Opening Stable Doors: Complexity and Stability in Nonlinear Systems

Inman Harvey

Evolutionary and Adaptive Systems Group
University of Sussex
inmanh@gmail.com

Abstract

Generic complex systems of many interacting parts can model both natural and artificial systems, and the conditions for their stability are of interest. Two influential papers (Gardner and Ashby, 1970; May, 1972) laid down a mathematical framework suggesting that, without some specific constraints on the interactions, such systems are very likely to be unstable as they increase in size and connectance. We draw attention to a programming error in the first paper and to flaws and omissions in reasoning in the second that discredit such conclusions when applied to nonlinear systems. With nonlinearity the connectance strength of an influence of any one variable upon any other will vary according to context, which May's analysis does not address. Further, in nonlinear systems there can be many equilibria, and global instability requires every relevant local equilibrium to be unstable; neglecting this invalidates the conclusions. We discuss the relevance of ambiguous circuits (Thomas and D'Ari, 1990) and consider simple classes of nonlinear functions that generate these, including the hat shaped viability functions that generate homeostasis in Daisyworld models. We demonstrate that the May results are unreliable even for the simplest families of nonlinear systems that model common biological, physical or artificial systems.

Introduction

An influential early paper (Gardner and Ashby, 1970) used computer simulations to assess the probability that a large system of interacting component parts that has been assembled at random, or has grown haphazardly, will be stable or unstable. They considered systems where the interactions between parts were linear, and looked at how the expectation of stability changed as the number of variables increased. This was a theoretical study, to be motivated by its possible application to both biological and man-made systems: brains (real or artificial), planetary climate systems, social or financial systems, ecosystems. The conclusion was the suggestion that all such large (random or haphazard) complex linear dynamic systems may be expected to show the property of being stable up to some critical, fairly small, level of connections; but above that phase transition value they are overwhelmingly likely to be unstable. From this it could be deduced that if one observed large complex linear systems that were indeed stable, there must be something exceptional and non-random about the way that the parts were connected.

The influence of this work stems primarily from its extension and development by Robert May, and the subsequent proliferation of a wide body of research in this area. He replicated a version of the results analytically rather than computationally (May, 1972), and claimed that their validity extended beyond the linear systems of Gardner and Ashby (hereafter: G&A) to systems "which in general may obey some quite nonlinear set of first-order differential equations". May's interest mainly focused on ecological systems, and a subsequent book (May, 1973) largely set the agenda for discussion of the relationship between complexity and stability in ecosystems ever since.

Before this work there was a common perception that the more diverse was the range of species in an ecosystem, the more robust and resilient to perturbations that system would be; and further, it was often assumed that this may well be due to some underlying law of large numbers that could apply very generally across all sorts of systems with many interacting components. But the work of G&A and May, apparently using very minimal mathematical assumptions, appeared to suggest that the opposite was true – at least, in the absence of further specific constraints. So subsequent argument and analysis have tended to focus on what further constraints, what limitations on the number, sign and size of interspecies interactions, might be necessary in order to make it likely that a complex ecosystem was stable. The mathematics, it has largely been assumed, is relatively simple and correct. Hence if we want to explain the existence of complex stable systems, it looks like we need to add further assumptions.

In this paper, we shall demonstrate that the reasoning within these two primary sources (Gardner and Ashby, 1970, May, 1972) is partially invalidated through omission and errors, and in particular should not be generalised in this way to nonlinear systems. Firstly, we draw attention to a programming error in the G&A paper, which has been noted previously (Solow et al., 1999). Secondly, we point out that May's attempted extension to nonlinear systems fails to specify the distribution from which the relevant connection strengths are drawn.

Thirdly, and fatally to May's reasoning, we point out a flaw where he claims to go beyond the purely linear systems of G&A towards a more general set of nonlinear systems. May considers local stability at just a single fixed point in the space of possible values for the system, a point that makes sense when considering linear systems with negative self-interactions. Unfortunately, when we move on to nonlinear systems there can be a large (and in some circumstances unlimited) number of points of potential stability to consider.

Global instability would require local instability at every one of those points. Hence the probability of global stability will be underestimated if one just considers local stability at a single fixed point, as May does.

These various criticisms are, as far as we are aware, all drawn together here for the first time. We present examples demonstrating that it is not merely exotic nonlinear functions that raise these issues. Even simple monotonic nonlinear functions such as sigmoids, or the simplest piecewise linear functions with a single change of slope, are sufficient to invalidate the reasoning. Hat shaped viability functions, as used in Daisyworld models, are discussed and it is shown how stability arises independently of the sign of the opposing effect. The 'ambiguous circuits' so produced are related to the multistationarity analysis of Thomas (Thomas and Kaufman, 2001a, 2001b).

