CSIRO Centre for Complex Systems Science

# Complex Systems Science – Can We Do Better Than Agent Based Modeling?

### Glynn Rogers 25/7/03

Faced with the rather daunting complexity of large scale multi-component systems, researchers have increasingly adopted the approach of treating the system components as semiautonomous agents acting according to a relatively small set of rules. Computer models enable numerical experimentation with different rule sets in an attempt to finding those general characteristics which fall into the category of self organizing or emergent behaviour. Artificial ant colonies are a favoured example.

The danger I see with this agent based approach is being seduced into playing amusing trial and error games with rule sets in the hope that something will pop out. Great fun - but is it science? Unless there's a conspiracy of silence here and all the answers really are there in Wolfram's book, it seems to me that, in spite of the intensive effort that has gone into the relatively simple world of Cellular Automata, there is no evidence of a substantial 'natural' theoretical structure to the rule sets. In short, there is nothing to find. I increasingly suspect that a concentration on rules will lead to the same type of blind alley that trapped AI (remember the Fifth Generation Computing debacle).

On the other hand the extensive development of the Synergetics concept by Herman Haken [1] and others has provided a wealth of theoretical insight into the self organization of continuous nonlinear systems. So it seems to me that we should not abandon these more traditional differential equation based methods in the belief that rule based agent systems are somehow more realistic.

The underlying problem, I suggest, is that rule based systems inhabit the world of finite state automata. Whereas the differential equations of continuum mechanics are based on a vector space richly endowed with topological and metrical properties, the state spaces of finite state machines are intellectually poverty stricken. They are in fact little more than an unstructured sets. There simply isn't enough underlying structure to support the intrinsic relationships that give rise to the 'emergent' properties of nonlinear dynamical systems. This suggests that, in agent based models, such properties would have to be installed by design, on a case by case basis, by hand crafting the rule sets. This may, like AI, lead to useful practical tools, but it's technology not science.

Let me present a couple of examples where agent based systems might appear, at first glance, to provide the appropriate model but where a closer examination suggests the opposite.

An ecosystem would seem superficially to be composed of a set of, at least semiautonomous agents running around interacting with each other and so generating the properties of the ecosystem. These interactions might seem to be governed by a finite set of simple rules so that, for example, the stability of the system would be a direct consequence of the rules – there are stable states in the state space and unstable states. Determining stability is then just a matter of analyzing the rule sets. Well maybe – but how to find general stability conditions? There doesn't seem to be even the hint of an answer. On the other hand a dynamical systems approach allows considerable progress to be made.

It is well known that Bob May's valuable early work in mathematical ecology, based on the Lotka Volterra equations, ran into the somewhat embarrassing difficulty that his important prediction of deceasing ecosystem stability with increasing diversity was simply wrong. As I understand it, the empirical evidence is now clear – the opposite is true – stability increases with diversity. Some quite recent work sheds considerable light on what went wrong – May's assumption that the 'community matrix' was essential structureless and his consequent adoption of a statistical method for determining the matrix components. The interesting paper by Neutel et al in Science [2] last year demonstrates, using empirical data, that trophic loops, particularly long loops, add considerable deterministic structure to the matrix. It also develops a mathematical argument to show how this structure accounts for system stability although this argument may be somewhat misdirected

A recent Russian paper by Pykh [3] uses the Lyapunov functional approach to investigate the stability of systems based on the Lotka-Volterra equations and again demonstrates the importance of structure although, in this case, in the interaction matrix. The fact that it is the interaction matrix rather than the community matrix underscores the significance of the approach. It deals with the question of stability from a global perspective without needing the detailed knowledge of the dynamics or the specifics of the equilibrium state that is required by the analysis of Neutel et al.

Indeed, applying Pykh's formulation of the original Volterra criterion to the three level system used as an example in Neutel et al shows immediately that stability depends on equality of the weights of the two long loops formed by the top predator preying on both the immediate as well as the bottom level species. However, as distinct from Neutel et al, the loop weights are calculated from the interaction rather than the community matrix ie without reference to the specific equilibrium state. Furthermore, this equality of loop weights leads directly to relationship between the three 'conversion efficiencies' ie "the efficiency with which food is converted into predator biomass".

