Benevolent Agents in Multiagent Systems

Abdulla M. Mohamed and Michael N. Huhns
University of South Carolina
Columbia, SC 29208 USA
{mohamed,huhns}@engr.sc.edu

Abstract

This short paper describes an analysis of benevolent agents in multiagent systems. We present a definition and motivation for benevolence that is appropriate for MAS. We construct a simulator, called Mattress In the Road (MIR), and use it to analyze benevolence. Finally, we suggest applications that are driven by benevolent agents.

1. Introduction and Background

Agents in a society might act alone or as part of a group. An agent's contribution to its group is controlled by its behavioral characteristics, such as cooperation and benevolence. This latter concept has been lightly studied in AI, but usually in terms of the mathematical utility for the agent acting benevolently, ignoring the origin of benevolence, whose long history explores virtue and moral duty [3], [8]. Herein we argue that benevolence should also have a classical basis that recognizes the moral goodness of an agent and includes social awareness.

Definitions of benevolence for agents are split into two different strands. Some researchers, such as Castelfranchi, Conte, Jennings, and Wooldridge, define benevolent agents as those that accept all other agents' requests for help [5]. Castelfranchi emphasizes the fact that a benevolent agent must adopt other agents' goals and interests without being asked, and even without the recipients' expectations [2]. Other researchers define benevolent agents in terms of the similarities of their goals, termed the benevolent agent assumption [7].

2. A Complete Definition of Benevolence

Based on the above, we define an agent as benevolent if

1. The agent voluntarily helps other agents without being commanded to do so.
2. Benevolent actions are intended to benefit the society to which the agent belongs.
3. The agent should not expect an immediate reward or benefit for its benevolent actions.
4. The benevolent action is taken while the agent is pursuing its own goals in such a way that it should neither prevent nor help the agent accomplish its goal.

According to our definition, benevolent actions should benefit the agent's society and will not stop it from reaching its goals. This is an indirect benefit, in that if the society is doing well, then all its members, including the benevolent agent, will be doing well too. Another motivation is the belief that the agent's benevolent actions may encourage others to act benevolently in the future, thereby providing compensation in the longer term. This relates to Blackmore's work on memes, where she states that altruism spreads altruism (meme-fountain) [1].

3. Mattress In the Road Simulation

The Mattress In the Road (MIR) simulator is a tool we have constructed to simulate agents' benevolent behavior. MIR consists of agents traveling on a road and encountering mattresses. Agents can be benevolent or non-benevolent. Benevolent agents will remove the mattress from the road at some delay cost, while non-benevolent agents will avoid the mattress at less cost.

MIR computes the average time t_{avg} for cars to complete the road to determine:

$$t_{avg} = f(d, m, b)$$

where d is the traffic density, m is the mattress probability, and b is the benevolence probability. t_{avg} is plotted versus d and m in Figures 1 and 2 for systems of just benevolent or just non-benevolent agents.

Figure 1 shows that the benevolent agents' performance is better than that of the non-benevolent agents, regardless

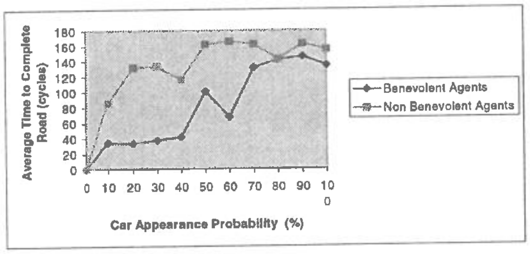

Figure 1. t_{avg} **vs. traffic density**

of traffic density. Above 80% traffic density, the difference is not significant.

In Figure 2 we can see that for mattress probabilities up to 40%, the benevolent agents' performance is better than that of the non-benevolent agents. After 40%, there is no significant difference.

