Agent Memory and Adaptation in Multi-Agent Systems

Kristina Lerman
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
lerman@isi.edu

Aram Galstyan
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
galstyan@isi.edu

ABSTRACT

We describe a general mechanism for adaptation in multi-agent systems in which agents modify their behavior based on their memory of past events. These behavior changes can be elicited by environmental dynamics or arise as response to the actions of other agents. The agents use memory to estimate the global state of the system from individual observations and adjust their actions accordingly. We also present a mathematical model of the dynamics of collective behavior in such systems and apply it to study adaptive coalition formation in electronic marketplaces. In adaptive coalition formation, the agents are more likely to join smaller coalitions than larger ones while there are many small coalitions. The rationale behind this is that smaller coalitions are necessary to nucleate larger ones. The agents remember the sizes of coalition they encountered and use them to estimate the mean coalition size. They decide whether to join a new coalition based on how close its size is to the mean coalition size. We show that the adaptive system displays most of the features of the non-adaptive one, but results in better long term system performance.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Theory

Keywords

multi-agent systems, analysis, coalition formation

1. INTRODUCTION

The study of artificial intelligent agents has received a great amount of interest in both academic and lay communities through an ever-growing roster of applications, which include personal digital assistants that filter email [18], recommend useful Web sites [22] and monitor travel plans [1].

AAMAS'03, July 14–18, 2003, Melbourne, Australia.

At the same time, there has been an explosion of interest in the study of multi-agent systems (MAS), *i.e.*, systems composed of many interacting artificial intelligent agents. Such systems can be used for distributed control (*e.g.*, network routing [4, 24]), resource management (*e.g.*, load balancing [23]), electronic marketplaces, and problem solving via teamwork [13]. Robotics [13, 7, 11] and embedded systems, such as sensor networks [5], are two important branches of MAS research. In these systems, limited on-board power and significant communications costs argue for distributed control solutions, in which interesting and useful collective behavior arises out of local interactions between many agents, whether robots or sensors in a sensor network. In fact, the so-called biological metaphor has inspired the design of several robotic multi-agent systems (see [20, 3, 19, 8, 9] among others). In addition to providing a mechanism for distributed control, the biological metaphor offers several other benefits for the design of agent-based systems, including (i) scalability: each agent has the same controller whether the group is composed of ten or 10,000 agents; (ii) robustness: group performance is robust to individual agent failure; (iii) flexibility: agents can be dynamically added or removed without significantly affecting the performance of the system; (iv) local sensing: in many cases the desired collective behavior can be achieved via local interactions only; (v) adaptability: allows for simple learning that enables agents to operate in uncertain hostile environments.

A distributed control mechanism based on local interactions offers yet another advantage: it can be mathematically modeled and analyzed. Mathematical analysis is an alternative to experiment and simulation, the two principal tools used by computer scientists. While experiments (*e.g.*, with robots) allow researchers an opportunity to study the systems's behavior in real environments, they are very costly and time-consuming, both to set-up and execute. Computer simulations, including sensor-based simulations used by roboticists [21, 6] that recreate the experiments under realistic conditions, and probabilistic microscopic simulations [11], also usually don't scale well. In many cases they are impractical tools for a detailed investigation of the properties of most multi-agent systems. As a result, many of the interesting questions about the behavior of agent-based systems, especially of large systems, remain unanswered. Using mathematical analysis, on the other hand, we can efficiently study dynamics of even very large agent groups, predict their long term behavior, gain insight into system design: *e.g.*, what parameters determine group behavior, optimize performance, prevent instabilities, *etc.* However, with the exception of work Huberman and Hogg [10] and Sugawara and coworkers [25, 26], as well as our own research [16, 15, 14], little mathematical analysis of multi-agent systems exists.

In our earlier work we have developed an approach [17] that allows us to create a mathematical model of collective dynamics of multi-agent systems. The approach is valid for agents that obey the Markov property — *i.e.*, only the present state determines the agent's future state. We have applied this approach to study coalition formation among intelligent software agents [16], as well as foraging [14] and collaboration [15] in groups of robots. The last two are compelling illustrations of the approach, because results of mathematical analysis agree with the results experiment and simulations, while at the same time allowing us to draw conclusions about the behavior of the system that would have been difficult to obtain through experiment and simulation alone.