These flaws in the two foundational papers by Gardner and Ashby (1970) and May (1972) suggest that a radical reappraisal is needed in the mathematical foundations of a substantial body of work that has built up over some 40 years. Rather than seeking a route to stability by adding further constraints to these abstract models, we need to open the doors to those possible locations of stability that have until now, through error or omission, been excluded. The significance goes beyond ecosystem theory to the study of all kinds of natural and artificial systems with complex nonlinear interactions, including financial systems (Haldane and May, 2011).

Gardner and Ashby on Linear Systems

Their short paper, a Letter of less than one page in Nature (Gardner and Ashby, 1970), was an early example of a computer simulation, using a Monte Carlo approach. They considered a very simplified formal model of any large system of many interacting parts. This could be traffic at an airport, or the neurons in a human brain. They asked the question: supposing one did not know all the details of the interactions between component parts, but modelled these as coming from some random distribution that gave the signs and sizes of these interactions, then what was the chance that such a large system will be stable? Although in the real world most of these large systems, perhaps biological or social, will be grossly nonlinear, they explicitly restricted themselves to considering only systems with linear interactions, as a first step towards a more general treatment. They were interested only in fixed point equilibria.

The model had n component parts. The intention was to investigate how the generic properties of such systems varied as n increases. The instantaneous state of the system can be expressed by a vector $\mathbf{x}$, where x_i represents the current value of the ith variable. In the very general case of nonlinear systems we would have, with different nonlinear functions for each i:

$$dx_i \big/ dt = NonLinFn_i(x_1, x_2, \ldots, x_n)$$

However in this restricted linear case this simplifies to a weighted sum of the current values of all the variables:

$$dx_i \big/ dt = \sum_j a_{ji} x_j \qquad \text{Eqns 1}$$

Because this is a linear system, there is a unique equilibrium point where for all i $dx_i/dt=0$. The issue will be: what is the probability that this unique equilibrium is stable, given the distribution from which the weights a_{ji} in the connection matrix $\boldsymbol{A}$ are drawn. $\boldsymbol{A}$ is the Jacobian matrix of the first-order partial derivatives, and in this case of a linear system these terms are all scalars, of fixed size and sign; when later we move on to nonlinear systems, these terms will be variable in both size and sign.

G&A chose to make this a partially connected system, with a proportion C of the off-diagonal weights being nonzero. These nonzero weights were distributed evenly between -1.0 and +1.0. Further, they ensured that all the weights a_{ii} in the main diagonal of the connection matrix (self-connections) were negative. They distributed these evenly between -1.0 and -0.1; in May's version that followed, May set all these to -1.0.

G&A are thus discussing a family of linear feedback systems, parameterised by these two values: n, the number of component parts, and C, the connectance or the proportion of possible interactions between parts that are non-zero. For any given values of n and C, their Monte Carlo approach involved testing many cases of such systems, with the connection weights drawn from the appropriate distributions, and finding out through computation what proportion of the systems were stable at their unique equilibrium point. For low values of the connectance, where the interactions are dominated by the stipulated negative values of self-connections, the probability of stability was close to 100% for all values of n tested. But as the connectance C increased, the probability of stability fell away. Using the limited computational facilities of their day (Gardner and Ashby, 1970), they tested examples where n equals 4, 7 or 10. Their conclusion, illustrated by a figure, was that as n increases the relationship between connectance and stability changes from (for n=4) a smooth falling away of probability of stability as connectance increases towards a step function for values of n of 10 or more. Their figure (partly replicated by the thin lines in Figure 1 here) suggests that for n=10 this phase transition from "almost certainly stable" to "almost certainly unstable" occurs at or around a connected value of 13%, C=0.13.

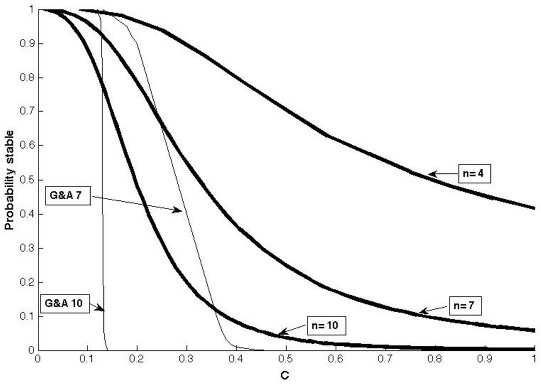

Figure 1: Thick lines give the correct results for G&A's examples, for n = 10,7,4 from left to right. Diagonal terms a_{ii} drawn from [-0.1,-1.0]; a proportion C of off-diagonal terms a_{ji} ($i \neq j$) drawn from [-1.0,1.0], with the remainder zero. Thin lines copy the incorrect results that G&A showed for n = 10,7 (Gardner and Ashby, 1970).

