

coordination in multi-agent systems

- definitions
- taxonomy
- applications (focus on implemented multi-robot systems)

multi-agent systems vs distributed systems

- distributed systems implies centralized control; multi-agent system implies autonomous control
- agents in distributed systems are assumed to be *benevolent* and *cooperative*
- agents in a multi-agent system are assumed to be *selfish*; they could be both (or either) *cooperative* and/or *competitive*
- in an MAS, cooperation is not governed; it is a result of *coordination*
- note that coordination is not necessarily a feature of every MAS

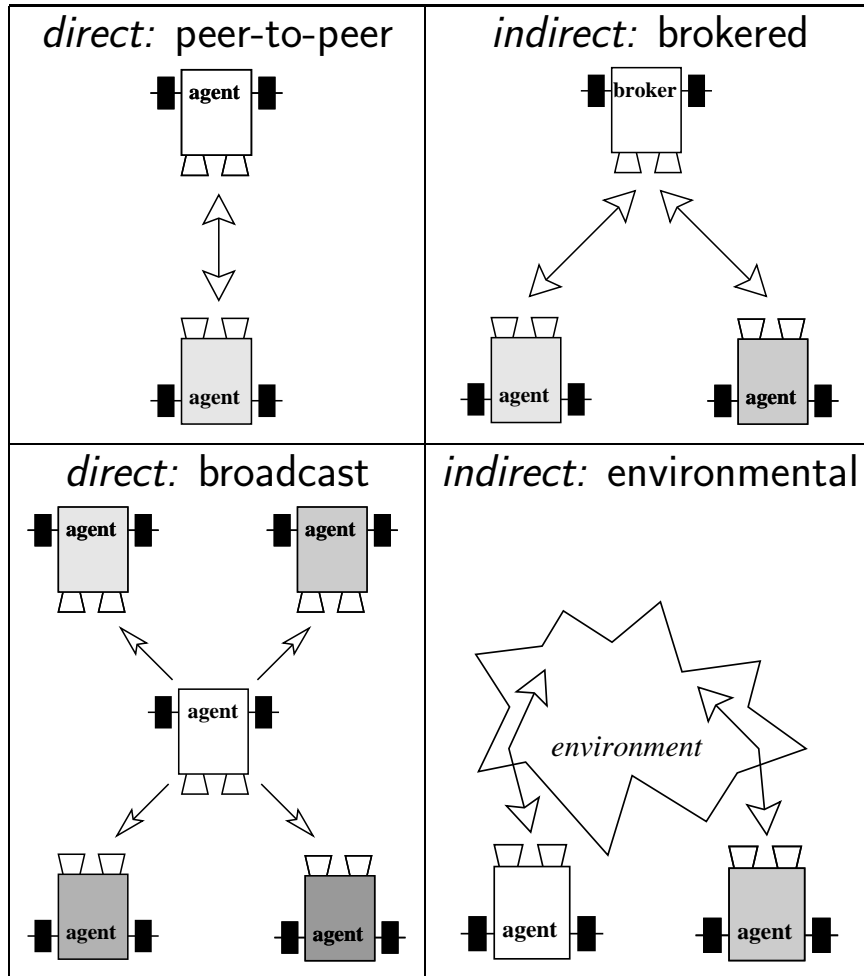
multi-agent system, society

- “sphere of influence” [Jennings, 2000]
when spheres overlap, one agent may interfere with the achievements of another
- dependent vs independent [Sichman and Demazeau, 1995]
 - *independent*:
when agents do not rely on each other at all; actions are chosen and performed separately.
If agent P wants to achieve goal state G_1 , it can do so without help from agent Q; and if agent Q wants to achieve goal state G_2 , it can do so without help from P. P and Q can each achieve their goals without interacting; and if they do interact, this will have no effect on their goals. This also means that Q’s actions cannot prevent P from achieving its goal (and vice versa).
 - *dependent*:
when agents rely on each other to achieve their goals.
agent P cannot, for example, achieve goal state G_1 without help or cooperation from agent Q.

