

Building Effective Teams

Gordon Beavers and Henry Hexmoor

University of Arkansas

Computer Science & Computer Engineering Department

Engineering Hall 313, Fayetteville AR 72701

{gordonb, hexmoor}@uark.edu

Abstract

Teams of agents can be made more effective by the introduction of team assistants to accomplish specific tasks. Team assistants may be *real* or *virtual* and may be first class team members or only team adjuncts. For example, a team assistant may be created as a virtual, first class team member to serve as team manager to coordinate the actions of team members. The need for a team assistant might be indicated by evaluations either internal to (subjective), or external to (objective) the team. Some example subjective measures for determining the desirability of invoking a team assistant are given.

Keywords

Agent Teams, Performance Metrics

1. Introduction

The nature of teams of agents and their mental states are considered in [Beavers and Hexmoor 2001], which contends that a team is an aggregation of agents with a common mental state that includes having a common intention, awareness of team membership and of the team's abilities, and a cooperative attitude. Furthermore, team members must have autonomy in order to form intentions.

The effectiveness of teams can be enhanced by the addition of agents specialized to a particular task, e.g., a team of soccer playing agents becomes more effective through the addition of a coach who develops strategy and coordinates the actions of the team members. The coach need not be an agent of the same type as the players; the coach might be a virtual agent. In the case of human agents, the coach is selected first and then the coach develops the team, but there is no reason that the process could not be reversed in the case of artificial agents. That is, a group of agents realizes that their performance could be enhanced with the services of a coach, and therefore decides to obtain or instantiate a coach. Note the asymmetry between natural and artificial agents. Natural agents can be recruited for and trained to a task, but artificial agents may be designed and instantiated in order to accomplish a task.

Autonomous agents endowed with social abilities and able to evaluate their own performance and the performance of their teammates, finding that they are less effective than they judge that they should be, will attempt to rectify the situation. Improvements in team effectiveness might be achieved by adopting a different plan, utilizing

available resources, e.g., a knowledge base, or reorganizing the existing team, or the team might determine that it could benefit from the assistance of an additional agent. This assistance could take many forms. An existing agent might be recruited as a new member of the team, or an agent might be created to join the team. Services might be obtained from a non-member agent, or a virtual agent might be instantiated solely to provide some service for the team. The social structure of the team will determine how the team decides to obtain assistance. If the team is egalitarian, consensus might be required, but if the team is under the direction of a manager, the manager alone might make the decision to seek assistance.

A team can possibly benefit from the services of multiple assistants, and thus might utilize separate assistants for separate tasks. These assistants need not be the same type, and may be organized into a hierarchy, with meta-assistants supervising sub-assistants. Assistants are facilitators, improving the effectiveness of teams. A distinction is made between the autonomous agent assistants that are considered here, and non-autonomous assistance as provided by, e.g., an expert system.

Motivations for invoking a virtual agent can include the need to distribute the workload. Effective teamwork requires considerable contribution of resources to the team effort. A plan must be discovered and agreed upon, and this activity can require extensive negotiation and communication. Cooperation requires monitoring the physical and intentional states of other team members. Before an agent delegates a task to another agent, the delegating agent must make an assessment of the trustworthiness, loyalty, and competence of the other agent. These activities often involve reduplication of efforts. In addition, different agents may be better suited to particular activities than other agents and an agent's abilities may wax and wane over time. In a sense/plan/act model, communicating the distributed sense data to a virtual planning agent who then communicates the action determined for each member, could greatly reduce the total amount of calculation, negotiation, and communication needed to determine a joint plan for the team.

The ideas, notation, and formalism found in this paper are constrained by the logic used to model autonomous agents, namely, multi-modal temporal BDI logic. In the remainder of this paper, we first review our account of teaming and the mental constituents required for teaming in section 2.

The requirements include team members to have nontrivial levels of ability, awareness, and benefit. Section 3 considers agent social abilities, while section 4 suggests augmenting teams with team assistants, virtual or otherwise. Section 4 further considers various subjective measures of team effectiveness such as efficiency, influence, dependence and redundancy. Section 5 addresses related work.

