Management of Complex Systems

R.S.MacKay, Mathematics Institute, University of Warwick, Coventry CV4 7AL, U.K.

1. The research question and its importance

A pervasive challenge of great socio-economic importance is to control or manage complex systems.

I will use the word “control”, though many social scientists feel it carries overtones of authoritarianism and prefer the more benevolent term “management". Nevertheless, mathematicians and engineers prefer the word “control” because “management” sounds too vague.

There is an important distinction to make, however, between two forms of control. In the strong form of control, the objective is to make the *trajectory* of the system follow some desired track or reach some target. In the weak form of control, the objective is to make the *probability distribution for the trajectories* of the system follow some desired track or reach some target.

The distinction is crucial. As soon as the dynamics of a deterministic system shows sensitive dependence on initial conditions (“chaos"), control of a trajectory requires feedback with higher gain than the maximal Lyapunov exponent, which may involve unrealistically high observational power and and unrealistically high actuator response. Similarly, control of a stochastic jump system requires control response time to be shorter than typical waiting times.

In contrast, chaos or stochasticity are good for making the probability distribution for the trajectories of a system relax rapidly to a unique one. Thus the probability distribution may be controlled by much slower observation/actuator feedback than the individual trajectories. Much less detailed observations and controls may suffice for this weaker form of control than the strong one, thus avoiding Orwell's “big brother" nightmare and being more acceptable to social applications.

Thus the aim of this proposal is to develop control theory for probability distributions for trajectories of systems with many interdependent components. I call these probability distributions “space-time phases".

2. The current state of research on the question

Control of trajectories for distributed systems is an established part of engineering. Control of probability distributions is also well developed for compact stochastic systems. But there is a huge need for results combining both aspects: control of *probability distributions* for *distributed systems*.

Barabasi made a splash in a recent issue of Nature with a paper on “Controllability of complex networks” but it was from the point of view of individual trajectories rather than probability distributions.

The simplest context in which to begin is that of probabilistic cellular automata. These consist of a fixed network of units whose states update in parallel in discrete time according to probability distributions that depend on the current state of the whole network but, conditional on the current state, the distributions for different units are independent.

For example, think of the units as being people with two possible states: healthy and infected. Each day a healthy individual has a chance of catching an infection from one of the people with whom they come into contact, which for simplicity we suppose is the same set of people each day; and each day an infected individual has a chance of recovering. If the infectivity rate is low enough and recovery rate high enough then any initial infection will die out in a relatively short time. But if the infectivity rate is high and recovery rate is low then any initial infection is likely to produce an epidemic, leading to a situation of endemic disease. Control theory is well developed in epidemiology, with various options like vaccination to reduce the infectivity rate and movement controls in case of an outbreak. Yet it is insufficiently well developed, witness the recent vast overspend on H1N1 influenza vaccine in several countries, and concerns in the USA about spread of disease in animal livestock (for which a new member of staff in my centre, Michael Tildesley, is working on a recently awarded grant funded by the US Department of Homeland Security).

Similar toy models can be made for the dynamics of opinions and for systemic risk.

In contrast to much of the literature, I do not assume that the network for a probabilistic cellular automaton is a regular lattice, nor that the units are identical, nor that the dynamics is autonomous, nor that the state spaces are finite, so the scope of application is much larger than might first appear.

Under suitable conditions, the operator representing the evolution of probability distributions for the state of the network is an eventual contraction in a suitable metric, and this leads to exponential convergence to a unique probability distribution for the trajectories. The resulting space-time phase depends smoothly on parameters of the model, thus its dependence on feedback control laws can be studied. Given design objectives, one could then seek feasible control laws to bring the statistical behaviour of the trajectories close to the objectives.

Although the above holds for all indecomposable systems with finitely many units and finite state spaces, a more appropriate approach for large systems with some strong interdependence of their units is to consider them as part of an infinite system, just as in equilibrium statistical mechanics. Then the possibility of non-unique space-time phase emerges.

As parameters are varied, the system may jump from one space-time phase to another one that is far away. This is a reflection of the popular notion of "tipping point". Even without parameter variation, the original finite system may best be described as making random transitions between two or more such phases. The ways the set of phases can depend on parameters is a fairly wide open question: some semi-continuity results hold, but there is a great need for an analogue of the bifurcation theory for simple attractors of deterministic dynamical systems, so that we could understand what are the typical qualitative changes in the set of phases.

