ARC CENTRE FOR
COMPLEX SYSTEMS

Technical Report

ACCS-TR-06-01

Control and Constraint: Cross Centre Insights from Modelling Cellular Morphogenesis

ACCS Workshop – 25th January 2006

T. Rudge and N. Geard

February 2006

ARC Centre for Complex Systems
School of ITEE, The University of Queensland
St Lucia Qld 4072 Australia
T +61 7 3365 1003
F +61 7 3365 1533
E admin@accs.edu.au
W www.accs.edu.au

Control and Constraint:
Cross Centre Insights from Modelling Cellular Morphogenesis

ACCS Workshop - 25th January 2006

Tim Rudge and Nic Geard

February 16, 2006

Complex systems, as a field of research, is predicated on the existence of a level of abstraction at which systems in different domains share common features and can be explored using common methodologies. The ACCS Workshop on Cross Centre Insights from Modelling Cellular Morphogenesis was organised to facilitate discussion between researchers and model builders interested in gene regulatory networks (GRNs), evolutionary economics (EE) and air traffic control (ATC). The focus of the workshop was the multi-agent cellular model developed to investigate the genetic control of morphogenesis and the insights obtained through the development and use of this model about the roles of control and constraint in understanding complex behaviour. The workshop was attended by 26 researchers representing all three ACCS streams (GRN, EE and ATC), as well as representatives from ITEE, ACMC, ACB, CMLR, CSIRO and Biomedical Sciences.

This document summarises the abstract framework presented at the workshop together with summaries of some of the main points raised in subsequent discussions.

1 The Constrained Multi-Agent Model

1.1 Model Structure

The modelling framework consists of four components:

Agents: containing both a controller of some nature (E.g., neural network, rule set) and a description of it's current state.

Interactions: describing some form of communication channel between agents (either direct or mediated by the environment).

Environment: in which the agents exist, and which may impose some form of constraints on the behaviour of the agent; either physical (E.g., spatial interactions) or non-physical (E.g., laws or regulations).

Adaptive Process: by which agent behaviour can be modified between instantiations of the system; this could be either automatic (E.g., an evolutionary algorithm) or a process of trial and error (E.g., manually reformulating a hypothesis).

1.2 The Role of Constraints

A key element of this conceptual model is the interaction between the environment and agents (figure 1). Agent controllers act in response to the current agent state to determine the state after some discrete time step. Controller output does not directly manipulate the agent's state but acts via the environment process, which encapsulates a set of biases or constraints on agent behaviour.

An agent's behaviour at a given point in time is therefore a function of both the actions of its control component and the effect of the constraints that act upon it. The performance of a multi-agent system may be defined in terms of the behaviour of all of the agents in the system. In the absence of any control, an agent's behaviour will be due to the effect of constraints alone. Therefore, in the absence of agent control, system performance will be a function of constraints alone.

Constraints therefore define the null behaviour of a system: what its behaviour will be if control is either absent, non-optimal, or fails. The term 'constraints' should not necessarily be interpreted as restricting the behaviour of a system in a negative way. Constraints can also operate in a positive fashion: if the null behaviour of a system (in the absence of control) can be configured in such a way that it is close to the optimal behaviour of the system, then locating that optimal behaviour may well be easier. In addition, if control subsequently fails, a certain amount of fail-safe robustness is built into the system in the sense that the null behaviour to which the system returns is known not to be catastrophic.

While the purely constraint-driven behaviour of a system may suffice in a predictable or static environment, in a more unpredictable or dynamic environment, adaptive control may be required to ensure adequate system flexibility. The optimum design of a system will therefore consist of a balance between roles for control and constraint, with the location of this balance depending on the context and specific details of the system.

