Using Decision Theoretic Techniques for Reasoning about Declarative Goals

Patricia H. Shaw (PhD Student)
Rafael H. Bordini (Supervisor)

University of Durham, UK

Abstract

The initial research to be done as part of this project is to look into existing methods for reasoning about goals, at the individual agent level, to improve the effectiveness of each individual agent in its ability to achieve its goals successfully. Some work has already been done on improving the efficiency of individual agents; however when extended to multi-agents systems the existing techniques require a global knowledge or control mechanism. The aim of the research is to produce a model for reasoning about goals that can be extended to a multi-agent system to improve the overall efficiency of teams. In particular, we aim at developing hybrid agents that have both representations of (declarative) goals in some symbolic language and use decision-theoretic techniques for reasoning about the feasibility of the plans available for achieving such goals.

1 Introduction

There has been much work on improving the efficiency of agent teams, including interaction between different types of agents such as humans, robots and software agents. Some of that work, in particular when involving only software agents, has required the use of global knowledge to help reduce the amount of communication required. However, this approach is not viable in most real-world, large-scale distributed applications so other techniques for improving the efficiency of the communication and coordination of agents within a team are preferred.

Our research will initially look into how reasoning about goals can be used within individual agents, in order to improve the efficiency and effectiveness of the agents in their ability to achieve the goals they commit themselves to achieving. Further, we aim at producing a model for reasoning about goals that can be extended to a multi-agent system to improve the overall efficiency of teams. We are particularly interested in hybrid models of agents, which include both symbolic (in particular BDI-based) as well as decision-theoretic techniques, along the lines of the work by Milind Tambe and colleagues [Tambe et al., 2005].

Intelligent agents should also be *rational* [Thangarajah et al., 2002]. This means they should not perform any actions that negatively affect their ability to achieve their own goals, whether as part of a team or individually. In order to ensure an agent performs rationally it needs to reason about any limited resources that are used, the actions it performs and how they interact, along with the effects of the interaction with other goals and, within teams, other agents to aid or prevent the success of a given goal. In the symbolic representation of declarative goals and available plans to achieve them, our work will follow a similar approach to the work of Thangarajah et al. [2002].

The main aim of our research is to consider how goal-based reasoning can be combined with decision-theoretic techniques. This is to improve the performance of agents and increase the chances of goals being successfully completed under extreme circumstances of goal interference and limited resources. This will involve looking at different methods of reasoning about goals within single agents, then distributing and expanding these techniques to see how they could be used for coordinating multiple agents. In terms of methodology, part of the work will be experimental, experimenting on how hybrids of existing multi-agent communication and coordination methods can be used to achieve the best performance in different environments.

A secondary aim is to make the use of the approach resulting from our research more applicable, by attempting to integrate this reasoning into an existing agent programming language, thus allowing future projects to immediately gain from the results of this research.

The specific problems to be addressed as part of this project are how to:

1. *improve the efficiency of a single agent*, for example a Mars rover, by reasoning more effectively about goals and plans with regard to potential conflicts from limited resources and positive or negative interference between the effects caused by executing the plans. This is in an attempt to try and ensure as many goals as possible are successfully completed, and if possible to try and find the most cost effective order and selection of plans to try and reduce the length of time used to achieve the goals and also make use of the least amount of resources necessary.

2. *improve the efficiency of teams of autonomous agents*, by extending the work from the first phase of the project to incorporate teams of agents that can cooperate to achieve common goals. This will incorporate reasoning about communication and coordination between the agents as well as attempting to reason throughout the team about the most efficient method of achieving their goals.

3. *embed the resulting techniques in an agent platform*, to make the use of the results from the previous stages of the work usable in practice.

A possible continuation of the proposed work would be to combine the resulting approach with various techniques in the Agents literature for the coordination of self-interested agents, in particular using game theory. The idea would be to take the work one step further to attempt to reason about how non-cooperative agents can attempt to work more efficiently to achieve each of their own goals in the context of our hybrid model. This would involve looking at issues of trust between agents, and also security issues where agents may not want to reveal everything they know.