The Programming Error

When we replicated their method[1] our results were similar for n=4, but noticeably different for n=7 or 10. The difference, shown in Figure 1, is striking, and in particular eliminates the sharp nature of the phase transition claimed for n=10. This was their main result, and May claimed on the basis of his analytical treatment (May, 1972) to have corroborated this: “The sharp transition from stability to instability, which was the essential feature of their [G&A] paper is confirmed”. Having failed to replicate this sharp transition, the first step was to check whether we had misinterpreted their methods. But eventually a colleague discovered a rarely cited 1999 reference (Solow et al., 1999) pointing out the same problem, with results agreeing with our own presented here. They attributed the problem to some unknown programming error in G&A’s code. Further, they comment that this nullifies one of May’s conclusions where he had assumed that the G&A phase transition was a real phenomenon. Correction of this programming error does not alter the conclusion that as n increases and C increases the probability of stability goes down; it does alter the conclusion that for values of n above some fairly small value the relationship between stability and connectance turns into something close to a step function.

For the purposes of this paper, this programming error is the least important of the errors and omissions to be discussed. Nevertheless, it is of note that it took nearly 30 years until this error was pointed out in print.

May’s analysis: linear systems

Whereas G&A explicitly limited themselves to the consideration of linear systems “merely as a first step towards a more general treatment” (Gardner and Ashby, 1970), May claims to be considering systems “which in general may obey some quite nonlinear set of first-order differential equations.” (May, 1972). His method is to focus on the behaviour of such nonlinear equations around “the equilibrium point”. Through making a Taylor expansion and ignoring the higher-order terms one can consider this locally as a linear system. Thereafter, May goes on to analyse the same kind of linear system as G&A, while still claiming that it generalises to nonlinear systems.

Insofar as May’s analysis is restricted to the linear version, he tackled analytically much the same class of systems that G&A had tackled computationally. To be precise, this was a slight variant with qualitatively the same behaviour; in place of just C or connectance he considers a term α that is the mean square value of the distribution of all off-diagonal elements, described as expressing the average interaction “strength” (measured on a scale that rates the negative self-feedbacks on the diagonal of the matrix at -1). May’s results were broadly similar, claiming that the central feature of the results for large systems is “the very sharp transition” from stable to unstable behaviour above a critical value that “accords with Gardner and Ashby’s conjecture”. As we have pointed out above, in fact the transition is not as sharp as G&A indicated; however the analytical results do agree with a correctly coded computational Monte Carlo approach. The influential take-home message from both the computational and analytical results has been: in any such system of many interacting parts, as soon as the average interaction strength (interactions between different component parts) rises above some small value, the probability that such a system will be stable drops to near zero. This limitation on stability becomes worse as n, the number of parts, increases. In the context of ecosystems, such a result challenges the commonly held assumption that the more diverse an ecosystem is, the better it is able to remain stable in the face of perturbations.

Picturing Stability

In preparation for understanding nonlinear systems, we first present in some detail a sketch of how to analyse and visualise stability in linear systems. This is basic textbook material, but that is the level of the flaws that we are going to exhibit when we move on later to nonlinear systems. For a simple system of two variables, we can graphically sketch the nullclines (where dx/dt=0 and dy/dt=0) and, by plotting the consequences of perturbations, analyse for stability. We start with two linear examples, Equations 2 and 3, sketched and analysed in Figures 2 and 3:

$$dx/dt = -x + y/2 + 1 \qquad dy/dt = x/2 - y + 1 \qquad \text{Eqns 2}$$

$$\begin{pmatrix} -1 & 0.5 \\ 0.5 & -1 \end{pmatrix}$$

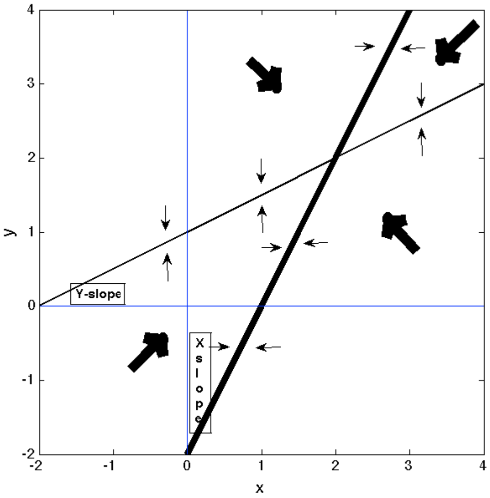

Figure 2: The nullclines for Eqns 2. Thick line for dx/dt=0, with horizontal small arrows indicating responses to x-perturbations. Thin line for dy/dt=0, with vertical small arrows for response to y-perturbations. The heavy arrows sum these responses, giving a stable equilibrium at the intersection (2,2).

