NSF/ESRC Agenda Setting Workshop on Agent-Based Modeling of Complex Spatial Systems: April 14-16, 2007

Distributed Simulation of Agent-based Models

Georgios Theodoropoulos

School of Computer Science
University of Birmingham

There has been considerable recent interest in complex systems, which involve dynamic and unpredictable interactions between large numbers of components including software, hardware devices (such as sensors), and social entities (people or collective bodies). Examples of such systems include from traditional embedded systems, to systems controlling critical infrastructures, such as defence, energy, health, transport and telecommunications, to biological systems, to business applications with decision-making capabilities, to social systems and services, such as e-government, e-learning etc. The complexity of such systems renders simulation modelling the only viable method to study their properties and analyse their emergent behaviour. Multi-agent systems (MAS) have emerged as a particularly suitable paradigm for modelling complex systems. When embedded in a real system, a MAS is itself a complex system whose properties and emergent behaviour have to be analysed via simulation. An *agent* can be viewed as a self-contained, concurrently executing thread of control that encapsulates some state and communicates with its environment and possibly other agents via some sort of message passing.

While agents offer great promise, adoption of this technology has been hampered by the limitations of current development tools and methodologies. Multi-agent systems are often extremely complex and it can be difficult to formally verify their properties. As a result, design and implementation remains largely experimental, and experimental approaches are likely to remain important for the foreseeable future. Over the last two decades, a wide range of MAS simulators and testbeds have been developed, and simulation has been applied to a wide range of MAS research and design problems, from models of complex individual agents employing sophisticated internal mechanisms to models of large scale societies of relatively simple agents which focus more on the interactions between agents. However, existing MAS simulations and simulators suffer from two key problems.

The first problem is lack of performance. The computational requirements of simulations of many multi-agent systems far exceed the capabilities of a single computer. Each agent may be a complex system in its own right (e.g., with sensing, planning, inference etc. capabilities), requiring considerable computational resources, and many agents may be required to investigate the behaviour of the system as a whole or even the behaviour of a single agent.

The second is lack of interoperability. The development of complex MAS simulation, usually requires collaborative effort from researchers with different domain knowledge and expertise, possibly at different locations. Furthermore, the effort required to develop a new simulation from scratch is considerable. There is therefore a strong incentive to reuse existing simulation components, toolkits and testbeds for a new problem. However, while many simulations have been developed, it is difficult to leverage this investment in the development of new agent simulations. Simulations developed for different agent simulators typically do not inter-operate, making it more difficult to re-use simulation components. This is particularly problematic in the case of spatial agent based models. Combining a simulation of an agent architecture developed for one simulator with a simulation of an environment developed for another typically involves re-implementation of one or both components. If the agent must be simulated in several different environments, the problem is compounded.

A solution to both these problems can be found in distributed simulation. The last decade has witnessed an explosion of interest in distributed simulation not only for speeding up simulations, but also as a strategic technology for linking simulation components of various types at multiple locations to create a common virtual environment. The culmination of this activity, has been the development of the High Level Architecture[1] (HLA), a framework for simulation reuse and interoperability developed by the US Defence Modelling and Simulation Office. Using HLA, a large-scale distributed simulation can be constructed by linking together a number of (geographically) distributed simulation components, or *federates*, into an over-all simulation, or *federation*. HLA, with minor revisions, has been adopted as an IEEE standard (IEEE 1516) and is likely to be increasingly widely adopted within the simulation community.

Distributed simulation and HLA offer an attractive potential solution to the problems of simulation and simulator reuse and simulation performance in MAS simulation. The development of HLA compliant agent simulators and simulation components would facilitate inter-operation with other simulations, allowing greater re-use of agent simulation components. In addition, the ability to distribute agent and other simulation components across multiple computers has the potential to increase the overall performance of a MAS simulation, given sufficient computational resources and favourable simulation characteristics. However most of the work in this area to date has employed various ad-hoc approaches to parallel simulation, e.g., distributing the agents over a network of processors interacting via some communication protocol, and has yielded relatively poor performance. MAS models present particular challenges for distributed simulation. An example key problem in the distributed simulation of spatial agent-based models is the efficient distribution of the agents' environment, namely the part of the world (or computational system) 'inhabited' by the agent. In simulations of situated MAS, the environment is represented by a large shared state space, which may

[1] https://www.dmso.mil/public/transition/hla/

be accessed by any of the agents frequently and in dynamic, non-deterministic patterns. It is therefore difficult to determine an appropriate simulation topology *a priori*. Encapsulating the shared state in a single process (e.g. via some centralised scheme) introduces a bottleneck, while distributing it all across the distributed resources (decentralised, event driven scheme) will typically result in frequent all-to-all communication and broadcasting.

While the HLA enables interoperability and the construction of large-scale distributed simulations using existing and possibly distributed simulation components, it does not provide support for collaborative development or configuration of simulation applications, nor does it provide any mechanism for managing the resources where the simulation is being executed. The emergence of Grid technologies provides exciting new opportunities for large scale distributed simulation of agent-based models, enabling collaboration and the use of distributed computing resources, while also facilitating access to geographically distributed data sets.

How should large-scale environmental data sets be distributed to enable efficient access by the agents in a distributed simulation? How can we develop efficient infrastructures to support load management, synchronisation, and query routing in large scale MAS simulations? How can we support collaborative model development and automated composition of simulation components? Is HLA the answer to interoperability? The workshop can discuss some of these challenges, which are at the heart of making agent-based simulation feasible at a scale necessary to master the challenges ahead of us.