2 Initial Reactions to the Model Overview

There was broad agreement that this abstract model of constrained multi-agent systems might be applicable to a wide range of problems tackled across the three streams of the ACCS. It was noted that it may even be possible to further specify

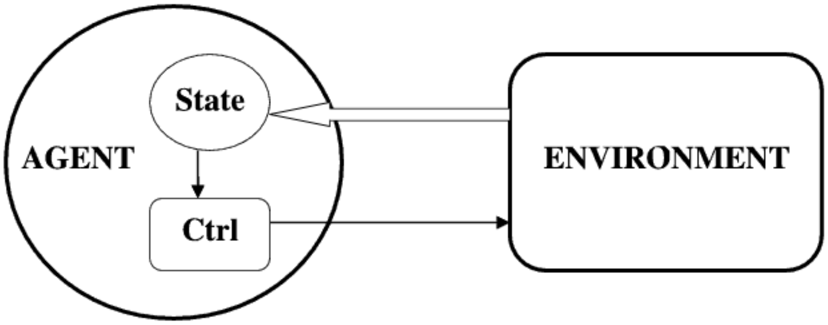

Figure 1: Agent controllers act via the environment process, which encapsulates a set of constraints or biases on agent behaviour and directly affects the agent states.

the nature of the environment process in order to capture common features of the relevant problem domains.

Incorporation of two types of processes was discussed: (1) Thermodynamic principles such as conservation of energy: these are being applied in economic models, and are obviously relevant to physically based systems. (2) Knowledge repositories would also be relevant in many areas, a prime example being DNA as a repository of available catalytic reactions.

Specification of the environment process in terms of these types of abstracted sub-processes could allow domain specific knowledge, techniques, and tools to be refactored in transferrable terms.

3 Summary of the Economics and ATC discussion

Discussion focused on the nature of environmental constraints in each of the two streams. Given a defined problem focus, the identity of each of the system components is trivial to deduce in these domains. It seems that the nature of environmental constraints is what will determine whether, and in what way, tools and methods can be reused across these domains.

Geographic considerations have not been widely considered in previous economic models. Discussion centred around the idea of geographic constraints in the broad sense, including transportation costs, regulations, international trade groups etc. Availability and location of resources were also considered to be environmental constraints, and are major influences on what types of businesses occur in which locations, where and with whom businesses must do trade, and ultimately the success of those businesses. Resource availability is not changed by agents, unless we include resource producers. In this case, the resource producers we include will ultimately depend on some other non-agent resource, which can be considered as an environmental constraint.

Direct agent-agent communications might be necessary to free-flight ATC (plane-plane communications) and economic systems (price negotiation, electronic markets, direct trading). However these cases may be conceptualised as specialisations of the general model, where all communications are via an identity process in the environment, rather than via a direct connection between two agents. In fact, electronic or direct trading is not direct agent-agent communication (i.e. is not simply an identity process in the environment) because the agent is not entirely free to choose any agent with which to deal. The environment defines with whom the agent may communicate, through geographic constraints, willingness to trade, and other factors.

Environmental communications as a message hub. If all agent-agent messages are mediated by the environment, this is analogous to a message hub in software engineering. The implications of this design are maximal separation of the communication channel from the agents. This means that the communication mode may be changed without requiring any corresponding change to the agents. This represents an easily modifiable and extensible design in which agents may be modified in any way whilst maintaining the capacity to interact with other agents in the system. In cellular morphogenesis we might think of multiple local message hubs, which are relatively uncoupled with respect to messaging.

The environment process is that which we consider to be fixed, but not necessarily static, within the time-frame of interest and with the agents as designated. In fact the environment is usually quite dynamic, but the nature of those dynamics is fixed with respect to the agents. For example, if we exclude petrol producers from an economic simulation then the price of petrol is part of the environment process. It may vary over time but cannot be directly influenced by any agent. If we include the petrol producer in the model as an agent, then the dynamics of petrol prices is no longer in the environment process, but is subject to agent control.

A common problem in Multi-Agent Systems design is the contention issue: Agents cannot share the same space. Space is used in the general sense of mutual exclusion where one resource cannot be occupied by more than one agent. The general definition of the environment used here includes the process which resolves contention issues. The cellular morphogenesis case of defining agents in terms of their boundary with other agents effectively solves the contention issues implicitly.

4 Summary of the GRN Discussion

Discussion began with a brief review of the different types of models that people are using (primarily in gene regulation/development, but also more generally). One distinction that became clear was between people involved in modelling the collective behaviour of a large group of agents, each individually modelled at a relatively simple level, and those modelling the behaviour of a single agent at a higher level of detail, possibly with the long term view of embedding these detailed models into a multi-agent context.