---

[1] Matlab code at www.informatics.sussex.ac.uk/users/inmanh/stable

The Jacobian matrix restates the fact that the self-connections are -1, and the cross-interactions are 0.5. These latter correspond to *tan*(X-slope) and *tan*(Y-slope) as those angles are indicated in the figure. In contrast, consider this example with the same nullclines, though swapped around:

$$dx/dt = -x + 2y - 2 \qquad dy/dt = 2x - y - 2 \qquad \text{Eqns 3}$$

$$\begin{pmatrix} -1 & 2 \\ 2 & -1 \end{pmatrix}$$

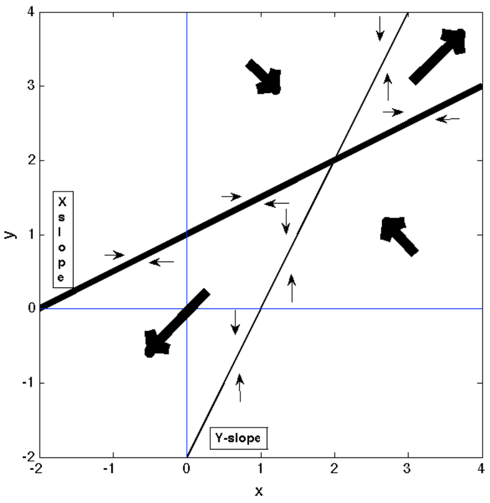

Figure 3: Nullclines for Eqns 3. Thick line (*dx/dt*=0) has now swapped places with thin line (*dy/dt*=0). Response arrows also differ from Figure 2, equilibrium at (2,2) is now unstable.

Here we can see that the equilibrium is unstable. We can note that the connection strengths, the off-diagonal terms in the matrix, also here *tan*(X-slope) and *tan*(Y-slope), are now 2 rather than 0.5. So anecdotally this conforms to a general picture that larger connection strengths are more conducive to instability; though we should also note that if these connection strengths had been of opposite sign, of whatever strength, stability would have been the consequence. We can now see how this analysis extends to the nonlinear picture.

May's analysis: nonlinear systems

May (1972) does not lay down any constraints on the very general class of nonlinear systems, bar implicitly that they should be smooth and differentiable so that they can be approximated by a linear system around any equilibrium point under investigation. For simplicity we start by restricting ourselves to systems of the form:

$$dx_i/dt = \sum_j Fn_{ji}(x_j) \qquad \text{Eqns 4}$$

and further restrict the classes of functions to just linear and sigmoid. We can demonstrate our essential points with a two-variable system:

$$dx/dt = -x + y \qquad dy/dt = 20/(1 + e^{(10-x)}) \qquad \text{Eqns 5}$$

The numbers have been chosen to demonstrate that there are now several equilibria, as demonstrated by the intersection of nullclines in Figure 4. We can see that two of these equilibria conform to the pattern of Figure 2 (and are stable), whereas the central equilibrium conforms to the pattern of Figure 3 (and is unstable).

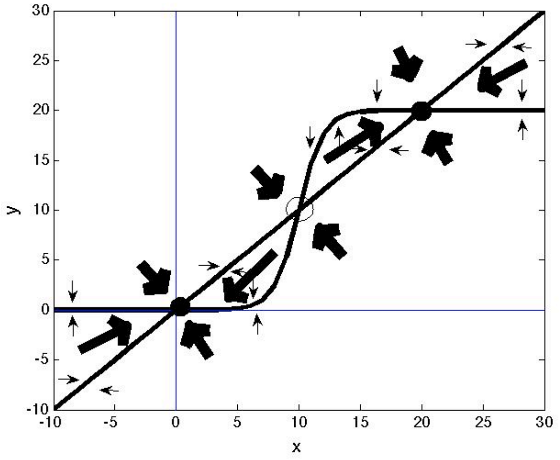

Figure 4: Nullclines for Eqns 5. Three equilibria are circled, the central one (open circle) is unstable, the other two (closed circles) are stable.

Which Distribution of Connection Strengths?

With the aid of this sketch we can make the trivial observations that the addition of even a single simple monotonic nonlinear function, such as this sigmoid, means that there can be several equilibrium points and that in general the slope of the nonlinear function, related to connection strength, varies from one equilibrium to another.

