Multi-Agent Systems for Autonomous Control

Tim Clarke & Peter Mendham, University of York [1]

1 Introduction

Autonomy has been of great interest to those developing space technology for some time. Future missions may well rely on autonomy technology to an unprecedented extent. Autonomy on board spacecraft can have a large impact on operating costs. There is the enormous potential to reduce the complexity of the command and control ground segment and lessen communication requirements. Autonomy is the key to supporting missions to explore the outer reaches of the Solar System - and beyond. We note an interesting link between the desirable properties of intelligent control systems for complex autonomous systems and the behaviour of agent-based systems.

- Agent-based approaches to software and algorithm development have received a great deal of research attention in recent years and are becoming widely utilised in the construction of complex systems.
- Agents use their own localised knowledge for decision-making, supplementing this with information gained by communication with other agents.
- Remaining independent of any kind of centralised control whilst taking a local view of decisions gives rise to a tendency for robust behaviour.
- The distributed nature of such an approach also provides a degree of tolerance to faults, both those originating in the software/hardware system itself and in the wider environment.

It is for these reasons that we consider an agent-based system to be a suitable model on which to base an intelligent control system for complex systems requiring a large degree of autonomy.

2 Agency

Although widely used, multi-agent systems research has also lead to a number of definitions of agency. Once again, in some cases, these definitions are inconsistent. In our context, the terms agent or intelligent agent refer to a software entity with one or more independent threads of execution, and which is entirely responsible for its own input and output from/to the environment in which it is situated. It is therefore autonomous. We assume that the agent has well-defined objectives or goals and exercises problem-solving behaviour in pursuit of these goals; reacting in a timely fashion.

It is this behaviour that allows us to refer to the agent as intelligent. Whilst being flexible problem solvers in their own right, the power of agents is only fully realised once multiple agents are combined and communicating. This is referred to as a multi-agent system (MAS). As agents are equipped with different abilities and different goals, each agent has a distinct sphere of influence within the environment in which all the agents are situated.

[1] Intelligent Systems Group, Department of Electronics, University of York, York, YO10 5DD

These spheres of influence may overlap, defining a fundamental relationship between agents. Further relationships may be superimposed through the use of communication channels.

A MAS, therefore, has all the basic properties of a complex system: entity, asynchronicity, concurrency and connectivity. However, the usefulness of a MAS to intelligent control lies in its application.

If the MAS is to mirror the real world, then two important decisions become immediately apparent: what and how much of the world should each agent represent? We call these two related questions the *representation* and *granularity* decisions respectively. In most cases the application of a MAS has the most to offer if it is physically representative, i.e., if the agents represent interrelated subsystems at some useful level of abstraction. The correct level of abstraction, and therefore the granularity of the agents, is likely to be problem dependent. Although we have carried out some research, particularly to begin addressing the issue of representation, a more general study would prove useful to further applications of MAS to physical, complex system.

3 Multi-Agent Control

The use of multi-agent systems for intelligent control has great potential. However, a number of important challenges remain before this potential may be fully realised. Here, we summarise key areas requiring research which have been brought to light by earlier discussions.

- General theories to guide designers in making representation and granularity decisions. This would ideally result in design methodologies for the application of agents to control.
- Methods for allowing the distribution of the dynamic control process across multiple agents.Whilst no existing method is suitable, per se, opportunities may exist to adapt techniques to fit the agent paradigm.
- Representational issues governing how agents store their local knowledge and how it may be interchanged with other agents in a commonly understood format. This is essential if the methods used in multi-agent hybrid intelligent frameworks are to be coupled to dynamic control without recourse to layered architectures.
- Techniques which detect aberrant behaviour, both within the agent system and within the plant itself. The distributed nature of multi-agent systems has the potential to offer a greater degree of fault tolerance and robustness. If the operation of these techniques could be proven then the applicability of multi-agent control could extend to safety-critical systems.

There is great scope for research. Nevertheless, the issues highlighted here need not be wholly solved for multi-agent control to be useful in practical engineering situations.

Multi-agent control has the potential to exhibit the necessary characteristics of a system designed for autonomy: intelligence, robustness, adaptivity and fault tolerance. The solution to each of these issues emerges directly from the dynamic reconfigurability inherent in the use of communicating software agents.