Agents were generally cells in the developmental systems. At a lower level of detail, individual genes or chemical species could be considered agents. Conversely, L-systems models often define segments or modules as the primary unit of definition. Interactions between agents took the form of physical interactions with the corresponding impact on cell shape and position, transfer of chemical signals with effects on cell behaviour, or, in the more low-level models, chemical reactions that could change the identity of species (E.g., dimerisation).

During discussion of individual models, three main issues emerged as being either common features of or points of distinction between the models: the issue of defining and distinguishing between control and constraint in a system; the aim of a modelling project – design or analysis; and the significance of constraints in configuring the default behaviour of a system in the absence of control.

Control and constraint. One issue that arose in discussion was the precise definitions of control and constraint in a multi-agent system. Because both control and constraint influence the behaviour of an agent, and an agent can have an impact on the environment (E.g., by signalling), it wasn't immediately clear where the boundary should be drawn between the two concepts. Candidate definitions were:

Control : the repertoire of actions that an agent could take to influence the behaviour of the system, including other agents and the environment.

Constraint : those aspects of the system which affect the agent, but over which the agent had little or no direct control.

It was noted that in biological evolution, these categories were not necessarily fixed: agents could evolve control over aspects of their behaviour or environment that were previously constrained. For example, while division orientation may be physically specified or stochastic in early multicellular systems, evolution may later produce organisms capable of controlling the orientation of cell division at a genetic level.

Furthermore, at least two possible sources of constraint could be identified. First, there is the external environment: in a developmental system, the laws of physics that constrain the manner in which cells interact in a physical fashion. In addition, there are constraints due to the design of the controller, the aspects of development which it can influence. At one level, this may be as simple as whether an agent does or does not have control over the direction of division (for example); at a deeper level, the choice of representation for an agent's controller – for example, a rule-based system versus a neural network – may impose different sets of biases on the default behaviour.

Modelling objectives: analysis and design. It became apparent that both types of approach to modelling were represented by the instances discussed, and that the requirements for model design could be different depending on the research question. For example, testing specific hypotheses required a means of easily expressing and modifying a hypothesis, whereas attempting to locate systems capable of displaying a desired behaviour requires some sort of automatic search process (E.g., evolutionary algorithms).

Constraints define default behaviour. The discussion also touched on the importance of the fact that constraints may define a default system behaviour that occurs in the absence of any explicit control. For any simulation system, removing the control aspect of the process will provide a potentially useful indication of the null behaviour of the system. For example, in evolution, removing any active selection leaves just passive drift. The exact nature of this drift will be further constrained by the types of mutation operator that are employed (E.g., point mutation, gene duplication, etc).

In the case of the air traffic control problem reviewed by PAL, base system behaviour could be investigated by assuming no air traffic control, and simply modelling the proposed path of each plane together with the Poisson distributed "noise" resulting from early/late transitions between zones of control. This experiment would provide a base level of system performance against which the introduction of a controller could then be compared/evaluated. The aim would then be to design a controller that increased performance above this base level – and minimised the possibility of degrading performance that would emerge from optimising at a local rather than a global scale.

Multiple/dynamic types of agent. One suggestion for how the CM-DRGN model system could be adapted for the purposes of air traffic control was to allow multiple types of agents with distinct controllers. For example, if both air traffic controllers and planes could be simulated as agents, the effect of transferring control between zones and planes in a dynamic fashion could be explored. It was recognised that, while there is nothing in the system that would prevent this being added, it did represent an extension to current functionality.

Software engineering issues. The discussion also raised the more general issue of the design and construction of models from a software engineering perspective. Some of the points that were mentioned included:

- the feasibility of designing systems that were modular – allowing different controller / agent definitions to be substituted.
- the possibility of recognising the flexibility of the control / constraint distinction. For example, enabling the type of evolutionary transition

from a property being a constraint to a being under agent control raised above.

- the degree to which real data could be integrated into models.

5 Future Directions

It was suggested that a future workshop focusing on modelling methodologies, tools and software engineering aspects of modelling would be useful. Such a workshop would include presentations by a number of people both on the tools and models they use as well as the processes that they use to develop these tools. Furthermore, discussions on the potential of transferring tools and methodologies that were initiated at this workshop could be continued and deepened.