May wishes to extend the conclusions of the linear analysis – where the probability of stability depends on the nature of the distribution from which connection strengths are drawn – to a nonlinear case with an undefined distribution of nonlinear functions. But this could only be done systematically by firstly specifying the distribution of parameters in the specified class or ensemble of nonlinear functions; and secondly, determining where on such functions one measures the slope. Since there can be several equilibria, this gives several possible values for the connection strength. Given that low connection strengths tend to be conducive to stability in the linear case, it can be noted that many nonlinear functions including these sigmoids have regions where the slope is low.

May would need to do all this to complete his project of generalizing to nonlinear systems. One could then in principle find the distribution of connection strengths over all the equilibria, and perhaps give an estimate of the proportions of these that were stable or unstable. But then further work would need to be done to assess whether the system as a whole was stable or not, since that is a global property.

Global versus Local Stability

For a system to be globally unstable, every single equilibrium point must be unstable. But for stability it is sufficient for there to be just a single stable equilibrium point within the region of interest. In the case of linear systems, global stability and local stability are one and the same, but May's analysis fails to take account of the fact that nonlinear systems are different. Even if we had an estimate of the probability of any specific equilibrium point being stable, this may well be a gross under-estimate of the chance of there being stability somewhere within the system as a whole.

In some classes of nonlinear functions, e.g. sinusoidal, there is the potential for an unlimited number of intersections with a straight line, corresponding to an unbounded number of equilibria in the two-variable system. For well-behaved curves, as we can see in Figure 4, stable and unstable equilibria alternate so that as long as we have more than one equilibrium we are guaranteed a stable one.

For simplicity, in order to get the main points across, the examples above are restricted to systems of just two variables. Extending this to an n-variable system with $n>2$ requires more analysis. But in summary, the May analysis simply ignores these crucial differences between nonlinear and linear systems, and in doing so typically underestimates, perhaps grossly, the probability of stability in nonlinear systems.

Ambiguous Circuits

Thomas and colleagues (Thomas and D'Ari, 1990; Thomas and Kaufman, 2001a, 2001b) discuss the roles of positive and negative feedback in nonlinear biological systems. It so happens that their main interest is in the positive feedback circuits that lead to multistationarity, or switching, in genetic regulatory circuits. Nevertheless, much of their analysis can be applied to investigating issues of negative feedback circuits leading to homeostasis or stability. As with May, they are considering a dynamic system of n variables where many (but typically not all) pairwise interactions are present. This leads to the same connectance or Jacobian matrix. But unlike May they explicitly note that in the general nonlinear case the strengths (and indeed possibly the signs) of these interactions will vary throughout phase space.

Following their analysis, we note that any connectance matrix $\boldsymbol{A}$ can be considered as composed of multiple overlapping feedback circuits. For any such circuit, the indices are circular permutations of each other. For instance in a 3-variable system as sketched in Figure 5, the full list of potential circuits is: $<\boldsymbol{a}_{11}>$, $<\boldsymbol{a}_{22}>$, $<\boldsymbol{a}_{33}>$, $<\boldsymbol{a}_{12}\boldsymbol{a}_{21}>$, $<\boldsymbol{a}_{23}\boldsymbol{a}_{32}>$, $<\boldsymbol{a}_{31}\boldsymbol{a}_{13}>$, $<\boldsymbol{a}_{12}\boldsymbol{a}_{23}\boldsymbol{a}_{31}>$, $<\boldsymbol{a}_{21}\boldsymbol{a}_{13}\boldsymbol{a}_{32}>$. If one or more of the connections in such a circuit is zero, that circuit as a whole is non-functional; but otherwise, a count-up of the number of negative connection weights decides whether that individual feedback circuit constitutes a negative feedback (odd number of negatives) or positive feedback (even number). The limiting case of such a circuit is that constituted by self-feedback, given by the term $\boldsymbol{a}_{ii}$ on the main diagonal; that minimal circuit will be non-functional, negative-feedback or positive-feedback depending on whether its value is zero, or its sign is negative or positive.

