

Management of Mobile Agent Systems using Social Insect Metaphors

Tony White¹, Bernard Pagurek², Dwight Deugo¹

¹School of Computer Science, Carleton University

²Department of Systems and Computer Engineering, Carleton University

1125 Colonel By Drive, Ottawa, ON K1S 5B6 Canada

{arpwhite@scs.carleton.ca, bernie@sce.carleton.ca, deugo@scs.carleton.ca}

Abstract

The management of mobile agent systems that solve problems in a network is an issue that must be addressed if mobile agents are to be deployed industrially. It is clear that insufficient or excessive numbers of agents can cause the problem solving capabilities of an agent-based system to be impaired. Also, agents being software entities are almost always flawed therefore requiring the upgrade problem to be solved. This paper presents distributed algorithms based upon ant social behaviour that solve the problems of agent density maintenance and the agent upgrade problem.

Keywords: mobile agents, self-organizing agents, agent system management

1. Introduction

This paper concerns itself with the management of software agents moving throughout a network for the purpose of solving problems using distributed computation. In any system that supports distributed computation in an unreliable network, there is a need to address issues of agent density and upgrading. The application of mobile agents to Network Management [1] requires that these problems be solved in order to avoid the issues described in [6]; namely, that the solution becomes a problem itself.

An excessive number of agents in a network can significantly degrade the functioning of that network, while too few may also compromise performance [1], [5]. Generally, it is sufficient to maintain agent density within a range; a single point value is unnecessarily restrictive. In fact, point control can be the cause of oscillatory behaviour in a controlled system. In an unreliable network we find that links or network nodes may fail with a resulting loss of any agents executing on the node or in transit between two nodes. Some might argue that the computing infrastructure should support reliable computation and transport. However, this adds significantly to the complexity of the mobile agent infrastructure required and seems unnecessary

when we observe that naturally occurring agent systems (for example ants and wasps) are extremely tolerant to individual agent loss. Maintaining accurate statistics on agent numbers and position is really unnecessary for solving this problem if we view it from a decentralized viewpoint. In fact, agent populations should self-organize, as is clearly seen in nature, [8].

It is also too commonly the case that software, and agents are not likely to be exceptional here, is flawed either logically or functionally. When software flaws are discovered, software modification of individual agents or complete replacement of the defective agent must occur. A computing infrastructure supporting agent versioning is a more challenging problem than density maintenance, in that we know neither the positions of individual agents nor the versions actively moving through the network. This presents a serious halting problem in that we do not know the numbers of particular versions of the agent that are active in the network. It is not possible, then, to know when the upgrade process is complete. In other words, any algorithm or solution technique should be capable of upgrading older versions of agents for all time.

It is difficult to conceive of an environment that automatically upgrades agent software when changes are available that does not rely on some form of global information. For example, the current mechanism for software upgrade used on personal computers is to check periodically with the supplier of the software and download improvements when available. While mirror sites partially solve the load balancing problem, the solution is still a central one, one in which global information is held on the supplier's web site. Should the supplier move or disappear completely, the upgrade process fails. This, obviously, is an inferior solution. A decentralized solution to the versioning problem would have no such limitation, relying instead on only local information.

This paper proposes the use of algorithms that exploit ideas inspired by ants; relying exclusively on local information and the emergent behavior of large numbers of agents. The paper consists of four further

sections. Section 2 provides a brief description of the motivating ideas. Section 3 provides a description of the agent density control problem and how it can be solved using ant-like communication. Section 4 addresses the problem of upgrading software agents in a network and provides algorithms for the upgrade problem. The paper concludes with a section that summarizes the key contributions of the paper.

2. Motivating Ideas

It is difficult to argue against the effectiveness of many naturally occurring multi-agent systems and, in particular, systems exhibiting mobility. Societies of simple agents are capable of complex problem solving while possessing limited individual abilities. Many algorithms inspired by the social behaviour of insects have recently been documented [8].

Problem solving by societies of simple agents has a number of common characteristics. Inter-agent communication is local; no single agent has a global view of the world. Communication is also achieved using simple signals and these signals dissipate with time. Signal levels and gradients provide the driving force for migration patterns. Individual agents sense and contribute signal energy to the environment. In this description of the problem solving process, there are two distinct and important agent characteristics. First, there is the role of the agent within the problem solving process; i.e., how the work of problem solving is distributed to a *diverse* set of agents. Second, the degree to which the actions of one agent reinforce the actions of other agents in the society of problem solvers is significant.