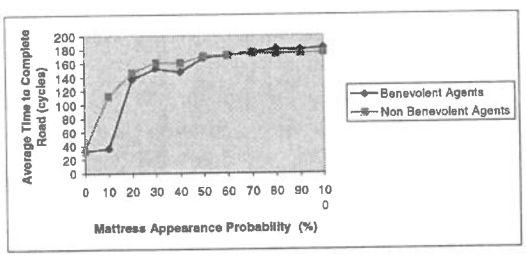

Figure 2. t_{avg} **vs. mattress probability**

Figure 3 shows that once the percentage of benevolent agents exceeds 20%, the MAS's performance increases dramatically. However, beyond 50%, the performance does not change significantly. Benevolent agents help the MAS to perform better, but we do not need every agent to be benevolent—only about half.

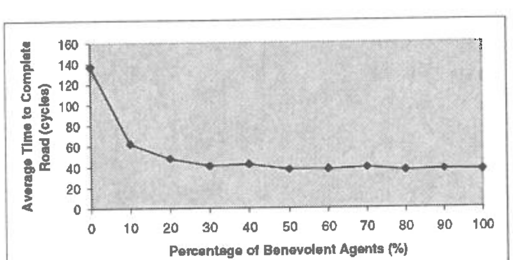

Figure 3. t_{avg} **vs. % benevolence**

4. Future Applications of Benevolent Agents

We predict that benevolence will play a major role in Internet computing. Many agents will soon be populating the Web, and if they behave benevolently by helping each other search for information, for example, it will greatly reduce Internet traffic. When asked, a benevolent query agent would freely share its results with other agents, even though it may have consumed substantial resources to get this knowledge and might have to consume more to share it. Through one agent's benevolence, other agents charged with similar queries would not have to explore all the sites or databases the first explored: they can simply use its results, thus reducing Internet traffic for all [4].

A second application is a collective store [6] consisting of query results and information. Benevolent query agents contribute search results to this store, and then when heavy traffic degrades their performance, the agents could find what they need in the store. Agents making greater contributions to a collective store might be given higher priorities in its subsequent use.

References

[1] Susan Blackmore, *The Meme Machine*, Oxford University Press Inc., New York, 1999.

[2] Rosaria Conte and Cristiano Castelfranchi, *Cognitive and Social Action*, UCL Press, London, UK, 1995.

[3] Charles Darwin, *The Descent of Man and Selection in Relation to Sex*, John Murray, London, 1871.

[4] Michael N. Huhns and Abdulla M. Mohamed, "Benevolent Agents," *IEEE Internet Computing*, Vol. 3, No. 2, March–April 1999, pp. 96-98.

[5] N. Jennings and M. Wooldridge, "Agent Theories, Architectures, and Languages: A Survey," *Intelligent Agents: ECAI-94 Workshop on Agent Theories, Architectures, and Languages*, Nicholas Jennings and Michael Wooldridge (eds.), Springer-Verlag, Berlin, 1995.

[6] Domenico Parisi, "What to Do with a Surplus," *Simulating Social Phenomena*, Conte, R., Hegselmann, R., and Terno, P. (eds.), Springer-Verlag, Berlin, 1997.

[7] Jeffrey S. Rosenschein, *Rational Interaction: Cooperation Among Intelligent Agents*, Ph.D. Dissertation, Stanford University, Stanford, CA, 1985.

[8] R. L. Trivers, "The Evolution of Reciprocal Altruism," *Quarterly review of Biology*, Vol.46, 1971, pp. 35–56.

Proceedings

Fourth International Conference on MultiAgent Systems

Boston, Massachusettes, USA
July 10 – 12, 2000

Sponsored by
IFMAS, Charles River Analytics, Nokia, DARPA, NSF

In cooperation with
AAAI

IEEE COMPUTER SOCIETY

Los Alamitos, California

Washington • Brussels • Tokyo

IEEE Computer Society Order Number PR00625
ISBN 0-7695-0625-9
ISBN 0-7695-0626-7 (case)
ISBN 0-7695-0627-5 (microfiche)
Library of Congress Number 00-101165

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1-714-821-8380
Fax: + 1-714-821-4641
E-mail: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1-732-981-0060
Fax: + 1-732-981-9667
http://shop.ieee.org/store/
customer-service@ieee.org

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Editorial production by Bob Werner
Cover art production by Alex Torres
Printed in the United States of America by Technical Communication Services

IEEE
COMPUTER
SOCIETY