Experimental Analysis for Large Agent Systems

Dave Robertson
Division of Informatics, University of Edinburgh
D.Robertson@ed.ac.uk

This extended abstract gives a basic introduction to the aims of a new project, funded by the European Commission, on the experimental analysis of large multi-agent systems. The project involves the University of Edinburgh (Anderson, Fourman, Robertson, Sannella, Vasconcelos, Walton), the Instute for Artificial Intelligence in Barcelona (Agusti, Sabater, Sierra) and the University of Liverpool (Parsons, Wooldridge).

Engineers working on large, distributed, multi-agent systems face a problem which differs from conventional software engineering. They build software systems which must coexist with other systems about which little may be known, yet we wish the overall behaviour of the population of systems to be predictable in certain ways depending on the domain of application. There is a loose analogy to the biosciences, where the scientific response to the problem of understanding population behaviour has been to build mathematical models at various different levels of granularity of detail and use these to help form hypotheses about the driving forces in very complex ecological systems. We are beginning to follow a similar path in analysing multi-agent systems.

Two technical prerequisites for solving this problem are a framework within which to design and run experiments on models of large agent systems and clear software engineering methods which allow the results of experimental analyses to be related to agent design choices. These currently do not exist. There are numerous agent deployment systems but no convincing systems for modelling agent populations and their evolution. There are numerous software engineering methods but none of these translate easily to agent design. Providing a combination of analysis and design in this area must therefore be a ground-up exercise, drawing upon fragments of research from related areas.

Section 1 gives an overview of what we hope to achieve by the end of the project. Section 2 gives a flavour of the sort of analysis by working through a basic example constructed in the first weeks of the project.

1 What we Hope to Achieve

An overview of our proposed framework is shown in Figure 1. In the centre of the diagram is a laboratory system in which experiments are run on multi-agent models. Two major processes surround the laboratory. To the left is the design of agents which is constrained by design rules (determining key features agents can have) and hazard analyses (warnings of threats created by design choices). These controls on design are extended and revised in the light of analyses from laboratory experiments, allowing us to feed empirical results relating to system behaviour back into the definition of design controls - a "virtuous cycle" of design $\rightarrow$ experiment $\rightarrow$ redesign. The second major process (to the right) develops an overall theory, for our experimental systems, of the response of ecosystem properties to design decisions. This influences the conduct of subsequent experimental analyses, which are used to reinforce and extend the experimental theory. It is also compared to case studies from real-world systems which either confirm the experimentally observed behaviour or generate exceptions to it. These exceptions prompt theory revision, giving a second "virtuous cycle" of experiment $\rightarrow$ theorise $\rightarrow$ validate $\rightarrow$ re-theori se $\rightarrow$ re-experiment.

