

Literature Review

Bradford Heap

August 2010

1 Introduction

Efficient and effective teamwork amongst autonomous robots in a dynamic environment is a complex and open problem in Artificial Intelligence. During the last twenty to thirty years a number of techniques have been developed for addressing teamwork in closed environments. Other research has focussed on detailed areas of teamwork such as communication, coordination, task allocation, and planning.

In this review the specifications and differences between classifications of multi-agent and multi-robot systems are defined. Models of behaviour within teams are discussed focussing on collective, cooperative, collaborative and coordinative interaction. This is followed by a review of task and resource allocation methods focussing on the use of markets and auction based controllers. The review concludes by briefly outlining a proposed area of research.

2 Definitions

An *agent* as defined by Russell and Norvig [1] is “anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.” Expanding on this definition Russell and Norvig propose five models of rational agents:

- *Simple reflex agent.* These agents sense the world around them and act according to what they can see. They have no prior knowledge of the environment or their own actions.
- *Model-based reflex agent.* These agents have prior knowledge of the environment and their actions and combine this knowledge with what they can currently sense when reasoning about their next action.
- *Goal based agent.* These agents incorporate a model-based reflex agent with the added desire to achieve a certain goal. This form of agent will seek to execute actions that help it to achieve its goal.
- *Utility based agent.* When executing tasks to achieve a goal there will often be a cost associated with the execution of the task. Utility agents will take these costs into account when reasoning about actions to ensure that they achieve their goal in the minimum amount of cost to itself.
- *Learning agents.* The ability to learn and adapt to an environment or particular problem is desired in many problems within the field of Artificial Intelligence. The ability to learn can be incorporated into any of the above agents that have some form of prior knowledge. This learning can then be used in the reasoning process to select the best action based on the outcomes of prior actions.

In its simplest form a *Multi-Agent System* (MAS) is purely a collection of many agents. Within this system these agents can interact with each other either to cooperate or compete to achieve individual or collective tasks [2]. Interaction is typically facilitated by communication through the exchanging of messages across a network [3] but other methods such as passive observation and reaction to the behaviour of other agents allows for interaction without explicit communication.

Multi-Robot Systems (MRS) are built upon the ideas incorporated in MAS. Farinelli *et al.* [4] define a MRS as “a set of robots operating in the same environment.” They argue that a physical MRS is a lot more complex than that of a MAS. In particular the use of physical robots in a MRS introduces issues not found in MAS such as uncertainty and incompleteness of information acquired from the environment, errors in communication between robots and hardware failures. One way of thinking about a MRS is that each robot consists of its own internal MAS but acts as a singular entity interacting with the environment.

3 Classifications

In Doran *et al.* [2] Franklin presents a typology of cooperation in MAS (Figure 1). At the root level Franklin classifies MAS into two branches: *independent* “if each agent pursues its own agenda independently of others” and *cooperative* if the agenda of agents include “cooperating with other agents in some way”.

On the independent branch two leaves are specified: *discrete* “if the agendas of the agents bear no relation to one another” and *emergent cooperation* where “from an observer’s viewpoint, the agents appear to be working together, but from the agent’s viewpoint they are not”. For instance, emergent cooperation is seen

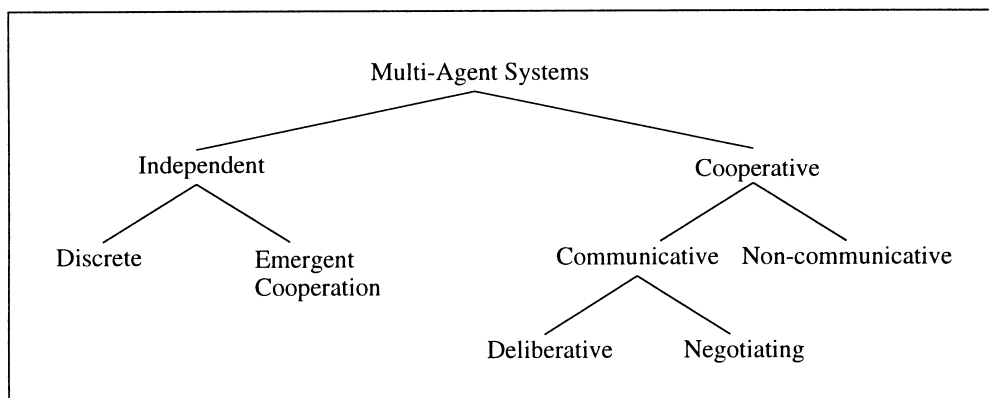


Figure 1: Franklin’s Cooperation Typology [2]