2 Reasoning about Goals

There are multiple types of conflicts that rational agents need to be aware of; these can be internal to the individual agent, or external between two or more agents [Tessier et al., 2001]. While conflicts can occur in social interactions, when attempting to delegate or collaborate over a set of given tasks the main focus of the research is to look at conflicts between goals, both within an individual agent and as part of multi-agent systems.

The conflicts arise within a single agent when it has taken on two or more goals that are not entirely compatible. The conflicts may be caused if there is a limited amount of resources available [Thangarajah et al., 2002, Raja and Lesser, 2004], or it may be due to the effects the actions involved in achieving the goals have on the environment between two concurrent goals [Thangarajah et al., 2003a,b].

Resources [Thangarajah et al., 2002] can be reusable or consumable; for example, a communications channel is a reusable resource, while energy or time are consumed so they cannot be reused. The interaction of effects can have both positive and negative impacts in relation to other goals, particularly where causal links exist between the plans for each goal. By looking at the pre-conditions, in-conditions (i.e., invariants during a plan execution), and post-conditions (also called effects), it is possible to determine more effective plan selection and scheduling so as to avoid conflict, or so as to gain from positive interactions [Thangarajah et al., 2003a,b].

An example of negative interaction used by Thangarajah *et al.* is a Mars rover having the goal of taking rock samples from locations A and B. If the rover starts by going to location A then tries to move to location B before having taken the rock sample this would cause conflict. The plan of going to location A sets up the pre-conditions for taking the rock sample at location A, and this must be protected until the rock sample has been taken. On the other hand, an example of positive interaction may be that the agent also has to take a soil sample as well as a rock sample at location A, so rather than returning to a base station between taking the two samples, the agent can merge the two goals and take both samples at the same time.

3 Hybrid Approaches

Research in the area of multi-agent systems has resulted in a wide range of different techniques for practical reasoning and agent coordination. These have developed into various diverging approaches, with different advantages and disadvantages associated with each, making them more applicable to specific sets of conditions and environments. Recently, attempts have been made to bring some of these approaches back together again in the form of hybrid techniques [Tambe et al., 2005], which make the most of the strong points from one approach and combine it with the strong points of another, to allow for more flexibility in a wider range of conditions and environments.

We aim to follow in that direction, i.e., combining different approaches to achieve efficient goal reasoning that can be used within a high-level agent language. On the BDI agents side, our work takes inspiration from the work on reasoning about goals within individual agents by Thangarajah *et al.* In their work, they have used "goal-plan trees" to represent the goals, and produce sets of summary information to reason about the three types of conflicts mentioned above. The use of summary information and hierarchies is very similar to work done by Clement and Durfee [2000] for reasoning about conflicts in a multi-agent system.

Alternative approaches to (single) agent decision making include Petri-nets [Seghrouchni and Haddad, 1996] and CSPs [Tessier et al., 2001, chap. 3], while techniques that are currently being used for multi-agent systems include DCOP [Modi et al., 2003], POMDPs [Varakantham et al., 2006], and Markov games. Hybrids of these approaches have been produced for coordination and communication problems in teamwork [Tambe et al., 2005], but the combination of these with the reasoning about interference between conflicting plans within an agent has been largely neglected.

Through this research, this problem will be addressed, adding additional reasoning about goal-plan trees to individual agents and teams of agents to improve the overall effectiveness and performance of teams.

In the research done so far, we have initially modelled the reasoning problem using Petri-nets, in particular defining a modular structure for representing plans as sequences of actions and goals, and goals with alternative plans to achieve them. The modular structure allows us to change the Petri-net representation on the fly, to accommodate new goal requests that an agent might, for example, receive from its owner. The Petri-nets allows an agent to check whether the addition of a new goal will cause other existing goals not to be achievable anymore, due to lack of resources, or due to conflicts between the effects caused by plan execution and the pre-conditions for the execution of other plans.