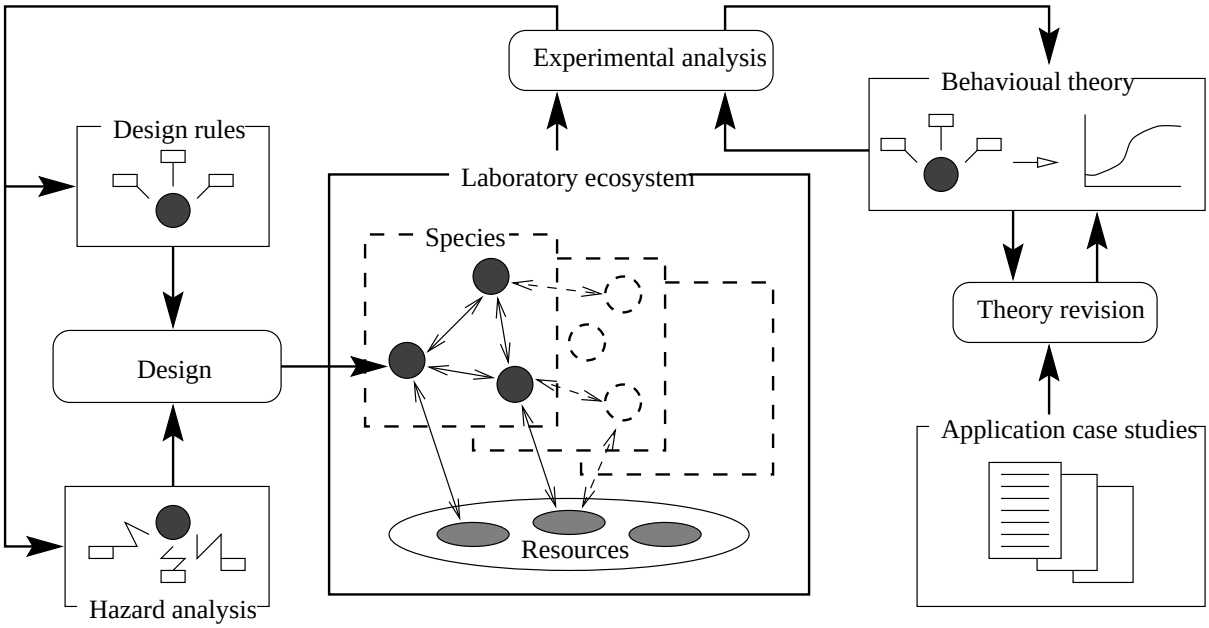

Figure 1: Conceptual overview of the project

Our aim is to innovate in the design lifecycle and validate with respect to ecosystem lifecycles. We are interested in relating the methods we use to design ecosystem inhabitants to aggregate behaviours observed at the ecosystem level.

2 Introductory Example

The purpose of this section is to give a straightforward example of analysing a rudimentary agent system. First we describe the system to be modelled; then describe the model itself; then perform some experiments with the model; and finally show how a potential threat occurring in the model may be countered.

2.1 The System to be Modelled

In the financial industries we now find systems where some centrally held resource cannot in practice be distributed directly from the centre but must be distributed via mediators. To allow consumers of the resource to be dealt with promptly, each mediator has its own system of bookkeeping and authority from the centre to decide which consumers to supply. Messages passed back to the centre allow the overall pattern of resource use to be monitored.

An example of this type of system is in bank account transactions. The resource in this case is the money in the bank. Its consumers are customers with bank accounts. The mediators are the points of contact for account withdrawals: at a branch, through a telephone call centre or via an internet site. Each of these points of contact may keep its own database of customer account details and these frequently will be reconciled with one or more central accounts databases.

2.2 The Model

The agents in the model and their potential interactions are illustrated in Figure 2. We can have any number of central resource agents (labelled R on the illustration). Each of these has some number of mediators (labelled M) and we shall assume for simplicity that this number is the same for each central resource. There can be any number of consumers (labelled C) and each may interact with one or more groups of mediators (in the illustration, C1 and C2 interact with the mediators for R1 while C3 interacts with the mediators for R1 and R2).

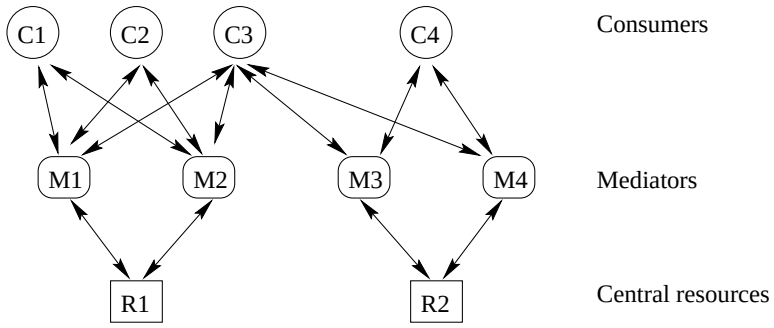

Figure 2: Overall architecture of the model

The only messages which are passed between agents (corresponding to the arcs in Figure 2) are as follows:

- A consumer can request a quantity of resource from a mediator.
- A mediator can credit a consumer with some quantity of resource.
- A mediator can inform its central resource that it has supplied a quantity of resource to a given consumer, along with the status of the resource account for that consumer at the time of supply.
- A central resource can give one of its mediators updated account information for a customer.

