Weakly Dependent Agents May Rely on Their Luck or Collaboration: A Case for Adaptation

Henry Hexmoor
Computer Science & Computer Engineering Department
Engineering Hall, Room 313, Fayetteville, AR 72701
Hexmoor@uark.edu

Abstract

We contrast strategies for agent interaction against several parameters when each agent has a set of acyclic ordered list of goals to perform. After developing a set of metrics for utilities of agent parameters, we suggest that such an agent must determine the most suitable interaction (adaptation) for itself in each domain.

1 Introduction

Two agents may socially depend on one another in one of 16 ways (Alonso, 1999). A type of social dependence he calls *weak dependence* is when two agents A and B are capable of achieving a goal individually but may prefer to have the other agent achieve it. If both agents prefer this, it is called *reciprocal or mutual, weak dependence,* which is the focus of this paper. Under the right circumstances, by sharing goal achievement, a weakly dependent agent's workload might be lowered. This is often at a temporal cost since agents may have to wait until dependencies are fulfilled. Sharing of goal achievement is central to producing agent teams (Cohen, Levesque, and Smith 1997). This paper extends our explorations of motivations for teaming (Hexmoor and Duchscherer, 2001). Agents who want to maximize their gain may consider negotiation prior to collaboration in order to lower their workload and temporal penalties (Beer, et al, 1999). Deals made in negotiation may in turn incur a cost and vary from goal to goal and agent to agent. We will argue that since newcomer agents may not know the overall outcomes of entering collaboration in terms of gains and losses they should consider adaptive techniques for collaboration. Finally, agents who cannot see other agents' goal structure but may still benefit from their goal execution, known as stigmergy (Beckers, et al , 1994), may become inclined toward collaboration or rely on serendipity of stigmergy. Agents should learn to predict the level of expected stimergy and when to abandon negotiation in favor of stimergy.

In this paper, we are making a case for how agents may benefit from collaboration. Our agents are self-interested and do not care about gain and loss in others (Brainov, 1996).

Here we will not favor any particular machine learning such as reinforcement learning (Kaelbling, et al. 1996) or logic-based approaches (Alonso and Kudenko, 1999). We do provide metrics for decision-making, which can be used in adapting a learning algorithm to adaptive deliberation over collaboration.

In general, adaptation as a beneficial change is broader than learning (Joshi, et al, 1997]). We provide scenarios and metrics that exhibit several common situations in multiagent systems. Other than learning, empirical adaptation methods can be used to detect the most applicable scenario and metric. For instance, agents can use our scenarios to decide whether to engage in negotiation as an adaptation. Norms as adaptive mechanisms can be designed using our scenarios and metrics (Fitoussi and Tennenholtz, 2000). Agents may lie and provide false information about their goal structure in order to exploit other agents. This is yet another reason adaptive collaboration.

In this paper we present our model of agents with their own goal dependency and then show utilities used by agents to deliberate about selecting partners to reduce their workload. The tradeoff between reduction in workload and temporal penalty is also discussed. We then turn to costs involved in deals and discuss how agents can benefit by reinforcement learning.

2 Weak Social Dependence

Let's consider a finite number of agents, each agreeing to achieve a partially ordered set of goals. The ordering among the goals of each agent is goal precedence such that each agent builds a directed a-cyclic graph (DAG) of its goal set. This is different than AND/OR goals trees used in (Alonso, 1999). In an AND/OR tree, only the leaves are actions that form a plan when they appear AND related. Alonso calls a goal directly associated with an action *bottom goal*. In our model, all goals are *bottom goals* and the partial ordering of our goals is a nonlinear plan, as in (Sacerdotti, 1977).

Agents who consider teaming can consider the potential payoff in reducing redundancy by considering composition of their goal DAGs. Consider two agents A and B and two goals 1, and 2 with the following respective DAGs: 1→ 2 and 2→ 1. These agents would not team since the combined DAG will have a cycle.