2. Properties of Teams

Here we review criteria set forth in [Beavers, Hexmoor 2001] for the existence of a team. In order for a team to exist the agents forming the team must have a shared mental state, which includes the following four components:

1. The team must have a joint intention, in the sense that there is at least one intention ψ such that each agent of the team intends ψ , and believes that the other team members also intend ψ . An intention is a way that the future could unfold. Intentionality requires autonomy.
2. Each individual team member must believe of itself, and of the team, that sufficient ability is available to achieve the joint intention of the team.
3. Each individual team member must be aware of being a team member.
4. Each individual team member must maintain a cooperative stance toward the recognized members of the team with respect to the shared goal.

Teams are individuated by their joint intention and membership. *Intention* depends upon *ability* and *choice*. In order to be a team member, an agent must possess an ability (in the sense of being able to perform an action), which contributes to the achievement of the team goal. Using the development of *Can* in [Wooldridge 2000] as inspiration, our notion of *ability* makes it explicit that it is the action (or series of actions) α , that is expected to bring about the state ψ . In the case that there are multiple actions which lead to similar states, an agent may be permitted to perform some of those actions, but not others. This leads to a quantified notion of ability. The agent i has ability to achieve the state ψ if and only if there is an action, or series of actions, α that i knows i can perform and i knows that α will achieve ψ in some possible world. Generalizing to the multi-agent case yields “the group g has the ability to achieve the state ψ if and only if there is an action, or series of actions, α that g collectively knows g can collectively perform and g collectively knows that α will achieve ψ in some possible world.”

Can stands for the ability of a single agent, whereas “J-*Can*” stands for the joint ability of a group of agents. In J-*Can*, the agents individually have a belief that they collectively have the ability to achieve the goal and they know of an action, or series of actions, that will lead to the achievement of the goal. The team members believe that they collectively have some ability to achieve a shared objective, that is, they believe the group to have adequate collective ability, including know-how. Each individual team member believes that it has some ability that could

contribute to the achievement of the shared goal and collectively the team must have some non-trivial probability of achieving its intention. An agent that has the ability to achieve some state ψ may also desire to have this be the case, may have permission to perform an action to bring that state about. Permissions are derived from social considerations. The combination of ability, desire, permission and autonomy enables the adoption of an intention.

The members of a team with the ability to achieve a goal have a joint intention to achieve the goal. Adopting an intention implies that the agent has some degree of autonomy. Autonomy is agent self-determination in a non-reactive manner. Intentionality distinguishes swarms from teams. Ants in a swarm are too simple to be able to form intentions, although they do individually make reactive decisions that collectively lead to coherent group behavior as judged by an external observer. Individual ants are unable to persuade other ants to pursue a goal other than the one reactively adopted.

We also distinguish between a team, which is formed to achieve a specific goal and a coalition, which is a work group formed to increase the mutual benefit of the individual group members. Members of a team share a joint persistent goal while members of a coalition share a joint utility and strive for mutual benefit. Teams are work groups that attempt to achieve a goal even though in doing so they may engage in activities that adversely affect their individual condition. The team persists so long as the joint goal persists. Coalitions, on the other hand, are work groups that pay constant attention to their individual benefit. Stability, viewed as the longevity of the work group, is desirable in coalition formation, however, the sense of stability comes from the expected longevity of benefit and there is no notion of a determinate persistent intention as there is when a team is seeking a goal. In comparison, teams seem more suitable when agents must reason about goal achievement.

An agent’s autonomy with respect to an action is a relative measure of its liberty to exercise individual preference with respect to that action. This individual preference is constrained by social considerations in that it includes consideration of other agents, how the other agents contribute to the agent’s sense of freedom to choose and to the agent’s performance and how it prefers to work with others. These issues may be combined to yield an assessment of an agent’s *permission* to perform an action. The personal autonomy (with respect to a given goal) of an individual might vary from complete self-determination (internal, high individual autonomy) of every aspect of the agent’s relationship towards the goal, to complete other-determination (external, low individual autonomy) of those aspects. The notions of group autonomy and cooperation can be developed in a similar manner. Autonomy is a complex and subtle concept that is here treated as an unexplained primitive. Variants of autonomy might include quantified versions such as “ i is permitted to achieve ψ by any means possible” and “ i is permitted to achieve ψ by

some means.” Group autonomy is conceived in terms of individual autonomy.