in Dagaëff *et al.* [5] where agents are programmed to be deliberately antagonistic towards other agents however cooperation emerges. In Iba [6] emergent cooperation appears during the use of genetic algorithms to develop agents in a MAS. Shehory *et al.* [7] argues that in a large scale MAS the complexity of too much communication will cripple the system and therefore presents a framework for cooperative goal-satisfaction in a large-scale environment that shows emergent cooperation.

Within the cooperative systems branch, MAS are divided into two further sub-classifications. In *communicative* systems agents cooperate via intentional messages being sent and received. This communication can take one of two forms, *deliberative* or *negotiating*. Deliberative systems see agents jointly plan any actions they take, which may or may not entail cooperative joint actions.

Negotiating systems also involve communication in planning but agents can bid or compete against each other to achieve the best outcomes for themselves. In *non-communicative cooperative* systems agents cooperate through observing and reacting to the behaviour of others. This behaviour differs from that of independent emergent cooperation in that in non-communicative cooperation agents have the ability to anticipate and assess the actions of other agents. Typically, this ability to anticipate and assess operates on the principle of *stigmergy* where changes in the environment provides information indirectly to other agents. Katoh *et al.* [8] explore this type of MAS in detail using potential fields.

Farinelli *et al.* [4] present a Taxonomy of MRS (Figure 2). Their taxonomy follows a similar pattern of thought to that of Franklin's cooperative branch in classifying MAS. However, unlike Franklin's focus on communication, Farinelli *et al.* primary focus is on different forms of coordination amongst MRS. The authors present two different types of coordination, *strong coordination* represents MRS that rely on a coordination protocol, whereas *weak coordination* represents those systems that do not operate via a protocol. They argue that weak coordination approaches

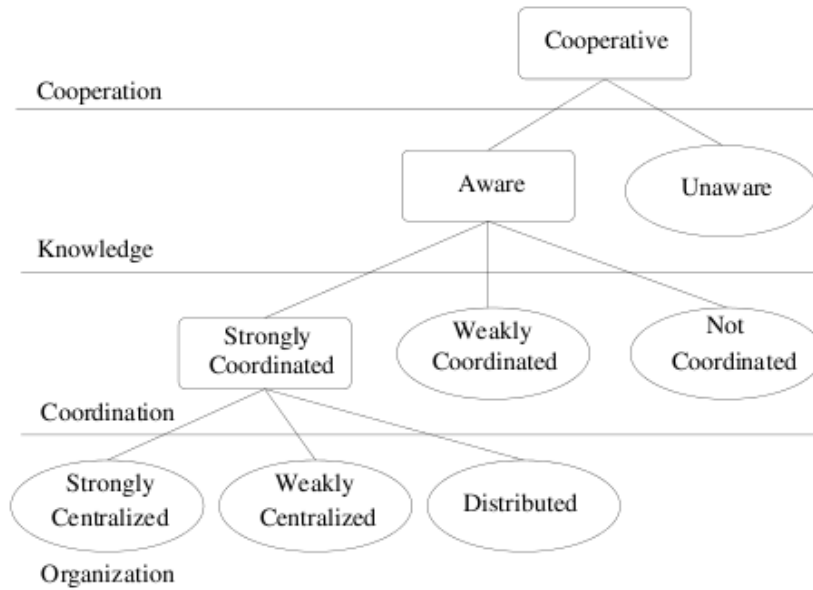


Figure 2: Farinelli *et al.* MRS Taxonomy [4]

are of more interest to MRS researchers because the use of strict coordination protocols is difficult and involves many overheads.

At the organisation level a distinction is made between the autonomy each agent has in the system’s decision process. A *centralised* system has a leader who is solely responsible for organising the entire team of robots. This centralised decision process can be classified as *strong* where one leader has complete command during the entire mission, or *weak* where a number of different leaders may have command during different periods of the mission. In a *distributed* system agents are completely autonomous in their decision making and no sole team leader exists.