We make the following assumptions:
1. Goals are equal in cost and the cost is nontrivial.
2. In a group of agents, goal achievement is uniform and the quality of goal achievement does not change from one agent to another agent.
3. Each goal takes one unit of time for achievement and during that period, agents do not see one another's activity.
4. In a group of agents, an accomplished goal is visible to all agents and performing the goal once is as good as multiple times provided the agent makes available its results and the goal's precedence is correct. This goal sharing can be accomplished as being part of a team (i.e., explicit agreement to share a goal) that share the goal. Another way for sharing is to make the results public then any agent may use it and sharing is implicit, Else, communication is by simple stigmergy. This assumption is not benevolence as the agent is indifferent whether to anyone uses the results of its accomplishment.
5 There are two options here:
a. Agents see one another's goals and their dependency.
b. Agents do not see another agent's goal or their goal dependency.
6. The agents are self-interested (indifferent) towards other agents; they are not benevolent or malevolent [Brainov, 1996].

In order to make these examples concrete, consider a number of satellites that orbit the Earth and are tasked with a sequence or pattern of pictures of geographic locations they need to collect. Each picture is a satellite goal. The satellites may either make the pictures publicly accessible or through an agreement as in team formation. In the remainder of this section we discuss a variety of goal patterns.

Example (shared second goal): Consider two agents A and B and three goals 1, 2, and 3 with the following respective DAGs: 1→ 2 and 3→ 2. Both these agents want goal 2 to be done. If they collaborated under assumption 5a, only one needs to do goal 2, whereas separately each must perform goal 2. There is a saving of one goal in the total number of goals executed. Temporally, if we assume goals 1 and 3 can be carried simultaneously, the goals are done in 2 units of time. If we extend this to n agents and n goals, the goal savings is n-1. There is no change in time with collaboration. Under assumption 5b, there is no savings.

Example (shared first goal): Consider two agents A and B and three goals 1, 2, and 3 with the following respective DAGs: 1→ 2 and 1→ 3. Both these agents want goal 1 to be done. Under assumption 5a, if they collaborated, only one needs to do goal 1, whereas separately each must perform goal 1. There is a savings of one goal in the total number of goals executed. Temporally, the goals are done in 2 units of time since goals 2 and 3 can be simultaneous goals. If we extend this to n agents and n goals, the savings is n-1. Under assumption 5b, there is no savings.

Example (goal dependency chain): Consider three agents A, B, and C and three goals 1, 2, and 3 with the following respective DAGs: 1→ 2, 3→ 2, and 1→ 3. An alterative but equivalent DAG is : 1→ 2, 2→ 3, and 1→ 3. In the 3 way combined DAG, goals have to be done in the order 1-3-2 (or 1-2-3 in the alternative DAG). Under assumption 5b, this example has no savings. Under assumption 5a, the team will save 3 goals but suffer a time lag of one unit. Each agent's average job is reduced to 1. However, the agents may have uneven split of jobs. For example, it is possible for two agents to carry out the goals and for one agent to do nothing. One agent can be redundant.

Let's extend this example to n agents (A1 … An) and n jobs (1 … n) with the following DAGs: 1→ 2, 2→ 3, …, n-1 → n, 1→ n. The combined DAG is rooted at 1 and the edges form is a path from 1 to n. This team will save n goals since average goal for each agent is one. The team suffers a time lag of n-2, which is large. Among agents A1, A2 and A3, agent A2 could rest and benefit from agents A1 and A3 doing their collective goals1, 2 and 3. Therefore A2 could be redundant. At most, A2, A4, …, An-2 could be redundant. In general, agent redundancy is (n/2)-1, almost half. Although reality of whether agents choose to be redundant or overactive in a given team may not be known at the outset, agents know that they could *on the average do only one job instead of two*. This fact can influence their autonomy consideration of whether they proceed with their goal by themselves or become part of a team.

Example (maximal dependency): Consider 4 agents A, B, C, D and 4 goals 1, 2, 3, and 4 with the following respective DAGs: A1: 1→ 2, 1→ 3, 1→ 4; A2: 2→ 3, 2→ 4, A3: 3 → 4, A4: 4. Agent A1 has 3 precedence relations, A2 has 2 precedence relations, A3 has one precedence relation, and A4 has a single goal and no precedence relation. Under assumption 5b, if each agent starts to execute its tasks, goals 1, 2, 3, and 4 will be completed in first time unit. In time unit 2, agent B will consider one of its dependent goals, for example goal 2.Although this goal is executed in time unit 1, it is of no use to A since it is executed too early and not after goal 1. If agent B performed task 3, and agent C performed goal 4, all goals will be compete in the second time unit and all in the right order. Agent A benefits from this stigmergy by having goals 3 and 4 performed by others. Agent B benefits from stigmergy by having goal 4 performed by agent C. Agents C and D do not see any benefit.

