Research Profile: Rob Kremer, Computer Science

Title: Agent Communication

The term *agent* refers to computer programs that are typically given goals and act autonomously to achieve their goals, as opposed to more traditional programs that require complete and detailed specifications for all of their actions. Agents are typically *social* in that they work together to accomplish their goals. This means they must communicate among themselves to co-ordinate their actions. Early work in agent communication involved specifying *protocols* or *conversational scripts*, but this proved unwieldy and brittle in that it was almost impossible to deal with conversational variants, on-the-fly changes, errors, failures, and other nuances typical of robust human conversation. Logical conversational models took over from protocols, the dominant model being the FIPA standard which uses a modal logic called BDI (Belief, Desire, Intention). There are several problems with the FIPA standard including it's cumbersome expressions (a typical one being along the lines of "I believe that you believe that I believe that you believe…"), and FIPA's assumptions such as "agents never lie" and that agents can have an omniscient knowledge of the internal states of all other agents. A more recent agent communication paradigm, called *social commit theory*, addresses some of the issues. In social commitment theory, a set of *social norms* or *policies* is established among a community of agents. These *policies* specify how messages can commit agents to certain *shared social commitments*, and how conversations (specific sequences of messages) can be used to negotiate other *shared social commitments*. A *shared social commitment* involves a debtor who is responsible for fulfilling the social commitment and creditor for whom the social commitment is owed and who is responsible for it's eventual acceptance. One of the advantages of social commitment theory is that agents aren't required to maintain a complex and omniscient view of all other agents, only a list of commitments between agents that can be easily inferred from observation of conversations (relative to the social norms).

A large part of Dr. Kremer's research is the development of the CASA (Collaborative Agent System Architecture) infrastructure which is designed to be both a practical multi-agent system development platform and an extremely flexible research tool that can support all three agent communication paradigms mentioned above. One of the biggest problems in development of CASA is achieving a balance between ease-of-use and flexibility. Ease-of-use is necessary to allow both practitioners and researchers to concentrate on their objectives without worrying about implementation details. Ease-of-use often entails rigidity of the system. On the other hand, flexibility is required to allow researchers to experiment with different (existing and new) agent communication paradigms and methods and run controlled experiments to compare them.