When considering MAS and MRS the classification of the complete system is important in determining the behaviour of individual agents. The classification of the system typically determines the form of interaction individual agents have in their communication and cooperation with other agents. This is explored in detail in the following section.

4 Interaction

Positive interaction between agents in both MAS and MRS plays a very important role in determining the overall success of a system. Parker [9] defines four common forms of interaction: *Collective*, *Cooperative*, *Collaborative*, and *Coordinative*.

The choice of the best type of interaction depends on the particular problem requiring a solution. Within a problem there are often a set of goals that need to be achieved. In MAS/MRS there are two types of goals, *individual* and *shared*. Individual goals can be homogeneous in that all agents have the same aim, or heterogeneous in which each agent has a different goal to satisfy.

When discussing interaction it is best to consider why interaction is required in the first place. It is possible to operate a MAS/MRS without any interaction between agents (as previously discussed in Section 3), however, as the system becomes more complex, the lack of interaction between agents will lead to bottlenecks, collisions and other problems as agents fail to account for each other. Jennings [10] highlights this problem when considering a group of people in a park, when it rains they all individually decide to seek shelter under the same tree. Jennings goes on to explain in detail the need for, and some models of, individual behaviour that also include communication and coordination amongst other individuals.

4.1 Collective Interaction

Returning to Parker's forms of interaction, *collective interaction* is defined as where "entities are not aware of other entities on the team, yet they do share goals, and their actions are beneficial to their team-mates". Systems using collective interactions often follow models that are biologically inspired such as flocking and herding. The most widely studied form of collective interaction is *swarming*.

Swarming is a form of independent MAS that shows emergent cooperation. In a swarm all agents share the same goal and act independently of each other, despite this they often appear to be working cooperatively together. Trianni *et al.* [11] bridge individual and collective interaction in their research which focuses on autonomous robots deciding if they should act as an individual or use the information that they have gathered from the environment to act with other robots as a robotic swarm.

In their research three robots are placed in a circle with two goals, only one of which can be achieved: individually leave the boundaries of the circle through a defined exit, or group together as a swarm in the middle of the circle. The achievement of either goal is considered a success and the best agents are those that completed either task in the shortest period of time. To improve the agents, the authors used genetic algorithms to evolve new controllers for the agents. While the use of genetic algorithms is not a focus of my research the ideas involved in the decision to switch from individual to collective and/or other forms of interaction are an open research area that can be further explored.

4.2 Cooperative Interaction

Parker's second type of interaction is called *cooperative interaction* and is specified by "entities are aware of other entities, they share goals, and their actions are beneficial to their team-mates". This type of interaction can take two forms: *joint task coordination* or *shared information coordination*. In joint task coordination agents work together to achieve a task that an individual agent could not achieve by itself, such as coordinated robotic construction [12].

Shared information coordination involves agents acting independently but sharing information. One of the major challenges in shared information coordination is determining the correct amount of information to share. If too little informa-

tion is shared between agents the system may fail or cease to perform adequately, conversely, if too much information is shared the system may be swamped in the communication process and the task is not achieved.

In Settembre *et al.* [13] the authors outline cooperative situation assessment for use in search and rescue “that balances the use of communication bandwidth with the need for good situation assessment”. The basis of their approach is for each robot to form a plan of action, once they have formed a plan they randomly send this plan to another robot in the system, if this robot agrees with the plan they then randomly forward it on to another robot. This process continues until a fixed number of robots share the same plan, subsequently the plan is executed. If the receiving robot does not agree with the plan proposed, it will send back additional data supporting its reasons why it does not agree to the proposed plan. The originating robot can then either decide to revise its own plan based on this new data or continue with its original plan.

There are two primary benefits of this approach to distributed planning. The first is low levels of communication overhead as there is only one-on-one communication between agents. The second is that by reasoning across multiple robots errors in a single robot’s sensors are filtered out. The authors point out that one limitation in their approach is when the environment in which the robots are operating is sparse, the information and views of each robot can be quite contradictory which degrades the performance of the algorithm.

As an aside, it should be noted that the authors chose to randomly select another robot to send their plan to. Because there is no requirement that this receiving robot is in close proximity to the sending robot it can be reasoned that communication and plans between robots that are far apart will often conflict and therefore slow down the negotiation process. As an improvement to this one could apply rules of an Ising or Sznajd model [14, 15]. In an Ising model opinions flow inwards, typically if a surrounding group of agents share an opinion the agents enclosed by this group

will change their opinion to match those of the outer group. Sznajd models show the propagation of an opinion from a small dense cluster of agents outwards. In the past it has been used to model the emergence of new ideas [16]. The application of this to robotic plan negotiation should see clusters of agreed opinion form within the system which may or may not converge on a solution faster than a randomised approach.