The simple models we investigated so far cannot describe systems composed of agents with memory, learning or deliberative capabilities. This is a major shortcoming, since majority of the research in the agents community is focused on deliberative agents that can use past experience to guide future actions. In this paper we describe an extension of the modeling approach that allows us to incorporate history or memory of agents' actions into the mathematical model of a multi-agent system. This allows us to model agents that make decisions about future actions based on past experience, and thus study adaptation and learning in multi-agent systems. Although learning has been one of the most important topics in computer science, few mathematical descriptions of MAS composed of concurrent learners exist [27]. This is an important step towards creating a *science* of multi-agent systems that will allow quantitative understanding of the behavior of a collective of agents in a given environment. Such a science of MAS will enable researchers design better and more robust MAS control mechanisms.

The rest of the paper is organized as follows. In Section 2 we present the model of the collective behavior of agents with memory. In Section 3 we illustrate our approach by studying a system in which agents use memory to adjust individual actions and at the same time improve collective performance. We extend the model of coalition formation in electronic marketplaces introduced in an earlier work [16] to an adaptive scenario where agents rely on the history of past states to choose what action to take (join a coalition or not). We show that adaptive systems have better collective performance than systems composed of memoryless nonadaptive agents.

2. MODELING RETROSPECTIVE AGENTS

In earlier works [16, 15, 14, 17] we showed that the dynamics of collective behavior in a system of agents obeying the Markov property is captured by the Rate Equations that describe how the average number of agents in each state changes in time. *State* labels a set of related agent behaviors required to accomplish a task. As an example, consider a robot engaged in a foraging task, whose goal is to collect objects scattered around an arena and deliver them to a prespecified home location. The foraging task consists of the following high-level behavioral requirements [2] or states: (i) *homing*, (ii) *puck pickup* and (iii) *searching*. Each of these states corresponds to a single action or behavior; however, it is often useful to coarse-grain the system by grouping related behaviors into a single state. Such coarse-graining not only helps in conceptualizing the system, it also keeps the mathematical model compact and tractable by reducing the number of states. For example, when a robot is said to be in the *searching* state, it is wandering around the arena, detecting objects and avoiding obstacles. In the course of accomplishing a task, the robot will transition from the *searching* to *pickup* and finally to *homing* states. It is clear that during a sufficiently short time interval each agent in a multi-agent system is in exactly one of a finite number of states.

Let $p(n,t)$ be the probability an agent is in state n at time t. Agents that use a finite memory of the past of length m in order to make decisions about future actions can be represented by a generalized Markov processes of order m. This means that the state of an agent at time $t+\Delta t$ depends not only on the configuration of the system at time t (as in simple Markov systems), but also on configurations at times $t-\Delta t,\ t-2\Delta t,\ \ldots,\ t-(m-1)\Delta t$, which we refer to as history h of the system. In the derivation below we will employ the following identities:

$$\begin{aligned} p(n,t+\Delta t|h) &= \sum_{n'} p(n,t+\Delta t|n',t;h)p(n',t|h) \\ 1 &= \sum_{n} p(n,t+\Delta t|n',t;h). \end{aligned}$$

The change in probability density Δp is:

$$\begin{aligned} \Delta p(n,t) &= p(n,t+\Delta t) - p(n,t) && (1)\\ &= \sum_h \big[p(n,t+\Delta t|h) - p(n,t|h)\big]p(h) && (2)\\ &= \sum_h \sum_{n'} p(n,t+\Delta t|n',t;h)p(n',t|h)p(h) && (3)\\ &\quad -\sum_h \sum_{n'} p(n',t+\Delta t|n,t;h)p(n,t|h)p(h) \end{aligned}$$

In the continuum limit, as $\Delta t \to 0$, $\Delta p/\Delta t$ can be written as

$$\begin{aligned} \frac{dp(n,t)}{dt} &= \sum_h \sum_{n'} W(n|n';h)p(n',t|h)p(h) && (4)\\ &\quad -\sum_h \sum_{n'} W(n'|n;h)p(n,t|h)p(h)\,, && (5) \end{aligned}$$

with transition rates

$$W(n|n';h) = \lim_{\Delta t \to 0} \frac{p(n,t+\Delta t|n',t;h)\Delta t}{.} \qquad (6)$$