- types of dependencies [Sichman and Demazeau, 1995]
 - *direct dependency*:
must be resolved for P to achieve G_1
 - *indirect dependency*:
less critical; P could achieve G_1 without Q (maybe less efficiently or quickly)
 - *unilateral*:
one agent is dependent on another, but not vice versa;
P needs Q to achieve G_1 , but Q does not need P to achieve G_2
 - *reciprocal*:
two agents depend on each other to achieve different goals;
P needs Q to achieve G_1 , and Q needs P to achieve G_2
 - *mutual*:
two agents depend on each other to achieve same goals
P needs Q to achieve G_1 , and Q needs P to achieve G_1 (same goal)

interactions in multi-agent systems

- agent interactions can be *direct* or *indirect*
- “real” systems often use a combination
- *direct*:
agents formulate messages and deliberately send them to other(s)
e.g., *peer-to-peer* or *broadcast*
- *indirect*:
messages are not sent explicitly to particular agents
e.g., *environmental (tacit agreements or stigmergy)* or *brokered*



process of interaction

- interaction = *communication* + *decision-making*
- three phases:
 1. agents decide individually what to do
 2. send messages reflecting individual decisions (proposals)
 3. discuss proposals amongst (other) agents
 4. may require additional communication
 5. (possibly) revise initial decisions about what to do
- *plan consideration* — deciding what to achieve
- *plan formation* — deciding how to achieve it
- in a *centralized* system: “leader” agent decides plans for all agents and transmits plans; there is no individual plan consideration or formation
- in a *decentralized* system: plan consideration and formation occurs at the level of the individual agent
- in a *hybrid* system: some combination of local/central plan formation/consideration occurs

message content

- binary / Boolean
- scalar
- object
- vector of scalars
- vector of objects

tacit agreement

- minimal interaction mechanism
- no explicit communication
- *social norms* dictate behavior
- can control access based on which agent(s) are privvy to the agreement
- e.g., pedestrian traffic

environmental cues or “stigmergy”

- agents signal to each other by modifying their (shared) environment
- message content usually simple
- binary = presence/absence of a signal
- scalar = strength of signal
- no control over access — all agents have potential to receive message
- but access is limited to agent(s) in physical vicinity of message
- message has a “lifetime”; message/signal may decay and disappear over time
- e.g., ants and pheromone trails

signal broadcasting

- agents transmit messages to anyone within range; there is no specific, designated recipient
- access limited to agents within transmission range
- can reach a wider set of agents than stigmergy, depending on physical distribution of agents and communication means
- can have a lifetime, if retransmission does not occur
- e.g., announcements on the subway

simple auction

- simplest market mechanism
- agents “bid” on single items, typically tasks or resources
- bid consists of a number (“how many”) and an indication of which item (or task or resource) bid applies to
- often used for “role allocation” or “task allocation” in a multi-agent system
- hybrid approach: local plan consideration, global task assignment
- e.g., 3 red dogs

combinatorial auction

- more complex version of simple auction since agents can bid on “packages” of items
- bid consists of a vector of simple bids (see simple auction, above)
- vector represents the package
- hybrid approach: local plan consideration, global task assignment
- e.g., 3 red dogs and 4 blue cats

negotiation

- iterative interaction in which agents can go back and forth with bids
- frequently brokered
- e.g.:
 1. P's bid: I would like 3 red dogs
 2. Q's bid: I would like 5 red dogs
 3. BROKER'S internal deliberation: P has priority, so I'll grant P's bid but since I only have 5 dogs, I'll have to short-change Q
 4. BROKER'S RESPONSE TO P: will you accept 3 red dogs?
 5. BROKER'S RESPONSE TO Q: will you accept 2 red dogs?
 6. P: okay
 7. Q: no!!
 8. BROKER'S DELIBERATION: I could make Q happy by giving all the dogs to Q, but then P will be unhappy. Let me try a modified response in which both agents are offered less than what they requested.
 9. BROKER'S RESPONSE TO P: will you accept 2 red dogs?

10.BROKER'S RESPONSE TO Q: will you accept 3 red dogs?

11.P: well, it's not optimal but I guess I don't mind

12.Q: no, I wanted 5 dogs!! (pout pout)

13.BROKER'S RESPONSE TO P: will you accept 1 red dog?

14.BROKER'S RESPONSE TO Q: will you accept 4 red dogs?

15.P: okay

16.Q: okay

argumentation

- also iterative, like negotiation
- but agents (can) offer reasons for their bids
 1. P's bid: I would like 3 red dogs
 2. Q's bid: I would like 5 red dogs
 3. BROKER'S internal deliberation: P has priority, so I'll grant P's bid but since I only have 5 dogs, I'll have to short-change Q
 4. BROKER'S RESPONSE TO P: will you accept 3 red dogs?
 5. BROKER'S RESPONSE TO Q: will you accept 2 red dogs because I only have 5?
 6. P: okay
 7. Q: no, I wanted 5 dogs!! (pout pout)
 8. BROKER'S DELIBERATION: I could make Q happy by giving all the dogs to Q, but then P will be unhappy. Let me try a modified response in which both agents are offered less than what they requested.
 9. BROKER'S RESPONSE TO P: will you accept 2 red dogs?
 10. BROKER'S RESPONSE TO Q: will you accept 3 red dogs?