Now let me skate out onto the really thin ice. The Lyapunov functional approach to the three level example suggests two elements of speculation. Firstly, calculating loop weights in terms of the interaction matrix may provide a more fundamental view of ecosystem stability than the community matrix approach of Neutel et al. This would have the potential advantage of providing stability relationships without having to determine a particular state and perform the standard linearisation of the equations at that state. (Oops - what's that cracking sound?). Secondly, Neutel et al have confused cause and effect in seeing stability requirements as imposing constraints on the 'intraspecific' (diagonal) terms of the community matrix. In the Lyapunov functional approach, stability is determined by the off diagonal terms of the interaction matrix only. This suggests that the intraspecific parameters (eg non predator caused death rate) determine the actual population levels at the stable point but not the stability itself. The stability analysis is simplified by the removal of a class of parameters (and an arbitrary constant that Neutel et al had to introduce). Ok – the cracking sounds are now deafening and I'd better retreat before I go plunging in completely out of my depth!

The essential point I am trying to make here is that this approach provides a 'generalised energy surface' picture of the system with the fixed points at the maxima or minima of the surface. This, in turn, may lead to insights, or at least directions of inquiry, into the global structure of the system which would be difficult to see down amongst the trees. Furthermore it generalises the notion of fitness functions, links to engineering optimization via the concept of utility functions and provides a conceptual and mathematical link to molecular science, in all its glory, via the Potential Energy Surface concept.

Which brings me to the second example – protein folding. Having read a paper on cytoplasmic inheritance, I got into an extensive discussion on this with a couple of cell biologists and a protein chemist from Sydney Uni at a social gathering. The 'straw man' I'm setting up here is the simplistic view that rule based agent systems are powerful because genomics is a rule based system that has sufficient power to produce the biological world. Well rule based maybe but, of course, all the rules do is specify the manufacture of amino acid chains - the next step, the production of proteins, is the big one and that appears to involve a great deal of physics and cellular machinery.

So the question was – does the genome contain sufficient information to specify both the amino acid and the manufacturing plant (analogously, the computer as well as the programme) or is there something else going on here (the subject of the paper). The eventual answer was that the genome was sufficient but it took a while to get there partly because the assembled experts found it difficult to free themselves from the background assumption that the cellular machinery was just there. Even then the conclusion was tentative so clearly there is still a lot of uncertainty about what is really going on there. One thing does seem clear – the 'programme' is not written in some general high level language but in the very specialised machine code of the cellular 'computer' and may well depend on some basic hard-wired functions.

Well, what else might be going on here. Some quick reading revealed that there is a lot of physics being invoked to explain protein folding – not just Potential Energy Surface ideas but Potential Energy Surfaces with very specific funnel structures. That is, we very quickly leave the rules behind and get back to more of a dynamical systems context capturing thermodynamics on the way.

So where does this leave genetic algorithms? I suspect very much in the same situation as neural networks (as distinct from neuronal networks) - simple caricatures of a much more complex reality whose main value turns out to be as concrete implementations of more abstract mathematical operations that have little to do with the original question. In the genetic algorithm case this appears to be nonlinear optimization. Then the really interesting questions are not to do with the nature of the algorithm but with the origin and structure of the functions (functionals?) that are the basis of the optimization.

So it seems that which ever way I look at what is at least an important class of complex systems, I eventually see a process which can be understood as seeking a maximum or minimum of some energy surface, a fitness function, a utility function whatever – a Lyapunov functional. So my central argument is that at the heart of complex systems science is the identification and understanding of these Lyapunov functionals.

Now I sense at this point a rising chorus of dissent asserting that I have missed one of the main points of complex systems science, namely adaptation. I have sketched a picture of a largely static universe where the only dynamics are associated with perturbations from equilibrium except in the case of non-static attractors. Even there the only potential interest is in quasi-periodic or outright chaotic attractors and I agree that they don't appear to help greatly with many of the important questions.

Well, there's more to the story – in fact I think it's the most interesting, if speculative part – because there is a higher level of structure involved in the relationship between the Lyapunov 'energy' function and the parameters which control it. An example that helps illuminate the idea is another interesting paper in Science - this time by Wales [4] in 2001 where he uses Catastrophe Theory to describe the evolution of the molecular Potential Energy Surface with variation in control variables such as the interaction range between atoms. This is a highly structured process in which local minima, and the saddle points between them, emerge from degenerate critical points with very simple geometrical relationships between them. One consequence is the orderliness of transitions between configurations induced by varying the control parameters.