4.3 Collaborative interaction

Collaborative interaction (also called *coalition formation*) is defined by Parker as “robots have individual goals, they are aware of their team-mates, and their actions do help advance the goals of others”. The best example of this sort of interaction in a MRS is within the domain of robot soccer. The overall goal of a team will be to win the game, however within the team each robot will have its own individual goal: the attacker to score, the defender to take the ball off the opposing team’s attacker, and the goalkeeper to block any shots on goal. When each agent achieves their individual goal in this domain it helps to contribute to the achievement of the overall team goal.

Vig and Adams [17] explore coalition formation in a multi-robot domain. They argue that a discrepancy exists between coalition formation in MAS and its application to the MRS domain. Coalition formation in MAS assumes that all agents are available and communication is always possible. In a MRS it is possible that robots will be far enough apart that communication can not occur across the entire team but instead only subsets. To test coalition formation in MRS three scenarios were set up where teams of robots need to work together to achieve a task: coordinated box-pushing, clean-up, and sentry duty.

For the box-pushing task two robots are required to push a large box. From a pool of 12 robots, two boxes needed pushing and two sets of two robots succeeded in

working together to move both boxes. The clean-up task involved a group of robots clearing an area of small boxes in which a single robot would individually move a single box at a time. The robots were able to successfully work together to clear the area of boxes. For the third task, sentry-duty, robots were required to navigate to a position and *collaborate information* to perform motion detection. Like the previous two tasks this one was performed successfully. For a final task a group of robots was given all three tasks to perform simultaneously, this required groups of robots within the MRS to form coalitions to achieve one of the tasks, in this scenario all three tasks were completed.

The authors discovered through the allocation of tasks that as more tasks are allocated and fewer robots remain, the number of messages exchanged to form a coalition decreases. This suggests that the algorithm they used in the experiment should scale well to larger systems. *Collaborate interaction* is the most interesting form of interaction amongst MAS and MRS as the combination of both individual and shared goals can provide an insight into the reasoning required for agents to determine when to act as an individual and when to act as part of a team.

4.4 Coordinative Interaction

The final type of interaction is *coordinative*. In this interaction “entities are aware of each other, but they do not share a common goal, and their actions are not helpful to other team members”. For instance, this form of interaction is seen in systems where agents operate in an enclosed environment and need to move from one area to another in the shortest distance and shortest time. In this scenario it is in an agent’s best interests to communicate with other agents to avoid a collision, but it would not be in its interest to give-way to many agents. This form of interaction typically sees agents acting against each other, which is in contrast to the previous three types of interaction.

5 Task and Resource Allocation

In Chevalyere et al. [18] multi-agent resource allocation is defined in its simplest form as “the process of distributing a number of items amongst a number of agents”. In the domain of MAS and MRS systems it is important to have well defined and efficient processes for distributing tasks and resources amongst agents to achieve their local and collective goals.

Within a large MAS the achievement of the global objective of the system will be dependent upon the successful completion of a number of sub-tasks. The requirements and dependencies of these tasks can vary widely. The simplest tasks are *independent*, meaning that they can be completed by one or more agents without any prior task needing to be achieved first. In other situations tasks are *dependent*, meaning prerequisite tasks must be successfully completed before the current task can begin. For instance, to achieve the task of driving a car, the task of starting the car must be achieved first. In some situations a group of tasks need to be sequentially achieved together. Take for example the installation of a car wheel onto a car, the car wheel must be placed onto the car before the bolts can be applied to the wheel to hold it in place.

5.1 Resources

To achieve some tasks the use of resources is required. There are broadly three forms of resources: *static*, *consumable* and *material*. *Static resources* are used to aid the completion of a task but are not consumed in the task and can be used later in a different task, or allocated to a different agent for use in completing their allocated task. *Consumable resources* are either partially or entirely used in the completion of the task but they do not form part of the final task output, for instance, the use of batteries to power robots completing a task with a limitation on the number of batteries available. Finally, *material resources* are used in the completion of a

task and form part of the final output. As an aside it should be noted that it is entirely possible for one task to create a material resource than can subsequently be consumed in a further task.