Equation 4 is similar in form to the stochastic Master Equation widely studied in statistical physics and chemistry [12]. It is a microscopic equation that describes the rate of change of probability density for an agent to be in state n at time t. The Master Equation has a simple interpretation: the probability density associated with state n will increase because of transitions from other states to state n, and it will decrease because of transitions from state n to the other states. It is more useful, however, to work with the macroscopic equation, the so-called Rate Equation, that describes how the configuration of the entire multi-agent system evolves in time. Following Van Kampen's indispensable text [12], we assume that agents are independent and indistinguishable. In this case, the configuration of the collective is described by occupation numbers $\{N\} = \{N_1, \ldots, N_L\}$, with N_n the number of agents in state n. Assuming that during a sufficiently short time interval only one of the agents in state n' will make a transition to state n, we can derive the macroscopic equation for the rate of change of $\langle N_n \rangle$, the average number of agents in state n:

$$\frac{d\langle N_n \rangle}{dt} = \sum_{h,n'} \left[W(n|n';h)\langle N_{n'} \rangle - W(n'|n;h)\langle N_n \rangle \right] P(h) \quad (7)$$

where $P(h)$ is the probability the system had a history h over the last m time steps. Using the mean-field approximation $\langle f(n) \rangle \approx f(\langle n \rangle)$ and the fact that N_n does not depend directly on h, we rewrite the above equation,

$$\frac{dN_n}{dt} = \sum_{n'} \left[\langle W(n|n') \rangle_h N_{n'} - \langle W(n'|n) \rangle_h N_n \right]. \quad (8)$$

Here for notational convenience $\langle \ldots \rangle_h$ denotes average over histories, and we have dropped angle brackets around N, although this variable still denotes an average quantity.

Equation 8 is very similar to the rate equation we used to study Markov-based agent systems [16, 15, 14], except that transition rates $W(n|n')$ are now replaced by their history-averaged values. We will use the above equation to study how agents can use histories, or memories of past events, to improve the collective behavior of the system. We illustrate the approach by re-examining coalition formation in electronic marketplaces, a memoryless version of which was introduced in [16]. Rather than blindly joining or leaving a coalition, the agent is now able to base its decision on what coalitions it has seen already. The agent designer can exploit the memory mechanism to allow each agent to adapt to the actions of other agents, resulting in an improved overall performance of the entire system.

3. COALITION FORMATION IN ELECTRONIC MARKETPLACES

In [16] we described a hypothetical electronic marketplace composed of mobile purchasing agents, in which each agent is given a task to obtain goods with the goal of minimizing the price paid for the goods. This market has the wholesale property, *i.e.*, the sellers can reduce their manufacturing, advertising and distribution costs by selling large quantities of the product in bulk and choose to pass some of the savings to the buyers. It is, therefore, in the buyers interest to form coalitions with other buyers allowing them entry into the wholesale market, thus reducing the price per unit.

The agents join coalitions by placing an order to purchase a product, and they leave coalitions by withdrawing the order. The orders remain open for some period of time to allow new orders to come in and coalitions to grow. At the end of the specified time, the orders are filled, and the price each agent pays for the product is based on the size of the final purchasing coalition.

In the same work we proposed and quantitatively studied a low-complexity mechanism for coalition formation, which is summarized below. Flowchart in Fig. 1 is a schematic of an individual agent controller. We assume that, given no additional information, agents have no *a priori* preference among vendors; therefore, each agent chooses a vendor site randomly. If it encounters another agent or a coalition of agents at that site, it may, with some probability, join the coalition or form a new one with the single agent.[1] An agent may also leave a coalition with some probability, because further exploration of the marketplace may result in it encountering a more beneficial coalition to join. Though we have not yet specified a relationship between coalition size and the utility of belonging to the coalition, we can assume, without loss of generality, that the benefit of being a coalition member depends on the coalition size; therefore, the decision to join or leave the coalition will also depend on the coalition size. In other words, the agent is more (less) likely to join (leave) a larger coalition rather than a smaller one (given there are no costs associated with leaving a coalition).