In this paper, ant-inspired agents solve problems by moving over the nodes and links in a network and interacting with "chemical messages" deposited in that network. Chemical messages have two attributes, a *label* and a *concentration*. Chemical messages are used for communication rather than raw operational measurements from the network in order to provide a clean separation of measurement from reasoning. In addition, chemical messages drive the migration patterns of agents, the messages being intended to lead agents to areas of the network that may require attention. Chemical labels are digitally encoded, having an associated string pattern that uses the alphabet $\{1, 0, \#\}$. This encoding has been inspired by those used in Genetic Algorithms. The hash symbol in the alphabet allows for matching of both one and zero and is, therefore, the "don't care" symbol.

3. Density Control

The problem of resource control for mobile agents can be attributed to Tschudin [3] and has been studied

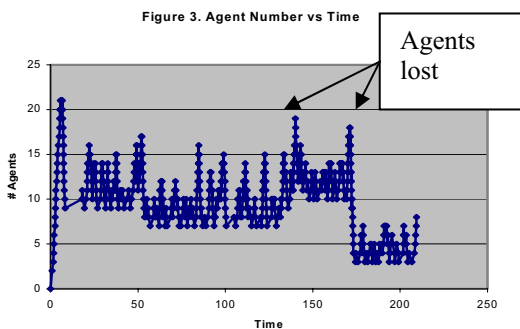
by several researchers [4], [5] and [7]. Shehory et al [5], for example, uses agent cloning to ensure appropriate agent densities for problem solving while acknowledging the relevance of the load balancing literature, while Bredin et al, [7] uses Market Based Control for resource allocation problem resolution. Clearly, mobile agent researchers have long recognized the importance of having an appropriate number of agents in a network performing a given task. This is the problem of density control. For example, in the routing problem described in [2], it was noted that if too few routing agents were sent out into the network, routes would not necessarily emerge. This is in complete agreement with Dorigo's work on AntNet [8].

The above scenario is a less interesting example of density control when compared to the general situation with agents moving through the network and never terminating their problem solving activity; for example fault detection [1]. It is this type of problem that motivates the research reported here.

Agent loss will inevitably occur in an unreliable network. In order to control agent density we propose the addition of a Density Control Agent (DCA) class. The purpose of the density control agent class is to circulate continuously in the network depositing chemical signals in that network such that agent classes whose density is being controlled will automatically adjust their numbers to fall within the target density range. The DCA class uses a random migration decision function in order to explore all parts of the network equally and is responsible for controlling its own density. It also remains at each node for a randomly generated period chosen from a uniform distribution in order to avoid correlations between agent actions. This is an important observation as, without it, significantly greater oscillations are observed with possible population extinction. Every agent class that is density controlled generates a visit chemical that is sensed by the DCA. The DCA controls its own density in order to solve the problem of managing management class agents. Therefore, the DCA also generates a visit chemical. Visit chemical concentrations are associated with the node. The visit chemical leaves a trail of activity for the density-controlled problem solving agents that is integrated across a number of network nodes by DCA agents for the purpose of generating birth or death signals that are in turn sensed by the density controlled problem solving agents. Birth or death signals are chemical in nature and these chemicals do not evaporate.

DCA agents generate birth signals when the aggregated visit chemical concentration for a particular density control problem solving agent class falls below a threshold value. Visit chemical

system is disturbed with agent injection. The oscillation can be further damped by choice of exponential averaging constant and by adjustment of the minimum to maximum visit chemical concentration threshold ratio. These experiments will be reported shortly. In the experiments charted in Figure 2 and Figure 3, a ratio of 5 was used to control the agent density. Figure 2 demonstrates the utility of having density control come online once the network has been colonized by problem solving agents. In these experiments, density control was disabled during the first 20 time units of the simulation. Interestingly, Figure 2 shows the system moving from one stable state to another at approximately 72 time units. Continuing the simulation beyond 200 time units saw no further changes in state. Contrasting the dynamics of this system with those displayed in Figure 1 clearly shows a much smaller transient and more rapid stabilization to the steady state network behaviour. The number of agents injected initially seems to make minor differences to the final stable network state. Experiments were run wherein up to 30 agents were injected initially; the system still converged to a mean number of agents of 6.