Thomas and Kaufman (2001a) defined a full-circuit as those circuits and unions of disjoint circuits that involve all the variables of a system. Hence in this 3-variable system, there are six possible full-circuits: $<\boldsymbol{a}_{11}\bullet\boldsymbol{a}_{22}\bullet\boldsymbol{a}_{33}>$, $<\boldsymbol{a}_{11}\bullet\boldsymbol{a}_{23}\boldsymbol{a}_{32}>$, $<\boldsymbol{a}_{22}\bullet\boldsymbol{a}_{31}\boldsymbol{a}_{13}>$, $<\boldsymbol{a}_{33}\bullet\boldsymbol{a}_{12}\boldsymbol{a}_{21}>$, $<\boldsymbol{a}_{12}\boldsymbol{a}_{23}\boldsymbol{a}_{31}>$, $<\boldsymbol{a}_{13}\boldsymbol{a}_{32}\boldsymbol{a}_{21}>$

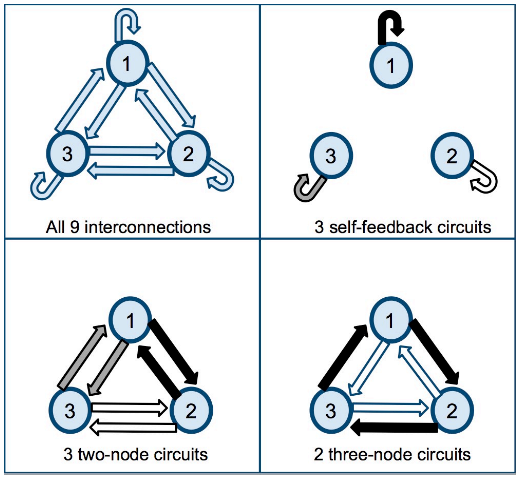

Figure 5: The eight potential circuits, differentiated by shading, within a system of 3 variables fully interconnected.

These correspond to the terms of the determinant of the Jacobian matrix. For any one such full-circuit, considered in isolation, the type of steady state this generates will be determined entirely by the signs, plus or minus, of the various component circuits that comprise this full-circuit. Given that in nonlinear systems any (or all) connection strengths can vary according to position in phase space, and given that the change of sign of any one connection will change the sign of any component circuit of which it is part, we can see that this will alter the type of steady state generated.

This highlights the significance of those connection strengths in a nonlinear system that change in sign as one moves through phase space. These arise from nonmonotonic functions that generate circuits that switch between negative and positive according to context – 'ambiguous' circuits – and thereby generate ambiguous full-circuits. Such changes in sign, in one or many such connections, carve up the phase space into different regions, and one can expect the properties of steady states to differ from one such region to the next. This gives the richness of possibilities to nonlinear systems that is missing from the linear ones.

Plausible nonlinearities

It might be argued that with some systems, although interactions are potentially nonlinear they are 'linear enough' for there to be only a single equilibrium. Here we present and discuss some simple nonlinear functions, to see where and how they generate multiple possible equilibria. If one was to analyse fully the probability of stability in some class of nonlinear systems, these might be appropriate simple classes to start on.

Sigmoids

Sigmoids are commonly used to model physical or biological systems, since they represent an effect that is monotonic yet with asymptotes at lower and upper bounds. A widespread example of where they are used in artificial systems would be Artificial Neural Networks. We have already seen above (Equations 5) that even a simple monotonic function such as a sigmoid is not 'linear enough' to avoid multiple equilibria. The ambiguous circuits discussed above, generating changes in stability through nonmonotonic functions, do not exhaust the ways in which multiple equilibria can exist. Figure 4 demonstrates how both stable and unstable equilibria can be generated merely by a change in strength of a connection without change in sign.

Piecewise linear with a single bend

Even simpler than a sigmoid, consider a piecewise linear function coupled with a linear function:

$$dx/dt = \max(0, 2 + y - 2x) \qquad dy/dt = x - y \qquad \text{Eqns 6}$$

These are both linear except that *dx/dt* is constrained not to go below zero. As can be seen from Figure 6, this is sufficient to generate a pair of equilibria, one stable and the other unstable.

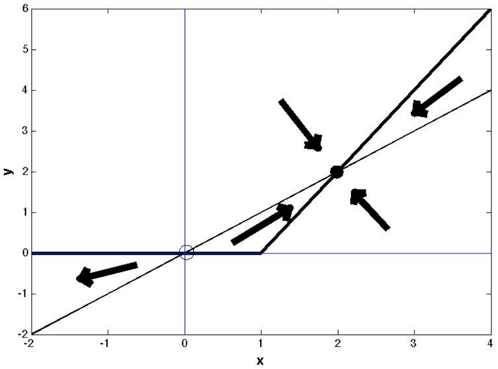

Figure 6: A perturbation analysis of Equations 6, using the same conventions as in Figure 4. There is a stable equilibrium at (2,2) and an unstable equilibrium at the origin (0,0).

