Defining the scope of an introductory MAS course: A balancing act
Position statement.

Roberto A. Flores & Philippe Pasquier
Université Laval, Département d'Informatique et de Génie Logiciel,
Sainte-Foy, Québec, Canada G1K 7P4
{flores, pasquier}@damas.ift.ulaval.ca

When designing a course anew (particularly in the case of new faculty), one is confronted with a myriad of basic questions in course-making, e.g., on the material (what is the state of the art?), on the teaching philosophy (what strategies are adequate to present this material?), on the results (how will students benefit from this course?), on the environment (how will this course fit the curricula of the institution?). These questions are not independent of each other, and the answer to one is likely to influence the others.

Of the above, perhaps the questions that may be more relevant to this workshop are on the material and the results, i.e., which material should be selected for an introductory course in multiagent systems (MAS) such that this material is representative of the area, and strikes a balance between theory and practice?

Just selecting representative topics is a challenge by itself: not only can MAS be seen as an intersecting arena of computer science areas (e.g., artificial intelligence, human-computer interaction, software engineering, internet computing) but it is no stranger to inter-disciplinarity either, as researchers commonly draw inspiration from sociology, linguistics, economics, and so on. It could be argued that computer science disciplines should be emphasized in a MAS course and that inter-disciplinarity should be covered through courses in their own faculties of study. This could perhaps be a motivation for a major in MAS as part of computer science curricula. (It is worth noticing that, even though inter-disciplinarity would naturally allow the development of MAS courses in faculties other than computer science, in this statement we assume a computer science perspective).

A more practical issue is the choice of computational tools that students should be exposed to during the course. On the one hand, the use of just one tool would allow for incremental assignments, where the expertise gained in one exercise could be passed onto the next. This would allow for depth but (likely) not for breadth, as there is hardly any one-size-fits-all computational tool for MAS. On the other hand, the use of several tools (incompatibility assumed) would increase breadth at the cost of depth. This issue is magnified by the experimental nature of existing MAS frameworks.

To summarize, there seems to be a contrast between the range of essential topics that should be covered in a MAS course, and the topics that could reasonably be covered in the time allotted. The absence of minimal topics expected in such a course may lead to (a valid but undesirable) *ad hoc* syllabi where the (research) preferences of the instructor would take precedence over a balanced overview, a distortion that could mar students' appreciation of the topics that are not covered. This problem is magnified by the narrow scope of current tools and their experimental nature; to use a range of tools would be helpful for breadth but would not likely allow for depth; and vice versa for the use of one tool only.

Hopefully, the experiences gathered from this workshop will shed light on these challenges.