On the Monitoring Selectivity Problem

Gal A. Kaminka
Computer Science Department
Carnegie Mellon University
galk@cs.cmu.edu

November 9, 2003

The field of Multi-Agent Systems (MAS) is seeing an explosion of research and applications bringing together ideas from AI and distributed systems, psychology and sociology, economics and strategic studies. Indeed, many are seeing agent-oriented systems, composed of many agents that interact with each other, as a new paradigm for system design and software engineering. However, this idea of bringing multiple agents to bear on a problem, while common, is surprising to me. It would seem to me that the multiple agents are either already there, before we (MAS researchers) arrive at the scene, or they are not. What has changed is not the nature of the problem, but our state of knowledge, as we discover that a naturally occurring phenomena (multiple agents interacting with each other) is apparently common not only among carbon-based agents, such as humans, but also among silicon-based agents, such as multiple servers running a distributed program.

Indeed, it is this view of the MAS field emerging as a scientific discipline, rather than another engineering discipline, that I wish to support. As scientists facing the exciting challenge of studying the behavior of a class of systems, we look for key characteristics that would enable us to classify different types of systems under study, and for rules that would help us predict, and later explain, the behavior of such systems. Yet, when we examin MAS literature, there are very few investigations aiming at such high-level goals. Most reports focus on either on successful, novel, designs of MAS for novel applications, or on prescriptive theories, which are to be used to guide the implementation or running behavior of such systems. I have begun to address these issues.

In this short abstract, I will attempt to demonstrate (a) that key characteristics of multi-agent systems exist; (b) that important sub-sets of MAS literature can be viewed as addressing rules of behavior arising from these characteristics; and, perhaps most important, (c) that appropriately, existing investigations, focusing on the design and prescription of behavior, can be used instead in a predictive role, to predict the behavior of new multi-agent systems that have similar characteristics.

First, I propose that there are (at least) two bases for characterizing MASs: The social relationships that are maintained by the systems (e.g., collaboration, adversarity, spatial formations), and by the amount of knowledge available to agents as they maintain these relationships (e.g., full knowledge of others' state vs. partial knowledge). The former characterization, though novel for classification, has to some extent been addressed in the literature, and has an intuitive appeal. The second one is novel, and can be explained briefly: There is a continuum that describes the amount of knowledge that an agent may have about its peers in a given multi-agent system. In one extreme end, an agent may have full knowledge (for instance, with unlimited bandwidth). In the other extreme, the agent may have very little knowledge about the other agents (for instance, if it can only observe them occasionally). The first extreme is often infeasible, if only because it requires much computational power to process all the incoming information (e.g., many humans have problem counting while monitoring a different count). The latter extreme is also problematic, as lack of knowledge causes uncertainty about other agents, which can severly limit the agents ability to function correctly in a given MAS. Multi-agent systems can be characterized by their placement on this contiuum, which can be a function of their design or evolution.

Taken together, these two characterizations define an instance of the *monitoring selectivity problem (MSP)* [2]: How can an agent maintain a *required relationship*, given its *placement on the continuum* discussed above. My claim is that much of MAS literature can be viewed as addressing MSP problem instances. I will focus on collaboration as an example. The Joint-Intentions collaboration theory [1] defines the correct behavior of a team-member towards its peers. In doing this, it puts forward requirements on the knowledge that the team-member will need of others. For instance, it requires that agents have mutual belief in a joint goal, thus requiring agents to know about the goals of others (this is already addressing a class of MSPs—the theory says that communicating about goals is important, as opposed to other communications). However, in building systems that implement this theory, it was found that in some cases, agents could not communicate continuously. Therefore, proposed solutions in the literature relied on agents' visual observations of each other to confirm agreement [2], and on decision-theoretic consideration of the cost of failure vs. cost of communciations [4].

Accepting that MAS literature can be viewed, to a large degree, as a collection of engineering solutions to monitoring selectivity problem, we can attempt to sythesize a predictive model of MASs from these instances of the phenomena under study. In particular, I have taken the prescriptive collaboration theories which is meant to be applied in guiding MAS design, and applyed it in a predictive role, to quantitative measure certain aspects of collaboration in actual real-world systems, with some success. For instance, in [2], I use such measures to distinguish between simulated soccer teams that rely on communications and teams that rely on improved visual sensing, and in a new, yet unpublished, study, I show that a correlation can be found between a teams soccer-playing ability, and its conformity to theoretical models of cooridnation, initially investigated for business processes [3].

I am currently pursuing two lines of work which are based in these novel ideas. First, I am exploring ways in which new engineering solutions can be guided by our understanding of the nature of the monitoring selectivity problems they are addressing. Specifically, I am interested in facilitating human teamwork by building teamwork-assistants, which help a user manage its interactions with team-members by relying on their knowledge of teams and humans. Second, I am continuing investigations of MASs, seeking to classify and measure them. I am now building the tools for measuring agent teamwork in several environments, including robots, software agents, and humans.

References

[1] Philip R. Cohen and Hector J. Levesque. Teamwork. *Nous*, 35, 1991.

[2] Gal A. Kaminka and Milind Tambe. Robust multi-agent teams via socially-attentive monitoring. *JAIR*, 12:105–147, 2000.

[3] Thomas W. Malone and Kevin Crowston. Toward an interdisciplinary theory of coordination. Technical Report CCS TR#120 SS WP# 3294-91-MSA, Massachusetts Institute of Technology, April, 1991 1991.

[4] Milind Tambe. Towards flexible teamwork. *JAIR*, 7:83–124, 1997.