An Artificial Agent Society is more than a Collection of "Social" Agents

Christian Gerber
German Research Center for Artificial Intelligence (DFKI GmbH)
Stuhlsatzenhausweg 3
66123 Saarbrücken, Germany
gerber@dfki.de

Abstract

This work focuses on scalability of multi-agent-systems (MAS) and on the development of methods to allow systems to configure themselves to any application scale and nature. I address this issue by referring to the assumption that social behavior of artificial agents cannot be achieved by simply designing a *social agent architecture* alone; social behavior needs *social control mechanisms* for agent groups or societies. Here, I shall outline such a mechanism, give an application example and put this work into relation to other fields of research.

Introduction

Whenever large scale societies of agents (artificial as well as natural) are brought together to achieve a certain goal, new aspects come into play. In large groups of humans (e.g., in companies) or animals (in particular insects) often *social dynamic* or *emerging functionality* effects occur. Such phenomena cannot be explained by observing group members independently.

In the world of artificial agents, increasing the number of members of a society may lead to serious problems: a system running efficiently in a small environment may fail in a large one if the system's algorithm is exponential in terms of environment size.

So, focus for programming in the large has to be put on two aspects: First, how to suppress negative properties of large scale artificial agent groups; second, how to achieve positive phenomena occuring in large natural societies, such as social dynamics. These aspects reflect the central problem: *How can a multi-agent society be organized to make it flexible enough to cope with applications of any scale?* Organizing such a society includes specifying the society structure, communication forms and agent architecture aspects. Furthermore, not only the sheer size of an application has to be taken into account, but also other properties characterizing its nature. Hence, the goal of this work is to provide a mechanism which adjusts a MAS to *any environment.*

An Adaptation Mechanism

I regard the task to organize a society of artificial agents as an optimization problem by characterizing a *search space* and an *objective function* to be optimized. The objective function has to denote the system's performance while a multi-dimensional search space must describe the system's set of possible configurations.

The *objective function* has to be defined for each application from scratch since it depends on several factors the application designer has to combine, such as operating time, quality of the result, etc.

Each modifiable property of the system reflects one dimension of the *search space.* The search space dimensions can be derived from scalable parameters of a multi-agent application on three levels: on the *agent society structure* level, parameters are for instance number of agents and organizational form of the society. On the *agent architecture* level, e.g., explicit resource distribution to the various modules are regarded as scalable parameters (see (Gerber & Jung 1997) for a unified approach based on resource distribution management). Finally, on the level of *agent communication/cooperation*, issues such as complexity of the communication process can be regarded as scalable parameters (for details see (Fischer *et al.* 1997)).

Figure 1 shows a simple, two-dimensional example: here performance depends only on the number of agents and the usage of a sophisticated knowledge representation (KR) component. The number of agents is represented by a discrete dimension, whose domain ranges from 0 to possibly infinity. KR usage is modeled as a continuous dimension, its domain ranging from 0 to 1. A number between those extremes indicates what percentage of an agent's computational time can be used in the KR component.

In many cases, multi-agent applications are intended to work ad infinitum. A system which was originally adjusted to work at a very high performance level, has to react dynamically to new inputs so that it might lose its performance over time as the environment changes.

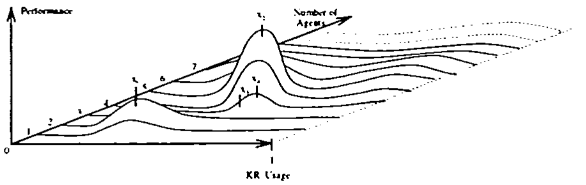

Figure 1: Example of a System Performance Relation

Maintenance of high performance is therefore another main issue. By applying optimality conditions to a current situation the system shall be able to detect sub-optimalities. In the following, I present a mechanism for achieving and maintaining high performance during the complete run of an application. As global optimality can hardly be achieved in a reasonable amount of time, but may be lost very easily, this mechanism bases on the steepest ascent method (Zoutendjk 1976) for finding *local* optima. The mechanism is enabled to include an expert's intuitive understanding on how to model an organization.

Method

1. The application designer produces an initial model of the application by instantiating scalable parameters at the *agent society structure, communication/cooperation* and on the *agent architecture* levels.

2. The system detects sub-optimalities in this layout by applying optimality conditions. The system modifies the layout in order to remove them. This step may be performed in interaction with a human expert.

3. The second step is repeated until the application is terminated.

In the above example[1] a starting point may be x_3. A maximal performance gain will be achieved by adding more resources to the KR component: x_4 may be achieved. By increasing the number of agents to five, the optimal configuration x_2 will be found.

For the realization of this algorithm two issues must be addressed: *How to retrieve optimality conditions mentioned in step 2* and *How to monitor the structure.*

(Local) optimality conditions may be retrieved through a theoretical examination of scalability quantities (e.g., through bottleneck analysis (Sie 1996)) prior to the run of the system, but also through on-line performance measuring based on a trial-and-error method. In addition, examination of related fields of research may lead to further insights: organizational forms defined in organization theory can be used to model substructures of the layout; heuristics about the perfect size of a group may be taken from psychology, etc.

Secondly, the structure has to be controlled or *monitored* in order to detect and remove sub-optimalities. It may not be reasonable to introduce only one such optimization procedure as too many scaling dimensions may occur. It therefore might be wiser to introduce one optimization process for scaling dimensions which characterize general system properties, several processes for dimensions characterizing smaller groups of agents and a process for each agent, addressing dimensions on the agent architecture level.

