Engineering Self-Organizing Multi-Agent Systems (Demo Paper)

Radhika Nagpal
Harvard University
Cambridge, MA, USA
rad@eecs.harvard.edu

Chih-han Yu
Harvard University
Cambridge, MA, USA
chyu@fas.harvard.edu

Daniel Yamins
Harvard University
Cambridge, MA, USA
yamins@fas.harvard.edu

ABSTRACT
In this demo session, we will present two examples of how one can systematically program self-organizing multi-agent systems, using inspiration from biology. The first system is a modular robot that autonomously adapts to satisfy complex environmentally-adaptive goals through the cooperation of multiple module agents. The second is a global-to-local compiler that can transform a user-specified pattern formation goal into a multi-agent program and reason about the agent resources required. These systems show (1) how biological design principles can be formally captured and theoretically analyzed, and (2) how global goals can be translated into local interactions amongst many simple agents. Both systems will be demonstrated in real-time and interactively with the audience.

Categories and Subject Descriptors
I.2 [**Computing Methodologies**]: Artificial Intelligence

General Terms
Algorithms, Experimentation

Keywords
Self-organization, global-to-local compilers, cellular automata

1. INTRODUCTION
Biological systems, from embryos to ant colonies, achieve tremendous mileage by using vast numbers of cheap and unreliable components to achieve complex goals reliably. We are rapidly building embedded systems with similar characteristics, from self-assembling modular robots to vast sensor networks. How do we engineer robust collective multi-agent behavior?

Our group investigates bio-inspired programming paradigms for achieving robust collective behavior in large-scale multi-agent and distributed systems. We draw inspiration mainly from developmental biology. Our goal is to build a framework for designing and reasoning about self-organizing systems: (1) to show how biological design principles can be formally captured and theoretically analyzed (2) to contribute to a fundamental understanding of how user-specified global goals can be translated into local interactions amongst many simple agents.

Cite as: Engineering Self-Organizing Multi-Agent Systems (Demo Paper), Nagpal, Yu, Yamins, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, p. 1717.

In this session, we will present two projects that demonstrate how one can design self-organizing systems that rely on simple decentralized agent interactions but are able to provably generate a wide variety of user-specified goals. Both systems are physical demonstrations of papers by our group at this AAMAS conference. While the papers focus on the theoretical aspects, this demonstration will show how these ideas work in practice.

The first system is a self-adaptive modular robot that can form complex environmentally-adaptive shapes, solely through the cooperation of locally-interacting module agents. This system is inspired by multi-cellular tissues where cells cooperate to adapt global shape and function. Underlying the adaptive behavior is a simple form of multi-agent interaction, that is closely related to bird flocking and firefly synchronization. We have shown that this can be generalized to more complex goals such as *homeostasis* (constraint-maintenance). In the demonstration we will interactively show how the physical robot adapts as the environment changes.

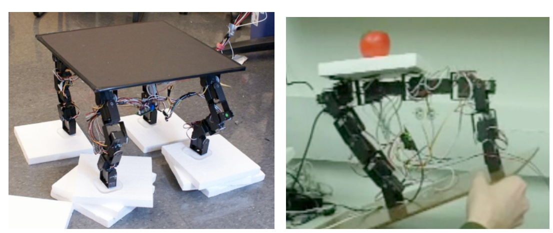

Figure 1: Images of the physical modular robot

The second system is a computer program inspired by pattern-formation in natural systems, where many cells/agents cooperate to achieve complex, reliable, but adaptable patterns. Given a user-specified 2D pattern, this program automatically generates a multi-agent (cellular automata) program that creates the pattern in a manner that is robust to number of agents, initial conditions, and asynchronous timing. In addition, the program can reason about the minimal state and minimal interaction radius required by the agents. Underlying this program is a theoretical understanding of the local-to-global relationship in spatial systems, that has many potential applications in swarm and modular robotics. A key aspect of both systems is that they demonstrate "global-to-local compilation": the agents rely on simple and local interactions that provably self-organize a wide class of user-specified global goals.