11.P: well, it's not optimal but I guess I don't mind

12.Q: no, I wanted 5 dogs!! (pout pout)

13.BROKER'S RESPONSE TO P: will you accept 1 red dog?

14.BROKER'S RESPONSE TO Q: will you accept 4 red dogs?

15.P: okay

16.Q: okay

taxonomy of interaction/coordination mechanisms

classification	interaction medium	message content	decision mechanism	messaging sequence	synchronization
tacit agreement	none	none	distributed	n/a	none
environmental cues	environmental	scalar	distributed	continual	none
signal broadcasting	broadcast	binary / scalar	distributed	continual	none
simple auction	brokered	scalar	hybrid	once (or in rounds)	turn-taking (with broker)
combinatorial auction	brokered	vector of scalars	hybrid	in rounds	turn-taking (with broker)
negotiation	peer-to-peer / brokered	vector of scalars or object	distributed or hybrid	in rounds	turn-taking (with peers or broker)
argumentation	peer-to-peer	object	distributed	in rounds	turn-taking (with peers)

applications: tacit agreements — formation control

- formation control methodologies
 - *leader-follower* —
one agent is designated as the leader and the rest follow it
 - *virtual structure* —
centralized controller synchronizes actions of all agents to maintain a particular shape
 - *behavioral* —
individual agents are given rules such as avoid others or approach others; result is desired formation
- [Fierro et al., 2002] experimented with three different control algorithms in simulation:
 - *separation-bearing* —
agent follows another while keeping a specified distance away at a certain relative bearing
 - *separation-distance-to-obstacle* —
agent follows another while keeping away from obstacles
 - *separation-separation* —
agents follows two others while staying specified distances away from each

- [Balch and Arkin, 1995] and [Balch and Arkin, 1998] experimented with formation control algorithms in simulation and robots:
 - formations: lines, columns, diamonds, wedges (V-shaped)
 - two-step algorithm: agent perceives where it should be in the formation (relative to its current location in relation to other agents), then it moves into position to maintain formation
 - hybrid approach: robots (sometimes) broadcast their position information
 - results provided recommendations on different formations for achieving different tasks

applications: tacit agreements — coordinated map building

- [Yamauchi, 1998]:
 - agents independently determine “frontier” cells and move to the closest ones
 - no explicit coordination mechanism
- Centibots [Konolige et al., 2004]
 - agents explored independently
 - then, agents merged maps

applications: tacit agreements — other methods

- *locker-room agreements*

- in RoboCup soccer simulator [Stone and Veloso, 1998]
- agents agree *a priori* on a set of roles and strategies for team members and changes in environment that would signal which roles/strategies to employ

- *social norms*

- [Shoham and Tennenholtz, 1992a, Shoham and Tennenholtz, 1992b]
- system contains inherent motivation for agents to conform to “norms”
- agents can evaluate conformity of other agents, e.g., tally how many agents have made particular choices
- norms are determined dynamically instead of *a priori* (like locker-room agreements)

applications: electronic institutions

- [Esteva et al., 2001] present a structure for organizing a multi-agent system
- includes software packages for designing, testing and model-checking (EIDE, ISLANDER, AMELIE,...) [Esteva et al., 2002]
- *electronic institution* consists of:
 - *roles* —
agents, defined according to characteristic (typically task-oriented) behavior categories
 - *dialogic framework* —
communication language, ontology and locution rules
 - *scene* —
series of locutions
 - *performative structure* —
series of scenes
 - *norms* —
“commitments, obligations and rights” of agents

applications: environmental cues — ant systems

- classic example: [Dorigo et al., 1996]
- “simulated agents” (ants) solve traveling salesman problem (TSP)
- cities are represented as nodes in a graph; roads between cities are links between nodes
- ants traverse links and leave “pheromone” (chemical) trails
- more-travelled links have more pheromone
- ants are attracted to pheromone
- solution emerges
- pheromone decays over time; old trails essentially disappear
- algorithmic parameters/variations allow modification of decay rate, amount of attraction, etc
- [McLurkin, 1995] constructed small robots but never implemented pheromone trails
- [Svennebring and Koenig, 2004] built single ant-robot that left a trail

applications: environmental cues — other methods

- ALN

- [Kube and Zhang, 1992] experimented in simulation and robots
- *adaptive logic network* (ALN): neural network architecture that recognizes “perceptual cues” indicating changes in the environment
- recognized states triggers behavioral responses that result in coordinated activity