4 Planned Evaluation

Various different hybrid approaches (i.e., different combinations of the techniques mentioned in the previous section) will be modelled and prototyped in an agent language such as AgentSpeak [Bordini et al., 2005] and then tested using simulations of different environments to provide a range of operating conditions. A suitable performance evaluation testbed for a single agent would be the Mars rover scenario mentioned earlier, while the Robocup Disaster Rescue Simulation testbed would be more suitable for teams of agents. This testbed is becoming increasingly used as a benchmark for teamwork in multi-agent systems. The testbeds will be used to obtain experimental results regarding the performance of the proposed techniques under a number of different conditions.

From those results, the various prototypes can then be refined to improve their performance and evaluate the effects of different environments on each of them. We will then be able to understand how the best performance can be gained by a set of agents in certain situations, and which techniques are more suitable for each of the specific situations.

Existing work provides experimental results for their research so by using a consistent set of conditions it will be possible to compare the different approaches we shall propose with existing approaches. Ideally, we would be able to find an approach that consistently gives better results than existing approaches for the Robocup testbed in particular. Naturally, it is expected that the overall performance will be a trade-off between the improvement in agents' effectiveness and efficiency in achieving their goals due to goal reasoning, and the costs of the additional computation for reasoning about conflicts and interference, which also needs to be sufficiently flexible to allow dynamic updating of the goal model during runtime.

References

R. H. Bordini, J. F. Hübner, and R. Vieira. ***Jason*** and the golden fleece of agent-oriented programming. In *Multi-Agent Programming*. Springer, 2005.

B. J. Clement and E. H. Durfee. Performance of coordinating concurrent hierarchical planning agents using summary information. In *Proceedings of ICMAS*, pages 373–374. Springer Verlag, July 2000.

P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. An asynchronous complete method for distributed constraint optimization. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, July 2003.

A. Raja and V. Lesser. Reasoning about coordination costs in resource-bounded multi-agent systems. *Proceedings of AAAI 2004 Spring Symposium on Bridging the multiagent and multi robotic research gap*, pages 25–40, March 2004.

A. E. F. Seghrouchni and S. Haddad. A recursive model for distributed planning. In *Second International Conference on Multi-Agent Systems (ICMAS)*. IEEE press, December 1996.

M. Tambe, E. Bowring, H. Jung, G. Kaminka, R. T. Maheswaran, J. Marecki, P. J. Modi, R. Nair, S. Okamoto, J. P. Pearce, P. Paruchuri, D. Pynadath, P. Scerri, N. Schurr, and P. Varakantham. Conflicts in teamwork: Hybrids to the rescue. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 3–12, New York, 2005. ACM Press.

C. Tessier, L. Chaudron, and H.-J. Müller, editors. *Conflicting Agents: Conflict Management in Multiagent Systems*. Multiagent systems, Artificial societies, and Simulated organizations. Kluwer Academic Publishers, 2001.

J. Thangarajah, M. Winikoff, and L. Padgham. Avoiding resource conflicts in intelligent agents. In F. van Harmelen, editor, *15th European Conference on Artifical Intelligence 2002 (ECAI 2002)*, Amsterdam, 2002. IOS Press.

J. Thangarajah, L. Padgham, and M. Winikoff. Detecting and avoiding interference between goals in intelligent agents. In *Proceedings of IJCAI*, pages 721–726, 2003a.

J. Thangarajah, L. Padgham, and M. Winikoff. Detecting and exploiting positive goal interaction in intelligent agents. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 401–408, New York, NY, USA, 2003b. ACM Press.

P. Varakantham, R. T. Maheswaran, and M. Tambe. Implementation techniques for solving POMDPs in personal assistant domains. In *Proceedings of Programming Multiagent Systems (PROMAS)*. Springer Verlag, 2006.