

Design and Mathematical Analysis of Agent-based Systems

Kristina Lerman

Information Sciences Institute
University of Southern California
Marina del Rey, CA 90292, USA
lerman@isi.edu

Abstract. Agent-based systems that are composed of simple locally interacting agents but which demonstrate complex group behavior offer several advantages over traditional multi-agent systems. A well-designed complex agent-based system is an efficient, robust, adaptive and stable system. It has very low communication and computational requirements, meaning that there are virtually no constraints on the system size. The simplicity of agent interactions also makes it amenable to quantitative mathematical analysis. In addition to offering predictive power, mathematical analysis enables the system designer to optimize system performance.

To date, there have been relatively few implementations of complex agent-based systems, mainly because of the difficulty of determining what simple agent strategies will lead to desirable collective behavior in a large system. We claim that there exists a set of primitive agent strategies, similar to the basis behaviors in behavior-based robotics, from which complex group behavior can be designed. Moreover, these simple primitive strategies naturally lend themselves to mathematical description, making a quantitative study of agent-based systems possible. We present a case study of coalition formation to show that two simple behaviors, *dispersion* and *aggregation*, can lead to coalition formation in a multi-agent system under some conditions. We use this system to illustrate the process by which a mathematical description of the agent-based system is created and analyzed, and discuss the insights the analysis provides for designing coalition forming agents.

Introduction

Recent years have witnessed an explosion of interest in the study of agent-based systems, *i.e.*, systems composed of many interacting autonomous, artificial intelligent agents. Such systems may be used for distributed control (*e.g.*, network routing [24, 4]), distributed resource management (*e.g.*, load balancing [23, 9]), optimization [11], and electronic marketplaces [3, 14]. Distributed robotics is a field of embodied agent-based systems that addresses the problems of coordinated action [12, 15, 16, 6] as well as learning [18] in a group of robots. Agent-based computing has been introduced as the next step in the development of

computation, in which autonomous proactive components replace the mostly reactive objects programmers create today [10].

Individual *vs* Emergent Complexity

The central problem in the design of a multi-agent system is how much intelligence to place in the system and at what level. Does a central authority create and direct group behavior, or does the global behavior emerge from interactions among many individual agents? How complex should each agent's behavior be? The vast majority of the work in this field has focused on making agents more knowledgeable and able. This has been achieved in several ways: by giving the deliberative agent a deeper knowledge base and ability to reason about data [21], giving it the ability to plan actions [26], negotiate with other agents [20], or change its strategies in response to actions of other agents [23, 8]. At the opposite end of the spectrum lie agent-based systems that demonstrate complex group behavior, but whose individual elements are very simple. Such systems have not received much attention in the agents community.

It has long been recognized in physics and biology that complicated global activity can result from very simple local interactions. Examples of this phenomenon pervade the natural world and include among others: pattern formation in thermal convection [13], Turing patterns [27] in chemical reaction-diffusion systems [19], the transition by which single cell slime mold amoeba aggregate to form a functioning multicellular organism [22].

Collective behavior emerging from local microscopic interactions has also been observed in many species of social insects. Though individual insects are arguably very simple creatures, having very limited memory, knowledge or reasoning facilities, an insect society demonstrates many complex behaviors. Hives, ant trails and swarms are examples of such robust, adaptable, seemingly "organized" collective behavior that does not have any central control [2].

Design of Agent-Based Systems

Complex agent-based systems (CAS) composed of simple agents that demonstrate complex collective behavior offer several advantages over traditional multi-agent systems that rely on deliberative agents. Though some problems are best suited for knowledgeable and able agents, traditional multi-agent systems (MAS) that utilize them have significant shortcomings in at least one of the following areas: robustness, adaptability, stability and scalability. Complex agents may fail, and if a central controller is involved in directing actions of agents, it has to be able to recover in the event of agent failure. Systems in which agents change their strategies in response to actions by other agents can quickly adapt to environmental changes; however, this feature is usually achieved at the expense of global stability [8]. The high communication and computational cost required to coordinate agent behavior constrain the size of the traditional MAS to at most a few dozen agents. Yet another disadvantage is that the complexity of the

agent’s internal states and its interactions with other agents make these systems ill suited for rigorous quantitative analysis.