Under assumption 5a, such a team will save 10-4 = 6 goals. The redundancy in agents is non-uniform. If A1 does all its goals, others can be redundant. Agent A4 can rely on agents A1, A2, and A3 to perform its goal. A4 may enter a team and not perform any goals. Whereas, A1 enters the team having to do at least one goal. 3 agents might be redundant.

Let's extend this example to n agents A1 … An and n goals 1 … n with the following respective DAGs. A1: 1→ 2, 1→ 3, …, 1→ n; A2: 2→ 3, 2→ 4, …, 2→ n, …, An: n. Agent A1 has n-1 precedence relations, A2 has n-2 precedence relations, and on until An has a single goal with no precedence relation. Under assumption 5a, goal latency for agents is uneven. An must wait n-1 time units. In general, agents wait n-1 units of time. At most, n-1 agents are redundant if A1 does all their goals. Such a team will save n(n+1)/2 – n goals with an average per agent savings of n/2 –1. Using this, the savings of large teams are far more than savings of smaller team. Average savings may be meaningful for A1 but not for An. If n = 2, A2 will have no benefit from teaming. Under assumption 5b, agent A1 can at most benefit n-2 goals performed by other agents, agent A2 will at most benefits from n-3 goals being done by other agents, and so forth.

Example (goal bi-partite graph): Consider two agents A and B and four goals 1, 2, 3, and 4 with the following respective DAGs: A: 1→ 2, 1→ 3, B: 4 → 2, 4 → 3. The combined graph is a bi-partite graph. Both these agents want goal 2 and 3 to be done. Under assumption 5a, if they collaborate, only one needs to do goal 2 and 3, whereas separately each must perform goal 2 and goal 3. There is a saving of two goals in the total number of goals executed. Temporally, there can be a huge penalty. If either agent performs its first goals, it must wait until the other agent executes its first as well before they can proceed to the dependent goals. This team needs to coordinate their goal to minimize waiting for one another. If we assume simultaneous execution of their first goals, then there is no temporal penalty. Under assumption 5b, at best they will do a different second goal and save one goal each.

Let's extend this example to n agents and n goals. Under assumption 5a, such a team will perform $n^2/4 + n/2$ goals as opposed to n goals. The goal savings is $n^2/4 – n/2$, which is good but not as good as the maximal example. Under assumption 5b, each agent at most can save n/2 - 1 goals.

3 Metrics and Utilities

In contrast to global savings, here we will take an individualistic and utilitarian view of motivation for agents. Each agent must examine its own gains in its consideration of teaming. In this section we will define concepts to measure an agent's utility of interaction with other agents.

If multiple agents might perform a given goal, agents might have non-uniform chance of performing the goal. Let L_i^j denote agent i's likelihood of performing goal j. For example, in the example of maximal dependency with 4 agents and 4 goals under assumption 5a, if we assume each agent who wishes an action to be executed will have equal chance of performing a goal, agent 3 considering goal 4 will have 1 in 4 chance of performing it, $L_3^4 = 0.25$. If we averaged an agent's likelihood of goal performance over all its goals in a given group, we will have our first metric that reflects the agent's average likelihood of performing a goal in the group. We will denote this metric for agent i by Λ_i. I.e., $\Lambda = 1/n\ \Sigma_1^n\ (L_i)$. We suggest that Λ can be used as a parameter to be used in motivation value for teaming. For agents that must perform their own goal, Λ = 1. An agent's utility (U) can be defined as the gain it derives, i.e., compliment of having to absolutely do the job, $U^1 = 1 - \Lambda$.

Let's consider a variation of the U^1 metric. Instead of workload, let's consider expected cost C of collaboration with equal probabilities of task preference as before (or its compliment, expected savings). This metric considers all options with their respective probability and costs. We will use the same example of maximal dependency with 4 agents and 4 goals under assumption 5a. We will assume each task takes one unit of cost. For example:
$C_1 = 1 + (1/2)(1/3)(1/4)3 + (1/2)(2/3)(1/4)2 + (1/2)(2/3)(3/4) + (1/2)(1/3)(1/4)2 + (1/2)(2/3)(3/4)3 + (1/2)(2/3)(1/4) = 53/24$.
Along the same lines, $C_2 = 11/12$, $C_3 = 6/12$, $C_4 = ¼$.
The alternative metric is derived by subtracting the cost without collaboration from the expected cost, $U^{1'}$.