Agents with group autonomy could join a team. In addition to shared autonomy, agents should have a joint cooperative attitude. Agents must adopt the principle of social rationality in order to be cooperative [Jennings, Campos 1997]. Social rationality is used to develop a theory of joint responsibility. We assume that an agent adopts an intention only after consideration of its desires and its autonomy with respect to those desires. We assume that there is a measure of individual benefit that partially orders the possible worlds with respect to desire. An action is chosen by the agent to cause a transition to a desired state. We use the notation ($\text{Choice } i \alpha$) to indicate that the application of the choice function (which has performed a cost/benefit analysis) by the agent i yields the action α . The state at which the choice function is implemented will affect the result. An individual i has an intention to attain the state ψ through the action α if the individual has the autonomy to do so and so chooses. Commitment to the joint intention is part of being a team member and this commitment will demonstrate itself in the persistence of the intention. Joint intention is defined in the pattern of individual intention, namely, as joint autonomy and choice.

In addition to having a joint intention and joint ability the members of a team must be aware of their membership in the group and maintain a cooperative attitude towards the group with respect to the joint intention. Awareness captures an agent’s cognizance of at least some of its teammates’ intentions. We leave for a later paper an explanation of the means for becoming aware and an account of ongoing awareness maintenance (see [Kaminka, Pyandeth, Tambe, 2001]).

A part of the requirement for being a team is that the agents forming the team must individually intend that the team achieve the goal which requires coordinated action of the team members and this coordination in turn requires that the individual agents be able to determine when to start or stop an activity, when to switch activities, and when to adjust their contribution. A potential team member must have the ability to accept or reject a goal, that is, revise its intentions. Part of this ability is the autonomy to make this decision, that is, without the autonomy to adopt an individual intention an agent cannot become a team member.

Coordinated action generally requires an agent to have some knowledge of the activities of other team members. The agents forming the team must individually be aware of the existence of the team and intend to coordinate their actions with other perceived team members. Agents must know what the joint intention is and that they are part of the team with that particular intention.

An agent is aware of being a member of a team g that intends to achieve a state ψ if and only if it has a true belief that it is a member of the team, and that the team has both joint ability and joint intention to achieve ψ . Note that awareness requires only a belief that the agent is a member of the group, and does not require that the agent know who the other members of the group are. A group g of agents is

jointly aware of its intention to bring about the state ψ if and only if each agent of the group g is individually aware.

An analysis of the notion of cooperative attitude has not yet been performed, and thus we include it as a primitive component of teaming. The definition of a team can now be given: “A group of agents is a team if and only if the group has a joint intention, group autonomy, joint benefit, and joint awareness.” We have formally stated the notions in this section in [Beavers and Hexmoor, 2001]. Space limits us from using more formal notations in this paper.

3. Social Abilities

Teams of social agents have other criteria beyond those addressed in section 2. We have been working on a conceptual model that adds mental states capable of dealing with social elements to BDI architectures. It has long been recognized that social elements enhance multi-agent performance. Extending BDI architectures to incorporate norms, values and obligations are presented has been considered. Norms and values endow multi-agent systems with social attributes thereby facilitating the cooperative behavior of agents.

Awareness of norms and values among agents can be used by agents to adapt to changing conditions. One agent must recognize the values and norms being observed by other agents in order for the agent to respond appropriately to the actions of the other agents. Without this recognition an agent, for example, might be benevolent and have the well-being of other agents in mind, but work at cross purposes to those agents because it does not recognize the plan being followed by the other agents, or the roles that those agents have adopted or the norms that are associated with those roles.

To see the importance of values, roles, and norms in cooperative activity consider the following concrete example. Suppose that three agents have formed a team to fulfill the duties at the framastat installation station on a widget assembly line. They first agree to adopt the values “be a team player”, “work efficiently”, “be independent”, and “don’t interfere with the activities of other agents on your team” since these values are seen to be to their mutual benefit. Note that the value “be independent” and “be a team player” are in some conflict while the value “be a team player” seems to subsume the value “don’t interfere”. The agents then agree to assume the roles required to fulfill their duties. Suppose that their duties include (i) two agents guide and push the widget to the framastat installation station, (ii) one agent retrieves the framastat and inspects it while another agent obtains the fasteners, (iii) one agent positions the framastat and holds it in place while the second and third agents place and secure the fasteners, (iv) the widget/framastat assembly is inspected, and finally, (v) the widget is guided to the next station on the assembly line. The roles seem to be ‘pusher’, ‘navigator’, ‘framastat fetcher’, ‘fastener fetcher’, ‘framastat inspector’, ‘framastat positioner’, ‘framastat fastener’, and ‘widget/framastat inspector’. It is apparent that the adopted roles have dependency among them. Furthermore, it is possible that the agents will undertake the completion of their respective

tasks in such a manner that there is redundancy among them. Redundancy is between tasks. For example, multiple agents might duplicate the inspection of the widget on the assembly line. The agents attempt to remedy redundancy and dependence illustrates adherence to two of the values given above.