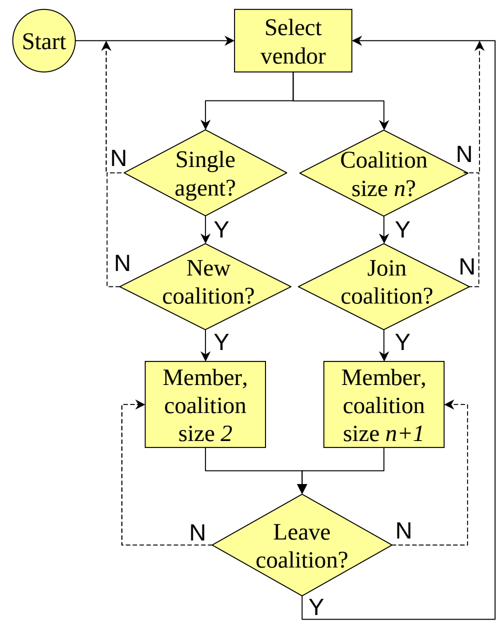

Figure 1: Single agent controller.

This mechanism, outlined in Fig. 1, requires minimal communication between agents, and because their decision depends solely on local conditions, it also requires no global knowledge. We present a mathematical model that describes the dynamics of the coalition formation process, *i.e.*, how the number and size of coalitions changes in time. In [16] we investigated a homogeneous multi-agent system, in which the agent's probability to join or leave a coalition was independent of its size. In this paper we examine an adaptive system where agents can modify their probabilities to join or leave a coalition based on their past experience and compare it to a non-adaptive system.

[1]We assume that some (*e.g.*, manufacturing and distribution) constraints limit the size of the bulk order, thus limiting the maximum size that can be attained by a coalition. The agents will not be able to join coalitions of maximum size. This assumption can be relaxed by essentially raising the maximum size of the coalition until it is equal to the number of agents in the system.

3.1 A Mathematical Model of Adaptive Coalition Formation

The variables of the mathematical model of coalition formation are the number of coalitions of a given size. Let $r_1(t)$ denote the number of searching agents unaffiliated with any coalition at time t, and $r_n(t)$ the number of coalitions of size n at time t, $2 \leq n \leq m$, where m is maximum coalition size.

The mathematical model of coalition formation consists of a series of coupled rate equations Eq. 8, each describing how the dynamic variables, $r_1, r_2, \ldots, r_m$, change in time. Solving the equations, subject to the condition that initially (at $t = 0$) the system consists of N_0 agents and no coalitions, yields the coalition distribution at any time. The most difficult part of constructing the model is finding an expression for transition rates. We make a simplifying assumption that coalitions and agents are uniformly distributed in vendor space. Under this approximation, the transition rate $W(n+1|n)$ is a product of the rate at which an agent encounters coalitions of size n ($\propto r_n$) and the probability of joining a coalition of size n. $W(n|n+1)$ is simply proportional to the probability of an agent leaving a coalition of size n. In some systems these proportionality factors can be calculated from first principles. In others, they can be estimated (from *e.g.*, how many coalitions an agent visits in a given period of time), and in still others it may be expedient to leave them as parameters to be estimated by fitting the model to experimental data or simulations. Note that in adaptive systems transition rates are averaged over agent histories.

The rate equations are written as follows (*cf.* Eq. 8):

$$\frac{dr_1}{dt} = -2D_1 {r_1}^2(t) - \sum_{n=2}^{m-1} D_n r_1(t) r_n(t) + 2B_2 r_2(t) + \sum_{n=3}^{m} B_n r_n(t)\,, \quad (9)$$

$$\frac{dr_n}{dt} = r_1(t)\left(D_{n-1} r_{n-1}(t) - D_n r_n(t)\right) - B_n r_n(t) + B_{n+1} r_{n+1}(t)\,,\ 1 < n < m, \quad (10)$$

$$\frac{dr_m}{dt} = D_{m-1} r_1(t) r_{m-1}(t) - B_m r_m(t) \quad (11)$$

Parameter D_n, the attachment rate, controls the rate at which agents join coalitions of size n. B_n, the detachment rate, gives the rate at which agents leave coalitions of size n. These parameters are averaged over agent histories. Note that one of the equations above is superfluous because of the conservation of agents condition, $\sum_{n=1}^{m} n r_n = N_0$.