Other than workload, timely performance is an important consideration. To extend Λ, we will define T as a temporal motivation. If an agent may have to put off its goal achievement, it may pay a penalty of latency. Let's assume each goal takes one unit of time for achievement. Then an agent who is delayed N units of time for a goal, will incur a penalty of Π. For example we may define $\Pi = N/(N+1)$ for that goal. An agent with a single goal without considerations of other agents can be most timely by achieving its goal immediately and its penalty is 0. For agents with multiple goals, we can average their temporal latencies. We will now consider a utility parameter that is a linear combination of workload and temporal penalty: $U^2 = C1 * (1 - \Lambda) - C2 * \Pi$. C1 and C2 are relative importance of workload versus time and C1 <= 1.0, C2 <= 1.0.

We will consider that deals incur a cost. In fact, deals for each agent and each goal are different. Let's extend U^2 to include cost of deals. Let's denote the actual cost of agent i's deal with agent j about goal k as δ_{ij}^k. But since in general agents won't have access to this actual value they might have an expected value that might be based on prior experiences. If we assume $\delta_{ij}^k = D/(D+1)$ where D is the number of deals. Let $\delta\pi$ denote an agent's average projected costs of its deals. The new utility function has to account for the projected cost of deals: $U^3 = U^2 - C3 * \delta\pi$. Here, C1 <= 1.0, C2 <= 1.0, and C3 <= 1.0.

Let's consider assumption 5b where agents do not see one another but their actions may benefit one another. There are certain circumstances when agents who could ordinarily deal and negotiate with other agents might find it more advantageous to focus on their own goals without regard or attempt to communicate with other agents and instead use timely executed goals by others. This is the basis of relying on stigmergy. Stigmergy is not dependable and all an agent can do is to determine experientially the most it has benefited from other agent's past actions. If an agent can view other agent's goal structure it can determine the most fortuitous situation when other agents' goal performance can benefit them. Let's denote average workload of an agent assuming most fortuitous stigmergy by Γ. A utility based on this situation can be defined as $U^4 = 1 - \Gamma$.

4 Strategies for Interaction

Interaction is consideration of interface with other deliberating agents when making a decision to act on a goal. The need for modeling and understanding types of interaction can be found in e-commerce, auction portals with multiple buyers and sellers, across organizations as well as among individuals, and large space missions with scarce resources. Electronic negotiation in such domains can be counterproductive and the agents must economize (Faratin, Sierra, and Jennings, 1998). The scenarios in this section are a large class of problems in multiagency. In these strategies consider maximal dependency with 4 agents and 4 goals and assume 5a. We use this example for illustration but the ideas apply in general.

4.1 Non-negotiation

If several agents wish for a goal to be performed, each will be equally likely to perform it without negotiation. $\Lambda_1 = 25/48$, $\Lambda_2 = 13/36$, $\Lambda_3 = 7/24$, and $\Lambda_4 = ¼$ and $U_1^1 = 23/48$, $U_2^1 = 23/36$, $U_3^1 = 17/24$, and $U_4^1 = ¾$. Agent 1 gains the least while agent 4 gains the most. If we consider this example with a variation of $U^{1'}_1 = 43/24$, $U^{1'}_2 = 25/12$, $U^{1'}_3 = 18/12$, $U^{1'}_4 = ¾$. With this modified metric, agent 2 clearly gains the most while agent 4 gains the least.

4.2 Non-negotiation and temporal penalty

Let's consider temporal penalties with C1 = C2 = 1.0 and considerations of section 4.1. Then $\Pi1 = 0$, $\Pi2 = ½$, $\Pi3 = 2/3$, $\Pi4 = ¾$, $U_1^2 = 23/48$, $U_2^2 = 5/36$, $U_3^2 = -1/24$, and $U_4^2 = 0$. Now agents 1 and 2 still have some gains whereas agent 3 will sustain a loss and agent 4 is even.

4.3 Negotiation

An agent facing this disparity in motivation might consider that is fairer if all team members had an equal average likelihood of performing goals. In this case, agents need to enter a negotiation bargaining for different probabilities of considering goals in order to attain an equal Λ for each of them. For example, in the example of maximal dependency with 4 agents and 4 goals, if we desire the agents to have an equal Λ_i, a solution is derived by a system of quadratic equations that assume equal average workload for all agents (i.e., equal Λ) with the following results (recall that L_i^j is the likelihood of agent i attempting goal j): $L_1^1 = 1.0$, $L_1^2 = L_1^3 = L_1^4 = 2/15$, $L_2^2 = 13/15$, $L_2^3 = L_2^4 = 3/15$, $L_3^3 = 2/3$, $L_3^4 = 31/180$, and $\Lambda_i = 6/15$ and $U_i^1 = 9/15$. In general, the system of equations for producing equal Λ is under constrained. This makes it a suitable for agents to negotiate and to arrive at equitable task decomposition.

