

Modelling a Society of Simple Agents: from Conceptual Specification to Experimentation¹

F.M.T. Brazier, P.A.T. van Eck, J. Treur

Faculty of Mathematics and Computer Science, Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
email: {frances, patveck, treur}@cs.vu.nl URL: <http://www.cs.vu.nl/>

Abstract. In this paper, the compositional multi-agent modelling framework DESIRE is not only successfully used to develop a conceptual specification of the simple agents discussed in (Cesta et al. 1996), but also to simulate the behaviour in a dynamic environment. In the DESIRE framework, a conceptual specification, which provides a high-level view of an agent, has enough detail for automatic prototype generation. The prototype implementation of the conceptual specification of the simple agents has been used to replicate and extend one of the experiments reported in (Cesta et al. 1996).

1 Introduction

Although much research within the multi-agent community has focused on the design of individual agents and their interaction, other research has addressed emergent behaviour within societies of agents (see for instance the three papers in the chapter on emergence in (Van de Velde & Perram 1996)). The behaviour of an individual agent can often be conceptually specified as can the interaction between individual agents. The result of interaction between larger numbers of agents in a dynamic environment is often not easy to predict. Experimental research, in which interaction between agents is studied in a simulated dynamic environment, provides a means to actually test and compare results of conceptual specification (Hanks et al. 1993).

In this paper, the compositional multi-agent modelling framework DESIRE (Brazier et al. 1997a, Brazier et al. 1995) is used to conceptually specify individual agents and to examine the behaviour of individual agents within a large group of agents. This framework includes tools with which detailed specifications (with which the behaviour of individual agents and their interactions is defined) can be automatically translated into prototype implementations. To examine such behaviour, experiments reported by (Cesta et al. 1996) with which social theories are tested by simulating interaction between different types of simple agents (i.e., agents with limited knowledge and capabilities), have been repeated and extended.

¹ This paper is a short version of (Brazier et al. 1997d).

Due to the nature of the environment in which the experimentation of (Cesta et al. 1996) was originally performed (using the MICE testbed (Montgomery & Durfee 1990)), most information about agent characteristics and behaviour is implicitly defined by the implementation and simulation environment. (Agent functionality is specified by LISP functions in MICE.) On the basis of the informal, textual descriptions provided by (Cesta et al. 1996), within the DESIRE framework a generic model of a simple agent is defined and refined for each of the four types of agents (Cesta et al. 1996) distinguished: social, solitary, selfish and parasite. One of the aims of this paper is to show how this approach leads to a flexible, conceptual-level specification, from which prototypes can be generated automatically for experimentation.

The second aim of this paper is to investigate the flexibility of the resulting specification with respect to adaptability to different agent behaviour. One of the advantages of a conceptual description of an agent and its behaviour is that it can be easily adapted at a conceptual level (without having to rewrite low-level code for each agent). In this paper, not only are these experiments repeated with agents automatically implemented from the conceptual DESIRE specifications, additional experimentation was performed to examine the influence of an increase in the number of directions (8 instead of 4) in which agents are allowed to move.

2 The MA modelling framework DESIRE

The multi-agent compositional modelling framework DESIRE provides support for the design of a conceptual model of the agents described in (Cesta et al. 1996) and the simulation of agent behaviour. Compositional agent models define the structure of the architectures: components in a compositional model are directly related to agent tasks. Existing generic agent models can be used to design a specific agent model. During analysis and design, relevant components in a generic model are refined by (1) more detailed analysis of the tasks of which such components are comprised and/or (2) inclusion of specific domain knowledge. Within the DESIRE framework, five types of knowledge are represented at a conceptual level, detailed level and at an operational level:

- The compositional structure of agents and their tasks. Tasks can be composed or primitive and are characterised by their input and output knowledge structures, which can be modelled at an object-, a meta-, a meta-meta-level, etc.
- Interaction within and between agents and tasks;
- Temporal relations between tasks, represented rules in a temporal logic;
- Delegation of tasks to agents;
- Knowledge structures, represented by knowledge rules in a three-valued, order-sorted predicate logic.

The representation at the operational level is automatically generated from the representation at the detailed level.

In DESIRE, models and specifications can be developed (Brazier et al. 1997a, Brazier et al. 1997b) in which all agents have specific knowledge of other agents and of their needs with respect to information exchange with these other agents.

The desired behaviour of individual agents and their interaction capabilities is the basis for the design of the system. For agents with a complex structure, such as in the model presented in (Brazier et al. 1997c), which can show complex behaviour and interactions themselves, it is complicated to perform experiments in a systematic manner. For simple agents in a dynamic environment with only a restricted repertoire of possible behaviours and interactions, systematic experimentation is feasible: the number of parameters involved is relatively small.

In models of agents with a more complex structure, such as the model presented in (Brazier et al. 1997c), eight agent tasks are performed by the eight components distinguished: control of an agent's own processes (Own Process Control), interaction with other agents (Agent Interaction Management), maintaining knowledge of other agents' characteristics (Maintain Agent Information), interaction with the external world (World Interaction Management), maintaining knowledge of the external world (Maintain World Information), maintaining information regarding past observations and interactions (Maintain History), managing cooperativeness (Cooperation Management), and performance of agent specific tasks (Agent Specific Tasks).