- Tron

- [Funes et al., 1998] built video game in which agents left “light trails” (ala Tron movie)
- agents were controlled by genetic programs
- agents played humans, behaviors *co-evolved*
- [Sklar et al., 2001] conducted follow-up experiments in which agents were trained using database of human interactions; agents were controlled by neural networks

applications: signal broadcasting

- [Yanco and Stein, 1993]
 - experimented with simulation (3 agents) and robots (2 agents)
 - implemented *leader-follower* behavior
 - leader sends signals to follower(s)
 - leader interfaces with human “instructor” who provides “rewards” for desirable behavior
- ALLIANCE [Parker, 1998, Parker, 2000]
 - experiments with robots
 - dynamic task allocation, emphasized robustness
 - team can recover from partial (or total) robot or communication failures
 - two behaviors:
 - * *impatience* — take over others’ tasks when they are not being achieved
 - * *acquiescence* — release their own tasks when they are not being achieved

- *embodied evolution* [Watson et al., 1999]
 - experiments with robots
 - devised special arena with continuous power supply
 - evolution occurs on-board the robots
 - robots seek a light source
 - robots are controlled by a perceptron
 - robots broadcast their weights: the more light they sense, the more frequently they broadcast
- *broadcast of local eligibility* (BLE) [Werger and Mataric, 2000]
 - cooperative multi-robot observation of multiple moving targets (CMOMMT)
 - robots exchange “fitness” information, via broadcast, to determine which agent is most fit to accomplish given task
 - less fit robots inhibit their tracking behavior

- honey bee emulation [Vaughan et al., 2000]
 - robots broadcast “global crumb list” and maintain “private crumb list”
 - crumb list is a time-sequenced list of headings and positions
 - emulates bees’ “waggle dance” by broadcasting private crumb list when a resource is found
- zone surveillance [Saffiotti et al., 2000]
 - team of robots is viewed like one very flexible and capable agent
 - team tracks and covers a target by distributing “zones” of responsibility
 - robots transmit location and estimated velocity of target (when they can see it)
 - “desirability” function assigns quantitative preferences to robots’ actions
 - experiments used three schemes:
 - * no communication
 - * local communication (*works best)
 - * global communication (sensitive to failure and also “little league” effect)

- PC-MVERT [Kalra et al., 2004]
 - “passive coordination”
 - robots develop plans using look-ahead of several steps
 - robots broadcast their current plan
 - robots coordinate internally by comparing/evaluating their plan with others’ received
 - works better than not broadcasting plans
 - other experiments included: “tight”, “planned”, “computationally feasible” coordination mechanisms

applications: auction mechanisms

- distinction between *simple* and *combinatorial* auctions is (often) blurred in robot applications
- Contract Net [Smith, 1977, Smith and Davis, 1980]
 - most widely used auction-based protocol in robot systems
 - TRACONET [Sandholm, 1993, Sandholm, 1998a, Sandholm, 1998b]: vehicle routing problem; proved that exchanging sets of tasks can help avoid local minima in solution space
 - Traderbots [Dias and Stentz, 2000, Dias and Stentz, 2002]: distributed traveling salesman problem (DTSP) explored; multiple agents share task of visiting multiple cities, bidding against each other for assignment to particular cities
 - extended work: [Dias et al., 2004]
- MURDOCH [Gerkey and Mataric, 2002]: experiments with auction mechanisms for robot team coordination, encompassing dynamic task allocation; focused on task allocation optimization and robustness

applications: negotiation

- *monotonic concession protocol* [Rosenschein and Zlotkin, 1994]: agreements over task allocations can be guaranteed if each agent makes a concession at each stage in the negotiation by offering a deal that is better for the other agent
- [Faratin, 2000, Faratin et al., 2000] extended this by ensuring that agents' concessions, while better for others, are not worse for themselves; takes advantage of different agents placing different utilities on different aspects of negotiated good
- Teamcore [Tambe, 2004]
 - general assignment problem (GAP)
 - “Machinetta” strategy, based on principle of “teamwork”, in which complex tasks (“roles”) are allocated to best-suited team members in order to optimize team’s objectives (NP-hard problem)
 - approximates analytically using DCOP (distributed constraint optimization problem) technique
 - roles (complex tasks) are represented as “tokens”; these are passed around to agents; agents accept tokens when they decide that they have the capabilities to complete tasks associated with the token

applications: argumentation

- theoretical work surveyed in [Rahwan et al., 2003]
- to date, not implemented in multi-robot system
- implemented in mixed-initiative system [Ferguson, 1995]
- decision-making and medical applications [Fox and Das, 2000, Fox et al., 1997]