In [16] we studied the uniform non-adaptive case where attachment and detachment rates were independent of coalition size and independent of the history of agent actions. A more realistic scenario is to design the agents to be *more* (*less*) likely to join (leave) larger coalitions than smaller ones, since the benefit of belonging to a larger coalition is greater than the benefit of belonging to a smaller one. However, this principle may not be useful at all times. When most of the agents are still in the searching mode and there are few large coalitions, it may be more beneficial for agents to nucleate many small coalitions in order to enable larger ones to grow more quickly. This type of adaptive behavior can be achieved when agents can adjust their preferences to join or leave coalitions based on what coalitions are already present in the marketplace. One way to design this functionality is for a central authority to measure the mean coalition size r_{mean} and broadcast it to every agent. The agents can then base their decision to join a coalition n based on how different its size is from r_{mean}, using a gaussian with mean r_{mean} and width σ, for example. In an alternative distributed design, agents use memory to estimate the mean coalition size from the sizes of coalitions they have encountered in the past. As an illustration, consider an agent with $m = 1$ memory who is trying to decide whether to join a coalition of size 4 and remembers seeing a coalition of size 2 at the last vendor site it visited. It will use a gaussian of a given width centered at 2 to calculate the probability of joining a coalition of size 4. If this number is greater than a threshold, *e.g.*, 0.5, it will decide to join; otherwise, it will continue exploring the marketplace.

Equations 9–11 describe the dynamics of such an adaptive system. In the simplest adaptive scheme, we let each agent use a gaussian function to compute its transition rates, as described above. Using the Central Limit Theorem, we argue that the macroscopic attachment and detachment rates D_n and B_n, which contain contributions from many agents, are also gaussian and centered on r_{mean}.[2]

$$D_n = D_0 \exp^{-(r_{mean}-n)^2/2} / N_D$$
$$B_n = B_0 (1 - \exp^{-(r_{mean}-n)^2/2}) / N_B. \quad (12)$$

r_{mean} is changing in time; therefore, D and B are also varying in time. D_0 and B_0 set the scale of the parameters and N_D and N_B are normalization factors as described below. We will compare the performance of the adaptive system to that of a non-adaptive system in which an agent is c times more likely to join a coalition of size $n+1$ than a coalition of size n, and, for symmetry, c times less likely to leave that coalition as the next smaller one. This behavior is described by parameters $D_n = D_0 c^{n-1}$ and $B_n = B_0 c^{-(n-1)}$. It is clear that the case $c = 1.00$ is equivalent to the uniform rates scenario studied in [16] where the agent's decision to join or leave a coalition is independent of the coalition size. Equations 12 are normalized so that the total transition probability ($\sum_n D_n$ and $\sum_n B_n$) for the adaptive system is equal to the total transition probability for the nonadaptive system.

3.2 Results

We solved equations 9–10 subject to initial conditions $r_1(t = 0) = N_0$ and $r_n(t = 0) = 0$ for $2 \leq n \leq m$ for both adaptive and non-adaptive parameters. Figures 2–3 show how solutions evolve in time for the non-adaptive case with $c = 1.10$, and the adaptive case. The $c = 1.10$ case corresponds to a system in which agents are 10% more likely to join a coalition of size $n+1$ as size n. For both adaptive and non-adaptive systems, the maximum coalition size is $m = 6$ and $B_0/D_0N_0 = 0.001$. The solutions are normalized by the total number of agents (N_0) to give the density of coalitions of some size. The x-axis is in units of dimensionless time D_0N_0t.

In the non-adaptive case we see that after a long time, solutions evolve to a steady state in which the number of coalitions does not change, even though individual agents

[2] Another way to think about it is there is that coalition sizes are distributed around r_{mean}. The agents sample this distribution by their observations of coalitions. The collective observation of many agents yields a good estimate of r_{mean}.

continue to leave and join coalitions. The number of searching agents r_1 drops initially, but after some time grows to a small steady state value, which depends on B_0/D_0N_0. The steady state is composed mainly of coalitions of maximum size and a decreasing number of smaller coalitions. In the adaptive system, larger coalitions grow more quickly, as designed. The steady state is composed mostly of coalitions of maximum size and some searching agents, with a negligible number of coalitions of other sizes. The number of searching agents is larger than for the non-adaptive case. This effect can be understood from the following argument: as the mean coalition size grows, any small coalitions quickly dissolve into searching agents, but because the only coalitions that exist are those of maximum size, the searching agents are no longer able to find a coalition to join (since they are barred from joining maximum size coalitions). The adaptive system nevertheless appears to be more efficient than the non-adaptive system, since the former appears to have a bigger fraction of larger coalitions in the steady state.