Chevaleyre *et al.* [18] provide a discussion of some different properties of resources:

- *Continuous vs Discrete.* *Continuous resources* are typically regarded as being infinity divisible, for instance, the supply of energy. However, within the context of the use and trade of resources continuous resources are typically packaged into discrete sizes. *Discrete resources* are typically indivisible, for instance, a barrel of bricks in which each brick is a discrete and fixed size. Agents can use multiple discrete resources but cannot break apart a singular discrete resource.
- *Shareable.* *Shareable resources* can be allocated to multiple agents at the same time. This type of resource will usually also be a *continuous resource* that can be distributed over a large number of agents. However, resources do not have to be physical, for example a resource could be a set of data, this set of data could be distributed amongst agents through the copying of this data and therefore each agent would have its own version of the resource.
- *Consumability.* As mentioned above certain resources can be consumed in the completion of a task. It is also possible for some resources to perish or expire if they are not consumed within a certain period of time. For instance, if an agent needs access to a certain fixed resource that is only available during a certain period of time and the agent fails to adhere to these time restrictions then the resource would have expired and would no longer be available to the agent. For example, batteries discharge over time even if they are not in use.

- *Quantity.* In some situations there can be many individual discrete resource items that make up one resource. As given in the barrel of bricks example, it is possible to have many bricks. If an agent used a brick resource it would not matter what exact brick they used. However, when considering a collection of bricks it may still be necessary to distinguish in some form of representation each individual brick.

5.2 Markets

The most common approach for the allocation of tasks and resources amongst MAS and MRS systems is through the use of a market. In Dias *et al.* [19] the common requirements of a market are defined:

- A global objective with subcomponents (tasks) that are achievable by individuals or sub-teams.
- A “global objective function” which gives preference to a particular solution to the global objective.
- Access to resources to achieve the objective.
- Each agent has a global utility function which quantifies its individual preferences for certain tasks and resources.
- A mapping between the global objective and the subcomponents so that the achievement of subcomponents aids the advancement of the overall solution.
- A mechanism for the redistribution of resources and tasks amongst agents. This mechanism accepts offers from agents for tasks and resources and subsequently allocates the tasks and resources in a manner that maximises the mechanism-controller’s utility.

Auctions are the most common mechanism for the distribution of resources and tasks. In their simplest form auctions have three phases: offer, bidding, winner determination. In the offer phase the auction controller announces to bidders a description of the lot that is available for bidding. During the bidding phase agents submit bids for the lot typically up to the maximum they wish to spend on the particular lot. Bids are typically representative of an agent's estimate of the resources required to carry out a task. Finally the winner of the auction is determined resulting in the exchange of payment and goods between the auctioneer and the winner.

There are a number of different auction methods that can be applied to resource distribution. Wooldridge [3] explores some of the popular approaches:

- *English auctions.* These are the most common form of auction. Bidding starts at a reserve price and increases as agents outbid each other. The auction stops when no agent is prepared to bid any higher. The lot is then awarded to the highest bidder.
- *Dutch auctions.* These auctions start with the auctioneer announcing a bidding price that is higher than any agent is prepared to pay. The auctioneer will then continue to lower the price for the lot until an agent accepts the offer price. The lot is then awarded to the bidding agent.
- *Sealed-bid auctions.* This type of auction involves only one bid from each agent being accepted. The bid each agent makes is private and is not disclosed to the other agents partaking in the auction. The winner is the agent that places the highest bid for the lot.

In all three of these forms of auctions the highest bidder pays the highest amount for the lot. One problem with this is in the competitive nature of bidding is bids can be placed that are well above the true valuation of the lot. In this situation instead of maximising their utility agents inadvertently harm themselves and subsequently

in a MAS harm the entire system. To prevent this problem in some situations the rules of the auction are altered so that the highest bidder pays only the price of the second highest bidder. This second highest price is expected to be a better representation of the true value of the lot.

When there are many lots available multiple auctions need to be run. The easiest way to distribute the lots is to run many auctions, one for each lot. However, this approach is not optimal as it is time consuming. Another problem is agents may require a certain number of lots but they may win the auctions for only some of the lots they require. When this happens they will need to either coordinate with other agents to share resources, or lots will need to be re-auctioned until all agents are satisfied and are able to complete their tasks with the resources they have.