*

References

- [Balch and Arkin, 1995] Balch, T. and Arkin, R. (1995). Motor-schema based formation control for multiagent robot teams. In *Proceedings on the First International Conference on Multiagent systems*.
- [Balch and Arkin, 1998] Balch, T. and Arkin, R. (1998). Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*.
- [Dias and Stentz, 2000] Dias, M. B. and Stentz, A. (2000). A free market architecture for distributed control of a multirobot system. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems*.
- [Dias and Stentz, 2002] Dias, M. B. and Stentz, A. (2002). Opportunistic optimization for market-based multirobot control. In *International Conference on Intelligent Robotics and Systems (IROS-02)*.
- [Dias et al., 2004] Dias, M. B., Zlot, R., Zinck, M., Gonzalez, J. P., and Stentz, A. (2004). A versatile implementation of the traderbots approach for multirobot coordination. In *Proceedings of the 8th International Conference on Intelligent Autonomous Systems*.
- [Dorigo et al., 1996] Dorigo, M., Maniezzo, V., and Colorni, A. (1996). The Ant System: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics Part B: Cybernetics*, 26(1):29–41.
- [Esteva et al., 2002] Esteva, M., de la Cruz, D., and Sierra, C. (2002). ISLANDER: an electronic institutions editor. In *First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)*, pages 1045–1052, Bologna, Italy.
- [Esteva et al., 2001] Esteva, M., Rodriguez, J. A., Sierra, C., Garcia, P., and Arcos, J. L. (2001). On the formal specifications of electronic institutions. *Agent-mediated Electronic commerce (The European AgentLink Perspective)*, LNAI 1991, pages 126–147.
- [Faratin, 2000] Faratin, P. (2000). *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, University of London, Queen Mary College.
- [Faratin et al., 2000] Faratin, P., Sierra, C., and Jennings, N. R. (2000). Using similarity criteria to make negotiation trade-offs. In *Proceedings of the 4th International Conference on Multiagent Systems*, pages 119–126.
- [Ferguson, 1995] Ferguson, G. (1995). *Knowledge Representation and Reasoning for Mixed-Initiative Planning*. PhD thesis, Department of Computer Science, University of Rochester.
- [Fierro et al., 2002] Fierro, R., Song, P., Das, A., and Kumar, V. (2002). *Cooperative Control of Robot Formations*, volume 66, pages 73–93. Kluwer Academic Press.
- [Fox and Das, 2000] Fox, J. and Das, S. (2000). *Safe and sound: Artificial intelligence in hazardous applications*. AAAI Press/MIT Press.
- [Fox et al., 1997] Fox, J., Johns, N., Lyons, C., Rahmanzadeh, A., Thomson, R., and Wilson, P. (1997). Proforma: a general technology for clinical decision support systems. *Computer Methods and Programs in Biomedicine*, 54:59–67.
- [Funes et al., 1998] Funes, P., Sklar, E., Juillé, H., and Pollack, J. (1998). Animal-Animat Coevolution: Using the Animal Population as Fitness Function. In *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior (SAB)*, pages 525–533. MIT Press.
- [Gerkey and Mataric, 2002] Gerkey, B. and Mataric, M. (2002). Sold!: Auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation*, 18(5).
- [Jennings, 2000] Jennings, N. R. (2000). On agent-based software engineering. *Artificial Intelligence*, 117:277–296.
- [Kalra et al., 2004] Kalra, N., Stentz, A., and Ferguson, D. (2004). Hoplitess: A market framework for complex tight coordination in multi-agent teams. Technical Report CMU-RI-TR-04-41, Robotics Institute, Carnegie-Mellon University.