In the absence of consumers the central resource tends to increase (if the resource is money then this increase is due to returns on investments of the fund). We assume that this is proportional to the total amount in the fund. Consumers lose resource at what we assume is a fixed rate over time (if the resource is money this loss is through personal expenditure). If the funds remaining in a central resource or consumer reduce to zero that agent dies.

Consumers can employ different strategies for acquiring resources. We model two of these:

- A "careful" strategy in which a consumer applies to a single appropriate mediator each time it wants resource and waits until it has received credit for the resource before looking for more.
- A "greedy" strategy in which a consumer applies as often as possible for a given quantity of resource to all the mediators, with the aim of obtaining as much resource as it can as soon as it can.

This allows us to run the model with different settings for various parameters including the following (a range is given where we varied the parameter during analysis) :

Parameter	Setting
The initial number of resources.	4
The initial number of consumers.	100 - 400
The number of mediators per central resource.	2
The percentage of consumers allowed to interact with each central resource (the choice of individuals being a random subset of the total pool of consumers chosen according to this percentage).	25%
The percentage of greedy consumers in the total consumer population.	0-100%

2.3 Experimenting with the Model

A property of interest in this system is the ability of its central resources to be sustained under exploitation by different proportions of greedy consumers. We investigate this by fixing the initial level of central resource and running the model with different initial numbers of consumers. These runs are replicated with three different proportions of greedy consumers: 100%, 50% and 0%. The results are shown in Figure 3. The higher curves in each of the three diagrams are for runs with lower numbers of consumers, since these deplete the central resource less. The topmost curve of each graph, at 100 initial consumers, shows that at this population size the central resource recovers quickly from initial exploitation. The effects of having more greedy consumers appear most strongly at around 300 initial consumers, where the central resource diminishes to zero almost immediately when all or even half the agents are greedy, whereas it recovers after coming close to zero if no agents are greedy.

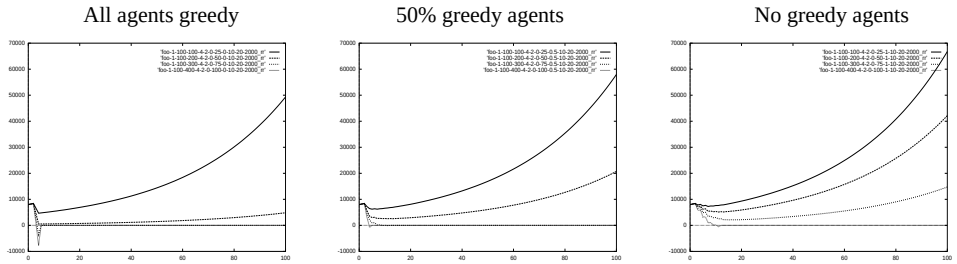

Figure 3: Central resource variation in response to proportion of greedy consumers.

The consumers in this model have a basic strategy of claiming resource whenever they can (modulo the greedy or careful strategy) which means that they exhaust their allocations around midway through the simulation and then slowly die of starvation. This gives the pattern of sharp resource increase followed by slow tail-off shown in the graphs of Figure 4 (each of which gives the consumers' view of the corresponding simulation run in Figurefig:basicexpt1). The peak is, unsurprisingly, most accentuated with greedy agents.

2.4 Counteracting a Threat

A surprising feature of the graph in Figurefig:basicexpt1)involving greedy agents is that the central resource can be depleted to below zero. This is a side effect of having more than one mediators autonomously distributing the same resource because greedy consumers will attempt to collect resource from all mediators simultaneously. Each mediator, however, has its own separate knowledge of how much resource each consumer is allocated. Although