A well-designed CAS, on the other hand, is an efficient, robust, adaptive and stable agent-based system. It lacks central control, meaning that the system can recover quickly from mistakes, agent failure and environmental change. Because it has very low communication and computational requirements, there are virtually no constraints on system size. This simplicity makes CAS amenable to mathematical analysis. Despite their numerous advantages, there have been relatively few implementations of CAS outside of distributed robotics. The scarcity is partially explained by the difficulty of designing a CAS. The designer, in a sense, has to reverse-engineer the problem, *i.e.*, determine what microscopic interactions, or basis behaviors, are necessary to produce the desired collective behavior [17].

Matarić [15] introduced basis behaviors as the fundamental components for behavior-based control in robots. A small set of primitive behaviors — collision avoidance, following, dispersion, aggregation and homing — is sufficient to synthesize complex behavior, such as foraging and flocking, in a single robot or a group of robots. We claim that a similar set of primitive agent strategies can be formulated for software agents, and they will serve as basic components for synthesizing collective behavior. Consider, for example, coalition formation. Coalition formation is a desirable behavior in systems where a group of agents can accomplish a task more effectively than a single agent can. The tasks may be very different – from collective block pushing, to commuter ride sharing, to consumers forming buying clubs to purchase products in bulk in order to save money – yet the underlying mechanism is always the same. We will demonstrate that coalition-formation in a system of software agents can result from two primitive agent strategies: dispersion and aggregation. Dispersion allows the agents to explore the environment in which they are situated and to encounter other agents and coalitions. Once an agent encounters a coalition, it makes a decision about whether to join it (aggregate). Other collective behaviors, such as distributed control and optimization (*e.g.*, task allocation), distributed resource management (*e.g.*, load balancing), collaborative information gathering, and cooperative transport in robots, may require the introduction of other primitive strategies.

Mathematical Analysis of Agent-based Systems

Decomposing a complex collective behavior into simple individual strategies allows us to create a mathematical model of the process and to analyze the system quantitatively. The analysis helps the agent designer determine what the important parameters of the problem are, and their effect on the global characteristics of the system. Analysis also helps the designer decide how the primitive strategies should be put together to optimize the system performance. We will focus on the *macroscopic* models that directly describe collective behavior. For coalition-formation, for example, the macroscopic model describes the number

and size of coalitions, not the behavior of individual agents. While some researchers have studied the behavior of agent-based systems quantitatively, they have almost exclusively focused on using *microscopic* simulations, such as molecular dynamics [5] and cellular automata, to model interactions between agents. Though microscopic simulations are an important tool for understanding the connection between the microscopic (agent) and the macroscopic (collective) behaviors, they offer only an indirect, empirical approach to studying the collective effects. Macroscopic models, on the other hand, directly describe collective behavior. Their other advantages are that they are more computationally efficient, because they use many fewer variables than the microscopic models; they are generic, meaning that the same model (with different parameters) is applicable to different systems, and they have better predictive power. Of course, microscopic and macroscopic theories are related, and understanding the connection between the two, *e.g.*, through simulation or by deriving the latter from the former, is an important goal of any complex systems research.

Coalition Formation in Agent-Based Systems

We present a case study of coalition formation in a multi-agent system to illustrate the process by which a mathematical description of the agent-based system is created, analyzed and what insights this analysis can provide for the system's design. The model we obtain is expressed mathematically as a series of differential equations that describe how the number and distribution of coalitions change with time. In addition to the dynamic variables, in this case coalitions, the equations contain variable parameters that determine how various agent strategies contribute to the behavior of the system. Many of the important global characteristics of the collective behavior, such as the existence of the steady state, the time it takes for the system to reach it, and the overall benefit of cooperation, depend on the values of these parameters. Information about the relationship between these parameters and the agent strategy is valuable to the designer of the system, who may want to control the global properties of a large-scale system.

Coalition formation is a valuable collective behavior in many systems. Consider a system where each agent is given a task to obtain goods at the lowest price [14]. We assume that agents know all the vendors that supply the requested goods and the retail (base) price for the product. Because bulk orders reduce manufacturer's costs, the vendors pass some of the savings to consumers. Therefore, the agents can lower the price they pay for the products by forming coalitions to buy them in bulk. Each agent moves among vendor sites. If it encounters other agents at a site, it can join the coalition or form a new one with a single agent. Manufacturing and other types of constraints limit the bulk order to a maximum size; therefore, vendors will not accept orders greater than this maximum.

The coalition formation mechanism outlined above is quite general and applies to any system in which it is beneficial for rational agents to form groups.