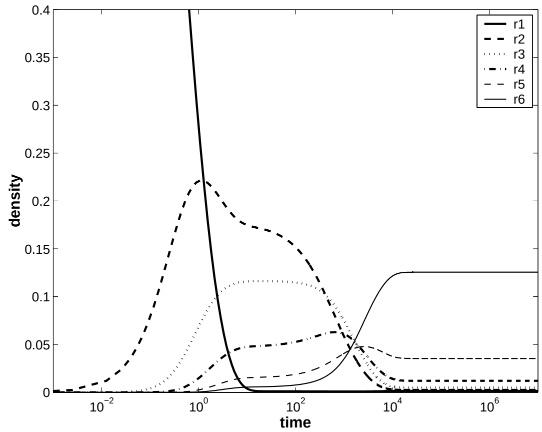

Figure 2: Time evolution of coalition densities for the non-adaptive system with $m = 6$**,** $B_0/D_0N_0 = 0.001$ **and** $c = 1.10$

We quantify the system's efficiency more precisely by the global utility gain. The utility gain is the price discount (reward) an agent receives for being a member of a buying coalition. If the retail price that an unaffiliated agent pays to the vendor for the product is p, and the coalition price each member of a coalition of size n pays is p_n, then $p_n < p$. In the simplest model we let $p_n = p - \Delta p(n-1)$, where Δp is some price decrement; therefore, the utility gain, or price discount, each coalition member receives is $G_n = \Delta p(n-1)$. The total utility gain measures the global performance of the system—the price discount all agents receive for joining coalitions. The value of this metric is expected to be high when there are many large coalitions, and conversely, it is low, when there is a large number of unaffiliated agents. The total total utility gain is:

$$G = \sum_{n=1}^{m} G_n n r_n = \Delta p \left(\sum_{n=1}^{m} n^2 r_n - N_0 \right) . \qquad (13)$$

The maximum utility gain, attained when all the agents belong to coalitions of size m, is $G_{max} = \Delta p N_0(m-1)$.

Figure 4 shows the evolution of the global utility gain per agent for two non-adaptive systems (with $c = 1.10$ and $c = 1.00$) and for an adaptive system. Non-adaptive system with size-dependent parameters leads to a bigger utility gain, especially at late times than the non-adaptive system with uniform size-independent parameters. However, adaptive system significantly outperforms non-adaptive systems both at early and very late times but underperforms them during intermediate times. The advantage of adaptive systems is most evident at very late times (in the steady state), when the utility gain is close to its maximum value. The slight decline at $t = 10^5$ is the result of an increase in the number of searching (unaffiliated) agents in Fig. 3. This behavior holds for different values of B_0 and D_0.

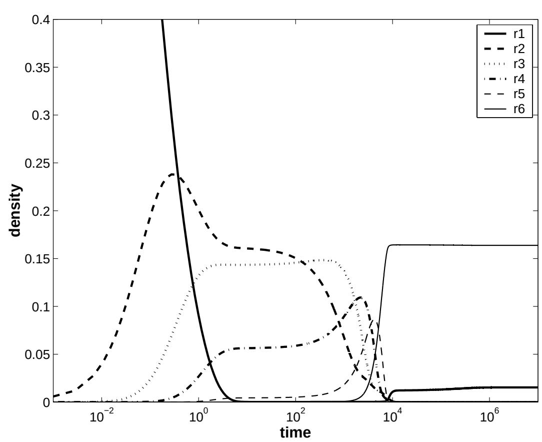

Figure 3: Time evolution of coalition densities for the adaptive system with $m = 6$ **and** $B_0/D_0N_0 = 0.001$

4. DISCUSSION

We have described a general mechanism for adaptation in multi-agent systems in which agents use memory of past events to change their behavior. Such behavior modifications can be elicited both by environmental changes or as a response to the actions of other agents. Specifically, the agents use their memory to estimate the global state of the system from the states of agents or environment they have