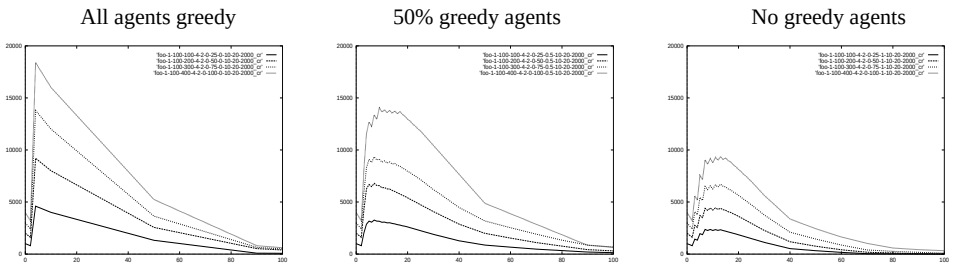

Figure 4: Consumer resource variation in response to agent mixture

this initially is consistent with the central resource it becomes inconsistent if more than one mediator simultaneously agrees to a release of resource for the same consumer, and this can cause more resource to be promised than the central resource has available. For instance if the resource allocated by R1 to C1 in Figure 2 is 10 units then M1 and M2 may simultaneously allow C1 to borrow 10 units. R1 now has to honour a commitment of 20 units although only 10 were intended. This phenomenon is known to occur in banking systems, where it is possible to withdraw more money than you have in your account by making multiple withdrawals through different mediators.

We would like to be able to counteract this threat but we cannot assume the easy solution of synchronising central resources and mediators each time a mediator is deciding whether to allocate resource to a consumer. An alternative is to create a new type of agent, call it an auditor, which at regular intervals asks a mediator and its central resource what they think is the unused allocation for a given consumer. If there is a discrepancy it broadcasts a warning message to all mediators who can then block further requests for resource from that consumer, ensuring that it starves to death. The arrangement is illustrated in Figure 5, where auditor A1 queries M1 and R1 about (say) C3 and finds a discrepancy so broadcasts a warning to M1, M2, M3 and M4.

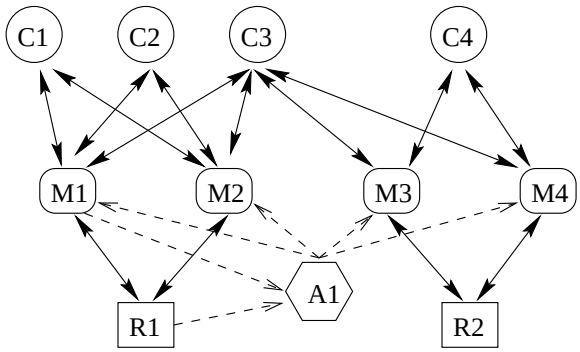

Figure 5: An auditor (A1) interacting with the resource system

To investigate whether this has a significant effect on the resilience of the resources against attack by greedy consumers simulations were run with 300 consumers, 50% of which were greedy, with first 0 auditors (as before) then 100 auditors; then 200 auditors. We can see the results of this in Figure 6. At 0 auditors the resource collapses to zero. With greater numbers of auditors it recovers quickly after an initial dip, as the warnings about greedy consumers exclude those consumers from the resources. We have not gone on to analyse whether there is an optimum number of auditors at which adding more gives lower increases in protection to the system but this seems likely to be the case.

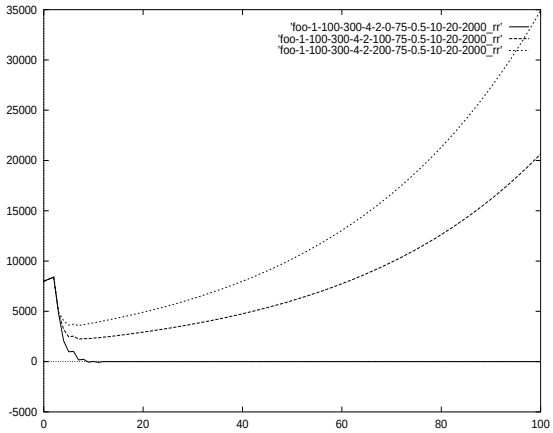

Figure 6: Shift in carrying capacity in response to auditing