Each role will have norms associated with it. For example roles requiring agents to work in close quarters might include “don’t come into contact with other agents”, and “always perform actions in the same way so that your actions are easily interpreted”.

4. Team Assistants

A team assistant is an agent recruited or created in order to improve the effectiveness of a team. Two types of agent assistants are manager agents and service agents. A team recognizing that it needs direction, guidance, coordination, or leadership may chose to invoke a manager agent to fulfill these needs. For an example of a service agent, consider a team that recognizes that individual team members are performing redundant calculations, and therefore chooses to invoke a virtual agent to perform the calculation for all agents and post the result to a blackboard. A team might also need a virtual agent to fulfill the role of social facilitator or councilor or provider of expert advice. Middle agents (as developed in [Wong, Sycara 2000]) constitute yet another type of team assistant. Team assistants will often be virtual, but whether or not the agents are *virtual* is not important. Teams might continue to function without these assistants. The team assistants are facilitators, improving the effectiveness of teams.

Efficiency, robustness (fault tolerance, ability to recover from error, etc.), cohesion, and redundancy affect the effectiveness of a team. Thus, the recognition that the team’s efficiency, robustness, cohesion, or redundancy needs improvement might cause a team to seek assistance or decide to invoke a virtual agent. Some subjective measures of efficiency, robustness, cohesion, and redundancy are given later in this section.

Goal achievement generally requires adoption of a plan, assignment of roles, adoption of norms (standards, and conventions). These tasks might be more efficiently handled by a virtual agent. Uses for virtual agents include maintaining a blackboard of common knowledge, coordinating communication, developing strategy, planning, role determination and assignment, monitoring of norm compliance, common mental state maintenance, monitoring of common intentions, and computation.

Much as with human agents, it is probable that effective teams will require that individual agent team members have confidence in a manager agent. Therefore, the types of agents constructed for teaming will probably need to be endowed with the ability to recognize the need for a virtual agent and the ability to evaluate that agent’s performance. These abilities would allow for the development of subjective or internal recognition of the need for an agent assistant. On the other hand, if the team is not completely autonomous, that need could be determined by an agent

outside the team. In this case, the determination is considered external or objective.

Effectiveness is considered as a measure of the likelihood that the team will achieve its goal. This measure might be subjectively determined by the team members, or objectively by an entity outside the team. With respect to effectiveness, measures for the following properties of teams are of interest: 1. efficiency, considered as a measure of the resources expected to be used by the team in its attempt to achieve the team goal, 2. robustness, considered as a measure of the team’s ability to recover from unanticipated circumstances, the failure of actions to achieve (sub)goals, and having team members withdraw from the team, 3. cohesion, considered as a measure of cooperation and synergy, and 4. redundancy, considered as a measure of the degree to which different agents each attempt to accomplish the same task.

These measures are not independent. Generally, an increase in the efficiency, robustness or cohesion of a team, or a decrease in redundancy, will yield an increase in the effectiveness of a team.

We now consider examples of these subjective measures that could cause a team to recognize the need to invoke a team assistant. The first measure is a subjective estimate by each team member of the perceived efficiency or ability of other team members in relation to the overall team objective. If an agent perceives a significant shortfall in the efficiency of another agent, it may announce this impairment so that the team can consider obtaining assistance. An agent’s assessment of its own efficiency is omitted in this example. Consider each agent’s estimation of another agent’s performance as a percent deviation from what is required. In the following table we show three agents. Each agent is considered to be fully effective at level 1.0. Each cell in the table from row i to column j contains agent i’s estimation of agent j’s efficiency. For instance, the table indicates that agent 2 believes agent 3 to be 30% inefficient. The first three entries in the last row are each agent’s averaged efficiency estimate. For instance, agent 2 is estimated to be 90% efficient by agents 1 and 3. Column 4 contains the sum of each agent’s perceptions about other agents. For example, agent 2 believes there is 40% ineffectiveness among agents 2 and 3. The values in column 4 are added to get an overall perception of the efficiency of the team, which is 20% ineffectiveness, in this case.