Experiments were conducted where agents were occasionally injected into the network in order to test the stability of the agent density management algorithm. As Figure 2 shows, injecting agents after the network (at 50 time units) has settled down merely causes the agent density to find a new stable point, which may, of course, be the same as the original point. This is to be expected in that many systems have several basins of stability. This characteristic is an attractive feature of the system in that we can alter the stable system trajectory by injection (or removal) of agents. In fact, as suggested earlier, we would propose the periodic injection of a density control agent in order to ensure that the system never remains locked at zero population for that agent. Agent injection also occurs at 90 time units. Figure 3 also shows the destruction of agents in the network after the settling period. This scenario represents the situation that initially inspired the density management

algorithm, namely the loss of agents as a consequence of network component failure. Figure 3 shows two failures, at approximately 140 and 175 time units, where multiple agents are lost. Clearly, the algorithm has performed well, with the natural trajectory of the system being quickly restored. Obviously a failure of *all* components in the network would cause the loss of all agents; however, the scenario demonstrated in Figure 3 actually represents a failure of 25% of the network that, in all likelihood, represents an extreme case. More general experiments, with random single node failures, provided equivalent support for the robustness of the density control algorithm and results are not included here as, it was felt, the scenario described in the previous paragraph provides a more dramatic illustration of the robustness of the algorithm.

4. Agent Upgrading

Software rarely has completely correct behaviour when first deployed. The problem of upgrading software in an operational environment is challenging and is currently the focus of considerable research [9], [10]. Hofmeister [9] describes three forms of dynamic reconfiguration: module replacement, structural change and geometric replacement. The software upgrade problem presents a unique challenge when the software is an agent and that agent is mobile, as we have no knowledge a priori of the location of any agent. The software upgrade problem is further complicated by no knowledge of the number of agents to be upgraded and their source of injection into the network. This latter piece of information is important as it implies that older, incorrect versions of a software agent may be injected into the network once the upgrade process has been supposedly completed. Together, these problems present a significant research problem, making geometric replacement the most attractive mechanism for upgrading agents. Again appealing to ant-like problem solving agents, and their tendency to be robust with respect to the failure of an individual agent, we view the agent upgrade problem as one of “failing” the faulty agent and injecting one with corrected behaviour.

Not knowing the source of the agent originally injected into the network leads to the disturbing and inevitable conclusion that the upgrade process is potentially never complete as an obsolete agent may be introduced from a source that has yet to receive the corrected agent.

Referring now to the previous section on density control, a potential mechanism for removal of the faulty agent is to exploit its response to the death signal and visit chemical concentrations. That is, increase the levels of the visit chemical for the faulty

agent such that an associated DCA agent generates the death signal for that agent throughout the network. This is achieved by having different visit chemicals for the faulty and corrected agents with the encodings so chosen that the faulty agent senses the visit chemical of the correct agent. Hence, fault agents will tend to see higher concentrations of visit chemical when compared to the corrected agent; the corrected agent will tend only to see its own. An example best illustrates this.

Consider two agent versions, v_f and v_c , representing faulty and corrected versions respectively having visit chemical encodings #####1 and ###11 respectively. Assuming a string length of 5, if we were to have concentrations c_f and c_c of the two chemicals, the faulty agent would sense $c_f + c_c$, whereas the corrected agent would see c_c . In other words, the faulty agent would see higher concentrations of visit chemicals and would be more likely to see the death signal as a result of exceeding the upper bound on visit chemical concentration. Similarly, the faulty agent would be less likely to see the birth signal as a result of higher visit chemical concentration.

This example can easily be extended. Consider a perfect agent that corrects faults in agent version v_c , say v_p . Let v_p have the visit chemical encoding ##111 and concentration c_p . Again assuming a string of length 5, the faulty agent would sense $c_f + c_c + c_p$, whereas the corrected agent would see $c_c + c_p$. In this case, both v_f and v_c would experience a reduced number of birth signals and elevated number of death signals. As the number of v_p agents increases, the tendency is for the number of v_f and v_c agents to decrease, eventually causing the imperfect agent types to disappear.


This mechanism works because of the encoding scheme used for the three agent types. It relies upon the masking of the existence of previous versions of the agent through the increasing number of bits being specified as we move to higher and higher versions of the agent. Obviously, this limits the number of versions that might be accommodated. However, a potentially infinite data structure such as can be provided by a tree removes this limitation. Obviously this is a more expensive solution from a computational viewpoint -- pattern matching by bit position versus matching by subtree -- but does solve the more general problem when an unbounded number of agent versions needs to be supported.

4.1 Agent Upgrading Results

This section presents experimental results that demonstrate the applicability of the agent upgrade algorithm. In the results presented below, the two networks used in the density control experiments were

also used as a simulation test bed for the agent upgrade algorithm. We used three agent versions in our experiments, called “flawed”, “corrected” and “perfected” with encodings #####1, ###11 and ##111 respectively.

Figure 4. Agent Number vs Time




In Figure 4, three flawed agents are injected into the network and are subject to the density control algorithm. Eventually the number of agents stabilizes around 7. At 78 time units, a single corrected agent is injected into the network and quickly establishes dominance over the flawed agents using the visit chemical masking mechanism introduced in the previous section. It is not possible for the population of flawed agents to recover once the number of flawed agents reaches zero, as the birth process is one of cloning. Only through injection of a new flawed agent into the network from some external source can the flawed population temporarily recover.