Moreover, this mechanism is local and requires minimal communication between agents. Agents learn indirectly about the presence and size of the coalition at a particular site by querying the vendor for the current price of the product. The agent knows the relationship between the size of the coalition and the expected benefit of joining it. This mechanism can be decomposed into two elements: dispersion and aggregation. Because these can be programmed directly into the agent's behavior, we call these elements primitive strategies or basis behaviors. The benefit of choosing these primitive strategies is that they are easy to describe mathematically, in addition to being relatively easy to implement. We list the primitive strategies below, together with additional simplifying assumptions that will make the construction of the initial mathematical model easier.

- *Dispersion*: Agents encounter other agents and coalitions randomly.
- *Aggregation*: Each agent's strategy is determined entirely by local conditions — the size of the coalition present at a particular site.
- Agents are homogeneous, in a sense that each agent has the same goal and follows the same strategy.
- The agents are spatially uniformly distributed, apart from the non-uniformities inherent in the coalitions. Even if this is not true initially, dispersion will tend to make the system uniform.
- There is no net change in the number of agents in the system.
- Agents are “mobile”, *i.e.*, free to choose vendors; coalitions are not.
- It is beneficial for agents to join a coalition; however, an agent cannot join a coalition already of maximum size.
- Agents are self-interested; therefore, given several alternatives, they will prefer and select ones most beneficial to their goals.
- An agent may leave a coalition and find a better one to join.

Though the assumptions above leave us with a severely idealized system, the model that we create using these assumptions still contains all the important ingredients of the coalition formation process and will correctly capture the behavior of the system. We can make the model more realistic by incrementally relaxing these assumptions and adding more realistic agent behaviors.

The Macroscopic Model

Using the axioms above we can construct a microscopic theory of the coalition formation process, which treats the individual agents as the fundamental units. This model would describe how agents make decisions to join coalitions. Alternatively, we can construct a macroscopic model that treats coalitions as fundamental units of the system. A macroscopic description offers several advantages, the most important is that such a model directly describes the global properties of the system we are interested in studying, namely the number and size of coalitions, and how these quantities change with time. Macroscopic theories tend to be more universal—the same mathematical description can be applied to other systems in which aggregation occurs. However, the macroscopic model

are usually phenomenological, and in some cases it may be necessary to derive the parameters of the model, and the model itself, from microscopic theory. This is not as vital in our application, because there is a simple connection between the microscopic behavior of the agents and the parameters of the model

We now present the macroscopic model that describes the time evolution of coalitions. The dynamic variables of the problem are the quantities we are interested in studying, namely coalitions, and they are labeled by their size. Let $r_1(t)$ denote the number of unaffiliated agents in the system at time t , $r_2(t)$ the number of coalitions of size two, *etc.*; $r_n(t)$ the number of coalitions of size n at time t , up to a maximum coalition size m .

Global Utility Gain

The global utility gain measures the efficiency of the system. For the e-commerce application, the total utility gain is the price discount all agents receive by being members of coalitions and was shown to be [14]

$$G = N \Delta p \left(\sum_{n=1}^m \frac{n^2 r_n}{N} - 1 \right).$$

where Δp measures how steeply the vendors decrease the price for each new member of coalition. We derived this expression from the discount $\Delta p(n-1)$ received by each member of the coalition of size n . Note that this form of the utility gain applies only to the e-commerce application. For other problem domains another expression for the utility gain function may be necessary.

Dynamic Equations

Initially (at $t = 0$) the system consists of N agents and no coalitions. We assume that there is no spatial dependence in the agent distribution, apart from coalition-based aggregation; therefore, the variables are functions of time only. A series of coupled ordinary differential equations describe how the number of coalitions of different size changes in time; therefore, the solutions of the equations yield the coalition distribution at any given time:

$$\frac{dr_1(t)}{dt} = -2D_1 r_1^2(t) - \sum_{n=2}^{m-1} D_n r_n(t) + 2B_2 r_2(t) + \sum_{n=3}^m B_n r_n(t), \quad (1)$$

$$\frac{dr_n(t)}{dt} = r_1(t) (D_{n-1} r_{n-1}(t) - D_n r_n(t)) - B_n r_n(t) + B_{n+1} r_{n+1}(t), \quad (2)$$

$$\frac{dr_m(t)}{dt} = D_{m-1} r_1(t) r_{m-1}(t) - B_m r_m(t) \quad (3)$$

Here $r_n(t)$ is the number of coalitions of size n at time t , and $\frac{dr_n}{dt}$ is the rate of change of this number. Parameter D_n , the attachment rate, controls the rate at which unaffiliated agents join coalitions of size n . This parameter includes

contributions from two factors: the rate at which agents encounter n -mers ($\propto r_1 r_n$, where the proportionality factor determines how many vendor sites an agent visits in a given period of time), and the probability of joining the coalition of size n . B_n , the detachment rate, gives the rate at which agents leave coalitions of size n . The solutions are subject to the initial conditions: $r_1(t=0) = N$ and $r_n(t=0) = 0$ for all $n > 1$.