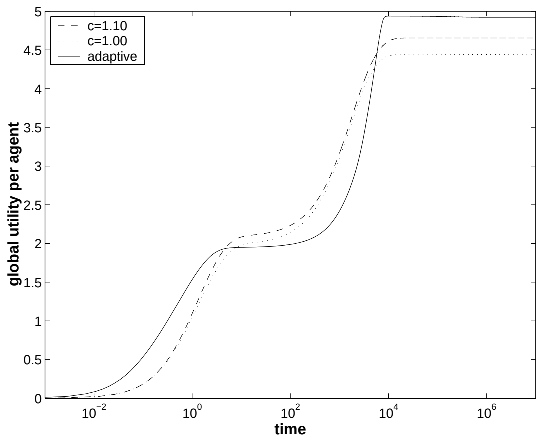

Figure 4: Time evolution of the global utility gain per agent G/N_0 for non-adaptive cases with $c = 1.10$ and $c = 1.00$ and for the adaptive case

encountered in the past, and use this estimate to adjust their own behaviors by modifying parameters that govern their actions. More importantly, we have presented a general mathematical model of the dynamics of collective behavior in such systems and applied it to study adaptive coalition formation in electronic marketplaces. In our illustration the agents form coalitions because they receive a benefit for doing so — specifically, the lower price they pay for a product when they buy it in bulk with other coalition members. We described a low complexity mechanism for coalition formation in which agents search the electronic marketplace for coalitions and decide whether to join a coalition or not. The agents may also decide to leave a coalition. In general, the decision to join or leave a coalition may depend on the coalition size. Starting only with searching agents unaffiliated with any coalition, the system evolves to a high utility state composed mainly of maximum size coalitions, but also of smaller size coalitions.

In an adaptive case, the agents use the memory of coalitions they encountered at previous times to estimate the mean coalition size. The agents are more likely to join a new coalition if its size is close to the mean coalition size, and less likely to leave such a coalition. The rationale behind such a scheme is that small coalitions have to form before larger ones can emerge; therefore, while there are few large coalitions, the agents should be more likely to join smaller ones. We showed that the adaptive system displays most of the features of the non-adaptive case (where the likelihood to join a coalition does not change in time), but is characterized by higher utility gains both at early and late times. We conclude that adaptation to the actions of other agents results in a better long term system performance.

We regard analytic results as still preliminary. Although its application to adaptive coalition formation was straightforward, several general questions remain: What is the role of history length m? What contributes to the error or variance of the agent's estimate of the global state of the system, and how does it influence collective behavior? Is it possible to reduce this variance with different memory lengths? Or do the time delay effects render adaptation meaningless? These and other questions are the focus of ongoing research activity.

Allowing agents to use memory of the past to adjust their behavior opens the door for us to study more complex and intelligent agents than the simple agents that have been studied in the past. In particular, we want to construct models that describe the behavior of decision-making and learning agents such as those described in the introduction. We believe that quantitative mathematical models will allow us to better understand and control the behavior of multi-agent systems composed of deliberative agents.

5. ACKNOWLEDGEMENTS

The research reported here was supported in part by the Defense Advanced Research Projects Agency (DARPA) under contracts number F30602-00-2-0573, in part by the National Science Foundation under Grant No. 0074790. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of any of the above organizations or any person connected with them.

6. REFERENCES

[1] Jose-Luis Ambite, Greg Barish, Craig A. Knoblock, Maria Muslea, Jean Oh, and Steven Minton. Getting from here to there: Interactive planning and agent execution for optimizing travel. In *The Fourteenth Innovative Applications of Artificial Intelligence Conference (IAAI-2002), Edmonton, Alberta, Canada, 2002*, 2002.

[2] Ronald C. Arkin. *Behavior-Based Robotics*. The MIT Press, Cambridge, MA, 1999.

[3] R. Beckers, O. E. Holland, and J. L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In Rodney A. Brooks and Pattie Maes, editors, *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems* $ArtificialLifeIV$, pages 181–189, Cambridge, MA, USA, July 1994. MIT Press.

[4] G. Di Caro and M. Dorigo. Antnet: Distributed stigmergetic control for communications networks. *J. of Artificial Intelligence Research*, 9:317–365, 1998.

