 v0.2.
- [33] Verma, D.C, Roberts, D., and Lock, G. Holistan – A Futuristic Coalition Scenario for International Coalition Operations. ITA Working paper. IBM Thomas J Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, 2007.
- [34] Wagner, T., Garvey, A., and Lesser, V. *Criteria-Directed Task Scheduling*. Technical Report 1997-16. Computer Science, University of Massachusetts, Amherst, MA, 1997
- [35] Yang, Q. Formalizing Planning Knowledge for Hierarchical Planning. *Computational Intelligence* 6(1): 12–24, 1990.