Results

Initial investigation of the model focused on the uniform attachment–uniform detachment case: $D_n = D$, $B_n = B$ for all n . The results have shown that mathematical analysis is not only feasible, but also yields non-obvious results. To simplify the analysis, we rewrite the equations in dimensionless form by making the following variable transformations: $\tilde{r}_n = r_n/N$ (density of coalitions of size n) $\tilde{t} = D N t$, $\tilde{B} = B/DN$ (dimensionless detachment rate). When the equations are written in dimensionless form, only a single variable — \tilde{B} , the dimensionless detachment rate — governs the behavior of solutions. The equations were integrated numerically using Mathematica for $m = 6$ and different values of \tilde{B} . In all cases solutions reach a steady state, in which the distribution of coalition densities no longer changes. Figure 1 shows how the steady state density of coalitions of each size changes. The leftmost unconnected set of points are for the no-detachment case $\tilde{B} = 0$. The steady state at this point consists mostly of coalitions of size two and three, a quickly decreasing number of larger size coalitions, and no unaffiliated agents. When \tilde{B} is small, \tilde{r}_1 is also small and the largest coalitions dominate. The number of unaffiliated agents, \tilde{r}_1 , increases with \tilde{B} . Coalitions start to “evaporate” quickly at $\tilde{B} \approx 1$; as a result, the number of larger coalitions drops precipitously. Note that there is a discontinuity at $\tilde{B} = 0$: the steady state solutions are qualitatively different for $\tilde{B} \rightarrow 0$ than at $\tilde{B} = 0$. Moreover, for $\tilde{B} \neq 0$, the steady state is an equilibrium state: even though agents are continuously joining and leaving coalitions, the overall distribution of coalitions does not change. For $\tilde{B} = 0$, the system gets trapped in a non-equilibrium state before it is able to form larger coalitions.

The global utility gain (per agent) in the steady state, calculated according to Eq. 1, is shown in the lower half of Figure 2. The utility gain is largest for small non-zero \tilde{B} . Its value for $\tilde{B} = 10^{-6}$ is $G/N = 4.87\Delta p$ — a substantial increase over the no-detachment case value of $G/N = 2.00\Delta p$. For large detachment rates there is virtually no utility gain, as the system is composed mainly of unaffiliated agents. The large increase in the utility gain for small \tilde{B} comes at a price, namely the time required to reach the steady state, plotted in the top half of Figure 2. While it takes $\tilde{t} \approx 10$ for solutions to reach the final state for $\tilde{B} = 10$, it takes $\tilde{t} \approx 10^9$ for the solutions to equilibrate for $\tilde{B} = 10^{-6}$.

The no-detachment ($\tilde{B} = 0$) case is especially interesting because mathematical analysis shows that global system properties, specifically, the steady state coalition distribution and the global utility gain, are outside of the agent designer’s control. We have already shown using dimensional analysis, that coalition distribution is independent of the attachment rate, D . We claim that for

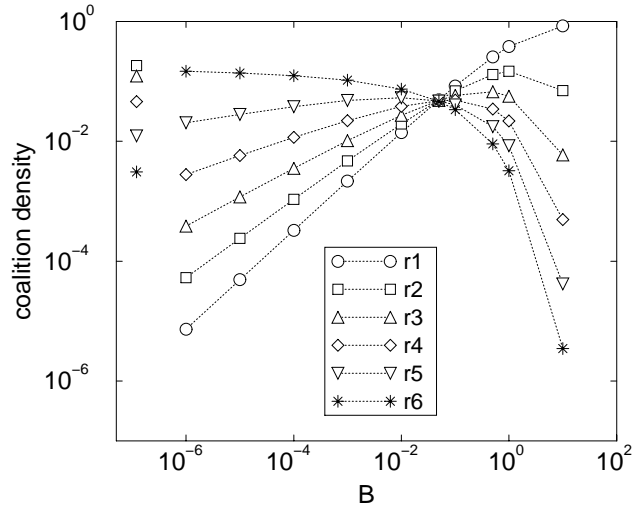


Fig.1. Steady state distribution of coalition densities *vs.* the dimensionless detachment rate. The open symbols are coalition densities r_1 (unaffiliated agents) through r_6 (coalitions of 6 agents).