A second approach to allocating the lots is to group lots together to form larger lots and run fewer auctions, for instance the grouping of a particular task and the resources required to complete the task. This is called a *multiple-item auction*. Multiple-item auctions allow the initial distribution of lots to occur in a much shorter period of time as there are fewer auctions held. However, if the individual items in the lots are not entirely wanted or required by the winning agent they may want to split the contents of their lots into separate lots and re-auction these new lots. This re-auctioning process can generate in total many more auctions than many single-item auctions.

Despite both *single-item* and *multiple-item* auctions not being an optimal method for the distribution of lots because of their simplicity they are still an active area of research. In a recent survey paper on agent coordination using repeated single-item auctions Koenig *et al.* [20] explore the use of auctions in real-time domains. They conclude that more work needs to be done comparing the use of auctions for both centralised and distributed coordination of agents, especially in the area of heterogeneous agents.

Nanjanath and Gini [21] used repeated single-item auctions to dynamically re-allocate tasks to rescue robots. The authors note that the auction algorithm is not optimal, however, because it is designed to work within a dynamic environment it performed well at adapting to changing conditions in the environment whereas a more optimal but slower algorithm may not have. In their algorithm whenever a robot achieved a task all uncompleted tasks were re-auctioned across all robots. This allowed robots that had encountered unexpected delays or failures to give up tasks that they were not able to complete and other robots who had achieved their tasks to continue with more tasks.

A third method of auctioning many lots is called *combinatorial auctions*. In combinatorial auctions all lots are offered to all agents at once. Each agent can bid on any combination of lots to suit their desires. This allows agents to bid for the specific lots they require and removes any need for repeated auctions to redistribute lots. This approach ensures that the auction results in an optimal solution. However, in order to achieve the optimal solution combinatorial auctions are NP-complete. Despite this Berhault *et al.* [22] used combinatorial auctions to allocate navigation tasks to robots. But as noted by Nanjanath and Gini [21] combinatorial auction algorithms do not scale well due to the computational complexity of generating and processing bids.

For auction algorithms to work there must be some form of currency exchanged and an accurate valuation of the lots. Agents bidding must also have enough, but a limited, supply of currency to bid and pay for lots. However, in many situations it is possible for the value of the currency to be merely arbitrary, when this occurs the entire market can be undermined and will cease to function with any efficiency. To avoid this problem altogether another form of market control exists involving bartering and trading without currency.

In a bargaining market agents will trade lots between themselves. Often one agent will value a certain set of lots at a higher value than another agent, and in reverse they may value their current lots less than another agent's view on the same lots. This discrepancy in the estimated valuation of lots allows for the trading of the lots between agents. To successfully trade agents negotiate to form agreements on the exchange of lots. Wooldridge [3] outlines three different forms of negotiation:

- *One-to-one negotiation.* One agent negotiates over the exchange of lots with only one other agent.
- *Many-to-one negotiation.* One agent offers a lot and negotiates with many other agents to determine the best deal for the lot.
- *Many-to-many negotiation.* Agents negotiate with many other agents in parallel. This is a decentralised approach to trading and allows agents to be negotiating on a number of different deals at the same time.

Auctions, markets, and trading is a very large research area in MAS. This section has barely touched on the key parts of the area. In Chevaleyre *et al.* [18] there is a discussion on the benefits and pitfalls of centralised versus distributed market controllers with regards to scalability. This discussion subsequently leads into the idea of agents operating in multiple commodity markets such as a market for task allocation and a second market for resource allocation. Wooldridge [3] devotes a number of chapters to discussion on various approaches to resource allocation. Shoham and Leyton-Brown [23] provide a chapter outlining many additional auction methods while Cramton *et al.* [24] have edited an entire book on combinatorial auctions methods.

6 Conclusion

This review has examined some of the prior research on communication and cooperation amongst agents. The chapter began by defining the concept of an agent and its relation to MAS and MRS. From this different classification types of these systems were reviewed, followed by forms of interaction amongst agents. Finally, methods for the distribution of tasks and resources were discussed.

