

Industry traction for MAS technology: Would a rose by any other name smell as sweet?

(Position paper submitted to the AAMAS-2008 session on the “Future of software engineering and multi-agent systems”)

Aditya Ghose, Decision Systems Lab, University of Wollongong, NSW 2522
Australia (aditya@uow.edu.au)

The multi-agent systems (MAS) community faces a crisis that many are unwilling to acknowledge. We are all acutely aware of how MAS concepts and technologies seem to have failed to gain significant industry traction (a few notable exceptions aside). Other research communities (e.g., the service-oriented computing community) have gained far greater industry acceptance for their outputs, in far shorter time. But the crisis stems from an even greater threat: the co-opting of agent/MAS concepts in other research communities, particularly the service-oriented computing community. The connections between agents and services run deep, and are to some degree recognized and addressed in the literature. Like agents, services can be viewed as autonomous, reactive components. Think of a dynamic service broker – itself a service – that re-computes service compositions on the fly (in response to changing service requirements and a dynamic operating context), in a manner akin to reactive agent programming, from a library of available services, just as one would compute composite agent plans from a plan library. Thus agent planning can manifest itself as service composition, agent negotiation as SLA (service level agreement) negotiation, and so on. I did an informal analysis of the papers presented at a recent major conference on service-oriented computing, and concluded that at least 60% of these had agent technology underpinnings in some form or the other. In other words, even as the MAS community worries about marginalization by industry, our research outcomes *are* finding useful and significant industry application, but under the banner of services.

It might be argued that this is not necessarily a bad thing. We might ask, for example, whether the nomenclature of our research (e.g., agents vs. services) should matter as long as we get to explore the really interesting and important questions. I submit that the nomenclature does indeed matter. We know from the sociology of research that different research cultures exist within different research communities. The MAS community offers a research culture that encourages the exploration of different questions, and in a different style, to those that the services community encourages. So should we worry about ending re-labelling our work as services research in our quest for industry relevance? Yes, we should. The label matters – it can influence research culture, style and content.

There are things, though, that we can learn from the success story of service-oriented computing in terms of gaining and retaining industry relevance:

- *Offer a simpler value proposition.* The services community offers a very simple value proposition: it is easier to model, design and deploy systems from distributed collections of components packaged as services (some of the recent discourse on service modelling refers to use of anthropomorphic constructs, further blurring the distinctions between our communities). The

MAS value proposition is far more complex. The notion of “agentification”, and the questions explored by much of the agent-oriented software engineering (AOSE) sub-community has a similar feel, but the MAS community has a myriad other technology offerings. These offerings often come with the baggage of legacy industry (mis)perceptions regarding “heavy” AI techniques that underpin a lot of MAS research.

- *Offer an incremental value proposition.* The MAS value proposition, in some ways, calls for radical changes to the state of industry practice. The service-oriented computing value proposition required far more incremental changes. Industry prefers incremental change to radical change.
- *Define a core agenda, while admitting a diversity of subsidiary themes.* The core research agenda for the services community is driven by software engineering concerns: service specification, discovery, composition and deployment. The AOSE research agenda has a similar feel, but the broader MAS research agenda is far more diffuse.

So what can we do now? Two strategies deserve our attention:

- *Mediate the deployment MAS research in industry through AOSE paradigms:* We have made comments that may appear critical of the MAS community: its complex value proposition for industry, its requirement for radical changes to the state of industry practice, and its diffuse research agenda. These should not be read as critiques. The breadth, depth and complexity of the ideas explored by the MAS community are its strengths and underpin the rich intellectual outputs that the community generates. AOSE research, by definition, addresses industry relevant concerns, and can form the packaging required to make the broader outputs of the MAS community more acceptable for industry.
- *Pro-actively engage the service-oriented computing community (both research and industry):* This reinforces a point made by Michael Georgeff in his panel submission. We need to highlight the substantial intersection between the services and MAS research agendas. We need to make explicit the agent technology antecedents of several key services concepts. We must use the connections with services concepts as yet another avenue for industry deployment of our research.

In the meantime, we must disagree with Shakespeare. He argues for the primacy of substance over form, of content over packaging when he says through the voice of Juliet: “that which we call a rose, by any other name would smell as sweet”. In our quest for industry relevance, form and packaging assume unexpected importance and impose on us additional obligations in interpreting our research results for industry.