Self Managed Systems - A Control Theory Perspective

Richard Taylor and Chris Tofts
Model Based Analysis Group
HP Laboratories
{r.taylor, chris.tofts}@hp.com

Abstract

Self-managed systems are essentially closed loop control systems. When designing such systems we should ensure that they do not allow fundamentally bad properties, too slow convergence, oscillation, chaotic behaviour, stuck modes; just as for any control system. How can we perform these checks in the presence of arbitrary (Turing complete) control functions? We argue that the space of control functions and compositions should be restricted to those with known 'good' properties and demonstrate such a space within cellular automata.

1 Introduction

As systems are growing more complex their control and configuration requirements frequently exceeds the comprehension and ability of their operators. Indeed, there are circumstances, the classic being modern fighter aircraft, where the system is designed *ab initio* to be uncontrollable by a human operator. The desire to move control from the human to the system is rational. Humans have a limited response rate, often make mistakes, are expensive to maintain, and difficult to train. In fact given all of these limitations it is unclear why we have retained humans is so many control loops for so long. The main reason for maintaining humans in management (or control) situations is their flexibility. Any fixed control system will only respond in the way in which it is designed, it cannot be aware of wider goals which can often be in conflict with its current control response. In almost all complex systems there are unforeseen situations to which a closed control system will not have a designed response. It is important to remember that self-managed systems are actually closed loop control systems.

They are by definition, systems that make observations about their current state[1], have the ability to control some (rarely all) aspects of their behaviour with the aim of reaching some 'good' (or even optimal state) and then repeat. The previous statement is simply the description of closed loop control. However, with such control systems there is a clear requirement to demonstrate that the control achieves the designers intent. This is of paramount importance when the intention is to leave the system running without human intervention for any extended period[2]. The rest of this paper will present a quick review of the observations from control theory of a 'good' control system, discuss how this might be demonstrated in a complex self-managed system and conclude with practical examples based on the exploitation of known cellular automata behaviour.

2 Control Theory Perspective

The structure and analysis of the feedback control system (figure 1) has been understood and applied for well over a hundred years, and control systems analysis and design is an integral part of engineering within the modern world.

There are a well known list of properties that are essential for a control system to be considered 'good'. These are the results of centuries of engineering practice and have generally been incorporated as the result of a disaster. A non-exhaustive list will certainly include:

1. convergence on the required solution over the whole input space;

---

[1] This obviously can include environmental inputs

[2] An extended period in this setting can be very short, for a fighter aircraft about 2 seconds

2. adequate responsiveness;

3. stability of solution on stable inputs, no unnecessary oscillations;

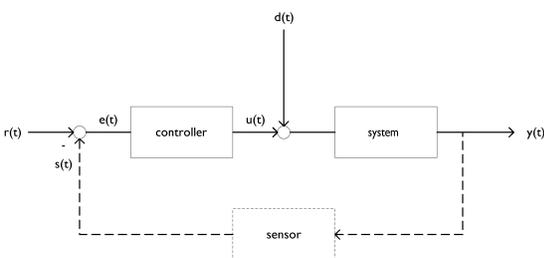4. no snap over on small changes of the inputs, 'chaos'.

Figure 1: There are two forms of control system - open and closed loop. We have a system which we wish to control (with a desired outcome of $y(t)$) based upon some reference trajectory, $r(t)$. In an open loop system, the controller bases the behaviour of the system on an idealised model, taking the reference trajectory and generating a control or actuation $u(t)$. In most systems however, disturbances to the system $(d(t))$ can cause the output or behaviour of the system to deviate significantly from the reference (or desired) trajectory. In addition, innacuracies in the model we use to predict the effect of $u(t)$ on the system can also cause significant deviations. In this case, a feedback loop (shown as a dotted line) is added to the system, enabling measurements taken by a sensor from the controlled system $s(t)$ to be subtracted from the reference trajectory $r(t)$ to provide error $e(t)$. The controller can then make decisions based upon the difference between the current state of the system and it's desired state (plus extensions of such direct measurement, including first and second derivatives). In practice, most complex systems make use of some form of feedback as well as multiple sensors and actuators.