Combining agent interaction models for group formation with market based approaches for task and resource allocation opens an interesting area of research bridging two fields. In addition to the already mentioned use of Ising and Sznajd models for the propagation of agreement, the use of groups of agents in bidding verses or in conjunction with individual bidding is a potential area of research. Counter to this it would also be interesting to explore how different auction mechanisms affect the formation of collaboration amongst agents.

References

- [1] Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. 2nd edn. Prentice Hall (2003)
- [2] Doran, J., Franklin, S., Jennings, N., Norman, T.: On cooperation in multi-agent systems. *The Knowledge Engineering Review* **12** (1997) 309–314
- [3] Wooldridge, M.: *An introduction to multiagent systems*. 2nd edn. Wiley (2009)
- [4] Farinelli, A., Iocchi, L., Nardi, D.: Multirobot systems: A classification focused on coordination. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* **34** (2004) 2015–2028
- [5] Dagaëff, T., Chantemargue, F., Hirsbrunner, B.: Emergence-based cooperation in a multi-agent system. In: *Proceedings of the Second European Conference on Cognitive Science (ECCS'97)*, Citeseer (1997) 91–96
- [6] Iba, H.: Emergent cooperation for multiple agents using genetic programming. *Parallel Problem Solving from Nature—PPSN IV* (1996) 32–41
- [7] Shehory, O., Kraus, S., Yadgar, O.: Emergent cooperative goal-satisfaction in large-scale automated-agent systems. *Artificial Intelligence* **110** (1999) 1–55
- [8] Katoh, T., Hoshi, K., Shiratori, N.: Cooperative behavior of agents based on potential field. *Multi-Agent Systems and Applications IV* (2005) 236–245
- [9] Parker, L.: Distributed intelligence: Overview of the field and its application in multi-robot systems. *Journal of Physical Agents* **2** (2008) 5–14
- [10] Jennings, N.: Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review* **8** (1993) 223–250
- [11] Trianni, V., Ampatzis, C., Christensen, A., Tuci, E., Dorigo, M., Nolfi, S.: From solitary to collective behaviours: Decision making and cooperation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **4648 LNAI** (2007) 575–584
- [12] Stroupe, A., Okon, A., Robinson, M., Huntsberger, T., Aghazarian, H., Baumgartner, E.: Sustainable cooperative robotic technologies for human and robotic outpost infrastructure construction and maintenance. *Autonomous Robots* **20** (2006) 113–123
- [13] Settembre, G., Scerri, P., Farinelli, A., Sycara, K., Nardi, D.: A decentralized approach to cooperative situation assessment in multi-robot systems. *Proceedings of the 7th international joint conference on autonomous agents and multiagent systems* **1** (2008) 31–38
- [14] Stauffer, D., Sousa, A., de Oliveira, S.: Generalization to square lattice of Sznajd sociophysics model. *International Journal of Modern Physics C-Physics and Computer* **11** (2000) 1239–1246

- [15] Stauffer, D.: Monte Carlo simulations of Sznajd models. *Journal of Artificial Societies and Social Simulation* **5** (2001)
- [16] Sznajd-Weron, K.: Sznajd model and its applications. *Acta Physica Polonica B* **36** (2005) 2537
- [17] Vig, L., Adams, J.: Multi-robot coalition formation. *Robotics, IEEE Transactions on* **22** (2006) 637–649
- [18] Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-aguilar, J.A., Sousa, P.: Issues in multiagent resource allocation. *Informatica* **30** (2006) 2006
- [19] Dias, M., Zlot, R., Kalra, N., Stentz, A.: Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE* **94** (2006) 1257–1270
- [20] Koenig, S., Keskinocak, P., Tovey, C.: Progress on agent coordination with cooperative auctions. In: *AAAI Conference on Artificial Intelligence*. (2010)
- [21] Nanjanath, M., Gini, M.: Repeated auctions for robust task execution by a robot team. *Robotics and Autonomous Systems* **58** (2010) 900–909 *Advances in Autonomous Robots for Service and Entertainment*.
- [22] Berhault, M., Huang, H., Keskinocak, P., Koenig, S., Elmaghraby, W., Griffin, P., Kleywegt, A.: Robot exploration with combinatorial auctions. In: *2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings*. (2003) 1957–1962
- [23] Shoham, Y., Leyton-Brown, K.: *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press (2009)
- [24] Cramton, P., Shoham, Y., Steinberg, R.: *Combinatorial auctions*. The MIT Press (2006)