4.4 Non-negotiation, temporal penalty, and deals

If all deals have equal weights: Agents 1 and 2 will perform 6 six deals each, agent 3 will do 5 deals, and agent 4 will do 3 deals. If we assume $\delta_{ij}^k = D/(D+1)$ where D is the number of deals and C1 = C2 = C3 = 1.0: $U_1^3 = 25/48 – 6/7$, $U_2^3 = 13/36 – 1/2 –6/7$, $U_3^3 = 17/24 – 2/3 - 5/6$, and $U_4^3 = 3/4 – 3/4 - 3/4$.

4.5 When serendipity is better/worse?

For example, consider the example of maximal dependency with 4 agents and 4 goals. Agent 1 can at most benefit to have 2 of its goal be done by other agents so under most fortunate circumstances, $\Gamma_1 = ½$, $\Gamma_2 = 2/3$, $\Gamma_3 = 1.0$, and $\Gamma_4 = 1.0$. $U_1^4 = 1/2$, $U_2^4 = 1/3$, $U_3^4 = 0.0$, and $U_4^4 = 0.0$. If we consider this example in the extended form with n agents and n task, at most fortunate $\Gamma_1 = 2/n$ and $U_1^4 = 1 – 2/n$. For agent A1, if n < 5, this is worse than negotiation. For agent A1, if n >= 5, this probability of workload is clearly as good or better than any outcome it could have secured by communication. Naturally, it is not clear that A1 will actually have $\Lambda_1 = 2/n$ since there is probability that stigmergy will not favorable. We conclude that agents must make observations in their environments and self-adjust to deal and communicate or rely on stigmergy. If they sense the presence of a high level of stigmergy, they should back off communication and merely use the fortuitous circumstances; otherwise, consider communication.

7 Related work

(Alonso, 1999) presents a model of society that focuses on an individual's agreements and negotiations. He discusses dependencies among agents and presents a calculus of negotiations that maximizes the agent's overall utility. We have adopted the notion of weak dependency from this work. Research with multi-robot experiments use the term team and society loosely as multiple robots that collectively perform a task. No analysis of teaming is offered. Agents are not endowed with abilities to reason about their team participation or team formation. (Jennings and Watts, 1998) suggest programming language constructs for teamwork written in Scheme, which is used to represent temporal structure of group plans. This system is applied to search and rescue experiments where multiple robots find and push large boxes but the agents are not endowed with abilities to reason about their own involvement. The programmer provides coordination schemes and plans for the group. (Mataric, 1998; Touzet 2000; Vaughan, et al 2000; Simmons, et al 2000) report on mobile robotic experiments that show cooperation among robots. The experiments describe cooperation algorithms. Although a central scheme governs coordination among agents, each robot does not independently enter and leave the group. (McCartney and Sun, 2000) describe a balanced binary tree data structure of tasks and their status to facilitate task allocation. This system assumes a globally shared plan with agents that are omnipotent and may take on any task. Unlike their claim to the contrary, since each agents must inform other agents of its view of task allocation and to broadcast its progress on its task, communication is high.

8 Conclusion

We have presented a scenario that agents who can ordinarily perform their own goals may either actively look into negotiation with others who plan to attempt the same goal or rely on luck of having their goals executed by others. We presented metrics and analysis of various cases. In realistic domains, agents may not have access to costs of deals as well as they may change their mind about what parameters such as time, workload, expected cost they care about. Such agents equipped with our metrics may compute the utility of their reliance on others and adapt their preference on goal execution.

Our assumption of likelihood of goal performance is not completely realistic since assuming equal probabilities is too simplistic. One way to address this to rely on agent's prior experiences of interactions to derive a estimation. Without making the estimation, computationally recursive, negotiations and agreements are useful. We will explore these alternatives future work.

Acknowledgements

This work is supported by AFOSR grant F49620-00-1-0302. I appreciate comments by Dr. Svet Brainov, Dr. Gordon Beavers and two anonymous reviewers.