For small numbers of inputs and response controls this is a well understood problem, all-be it when the control system is describable as a set of coupled differential equations. However, the question arises as to what should replace this approach to deriving control system properties in a complex setting, in particular where the control system uses coupling that cannot be easily analysed using differential equations or indeed computations that cannot be described in

this setting. There is a fundamental question as to whether we should employ control systems that cannot be analysed and should we restrict ourselves to mathematically analysable settings? There are further problems of potential higher order neglected factors in original model and of the consequences of approximation. For instance the Lotka-Volterra equations of population dynamics are not stable as a discrete probabilistic system and require the introduction of refuges. A presentation of Lotka-Volterra in this style would look like the following.

$$\Delta prey = bin(prey, birth) - bin(prey*pred, encounter)$$

$$\Delta pred = bin(prey*pred, encounter) - bin(pred, death)$$

Where $bin(n,p)$ is a sample from the binomial distribution with n samplings and p probability of success. With the above set of equations the only *certainly* stable point is $(0,0)$. The cycle point is *expectedly* stable, but is not guaranteed to be so, consequently eventually this system will end up at $(0,0)$.

There are two basic solution approaches to this problem. The first is to verify that any particular control system has the required properties by employing the appropriate mathematics. In the case of control systems whose nodes employ Turing complete levels of computation this is well known to be a difficult problem. The second is to limit the set of measure-response functions in the control system so that they are known to have the required properties in any setting. In the next section we outline how this can be achieved for cellular automata. As an example of control on a cellular automata, consider the problem of maintaining a pattern within the game of Life, given an arbitrary starting position and the ability either to set a limited number of cells, or only to set cells occasionally.

3   Examples - the basic cellular automata

Cellular automata make fascinating abstract, as well as practical example of systems that might be considered to have 'emergent' properties. Cellular automata (CA) are discrete dynamic sys-

tems (space, time and state values), with determinism and with local interaction. A CA is a finite dimensional lattice of sites whose values are restricted to a finite set of integers. The value of any site at any time step is determined as a function of the values of neighbouring sites at the previous time step and itself. In one simple example of Life, space is represented by a uniform mesh of 'processing' cells. Each cell within that space is connected to its eight nearest (surrounding) neighbours. Every cell executes a program that exchanges a token (dead or alive) with these eight neighbours, and then alters its own token based upon the number of 'live' neighbours that surround it. Too many neighbours and the cell 'dies' (over crowding), too few and the cell also dies (loneliness), a critical number and a cell is 'born' (i.e. change token from dead to alive). Such a simple system, replicated across a large continuous surface gives rise to apparently complex behaviour from simple seeds, as groups of cells are born, expand and die. Regular patterns of creation and dispersal appear to occur spontaneously, and in many systems, stable recurring patterns become established. It is the observation that recurrent, stable 'self organised' structures have emerged from apparently trivial 'rules' of computation and communication that have made CA of great interest as one model of parallel and distributed computation. Stability through self organisation on a massive scale has excited many researchers with an interest in scalable, robust parallel systems. Indeed practical and robust implementations of load balancing and fault kernels for operating systems, as well as telecommunications switching systems have been based on simple, well understood automata designs. Observations on the rate of dispersion of system perturbations (through load or fault conditions), combined with the required behaviour of system boundaries, for example, make it possible to construct reliable statements about the upper and lower bounds of system performance and degradation.