Sinusoidal functions

We have seen how the single inflexion of a sigmoid allows the possibility of 3 intersections with a straight line and hence 3 equilibria. Crudely speaking, the more bends the more possibilities for intersections, and with oscillatory functions such as a sine wave the slope changes in sign repeatedly and indefinitely. The combination of a straight line and a sine wave can lead to an arbitrary number of equilibria that will alternate between stable and unstable. Going further, it can be shown (Kaufman and Thomas, 2002) that a system of 3 variables:

$$dx/dt = -bx + \sin(y) \qquad dy/dt = -by + \sin(z)$$
$$dz/dt = -bz + \sin(x) \qquad \text{Eqns 7}$$

can, depending on the parameter *b*, move from having a single steady state for *b*>1, through multiple steady states as b decreases, with the number of steady states tending to infinity as *b*→0. The dynamics change from simple to chaotic, with periodic or multiperiodic windows. The many changes of sign within the regions where nullclines intersect provide ambiguous circuits and increase the richness of possibilities.

Hat-shaped functions

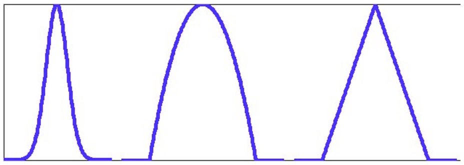

Figure 7: Three 'hat-functions' with broadly similar consequences: gaussian, truncated parabola, and witches hat.

Unimodal 'hat-shaped' functions whose slopes have a single change of sign are an important class of simple nonlinear functions that share some of the asymptotic properties of sigmoids. The examples in Figure 7 share the property of dropping to zero (or approaching zero in the case of a gaussian) each side of a central region. If we take any of these hat-functions as *y*=*H*(*x*), this could represent a viability function of an organism or species *y* that can only survive (in the case of the gaussian version: survive to any significant level) within some range of values of an environmental variable bounded above and below. These can be considered amongst the most basic of nonmonotonic functions, and it turns out that they do indeed play a crucial role in giving rise to homeostasis, or a particular form of stable equilibrium, in Daisyworld models. Those who use Daisyworld models (which are one class of nonlinear complex system) assert that homeostasis arises naturally in these, whereas many critics such as Kirchner (2002) consider the probability to be vanishingly small unless the parameters are fixed somehow. This controversy illustrates some of the archetypal contrasting viewpoints presented in the complexity-stability debate, and hence we shall review this at greater length.

Daisyworld

Lovelock introduced the Daisyworld model (Watson and Lovelock, 1983) as a possible explanation of how organisms coupled in mutual feedback with some environmental variable could form a homeostatic system, biotic-environmental, as is proposed in the Gaia Hypothesis (Lovelock, 1972). The Faint young Sun paradox (Sagan and Mullen, 1972) suggests that despite the heat output of the sun changing significantly over the last few billion years the planetary climate has maintained itself around the temperatures conducive for life. The Gaia Hypothesis suggests that this arises through homeostatic properties of the interactions between biota and environment. In the Daisyworld model the organisms (Daisies) have a viability whose dependence on temperature is given by a hat-function; the truncated parabola version is used in Watson and Lovelock (1983). In turn, through differential absorption or reflection of sunlight, these Black or White Daisies had a positive or negative affect on the same local temperature that influenced their viability. Such systems can be analysed for stability in the context of noise or perturbations at two levels.

In the first instance, any equilibrium state of such a system can be analysed for stability or instability in the presence of small levels of noise; only stable equilibria will persist, and only stable equilibria that have the biota (Daisies) within their viability zone are relevant. But the main interest of Daisyworld models is the extent to which such stable equilibria can persist in the face of major systemic external perturbations, such as major changes in heat output of the sun. It turns out that the Daisyworld temperature is maintained within the viability zone for significantly greater ranges of solar forcing *with* the biotic feedback to the local temperature, as compared to *without* such feedback. This homeostasis arises from the nonmonotonic nature of the hat-function.

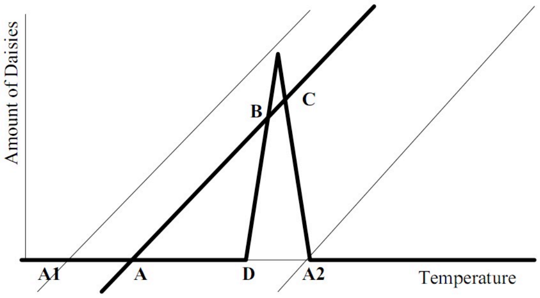

Figure 8: The witches hat-function represents the dependency of Black Daisies on local temperature.