References

E. Alonso, 1999. An individualistic approach to social action in Multi-Agent Systems. **Journal of Experimental and Theoretical Artificial Intelligence**, 11, 519-530.

E. Alonso and D. Kudenko, 1999. Machine Learning Techniques for Adaptive Logic-based Multi-Agent Systems. **The Second Workshop of the UK Special Interest Group on Multi-Agent Systems, UKMAS-99**. Hewlett-Packard, Bristol, UK.

R. Beckers, O.E. Holland, and J. L. Deneubourg 1994. From Local Actions to Global Tasks: Stigmergy in Collective Robotics, in R. Brooks and P. Maes eds. **Artificial Life IV**, Cambridge, Mass.: MIT Press, 181-9.

M. Beer, M. d'Inverno, N. Jennings, M. Luck, C. Preist and M. Schroeder, 1999. Negotiation in Multi-Agent Systems, **Knowledge Engineering Review**, 14(3), 285-289, 1999.

S. Brainov, 1996. Altruistic Cooperation Between Self-Interested Agents, In Proceedings of **ECAI-96**, p. 519-523.

P. Cohen, H. Levesque, I. Smith, 1997. On Team Formation, In J. Hintika and R. Tuomela, **Contemporary Action Theory**, Synthese.

F. Dignum, B. Dunin-Keplicz, and R. Verbrugge, 2000. Agent Theory for Team Formation by Dialogue, In Proceedings of **ATAL-2000**, Boston.

P. Faratin, C. Sierra, and N. Jennings, 1998. Negotiation decision functions for autonomous agents. **Robotics and Autonomous Systems** 24(3--4):159--182. Morgan Kaufmann Publishers

D. Fitoussi and M. Tennenholtz, 2000. Choosing social laws for multiagent systems: minimality and simplicity. **Artificial Intelligence**, 119(2000) 61-101.

H. Hexmoor and H. Duchscherer, 2001. Efficiency as a Motivation to Team, In **FLAIRS-2001**.

J. Jennings, C. Kirkwood-Watts 1998. Distributed Mobile Robotics by the Method of Dynamic Teams. In **Proceedings of the Conference on Distributed Autonomous Robot Systems (DARS 98),** May, 1998, Karlsruhe, Germany.

A. Joshi, T. Drashanksy, J.R. Rice, S. Weerawarana, and E.N. Houstis, 1997. Multiagent Simulation of Complex Heterogeneous Models in Scientific Computing, **IMACS Math. and Comp. In Simulation**, 44, pp 43-59, 1997.

L.P. Kaelbling, L. Littman, A. Moore, 1996. Reinforcement Learning: A survey, **Journal of Artificial Intelligence Research** 4:237-285.

M. Mataric, 1998. Using Communication to Reduce Locality in Distributed Multi-Agent Learning, **Journal of Experimental and Theoretical Artificial Intelligence**, special issue on Learning in DAI Systems, Gerhard Weiss, ed., 10(3), 357-369.

R. McCartney and H. Sun, 2000. Sampling and Estimation by Multiple Robots, In Proceedings of 4th **International Conference on Multiagent Systems,** Boston, MA.

E. Sacerdotti, 1977. **A Structure for Plans and Behavior**, American Elsevier, New York.

R. Simmons, D. Apfelbaum, D. Fox, R. P. Goldman, K. Zita Haigh, D. J. Musliner, M. Pelican, and S. Thrun. 2000. Coordinated Deployment of Multiple, Heterogeneous Robots, To appear in **Proceedings of the Conference on Intelligent Robots and Systems (IROS)**, Takamatsu Japan.

M. Singh, 1998. The Intentions of Teams: Team Structure, Endodeixis, and Exodeixis, **Proceedings of the 13th European Conference on Artificial Intelligence (ECAI)**, Wiley, pages 303-307.

C. Touzet, 2000. Robot Awareness in Cooperative Mobile Robot Learning, **Autonomous Robots**, 8 (1) p. 87-97.

R.T. Vaughan, K. Stoy, G.S. Sukhatme, and M.J. Mataric, 2000. Blazing a trail: insect-inspired resource transportation by a robot team. In **Proceedings of the 5th International Symposium on Distributed Autonomous Robotic Systems**, Knoxville, Tennesssee,

M. Wooldridge, 2000. **Reasoning about Rational Agents**, The MIT Press.