Going further, could controls be designed to collapse the set of phases into a desired unique one? This connects with another branch of deterministic dynamical systems theory called "ergodic optimisation" in which the aim is to stabilise an invariant probability distribution on an attractor that is not the one naturally chosen by the dynamics.

For some background, see the Technical Support Annex and the references [Mac08, Mac11, DM11], listed in the next section.

3. The proposer's qualifications for the project

Robert MacKay is a Professor of Mathematics at the University of Warwick, Director of its Centre for Complexity Science and Director of Mathematical Interdisciplinary Research at Warwick. He has made important contributions in Nonlinear Dynamics, such as the renormalisation understanding of the transition from order to chaos in classical mechanical systems and the resulting transport process, construction of a mechanical Anosov system (exhibiting the mathematically purest form of chaos that physicists had thought never occurred), results on existence, stability and interaction of "discrete breathers" (spatially localized time-periodic vibrations of networks of oscillators), and construction of examples of spatially extended dynamical systems exhibiting non-unique space-time phase. The last two topics led to his current emphasis on the mathematics of complex systems.

He has developed significant understanding of space-time phases, notably using them to put mathematical flesh on the concept of "emergence" [Mac08], producing examples with non-trivial time-dependence and other features [DM11], and promoting a metric for measuring the distance between the required multivariate probability distributions [Mac11]. He appears to be unique in pioneering this approach, though the tools he uses are the results of many deep mathematicians over the last 40 years.

He is well connected both to strong mathematicians in the relevant areas of probability theory and dynamical systems (e.g. C.Maes, Leuven; A.van Enter, Groningen; colleagues at Warwick like V.Kolokoltsov, R.Kotecky), and to people at the front edge of policy (a selection of such will serve as advisory board for this project, detailed in Section 8).

He has extensive research leadership and management experience, notably establishing and running the Nonlinear Systems Laboratory at Warwick with Rand (86-95), and the Nonlinear Centre in Cambridge (95-00) including providing scientific direction for a King's College Cambridge research centre project on Spatially Extended Dynamics (98-02), winning and coordinating an EC Research and Training Network on "Localisation and energy transfer" (00-04), taking over Directorship of Mathematical Interdisciplinary Research at Warwick in 00, and establishing the "Complexity Complex" in 06 (an association of research activities in Complexity Science at the University of Warwick).

His research has been recognized by the first Stephanos Pnevmatikos International Award for Research on Nonlinear Phenomena (93), a Junior Whitehead prize of the London Mathematical Society (94), election to Fellowship of the Royal Society (00), the Institute of Physics (2000) and the Institute of Mathematics and its Applications (03), and entry to the ISI highly cited authors list under Mathematics (08).

Relevant recent publications

[Mac08] RS MacKay, Nonlinearity in Complexity Science, Nonlinearity 21 (2008) T273-81.

[Mac11] RS MacKay, Robustness of Markov processes on large networks, J Difference Eqns & Applns 17 (2011) 1155-67.

[DM11] M Diakonova, RS MacKay, Mathematical examples of space-time phases, Int J Bif Chaos, in press.

4. The research methodology

The first step is to develop the theory of control of space-time phases for probabilistic cellular automata. The initial approach to be taken is to build on the results of [Mac11] in the regime of exponentially attracting stationary probability, where formulae for the change in stationary probability for a change in parameters are derived. A key question is about the spatio-temporal range of localised control. It is presumably short in this regime; I will seek to develop a tensor product version of [BM] to prove this and to provide realistic estimates.

It will be vital, however, to go beyond this simple regime. Numerical simulation shows that in the regime of non-unique phase, boundary control can have a strong effect, over an infinite range. Similarly, my former postdoctoral fellow Mobilia proved that in an opinion-copying model a zealot could force the whole system into its absorbing state. His model had a regular array of "voters" which can be in one of two possible states, who at each time step choose one of their neighbours at random and copy their current opinion. Into this array he substituted a single "zealot", a voter who never changes their opinion. In one and two dimensions, where the numbers of neighbours is two or four respectively, he showed that with probability one, after some time the whole array ends up voting the same way as the zealot.