3 The original experiment

(Cesta et al. 1996) examined the behaviour of different types of agents in interaction. Four types of agents are distinguished on the basis of their social characteristics: social agents, parasite agents, solitary agents and selfish agents. The effect of an agent's social characteristic on interaction with other agents is measured by simulating agent behaviour in a situation in which 30 agents try to survive on a 15 * 15 grid in which 60 pieces of food are continually available in random positions. An agent's welfare is measured on the basis of its energy level. The end result of a simulation is the number of agents that survive in a given society of agents, given the energetic value of the food available. Agents do not communicate explicitly but implicitly: a hungry agent changes colour, and this can be seen by other agents. An agents' social characteristics are assumed to be static. An agent does not change from being, for example, selfish to social. The implications of agents' social characteristics for its behaviour is as follows. A solitary will always search for food, regardless of its internal energy level. Likewise, a parasite agent will always look for help. A selfish agent will look for help only if it is in danger, otherwise it searches for food. A social agent will also look for help if it is in danger. If it is in a hungry state, it will search for food. If it is in a normal state, then it will search for food if no help-seeking agents are seen. Otherwise, the social agent will give food to one of the help-seeking agents nearby.

4 Conceptual model of simple agents

The generic agent model presented in Section 2 includes more functionality than required for the small agents described in (Cesta et al. 1996). Those small agents are not capable of communication and reasoning about other agents' knowledge,




Fig. 1: Generic structure of a small agent.

nor are they capable of reasoning about communication. Their only task is to stay alive in a dynamic environment. In fact, the only components within the generic agent model, applicable to these small agents, are the component Own Process Control and the component World Interaction Management. Figure 1 depicts not only the remaining composition of a small agent's tasks at the highest level of abstraction, it also shows the information links between the components.

The only information a small agent receives is the information it observes in the external world. This information is forwarded directly to the component World Interaction Management, which interprets this information. The result, information about the agent's position, about available food, and, if applicable, information about other needy agents is transferred to the component Own Process Control. The component Own Process Control determines which actions should be taken next, depending on the small agent's social characteristics and the agent's direct environment. This information is transferred to the component World Interaction Management that derives the information required to actually perform the action in the external world. This information is the only output a small agent provides to the external world, which is implemented as a C program. This program maintains a representation of the grid in which the agents live and is responsible for actually performing the agent's actions by changing the state of the grid with respect to agent's positions and appearance. This updated state is then observed by other agents. Other tasks of this program are growing new food at random locations if a piece of food is eaten and maintaining statistics with respect to the number of alive agents. A detailed design of simple agents is presented in (Brazier et al. 1997d).

4.1 The internal structure of component Own Process Control

The component Own Process Control is composed of four components: Own Resource Management, Own Characteristics, Goal Determination and Plan Determination. The component Own Resource Management receives information about its current energy level and the resources it has consumed, with which it determines its new energy level. On the basis of information the component Goal Determination receives about its own social characteristics and its own energy level, it determines the goals the agent is to pursue: for example to search for food, or to look for help. The component Own Characteristics receives information on the agent's

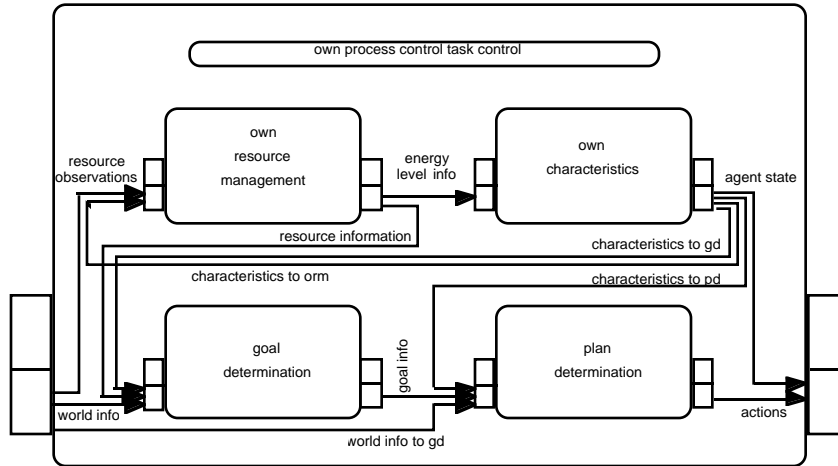


Fig. 2: Component Own Process Control.

energy level from the component Own Resource Management. This information is used to determine the agent's next state (e.g., hungry, normal or in danger). The component Plan Determination receives information (1) from the component Own Characteristics, namely the agent's current state, (2) from the component Goal Determination, namely which goals are to be pursued and (3) from outside the component, namely the current state of the world. With this information the component Plan Determination determines which actions to take in the external world.

4.2 Internal structure of the component World Interaction Management

The component World Interaction Management interprets information it receives from the external world, and it transforms information about actions to be taken in the external world into specifications for actions which the external world can execute. Two components are defined to perform these tasks: Observation Information Interpretation and Execution Preparation.