Unfortunately, life is not as rosy as one might think. If conventional computing and communication systems are used to construct CA and their resultant behaviour examined, a simple qualitative distinction can be made between them (sometimes referred within the research community as the *class* of the automata)

1. evolve to an homogenous state

2. evolve to simple, separated periodic structures

3. generate 'chaotic' aperiodic patterns

4. generate complex patterns of localised structures

While the forward problem - the description of properties of a given CA such as reversibility, invariants and criticality are well worn with observation based analysis, the reverse problem, the development of a set of techniques that that will generate a rule or a set of rules that can reproduce quantitative observations is primitive in all but the simplest subset of systems (mainly, but not exclusively a subset of class 1 and 2 automata). The bulk of 'successful' reverse analyses make use of a combination of probabilistic search strategies with forward experimental checking. For large systems, especially those of types 3 and 4, the limitations are obvious. Validation of design through simulation and/or construction is not practicable for large (and in the real world potentially expensive in both construction costs and the financial and social implications of error) systems. Successful commercial systems that do rely on CA stability typically live in the class 1 and 2 domains[4], where stability may often (but not always) be predicted analytically.

In practice the implications of these observations are stark. The consequences of introducing 'complex' behaviour into networks of simple, communicating devices can be catastrophic. Domino style failure modes in a number of high profile systems, including large scale telecommunications failure in the US East coast network, and widespread power failure in the US North East can be attributed in part to the failure of the designers to understand for (reasonable simple) communication systems the instabilities of these self organised networks. Modern communications construction (at least for fixed infrastructure) mandates that designs can be demonstrated to dampen the oscillations that invariably follow failure or excessive load conditions.

4   Example - Ant algorithms

As an example of the difficulty of validating particular algorithms see [1, 3, 5, 6]. These papers present verficiation of self organised control strategies in ant colonies for:

- task allocation;

- synchronisation;

- sorting

The work required to capture each of these algorithms is extensive and proving the convergence points for each of these algorithms is, from personal experience, extremely time consuming. Whilst it was possible in each of these cases to prove some properties about the robustness of the stable states. In these examples it was still virtually impossible to prove anything about convergence rates, either in general of from specific cases. Consequently the application of these aglorithms could be advocated in settings where their slow convergnce would be of little economic consequence. However, it would still be difficult to advocate their use in either safety critical or commercially critical settings. An example demonstrating the difficulty of proving convergence rates for a sequential algorithm of this type can be found in [2].

5   Conclusions

Whilst it is tempting, given the 'emulate anything theorem', the presence of Turing complete systems in a closed loop control system is extremely dangerous. It is very difficult to demonstrate the absence of unsatisfactory control behaviour in systems that employ such measure-response functions. Restricted models of computation can give 'good' solutions by construction, even with distributed solutions as the example from cellular automata demonstrated. Given the cost of validating control systems that are constructed *ad hoc* from an arbitrary underlying computation framework, and the risk that such a validation may well be incomplete, should this be a domain of control? Given the need for large numbers of control systems, and the speed with which the demands on the designer vary, it is clearly imperative that we derive as wide a class as possible of compositions and control functions that are inherently 'good'; instead of relying on the skills of a human designer and the observation of successful behaviour on a limited number of simulation exercises.

References

[1] N.R. Franks and C. Tofts, *Doing the Right Thing: Ants, Bees and Naked Mole Rats*, Trends in Evolution and Ecology 7: 346-349, 1992.

[2] Leslie Ann Goldberg, William E. Hart and David Bruce Wilson, *Analysis of a simple learning algorithm: Learning foraging thresholds for lizards*, Journal of Theoretical Biology, 197:361-369, 1999 .

[3] M.J. Hatcher, N.R. Franks and C. Tofts, *The Autosynchronisation of Leptothorax Acevorum(Flavius): Theory, Experiment and Testability* Journal of Theoretical Biology 157: 71-82, 1992.

[4] R Taylor, *Non-linear structural replication in locally connected computing structures*, New Developments in Neural Computing ed. J.G. Taylor and C.L.T. Mannion, (Adam Hilger, 1989)

[5] C. Tofts *Algorithmns for Task Allocation in Ants: A Study of Temporal Polyethism: Theory*, Bulletin of Mathematical Biology, 5:891-918, 1993.

[6] C. Tofts, *Using Process Algebra to Describe Social Insect Behaviour*[3], Transactions on Simulation, 1994.

---

[3]This should really have been titled *using process algebra to prove properties of social insect behaviour*