[5] Deborah Estrin, Ramesh Govindan, and John S. Heidemann. Embedding the internet: Introduction. *CACM*, 43(5):38–41, 2000.

[6] Brian P. Gerkey, Richard T. Vaughan, Kasper Sty, Andrew Howard, Gaurav S. Sukhatme, and Maja J Matarić. Most valuable player: A robot device server for distributed control. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001), Wailea, Hawaii, October 29 - November 3, 2001*. 2001. http://www-robotics.usc.edu/player/.

[7] Dani Goldberg and Maja 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.

[8] Dani Goldberg and Maja J Matarić. Robust behavior-based control for distributed multi-robot

collection tasks. Technical Report IRIS-00-387, USC Institute for Robotics and Intelligent Systems, 2000.

[9] Owen Holland and Chris Melhuish. Stigmergy, self-organization and sorting in collective robotics. *Artificial Life*, 5(2), 2000.

[10] Bernardo A. Huberman and Tad Hogg. The behavior of computational ecologies. In B. A. Huberman, editor, *The Ecology of Computation*, pages 77–115, Amsterdam, 1988. Elsevier (North-Holland).

[11] A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella. Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, 11(2):149–171, 2001.

[12] N. G. Van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier Science, Amsterdam, revised and enlarged edition edition, 1992.

[13] 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.

[14] Kristina Lerman and Aram Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2), 2002.

[15] Kristina Lerman, Aram Galstyan, Alcherio Martinoli, and Auke Ijspeert. A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life Journal*, 7(4):375–393, 2001.

[16] Kristina Lerman and Onn Shehory. Coalition Formation for Large-Scale Electronic Markets. In *Proceedings of the International Conference on Multi-Agent Systems (ICMAS'2000), Boston, MA, 2000.*, pages 167–174, 2000.

[17] Kristina Lerman and Aram Galstyan. A methodology for mathematical analysis of multi-agent systems. *submitted to Journal of Agents and Multi-Agent Systems*, 2002.

[18] Pattie Maes. Agents that Reduce Work and Information Overload. *Communications of the ACM*, 37(7):31–40, July 1994.

[19] A. Martinoli, A. J. Ijspeert, and L. M. Gambardella. A probabilistic model for understanding and comparing collective aggregation mechanisms. In Dario Floreano, Jean-Daniel Nicoud, and Francesco Mondada, editors, *Proceedings of the 5th European Conference on Advances in Artificial Life (ECAL-99)*, volume 1674 of *LNAI*, pages 575–584, Berlin, September 13–17 1999. Springer.

[20] M. Matarić. Minimizing complexity in controlling a mobile robot population. In *Proceedings of the 1992 IEEE International Conference on Robo tics and Automation*, pages 830–835, Nice, France, 1992.

[21] O. Michel. Webots: Symbiosis between virtual and real mobile robots. In J.-C. Heudin, editor, *Proc. of the First Int. Conf. on Virtual Worlds, Paris, France,*, pages 254–263. Springer Verlag, 1998. See also http://www.cyberbotics.com/webots/.

[22] M. Pazzani, J. Muramatsu, and D. Billsus. Syskill & Webert: Identifying interesting Web sites. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-1996)*, 1996.

[23] Andrea Schaerf, Yoav Shoham, and Moshe Tennenholtz. Adaptive load balancing: a study in co-learning. In Sandip Sen, editor, *IJCAI-95 Workshop on Adaptation and Learning in Multi-agent Systems*, pages 78–83, 1995.

[24] Ruud Schoonderwoerd, Owen Holland, and Janet 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] Ken Sugawara and Masaki Sano. Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica*, D100:343–354, 1997.

[26] Ken Sugawara, Masaki Sano, and Ikuo Yoshihara. Cooperative acceleration of task performance: Analysis of foraging behavior by interacting multi-robots. In *Proc. IPSJ Int. Symp. on Information Systems and Technologies for Network Society*, pages 314–317, Fukuoka, Japan, September 1997.

[27] David H. Wolpert, Kevin R. Wheeler, and Kagan Tumer. General principles of learning-based multi-agent systems. In Oren Etzioni, Jörg P. Müller, and Jeffrey M. Bradshaw, editors, *Proceedings of the Third Annual Conference on Autonomous Agents (AGENTS-99)*, pages 77–83, New York, May 1–5 1999. ACM Press.