- [Konolige et al., 2004] Konolige, K., Fox, D., Ortiz, C., Agno, A., Eriksen, M., Limketkai, B., Ko, J., Morisset, B., Schultz, D., Stewart, B., and Vincent, R. (2004). Centibots: Very large scale distributed robotic teams. In *Proceedings of the 9th International Symposium on Experimental Robotics*.
- [Kube and Zhang, 1992] Kube, C. R. and Zhang, H. (1992). Collective robotic intelligence. In *1992 International Conference on Simulation of Adaptive Behaviour*, pages 460–468.
- [McLurkin, 1995] McLurkin, J. (1995). The ants: A community of microrobots. Bachelors Thesis, MIT.
- [Parker, 1998] Parker, L. E. (1998). Alliance: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*.
- [Parker, 2000] Parker, L. E. (2000). Lifelong adaptation in heterogeneous multi-robot teams: Response to continual variation in individual robot performance. *Autonomous Robots*, 8(3):239–267.
- [Rahwan et al., 2003] Rahwan, I., Ramchurn, S. D., Jennings, N. R., McBurney, P., Parsons, S., and Sonenberg, E. (2003). Argumentation-based negotiation. *Knowledge Engineering Review*, 18(4):343–375.
- [Rosenschein and Zlotkin, 1994] Rosenschein, J. S. and Zlotkin, G. (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press.
- [Saffiotti et al., 2000] Saffiotti, A., Zumel, N. B., and Ruspini, E. H. (2000). Multi-robot team coordination using desirabilities. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS)*, pages 107–114.
- [Sandholm, 1993] Sandholm, T. (1993). An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*.
- [Sandholm, 1998a] Sandholm, T. (1998a). Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium Series: Satisficing Models*.
- [Sandholm, 1998b] Sandholm, T. (1998b). Contract types for satisficing task allocation: II experimental results. In *AAAI Spring Symposium Series: Satisficing Models*.
- [Shoham and Tennenholtz, 1992a] Shoham, Y. and Tennenholtz, M. (1992a). Emergent conventions in multi-agent systems. In *Proceedings of Knowledge Representation and Reasoning (KR&R-92)*.
- [Shoham and Tennenholtz, 1992b] Shoham, Y. and Tennenholtz, M. (1992b). On the synthesis of useful social laws for artificial agent societies. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92)*.
- [Sichman and Demazeau, 1995] Sichman, J. and Demazeau, Y. (1995). Exploiting social reasoning to deal with agency level inconsistency. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 352–359.
- [Sklar et al., 2001] Sklar, E., Blair, A. D., and Pollack, J. B. (2001). *Training Intelligent Agents Using Human Data Collected on the Internet*, pages 201–226. World Scientific, Singapore.
- [Smith, 1977] Smith, R. G. (1977). The contract net: a formalism for the control of distributed problem solving. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI-77)*.
- [Smith and Davis, 1980] Smith, R. G. and Davis, R. (1980). Frameworks for cooperation in distributed problem solving. *IEEE Transactions on Systems, Man and Cybernetics*, 11(1).
- [Stone and Veloso, 1998] Stone, P. and Veloso, M. (1998). Towards collaborative and adversarial learning: A case study in robotic soccer. *International Journal of Human Computer Studies*, 48:83–104.
- [Svennebring and Koenig, 2004] Svennebring, J. and Koenig, S. (2004). Building terrain-covering ant robots: A feasibility study. *Autonomous Robotics*, 16(3):313–332.

- [Tambe, 2004] Tambe, M. (2004). *Evolution of a Teamwork Model*, chapter 1. Cambridge University Press.
- [Vaughan et al., 2000] Vaughan, R., Stoey, K., Sukhatme, G., and Mataric, M. (2000). Blazing a trail: insect-inspired resource transportation by a robot team. In *Proceedings, 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, pages 111–120.
- [Watson et al., 1999] Watson, R. A., Ficici, S. G., and Pollack, J. B. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *1999 Congress on Evolutionary Computation*, pages 335–342. IEEE Press.
- [Werger and Mataric, 2000] Werger, B. B. and Mataric, M. (2000). Broadcast of local eligibility for multi-target observation. *Distributed Autonomous Robotic Systems (DARS)*, pages 347–356.
- [Yamauchi, 1998] Yamauchi, B. (1998). Frontier-based exploration using multiple robots. In *Proceedings of the Second International Conference on Autonomous Agents*.
- [Yanco and Stein, 1993] Yanco, H. and Stein, L. (1993). An adaptive communication protocol for cooperating mobile robots. In *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 478–485. MIT Press.