Fig.2. The global utility gain per agent in the steady state *vs.* the dimensionless detachment rate and the time it takes for the system to reach the steady state.

the $\tilde{B} = 0$ case, it is also independent of m , the maximum allowed coalition size. We introduce a new variable — $\tilde{r} = \sum_{n=2}^{\infty} \tilde{r}_n$ — the total coalition density. As m becomes large, the first sum in Eq. 1 approaches \tilde{r} . Summing all equations for every coalition size and using the fact that the sum of derivatives is the derivative of the sum, allows us to rewrite the rate equations in terms of two variables only: \tilde{r}_1 and \tilde{r} .

$$\frac{d\tilde{r}_1}{dt} = -2\tilde{r}_1^2(t) - \tilde{r}_1(t)\tilde{r}(t), \quad (4)$$

$$\frac{d\tilde{r}}{dt} = \tilde{r}_1(t)\tilde{r}_1(t). \quad (5)$$

The first equation is very similar to Eq. 1, and while the second equation might seem counterintuitive at first glance, it has a very simple meaning: namely, the total number of coalitions changes only when two unaffiliated agents join to form a coalition of size two. Solutions of these equations behave exactly the same way as solutions for each coalition size. Figure 3(a) shows the time evolution of the solutions to the full equations with $B = 0$ and $m = 6$, while the dashed lines in Fig. 3(b) show the time evolution of solutions to Eqs. 4–5. For comparison, we also plot the sum of coalitions of size two through six from Fig. 3(a). Although equations Eq. 4 and Eq. 5 are valid for large m , we can see that there is no appreciable difference in the steady state solutions, and therefore, the utility gain, already for $m = 6$.

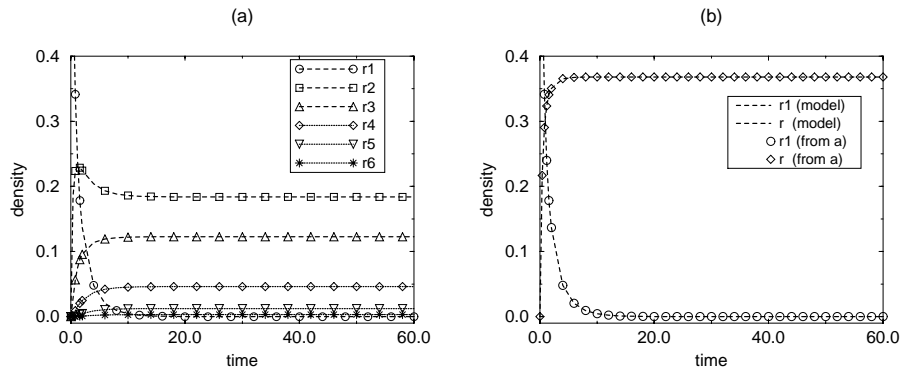


Fig. 3. (a) Coalition density *vs.* dimensionless time for the case where maximum coalition size is $m = 6$. (b) Solutions of equations for the single agent and total coalition densities *vs.* time (dashed lines). Single agent density (circles) and the sum of coalitions of size two through six (diamonds) in (a) are plotted for comparison.

We can try to repeat the analysis for the $\tilde{B} \neq 0$ case for large m by introducing a total coalition density variable and summing equations 1–3. Unfortunately, we can not eliminate every variable, and we are left with two equations and three

unknowns: \tilde{r}_1 , \tilde{r}_2 and \tilde{r} . This can be understood by considering that the total number of coalitions not only increases when two agents form a coalition of size two, but it can also decrease when a pair disintegrates. We can introduce a third equation for \tilde{r}_2 . However, this equation involves a new variable, \tilde{r}_3 , and to take *it* into account we need yet another equation, *etc.*, until we reproduce Eqs. 1–3. Fortunately, we can obtain analytic expressions for the steady state densities in terms of the monomer density from the rate equations by setting the left-hand side of Eqs. 1–3 to zero. We find that at late times the densities obey a simple relationship:

$$\tilde{r}_n = \tilde{B}^{-(n-1)} \tilde{r}_1^n. \quad (6)$$

By studying the behavior of solutions for different values of m , we empirically obtain a scaling law for the steady state monomer density,

$$\tilde{r}_1 \propto \tilde{B}^{\frac{m-2}{m-1}}. \quad (7)$$

This result is valid in the parameter range that we are interested in, namely where the utility gain is large and slowly varying. Equations Eq. 6 and Eq. 7, together, allow the agent designer to predict how the steady state density of coalitions of any size changes as the detachment rate or the maximum coalition size is changed. In particular, as maximum coalition size becomes large, the exponent of \tilde{B} approaches 1. In this case $\tilde{r}_n \propto \tilde{B}$, that is the number of coalitions of every size grows linearly with \tilde{B} .

Lessons for Agent Designers

We have shown that, at least for the uniform attachment–detachment cases, a steady equilibrium state is reached for all non-zero values of the control parameter \tilde{B} . It is as yet unresolved whether the equilibrium state is stable. If individual agents are not allowed to leave coalitions (corresponding to the case $\tilde{B} = 0$), the steady state coalition distribution is independent of any of the system parameters that are under the designer’s control. Even so, there is some utility gain in this system, reflecting the presence of small coalitions. Introducing even a very small detachment rate (or “shopping around” rate) to the basic coalition formation process allows the system to increase the global utility gain by more than a factor of two over the $\tilde{B} = 0$ scenario. As the relative strength of the detachment rate increases, the utility decreases until there is virtually no utility gain. (This fact is known to every shopper who has spent the day at the mall running from store to store comparing prices and not ending up buying anything.) The price for higher utility gain is that the time required to reach the steady state solution grows very large as \tilde{B} becomes small. However, utility gain remains large and decreases slowly over many orders of magnitude of \tilde{B} . The agent designer has much leeway in choosing parameter values that result in a substantial global benefit, while not requiring too long a wait for this benefit to be achieved. The agent designer can also predict the final distribution of coalitions, even for very large systems.

Related Research

Axelrod *et al.* [1] show how cooperative behavior can emerge among selfish autonomous agents. They use game dynamics to simulate interactions between two agents, in which the agents have to make decisions, each with a different pay-off or benefit to the agent. The agent's decision depends on choices made by other agents. Some strategies were shown to lead to stable cooperation, because it increases the overall pay-off to both cooperating agents. Others [23, 7] have applied game dynamics formalism to distributed control, where many agents adjust their strategies (a decision to compete or to cooperate) to grab a larger share of a finite resource. The focus of this work is adaptation in a distributed system, *i.e.*, how a group of agents can learn to cooperate to achieve a common goal without any central control. Some of the systems (see, for example, Huberman and Hogg's work on computational ecologies [8, 7]) are amenable to mathematical analysis, though most results about the stability of the system have been achieved through simulation. This line of inquiry is similar to ours, in a sense that it studies the global dynamics of a system of locally interacting agents. The aim of simulations, however, is to demonstrate the existence of evolutionary stable strategies that drive the system to the steady optimal solution, and how the agent's strategy evolves to maximize the benefit to itself. Our main goal, on the other hand, is understanding the global dynamics so that we can *control* the collective behavior of the system by manipulating individual agent's strategy.

Shehory *et al.* [25] have studied a large scale multi-agent system using a physics-based approach similar in spirit to ours. They suggest a low communication complexity coordination mechanism (though not coalition formation) for a large scale multi-agent system and use a physics-based microscopic model to analyze the system. Yet another physics-based approach used to study a multi-agent system is molecular dynamics simulations of swarm behavior [5]. In that work, the interactions among pairs of agents were modeled by a potential field. One advantage of microscopic simulations is that they allow agent behavior to be easily manipulated. However, simulations offer only an indirect way to study the collective behavior of the system, they are time consuming, and have little predictive power.

Acknowledgements

This research has been funded in part by the ISI-ISD Research Fund Award.

References

1. R. Axelrod and W. D. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.
2. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.