In fact this issue lies at the core of Haken's Synergetics and is a particular case of his Slaving Principle which demonstrates that phase transitions take place in a very limited dimensional subspace of the overall phase space and are controlled by a very few 'order parameters' with the remaining degrees of freedom 'slaved' to them. Indeed this subspace is spanned by eigenvectors of the linearisation of the dynamical system at the attractor which is approaching the transition. These eigenvectors are those whose corresponding eigenvalues have real parts that go to zero at the transition – those which, in a sense, cause the transition.

The essential point here is that, particularly from a design perspective, this strong high level structure provides considerable scope for exerting macroscopic influence on the system without having to control the microscopic details.

At the risk of encountering more thin ice, let me go back to the mathematical ecology example. While the Lyapunov functional method provides the stability analysis and the general relationships that need to be satisfied to ensure stability, understanding which parameters are critical for maintaining a particular stable state appears to require the Synergetics approach. Indeed, attempting to move a system to a particular state would seem to require the type of analysis that Wales employs, particularly if the new state is to be reached via a phase transition. In this picture, a system adapting to a change in external parameters, while maintaining stability, is changing its energy surface so that the minimum occurs at a point where the stability criteria can be satisfied with the new parameter values. To do this without encountering phase transitions, this move needs to be accomplished without sending eigenvalues to zero.

Understanding these mechanisms is, I suggest, an important issue for the engineering design of complex systems.

Now to return to the notions of self organization and emergence – can this approach provide and adequate account of them? Is it possible to explain how self organization and emergence arise? We need to proceed with some ontological caution here and not postulation the existence of new conceptual entities unnecessarily because, in this sense, it may be that self organisation and emergence are illusions - that there is nothing to explain. Stable states do not emerge in any real sense - they simply are the minimum energy states on some generalised energy surface and are arrived at via the normal dynamics of the system.

That this seems somehow unsatisfactory is, I suggest, simply because we approach the problem from the standpoint of statistical mechanics and expect to experience only ensemble averages, being surprised when we do not. We should heed the the warning of Carlson and Doyle in the context of HOT [5] that in areas such as biology, ecology, and engineering systems, we deal with states that are highly improbable from a purely statistical viewpoint. These states are the result of powerful optimisation mechanisms which run counter to the Second Law, rendering ensemble based methods inappropriate. The task then is to explain optimization and adaptation in the context of generalized energy surfaces.

Am I arguing that Complex Systems Science should be confined to the continuum mechanics arena? No, certainly not, but I am arguing that it provides a substantial intellectual base from which to develop a wider perspective whereas agent based modeling does not. That this wider perspective might eventually encompass agent based modeling concepts is suggested by a couple of lines of research. One is the investigation of hybrid continuous – discrete systems by the control community [6]. At a deeper level, the theoretical developments discussed in Section 5 of [7] promise a substantial conceptual connection between the two.

So the core of my argument is that a combination of the Lyapunov functional and Synergetics approaches is a key ingredient of Complex Systems Science. It promises a conceptual and mathematical framework within which to unify many of the problem areas that fall within Complex Systems Science – indeed to provide that underlying unification which I see as an essential attribute of Complex Systems Science. Developing this framework is a meta-level task which draws from and, in turn, illuminates the specific problem areas addressed by the Complex Systems Science Centre and should be part of the Centre’s activity.

[1] Herman Haken, “Advanced Synergetics”, Springer Verlag.

[2] Anje-Margriet Neutel et al, "Stability in Real Food Webs: Weak Links in Long Loops", Science Vol. 296, 10 May 2002, p1120. (also their ref 5)

[3] Yu. A. Pykh, “Lyapunov Functions for Lotka-Volterra Systems: An Overview and Problems”, Proceedings of the Fifth IFAC Symposium, “Nonlinear Control Systems”, 2001, p1655.

[4] David Wales, "a Microscopic Basis for the Global Appearance of Energy Landscapes", Science Vol 293 14 Sept 2001 p2067.

[5] J. M. Carlson and John Doyle, “Highly Optimised Tolerance: A Mechanism for Power Laws in Designed Systems”, Physical Review E, Vol. 60, No. 2, August 1999, p1412.

[6] Proc. IEEE, Special Issue on Hybrid Systems: Theory and Applications, Vol. 88, No. 7, July 2000.

[7] John C Doyle et al, “Robustness and the Internet: Theoretical Foundations” Draft,