Adaptive Supervisory Control of Multi-Agent Systems

Kiriakos Kiriakidis[1] and Diana Gordon[2]

[1] Department of Weapons and Systems Engineering, United States Naval Academy, Annapolis, MD 21402
kiriakid@novell.nadn.navy.mil,
WWW home page: http://wseweb.ew.usna.edu/wse/people/Kiriakidis
[2] AI Center, Naval Research Laboratory, Washington D.C. 20375
gordon@aic.nrl.navy.mil,
WWW home page: http://www.aic.nrl.navy.mil/~gordon

In the Proceedings of the First Goddard Workshop on Formal Approaches to Agent-Based Systems (FAABS'00).

Multi-Agent Systems (MAS) provide for the modeling of practical systems in the fields of communications, flexible manufacturing, and air-traffic management [1]. Herein, we treat the general MAS as an Interconnected Discrete Event System (IDES). By design or due to failure, the actual system is often subject to change and so is its IDES model. For example, at the highest level, an existing subsystem may fail or a new subsystem may be connected. At a lower level, the structure of a subsystem may change as well. In spite of these changes, we want the MAS to preserve important properties which describe correct overall behavior of the system. For example, we want agent actions to be non-conflicting. Formal verification followed by self-repair is a solution, albeit a complex one. The complexity is due to the credit assignment, i.e., if verification fails we must determine which agents' actions are responsible for the failure of the MAS to satisfy the properties, and what is the best way to resolve the conflicts. This could entail highly complex reasoning to resolve.

A simpler solution may be found by recasting the multi-agent paradigm in the context of supervisory control theory – a method well-known in the realm of discrete event systems that is applicable to the class of IDES. Unfortunately, structural changes in an IDES hamper the application of current supervisory control theory. At present, the literature offers but a few works on adaptive or robust supervisory control to tackle the problem of uncertainty in the IDES model [2]. Here we present a novel approach that tackles the issue of adaptive supervisory control in a multi-agent context. Two architectures are presented.

The first of our two MAS architectures is one where a single supervisor guarantees the system-wide properties of the MAS. The second is a distributed architecture in which supervisors are themselves agents. For this second architecture, each supervisor agent supervises a MAS subsystem. In this latter case, there exist interactions between agents via subsystem interconnections.

Our objective is to develop a method for *adaptive* control of the aforementioned MAS architectures. The proposed approach combines supervisory control,

learning, and verification as follows. Suppose that the supervisor controls the IDES so that it executes a desired language, which incorporates properties that the system needs to maintain regardless of structural changes. This desired language is in the form of a finite-state automaton (FSA). If a subsystem fails, the learning mechanism deletes from the desired language FSA events that pertain exclusively to the failed subsystem. Note that deleting events may fracture the desired language FSA into multiple isolated components. Therefore, to restore continuity in the FSA, an automaton repair algorithm is used. Then, model checking verifies that the resulting repaired FSA includes the aforementioned properties. If verification fails, the method applies a backup repair algorithm to the desired language FSA that requires no verification at the cost of yielding a more restricted desired language. In the end, the method automatically synthesizes a supervisor for the MAS from the new desired language, and closes loop around the IDES.

The advantage of this method is its simplicity. In particular, it obviates the need for a complex credit assignment process and instead substitutes a process of verification and repair to a single, typically small FSA – the desired language FSA. The complexity of translating repairs to the MAS is avoided. Instead we use the well-known process of automatic supervisor synthesis. Thus we obtain adaptability of the MAS via adaptive supervisory control.

1 Acknowledgements

This research is supported by ONR N0001499WR20010.

References

1. Jennings, N.,Sycara, K., and Wooldridge, M. (1998) A roadmap of agent research and development. Autonomous Agents and Multi-Agent Systems, **1**: 7–38.
2. Lin, F. (1993) Robust and adaptive supervisory control of discrete event systems. *IEEE Transactions on Automatic Control*, **38(12)**: 1842–1852.