3. A. Chavez and P. Maes. Kasbah: An agent marketplace for buying and selling goods. In *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 75–90, 1996.
4. G. Di Caro and M. Dorigo. Antnet: Distributed stigmergetic control for communications networks. *J. of Artificial Intelligence Research*, 9:317–365, 1998.
5. R. Duncan, T. D. McCarson, R. Stewart, P. M. Alsing, R. P. Kalle, and R. Robinett. Statistical paradigms for robotic swarm modeling. <http://www.mhpcc.edu/research/ab98/98ab41.html>, 1998.
6. Dani Goldberg and M. J. Matarić. Coordinating mobile robot group behavior using a model of interaction dynamics. In Oren Etzioni, Jörg P. Müller, and Jeffrey M. Bradshaw, editors, *Proceedings of the Third Annual Conference on Autonomous Agents (AGENTS-99)*, pages 100–107, New York, May 1–5 1999. ACM Press.
7. O. Guenter, T. Hogg, and B. A. Huberman. Learning in multiagent control of smart matter. In *AAAI Workshop on Multiagent Learning*, 1997.
8. T. Hogg and B. A. Huberman. The behavior of computational ecologies. In B. A. Huberman, editor, *The Ecology of Computation*, Amsterdam, 1988. North-Holland.
9. Bernardo A. Huberman and Scott H. Clearwater. A multiagent system for controlling building environments. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems*, pages 171–176, San Francisco, CA, 1995. MIT Press.
10. N. R. Jennings and M. Wooldridge. Applications of Agent Technology. In N. R. Jennings and M. Wooldridge, editors, *Agent Technology: Foundations, Applications, and Markets*. sv, March 1998.
11. J. Kennedy and R. Eberhart. Particle Swarm Optimization. In N. R. Jennings and M. Wooldridge, editors, *Fourth IEEE International Conference on Neural Networks*. IEEE Service Center, 1995.
12. H. Kitano, M. Asada, Y. Kuniyoshi, and I. Noda. RoboCup: A challenge problem for AI and robotics. *Lecture Notes in Computer Science*, 1395:1, 1998.
13. K. Lerman, D. S. Cannell, and G. Ahlers. Different convection dynamics in mixtures with the same separation ratio. *Physical Review*, E53:R2041, 1996.
14. Kristina Lerman and Onn Shehory. Coalition Formation for Large-Scale Electronic Markets. To appear in ICMAS'2000, 2000.
15. M. J. Matarić. Designing emergent behaviors: From local interactions to collective intelligence. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *Second International Conference on Simulation of Adaptive Behavior (SAB-92)*, pages 432–441, 1992.
16. M. J. Matarić. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4(1):50–81, December 1995.
17. M. J. Matarić. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16(2–4):321–331, December 1995.
18. M. J. Matarić. Learning in multi-robot systems. In Sandip Sen, editor, *IJCAI-95 Workshop on Adaptation and Learning in Multiagent Systems*, pages 32–37, 1995.
19. Z. Noszticzius, W. Horsthemke, W. D. McCormick, H. L. Swinney, and W. Y. Tam. Sustained chemical waves in an annular gel reactor: a chemical pinwheel. *Nature*, 329, 1987.
20. S. Parsons and N. R. Jennings. Negotiation through argumentation—a preliminary report. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems*. MIT Press, 1995.
21. S. Parsons, Carles Sierra, and Nick Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, June 1998.

22. W.-J. Rappel, A. Nicol, A. Sarkissian, H. Levine, and W. F. Loomis. Self-organized vortex state in two-dimensional dictyostelium dynamics. *Physical Review Letters*, 83:1247–1250, 1999.
23. A. Schaerf, Y. Shoham, and M. Tennenholtz. Adaptive load balancing: a study in co-learning. In Sandip Sen, editor, *IJCAI-95 Workshop on Adaptation and Learning in Multiagent Systems*, pages 78–83, 1995.
24. R. Schoonderwoerd, O. Holland, and J. Bruten. Ant-like agents for load balancing in telecommunications networks. In W. Lewis Johnson and Barbara Hayes-Roth, editors, *Proceedings of the 1st International Conference on Autonomous Agents*, pages 209–216, New York, February 5–8 1997. ACM Press.
25. O. Shehory, S. Kraus, and O. Yadgar. Emergent cooperative goal satisfaction in large-scale automated-agent systems. *Artificial Intelligence*, 110(1), 1999.
26. K. Sycara and A. S. Pannu. The RETSINA multiagent system: Towards integrating planning, execution and information gathering. In Katia P. Sycara and Michael Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (AGENTS-98)*, pages 350–351, New York, May 9–13 1998. ACM Press.
27. A. M. Turing. The chemical basis of morphogenesis. *Phil. Trans. Royal Soc. Lond.*, B327, 195.