Once the average efficiency of each agent is computed (the first three values in the last row), these values may be used to compute a overall subjective measure of efficiency by computing weighted average of the other agent’s estimates of the agent’s efficiency. These weights are the agent’s overall score (first three values in the last row). The idea is to differentially weight efficiency estimates of agents toward one another. The estimate of a more efficient agent will count more heavily than the estimate of a less efficient agent. Weighted averages are shown in the rightmost column. For instance, agent 3’s estimated efficiency is 0.91. By averaging the values in the last

column, we arrive at an overall estimate of efficiency, 84% in this example.

The second illustrative measure is a subjective measure of influence. Consider each agent's perception of influence with another agent as a deviation from 1.0. In the following table, each cell contains agent i's perceived influence with respect to agent j. For instance, agent 2 believes agent 3 has a 20% positive influence on it.

	1	2	3		
1	-	+0.1	+0.2	+0.3	0.88
2	-0.1	-	-0.3	-0.4	0.77
3	+0.1	-0.2	-	-0.1	0.91
	1.0	0.9	0.9	-0.2	0.84

Table 1 Efficiency

The last row contains the average influence produced by that agent. In this case, agent 3 has produced 40% positive influence, while agent 2 has produced a 10% negative influence. Column 4 shows, for each agent, the average influence that it has experienced from the other agents. In this example, all three agents have experienced 10% synergy. If we add the amount of influence produced and subtract it from the amount experienced, there is a net +0.2 influence. We can think of this group as being generally mutually supportive.

	1	2	3	
1	-	-0.1	+0.2	+0.1
2	-0.1	-	+0.2	+0.1
3	+0.3	-0.2	-	+0.1
	+0.2	-0.1	+0.4	+0.2

Table 2 Influence

Robustness is the ability to maintain team effectiveness in light of possible failures. Dependence of one agent on other agents detracts from robustness. Each team member may estimate its dependence on another member which is represented as a deviation from independence. The following table again shows three agents. Each agent is considered fully independent at level 0.0. Each cell in the table from row i to column j is agent i's dependence on agent j. For instance, agent 2 depends on agent 3 by 20%, which can be interpreted as agent 2 cannot finish 20% of its task without agent 3 completing its task. The first three entries in the last row are each agent's total dependence on others. For instance, agent 1 depends upon agent 2 by 10%. These values reflect the team's dependence on the agent. In this example, agent 2 is the most dispensable whereas agent 1 is the most needed.

	1	2	3	
1	-	0.1	0.3	0.6
2	0.7	-	0.2	1.3
3	0.1	0.0	-	0.2
	0.8	0.1	0.6	

Table 3 Dependence

In column 4, we show the sum of each agent's dependences on other agents. For example, agent 1 depends on agents 2 and agent 3 by 10% and 30% respectively. Since agent 2 depends on agent 3 by 20%, agent 1 indirectly depends on agent 3 by another 20%, which adds up to 60%. In this example, agent 1 is the most dependent on other members of the team, whereas agent 3 is least dependent.

Agents might duplicate tasks in part or in whole. For example, assembly inspection of a widget on an assembly line by all agents on the line is redundant and should be done only once. In the following table each cell from row i to column j contains agent i's overlap with agent j. For instance, agent 2 overlaps with agent 3 by 20%. Overlap is symmetric. In the fourth column we add the overlap with other agents and see the agent's total overlap. For example, agent 3 has 50% overlap with the agents 1 and 2.

	1	2	3	
1	-	0.1	0.3	0.4
2	0.1	-	0.2	0.3
3	0.3	0.2	-	0.5

Table 4 Overlap and Repetitions

Consider the largest overlap in table 4, which is between agent 1 and 3. A service agent might be useful in eliminating this overlap. This agent can either perform the duplicated action or keep track of which of the two agents 1 and 3 performs the overlapped task and alert the other agent of the completion of this action so it can be skipped. The service agent may also pay special attention to agent 3 since it has the most overlap with the other two agents. Consider a case where an agent has an overall negative influence on other agents and the group depends least on the agent. Such an agent might be the one to be evicted from the group. In the examples in this section, agent 2 comes closest to this case since the group depends on the agent by only 10% (table 3) and it has a 10% negative overall impact on the group (table 2).