On the other hand, Toom proved stability of some phases in some classes of example to small perturbation. By "stability" here I mean that there continues to be a nearby phase. Thus it is important to work out how widely his analysis applies and what happens at its limits. Part of the question is to explore what counts as a small perturbation; for example, even if a zealot looks small in space, its behaviour is very different from a normal voter, so the effect counts as a large change. I have the feeling the situation is closely analogous to that of uniformly hyperbolic attractors of deterministic dynamical systems, where $C^{1+\varepsilon}$-small perturbation makes only a small change to the SRB measure, but high gain control can move it onto a very different invariant measure, like that on a periodic orbit.

To gain some understanding of the range of control, I think that an aggregation procedure to group clusters of units into effective super-units with new effective interactions between them may be useful. This is an analogue of the real-space renormalisation group analysis of equilibrium statistical mechanics. An interesting feature of aggregation for social applications is that even if the individuals have complete preference orders (total orders, for mathematicians) over their actions, groups of individuals have at best an incomplete preference order (partial order, for mathematicians) over their group actions.

Background ideas and tools to achieve this part of the project are outlined in a technical support annex. They include the concepts of Gibbs phase, Dobrushin metric and renormalisation.

An interesting question is how one could test whether a control method for probability distributions is working. The problem is that one can observe only one realisation. Nevertheless, most probability distributions imply that some features are very likely and some other features very unlikely. As a simple example, in a sequence of N independent tosses of an unbiased coin, the probability of the proportion of heads being more than ε from N/2 is about $\exp(-2\varepsilon^2 N)$, miniscule as soon as N is significantly larger than ε^{-2}.

Once the theory for probabilistic cellular automata is well developed, it will be natural to seek to extend it to more realistic classes of system, for example continuous-time stochastic jump processes, systems of mobile units where the strength of interaction depends on distance in physical space, and deterministic systems with sufficiently chaotic dynamics. The extension to continuous-time should be relatively straightforward, though there will be technicalities. That to mobile units is much more challenging: the problem is that one may no longer wish to regard label space as the natural concept of space; instead real space should take over. For deterministic systems with chaotic dynamics, the way forward will be based on [GM].

An alternative approach would be to tackle a concrete problem area, formulate a toy model and seek to derive or simulate some results for it. This would have the advantage of proximity to a motivating problem, and could give great insight, but I think there is a need for fundamental theory of management of complex systems, to generate widely applicable principles.

The chief risk with my approach is that it will turn out to be very hard to say much rigorous. Already Toom's results proving examples with non-unique phase are highly non-trivial and have not been significantly improved upon since they were obtained 30 years ago. I hope that with a team of intelligent and quick postdocs we can achieve significant advances nonetheless.

The second chief risk is that the gap between what can be said mathematically and what is needed for socio-economic policy is too large to bridge with a single project like this. Nevertheless, I think what I am proposing is an essential first step and that its results will be able to feed into larger and longer duration projects, such as the proposed EC flagship FuturICT (Innovative ICT and science for a resilient and sustainable society) which is currently in its pilot year and to whose full proposal I am contributing. It is a 10-year project and is seeking a budget of 1 billion Euros.

The novelty of the approach is to take a dynamical systems view of the question of controlling probability distributions for spatially extended systems. Although the evolution of probability distributions is linear, the effect of interdependence between units is very similar to nonlinearity in deterministic dynamics, as I argued in [Mac08].

References cited

[BM] C Baesens, RS MacKay, Exponential localisation of linear response in networks with exponentially decaying coupling, Nonlinearity 10 (1997) 931-940.

[GM] G Gielis, RS MacKay, Coupled map lattices with phase transition, Nonlinearity 13 (2000) 867–888.

5. The outputs

The project will deliver some high quality publications on mathematical foundations for management of complex systems. It will also deliver some computer demonstrations of the key results, as the most effective way of communicating them to policy makers and the general public.

Success would constitute these publications and demonstrations being recognized as inspirational and their results being implemented in socio-economic contexts.

It would be heart-warming if the project is remembered years from now as the beginning of effective and structured methods for management of complex systems. A lot of complex systems are currently managed with very little understanding. Although many policy makers have great wisdom and there can be substantial collective wisdom, the results of relying on inadequate understanding can be disastrous, e.g. the 2008 financial crisis, the current national debt crises, the riots in England during August 2011, various recent animal and human epidemics, 9/11, the invasion of Iraq, the war in Afghanistan, Fukushima. There are enormous socio-economic challenges ahead, such as food supply, climate, demography, financial regulation, electricity pricing systems, social unrest, reorganisation of health services, epidemics, and management of the economy. It is essential to base their management on sound foundations that take into account factors like systemic risk and the differences between individuals' incentives and what is good for society as a whole.