The latter control structure can be integrated into the agent architecture. One approach to construct the other control mechanisms is to commission certain agents to control optimality conditions. Such *monitor agents* can be society members or additional agents whose sole functionality is to monitor. They monitor optimality conditions by applying them in a feedback loop, or by controlling them in a daemon-like fashion.

An Application Example

The MAS simulation environment SIF (*Social Interaction Framework* (Funk *et al.* 1997)) is currently being developed for the study of social interaction between artificial agents. Here, the environment is populated by agents that have some sort of "life energy", ranging from 0 (the agent is dead) up to 100 (the agent is perfectly healthy). Any activity an agent performs results in energy loss. The agents' main goal is to keep their life energy as high as possible. Increasing energy is achieved by consuming food which has be produced by processing various types of raw material. Cooperation between agents is reasonable because the process of producing food can hardly be achieved solitary.

Overall goal of the system is to derive an agent society as strong as possible, i.e., that consists of as many members as possible. A centralized control of agent activity may easily become intractable once the society has reached a certain size. Therefore, I incorporate a more decentralized and scalable approach: agents are enabled to found, join, or leave groups. In contrast to other group formation approaches (e.g., (Ketchpel 1993)), once evolved, agent groups are *explicitly* represented by a monitor agent equipped with the functionality described in the previous section. Agents groups again can cluster to larger units which again are represented explicitly.

[1]For the sake of clarity, I assume for this example that during the optimization process neither the search space changes, nor does the dependence of the performance function to scalable or non-scalable quantities. Of course, this assumption does not hold in general. However, the above search procedure can still be applied for the general case.

The main goal of a group (i.e., the *objective function* of the corresponding monitor agent) is to gain as much control over food as possible in order to enable the survival of group members. Agent individuals have unique skills to perform certain jobs. Hence, the monitor agent representing a group must determine how many agents, what type of agents, what patterns of command in the group, etc., are needed to optimize its objective function. All these parameters reflect *search space dimensions* in the optimization problem the monitor agent has to solve.

Agents, on the other hand, have to optimize their life energy. They have to reason whether or not to join or leave a certain group. So, not only group formation evolves naturally from goal optimization, but also agent characteristics such as *selfishness, social behavior* or *solitary behavior*, concepts well studied but mostly represented in an explicite manner.

Related Fields

This work addresses problems somewhere in the intersection of a wide range of research areas, mainly in the area of cognitive science, allowing to incorporate research results found in other fields. Furthermore, there is also the chance that this work will be useful for these disciplines, for instance, for simulating behavior of large human groups. Important points of contact to other fields are described below:

Psychology Group formation and development is one issue that social psychologists work on. Much research is carried out to find the optimal size of a certain group (see e.g., (Nasser 1988)). Furthermore, the strongly relevant question why people form or join certain groups is investigated. ((Tuckman & Jenson 1977) shows some aspects.)

Biology The examination of the sociology of animal societies, in particular of phenomena such as emerging functionality in insect societies, is important to how a society goal is split into subgoals, and then, how they are achieved. Furthermore, feedback loops regulating the size and structure of insect societies (see (Free 1987) as an example) are worth looking at.

Management Theory In business organization, consultants have the tough job of deciding how to structure very large companies. Thus, business administration has developed basic organizational forms for companies and rules designed for building company structures. (A survey can be found in (Wöhe 1981).)

Summary

In this paper I have presented an adaptation mechanism which enables multi-agent systems to configure themselves to any application scale and nature. This goal is motivated by the necessity to guarantee high performance of MAS of any scale and thus, to achieve scalability. Furthermore, I have sketched an application example to demonstrate the feasibility and practical relevance of the approach. Here, I have shown how social phenomena such as *the degree of social behavior* or *group formation* can evolve. Finally, related fields of research were addressed.

Acknowledgments

Part of this work has been supported by Siemens AG Germany.

References

Fischer, K.; Müller, J.; Pischel, M.; Gerber, C.; and Chaib-draa, B. 1997. A Simulation Approach based on Negotiation and Cooperation between Agents: A Case Study. *Submitted to: ACM Transaction on Modeling and Computer Simulation: Special Issue on Simulation of Scalable Systems.*

Free, J. 1987. *Pheromones of Social Bees.* London: Chapman and Hall.

Funk, P.; Gerber, C.; Lind, J.; and Schillo, M. 1997. Social Interaction Framework—A generic Testbed for Social Agents. Technical report, German Research Center for Artificial Intelligence. To Appear.

Gerber, C., and Jung, C. 1997. Towards the Bounded Optimal Agent Society. Working Notes of the Workshop on Distributed Cognitive Systems in conjunction with the 21st Annual German Conference on Artificial Intelligence (KI-97), To Appear.

Ketchpel, S. 1993. Coalition Forming Among Autonomous Agents. In *Proceedings of the 5th European Workshop on Modelling Autonomous Agents and Multi Agent Worlds (MAAMAW'93).*

Nasser, D. 1988. How to run a focus group. *Public Relations Journal* 44:33–34.

Siemens AG, Zentrabteilung SW2. 1996. *Enpaßanalyse -eine neue Methode zur Systembewertung.*

Tuckman, B., and Jenson, M. 1977. Stages of small-group development revisited. *Group and Organization Studies* 2:419–427.

Wöhe, G. 1981. *Einführung in die Allgemeine Betriebswirtschaftslehre.* Verlag Wahlen, 14th edition.

Zoutendjk, G. 1976. *Mathematical Programming Methods.* Amsterdam: North-Holland.