The component Observation Information Interpretation receives information from the external world, for example which pieces of food and which agents have been observed within a given range. This information, termed sensory information in (Cesta et al. 1996), is translated into information which can be used by the component Own Process Control to reason about new goals and plans.

As stated above, the component Action Execution Preparation receives information about actions to be taken in the external world from the component Own Process Control and translates these actions into specifications to be executed in the external world. These specifications are also the output of effectors in the terminology used in (Cesta et al. 1996): elementary actions to be performed in the external world.




Fig. 3: Component World Interaction Management.

5 Experimentation

The first goal of this exercise is to replicate some of the results presented in (Cesta et al. 1996), on the basis of a conceptual specification of agent behaviour, simulated in a DESIRE environment.

Method One of the simulation discussed in (Cesta et al. 1996) has 15 social agents and 15 parasite agents with varying food energetic values in one world, with 500 steps per agent per run. The same simulation was re-run for the same agents in an environment in which agents could move in more than 4 directions: they could move in 8 directions. For this new experiment, only a small adaptation to the detailed design had to be made. (See (Brazier et al. 1997d) for details.)

Results For graphs presenting the results, see (Brazier et al. 1997d).

Evaluation The results acquired in the DESIRE simulation are comparable to the results in (Cesta et al. 1996): social agents survive more often than parasite agents in situations with low food energetic values. The same holds for the experiment in which agents have more degrees of freedom. In our experiments, the chance of survival increases with an increase in the degree of freedom for both types of agents.

6 Discussion

Much research concerning the design of multi-agent systems (at a conceptual level) addresses complex agents which exhibit complex interaction patterns. Due to this complexity, it is difficult to perform rigorous experimentation. On the other hand, systematic experimental work regarding behaviour of societies of more simple agents, while reporting valuable results, often lacks conceptual specification of the system under consideration. In this paper, the compositional multi-agent modelling framework DESIRE is not only successfully used to develop a conceptual specification of the simple agents discussed in (Cesta et al. 1996), but also to simulate the behaviour in a dynamic environment. In the DESIRE framework, a conceptual specification, which provides a high-level view of an agent,

has enough detail for automatic prototype generation. The prototype implementation of the conceptual specification of the simple agents has been used to replicate and extend one of the experiments reported in (Cesta et al. 1996). One of the advantages of conceptual specification has been explored, namely the ease with which existing specifications can be modified.

In future research, slightly more complex agents will be designed: agents capable of adapting their own characteristics to increase their chances of survival. These agents must possess the capability to learn from the observed effects of their own behaviour and that of others. Explicit conceptual specification makes it possible to make such modifications at a conceptual level.

Acknowledgements The authors thank Amedeo Cesta, Maria Miceli and Paola Rizzo for answering our questions regarding details of their experiments.

References

- Brazier, F.M.T., Dunin-Keplicz, B.M., Jennings, N.R. & Treur, J. (1997a). DESIRE: modelling multi-agent systems in a compositional formal framework. In M. Huhns & M. Singh (eds.) *International Journal of Cooperative Information Systems, special issue on Formal Methods in Cooperative Information Systems: Multiagent Systems*, 6(1), 67-94.
- Brazier, F.M.T., Dunin-Keplicz, B.M., Treur, J. & Verbrugge, L.C. (1997b). Modelling internal dynamic behaviour of BDI agents. In *Proc. 3rd Workshop of the ModelAge Project*. (to appear).
- Brazier, F.M.T., Jonker, C.M. & Treur, J. (1997c). Formalization of a cooperation model based on joint intentions. In J. Müller, M.J. Wooldridge & N.R. Jennings (eds.) *Intelligent Agents III*. LNAI 1193, 141-155, Berlin: Springer-Verlag.
- Brazier, F.M.T., Van Eck, P.A.T. & Treur, J. (1997d). *Modelling a Society of Simple Agents: from Conceptual Specification to Experimentation*. Report, Faculty of Mathematics and Computer Science, Vrije Universiteit Amsterdam.
- Brazier, F.M.T., Treur, J., Wijngaards, N.J.E. & Willems, M. (1995). Formal specification of hierarchically (de)composed tasks. In B.R. Gaines & M. Musen (eds.), *Proc. of the 9th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, Vol. 2, 25/1-25/20, Department of Computer Science, University of Calgary: SRDG Publications.
- Cesta, A., Miceli, M. & Rizzo, P. (1996). *Effects of different interaction attitudes on a multi-agent system performance*. In (Van de Velde & Perram 1996), 128-138.
- Hanks, S., Pollack, M. & Cohen, P. (1993). Benchmarks, testbeds, controlled experimentation, and the design of agent architectures. *AI Magazine*, 13(4).
- Montgomery, T.A. & Durfee, E.H. (1990). Using MICE to study intelligent dynamic coordination. In *Proc. IEEE Conf. on Tools for Artificial Intelligence*, IEEE Computer Society Press.
- Van de Velde, W. & Perram, J.W. (eds.) (1996). *Agents Breaking Away*. Proc. 7th Eur. Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'96. LNAI 1038, Berlin: Springer-Verlag.