In Figure 4, the flawed agent type is re-injected into the network at 100 and 140 time units following the colonization of the network by the corrected agent. The flawed agent population is quickly removed from the network on both occasions despite large numbers of them being injected. While results are not included here, the number of flawed agents re-injected did not affect the final state of the network, merely the time to achieve it. In all cases, the flawed agents were eventually removed from the network, leaving only the corrected agent population. Clearly, this demonstrates that the corrected agent resists recolonization of the network by the flawed agent. The rejection of the flawed agents by the network is faster upon re-injection as a result of the established population of corrected agents.

In order to demonstrate the robustness of the algorithm in the presence of multiple agent versions, experiments were conducted wherein a third version -- the “perfected agent” -- was injected into the network after the corrected agent population had replaced the initial population of flawed agents. As can be seen in Figure 5, the perfected agent quickly established a dominant position in the network causing the corrected

Figure 5. Agent Number vs Time



agent population to vanish completely. Once again, subsequent re-injection of flawed or corrected agent populations did not affect the dominant position of the perfected agent. This is clearly shown in Figure 5 where corrected and flawed agents are re-injected into the network following the perfected agent establishing itself in the network. Reviewing the interval 25 to 45 time units in Figure 5 also shows that the algorithm can deal with multiple agent populations simultaneously trying to establish themselves. During this time interval, flawed, corrected and perfected agent types are moving throughout the network; however, the corrected agent population is quickly extinguished even before stabilization can occur. Despite the version control algorithm's simplicity, it seems to be remarkably effective in maintaining the correct dominant version in the network.

5. Conclusions

In a network managed completely by swarms of mobile agents, two important observations need to be made; namely, the number and positions of agents are unknown. These characteristics make management of the swarm populations extremely challenging. This paper has addressed two questions related to the management of mobile agent swarms.

The first question relates to the maintenance of population density and we have presented algorithms that maintain population density in a network having unreliable components. The algorithms presented rely on local knowledge only and we have shown by experiment that populations quickly settle to a mean population around which they oscillate. The algorithms presented are robust with respect to the introduction of agents after the network has stabilized as well as loss due to component failure.

The second question deals with the upgrading of agents over time. Agents, being software entities, rarely have correct behaviour when first introduced into service and often require upgrades. This paper presents an algorithm that solves the upgrade problem

by taking advantage of the density control algorithm in a form of parasitic behaviour where a corrected agent "fools" the flawed agent population into believing that there are more of them than there are, the flawed agent population quickly decaying to zero.

6. References

- [1] Bieszczad A., White, T., Pagurek, B., Mobile Agents for Network Management. In IEEE Communications Surveys, September 1998.
- [2] White T., Pagurek B. and Oppacher F., Connection Management using Adaptive Mobile Agents, Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '98), July 1998.
- [3] C. Tschudin. Open Resource Allocation for Mobile Code. Mobile Agents - First International Workshop, MA '97 (Berlin, Germany, April 7-8, 1997). Published as Kurth Roethermel and Radu Popescu-Zeletin, editors, Lecture Notes in Computer Science, 1219, Springer, 1997.
- [4] Jonathan Bredin, Rajiv T Maheswaran, Cagri Imer, Tamer Basar, David Kotz, and Daniela Rus; A Game-Theoretic Formulation of Multi-Agent Resource Allocation, Proceedings of the 2000 International Conference on Autonomous Agents, Barcelona, Spain, June 2000.
- [5] O. Shehory, K. Sycara P. Chalasani and S. Jha Agent cloning: an approach to agent mobility and resource allocation, IEEE Communications, pages 58-67, vol. 36 no. 7, July 1998.
- [6] Simoes P., Moura e Silva L., and Boavida F., Mobile Agent Infrastructures: a Solution for Management or a problem to Manage, in Proceedings of the 3rd IEEE Conference on Telecommunications, 23-24 April 2001, Figueira da Foz, Portugal.
- [7] Jonathan Bredin, David Kotz, and Daniela Rus. Market-based resource control for mobile agents. In Proceedings of Autonomous Agents '98, pages 197-204, 1998.
- [8] Bonabeau E. Dorigo M. and Theraulaz G, Swarm Intelligence, Oxford Press, 1999
- [9] Hofmeister C. Dynamic Reconfiguration of Distributed Applications, Ph.D. thesis, Computer Science Department, University of Maryland, 1993.
- [10] Noel de Palma, Dynamic reconfiguration of agent-based applications, Technical Report Project SIRAC, INRIA, 1999.