5. Related Work

[Pechoucek et al 2001] suggests the use of *meta-agents* to coordinate the flow of information across the agent community to facilitate project-oriented production. These meta-agents illustrate one type of agent assistant. [Sycara 2001] and [Wong, Sycara 2000] provide another type of agent assistant with *middle agents*. They are considering large open agent groups on the Internet. Much like public

services of a town, Sycara's middle agents provide general functions of playing a middle-agent between agents. Our agent assistants have a much more specific task of making sure the motivating factors for a team continue to exist and the team functions well. Middle agents need not be team members. We take teams much more seriously and conceive of teams of agents with highly coupled tasks.

Work in Milind Tambe's group has been very visible [Kaminka, Pyandeth, Tambe, 2001]. Recent developments from this group include a monitor for distributed teams that minimizes the amount of required communication through having the monitor infer the progress of team plans by observing the actions of member agents. Tambe's Overseer is a good example of agent assistants proposed in [Hexmoor, and Beavers, 2002]. Inferring agent states from team actions suffers from uncertainty and from growth of computational complexity that is exponential in the number of agents being monitored. By modeling the group as a whole, rather than modeling the individual agents constituting the group, Overseer is able to effectively predict team responses during normal and failed plan execution in linear time. The gain is accomplished through restricting attention to coherent monitoring hypotheses. The trade off is that the team is not accurately modeled in some failure modes because an inaccurate, but efficient, temporal model is used. Knowledge bases containing the plan hierarchy, and containing the group hierarchy enable *socially attentive* monitoring. Tambe's approach blends reasoning about teamwork and reasoning about uncertainty.

6. Conclusion

We have reviewed our formal account of teaming that relies on some joint understanding of a set of mental notions: intentions, autonomy, awareness, and benefit. In particular, nontrivial abilities of team members, inter-agent awareness, and benefit were introduced as additional criteria for a team. This paper has extended the treatment of teams to include team assistants, has considered motivations for obtaining assistants, and has suggested subjective measures for efficiency, influence, dependence, and redundancy, which all affect the effectiveness of a team. An estimate of effectiveness is to be used in evaluating the need for team assistants. We have shown preliminary motivations for assistant agents. Mechanisms are needed for agents to reason over team effectiveness measures and to propose team assistants. Subsequently, some protocols are needed for recruiting and enabling team assistants. Development of such protocols is part of our ongoing activity.

Acknowledgements

We thank Scott DeLoach for helpful comments. This work is supported by AFOSR grant F49620-00-1-0302.

References

- G. Beavers and H. Hexmoor, 2001. Teams of Agents, in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics* (SMC'2001).
- H. Hexmoor, and G. Beavers, 2002. Measuring Team Effectiveness, In proceedings of 20th IASTED International Multiconference: Applied Informatics (AI 2002), Innsbruck, Austria
- N. R. Jennings and J. R. Campos, 1997. Towards a Social Level Characterisation of Socially Responsible Agents, in: *IEEE Proceedings on Software Engineering*, 144 (1), 1997, 11-25.
- G. Kaminka, D. Pyandeth, M. Tambe, Monitoring Deployed Agent Teams, *International Conference on Autonomous Agents* (Agents '01), 2001.
- M. Pechoucek, F. Macurek, P. Tichy, O. Stepankova, and V. Marik, 1999. Meta-agent: A Workflow Mediator in Multi-Agent Systems, In McIntosh, A. Taylor R., editor(s), PAKeM '99: Second International Conference of Practical Applications of Knowledge Management, pages 115-133, London, May 1999.
- K. Sycara, 2001. Brokering and Matchmaking for Coordination of Agent Societies: A Survey. In *Coordination of Internet Agents*, A. Omicini et al. (eds.), Springer, 2001.
- R. Tuomela, 2000. *Cooperation: A Philosophical Study*, *Philosophical Studies Series*, Kluwer Academic Publishers.
- M. Wooldridge, 2000. *Reasoning about Rational Agents*, MIT Press.
- H.C. Wong and K. Sycara, 2000. A Taxonomy of Middle agents for the Internet, In *Proceedings of the Fourth International Conference on Multi-Agent Systems* (ICMAS'2000).