6. Justification for the resources requested

To allow MacKay to concentrate on this research project will require relieving him from administrative responsibilities. The model in the UK for research proposals is for investigators to bid to the funder for the fraction of their salary that the project will require. A realistic requirement is 50% of MacKay's salary for 2 years. With this time, he would be able to concentrate on formulating the ideas and research programme in depth. It is only he that can lead such a project and it is essential that his time be liberated for this purpose.

The project requires two 2-year postdoctoral research assistants. Ideally, one will be expert in mathematical analysis of probabilistic processes, and the other in numerical simulation and computational implementation of the theory to be developed.

It requires computers for the postdocs to use in their research. Each of them will need a desktop and a laptop. One of the desktops needs to have significant computational power. We may also need to pay for access to the University of Warwick high performance computing facility.

It requires travel funds to (i) attend relevant conferences to promote the results, gain feedback and keep up to date with cognate research, (ii) interact with relevant potential end-users, particularly the Sciteb consortium on financial regulation which is based around their offices in London and Boston, and a promising collaboration with National Grid (UK and New England) involving Michael Caramanis of Boston University, (iii) invite relevant experts so that the project would benefit from their knowledge (e.g. Christian Maes, Leuven BE; Aernout van Enter, Groningen NL).

7. Other sources of support

I have access to a large pool of research studentships, in particular ten 4-year MSc/PhD studentships per year for Sept 2011, 12 and 13 in Complexity Science, seven 2-year MSc studentships per year for Sept 2011, 12, 13 and 14 in Complex Systems Science, and six 3.5-year PhD studentships per year indefinitely in Mathematics or Interdisciplinary Mathematics. From these sources I could expect to attract one or two a year. Such research students would play an important role in testing ideas, working out details and moving the project along.

The project would interface nicely with a 4-week session at the Newton Institute, Cambridge UK on “Control and Management of global complex systems” proposed for

September 2013. The session would make a particularly appropriate context in which to promote our results.

One could consider this proposal suitable for submission to (or co-funding by) a national research council, such as the EPSRC in the UK. But they are in the process of conducting a major review of strategic priorities and it is particularly difficult to pitch a bid to them at this time, for an adventurous proposal that crosses traditional boundaries. On the other hand, a successful project funded now by the Sloan Foundation would put me in an ideal position to bid for UK funding in the future when the landscape has settled into its new form.

One might consider applying to the European Commission, but I don't think the project would fit either of their main funding streams. One is the Future Emerging Technologies scheme, but the grants have to involve collaboration between at least three European countries, whereas the nature of this project requires effort from a focused dedicated team in a single location (i.e. Warwick) if significant progress is to be made. The other is the European Research Council fellowships, which could well be suitable but the results from this initial Sloan Foundation project when obtained, digested and published, would provide a major contribution to the chances of success of such an application.

8. Advisory Board

The project will be overseen by a high-powered advisory board, consisting of people close to the front line in various areas of socio-economic policy. Those whose agreement I have sought are: N.Beale*, Executive Director of Sciteb, London and Boston, runs a consortium of banks, financial regulators, academics and insurance companies, with whom

MacKay currently co-supervises a PhD student on design of new financial regulation; N.Winser, Executive Director for Transmission at National Grid, the major electricity and gas distribution network in the UK and also a major player in New England; M.Caramanis, Professor of Electrical Engineering, Boston University, who devises electricity pricing systems; D.MacKay, Chief Scientific Advisor to the UK Department of Energy and Climate Change; N.Gaunt*, UK National Health Service Institute for Innovation and Improvement; A.Cohen*, Professor of Health Policy and Management in the Boston University School of Management and Executive Director of the Boston University Health Policy Institute; M.Keeling*, Epidemiology group at Warwick, who advises UK government on epidemics; Lord May*, former President of the Royal Society and former Chief Scientific Advisor to the UK government (* indicates they have already agreed).

They will confront MacKay with what is required for socio-economic policy, and bring to his attention any opportunities to contribute to concrete solutions.