Harvey (2004) showed how a simplification of the Daisyworld model produced the same effects, using a witches hat-function. A reduced version of such homeostasis can be shown with just one species of Daisies, e.g. Black ones. With Y black daisies, local temperature T, level of solar forcing S, then for suitable constants k_1 k_2 we have:

$$dY/_{dt} = H(T) - Y \qquad dT/_{dt} = S - k_1 T + k_2 Y \qquad \text{Eqns 8}$$

The equilibria are shown where the corresponding lines intersect in Figure 8. The different sloping lines, intersecting the temperature axis at A1, A, A2, correspond to different possible levels of solar forcing. It can be seen that, depending on the level of solar forcing, there is either one equilibrium (e.g. at A1 or lower temperatures, or at A2 and higher temperatures) or three (e.g. A, B, C). This latter case gives us: a possible stable equilibrium with zero Daisies at A; or an unstable equilibrium with Daisies at B, the instability being despite the temperature being viable; or a further stable equilibrium at C with Daisies present within their temperature viability-zone. This last stable equilibrium is the focus of interest, and we consider the range of solar forcing for which C exists; i.e., for which there is a stable population of Daisies within the local temperature viability zone. From inspection of Figure 8 we can see that the biotic feedback (from Black Daisies increasing local temperature) has given rise to viable local temperature over a wider range of solar forcing (corresponding to the range A1↔A2 in the figure) than in the absence of such feedback (corresponding to D↔A2, the unassisted viability range of the hat-function).

Thus the presence of Black Daisies extends the range of viability towards lower solar forcing (the 'faint young sun'); conversely, White Daisies (giving rise to a line ABC with a negative slope in contrast to the positive slope in Figure 8) would extend the range of viability towards higher solar forcing, a hotter sun. This increased range of homeostasis arises from the nonmonotonic nature of the hat-function generating extra possible equilibria.

Criticism of Daisyworld

This present analysis of the G&A and May papers was originally motivated by work on Daisyworld models (Harvey, 2004) that are one class of these nonlinear systems of a Gaian biota/environment. Such models display homeostasis under a wide range of conditions, yet critics frequently voice the suspicion that this must be because the parameters are carefully chosen from an improbable subset, biased towards negative feedback, in order to achieve stability. For instance Kirchner (2002) suggests that Gaian regulation depends on an implausible assumption that the influence of biota on the environment have a strong tendency to be environment-enhancing rather than environment-degrading. This, it is implied, suggests that such influence has been biased by the modeler to have the appropriate sign, positive or negative. Yet, as is shown in Harvey (2004), regardless of the sign of such a biota→environment effect, when combined with a hat-shaped viability function environment→biota, the resulting ambiguous circuit has the potential for both stable and unstable equilibria within the viability range. Stable equilibria will inevitably be 'selected' in preference to unstable, but since this is independent of the sign of the biota→environment effect it cannot be attributed to some biased choice of this sign. In either case the viable stable equilibrium gives a context that defines this effect as locally environment-enhancing.

This has inevitably been a limited review of the basics of Daisyworld models, missing out many layers of subtlety. For instance the role of hysteresis has not been mentioned, and the significance of those stable equilibria that are within the viability zone, as contrasted with stable equilibria corresponding to extinction, has been treated only briefly. But the main point to be emphasised here is that the interesting (and often counter-intuitive) properties of these models arise from exactly those features of nonlinear systems that May had omitted in his analysis.

Importantly in this context, the homeostasis of Daisyworld systems extends to those with large numbers of variables. Applying these lessons to the construction of artificial systems, it has been demonstrated (Harvey, 2004) that a simulated robot coupled with the environment via an arbitrarily large number of interactions comprising hat-functions (on sensory inputs) and linear functions (on consequent outputs) could find a homeostatic equilibrium. This is so even if the signs of the linear functions are set positively or negatively at random, and the relevant parameters are varied across some two orders of magnitude. Other examples of systems with multiple interacting component parts achieving equilibrium through the use of hat-functions can be found in Dyke et al. (2007) and McDonald-Gibson et al. (2008). In these cases there was a single environmental variable, and numerous biotic variables subject to hat-function viability limits. The Daisystat (Dyke, 2010) extends this approach to multiple environmental variables.

Conclusions

The core of this paper is the demonstration that May's (1972) generalization to nonlinear systems – of results that largely hold true in linear systems (Gardner and Ashby, 1970) – is flawed. The method, through linearization around an assumed single equilibrium point, will at best give local stability; there may be many equilibria, and global stability can arise through stability at just one of these. With nonlinear interactions the size, and potentially also the sign, of the connection strengths varies according to position in phase space, and there is no attempt to account for this. To be rigorous, the probability of global stability would depend on assessing the (differing) probabilities of local equilibria, and combining these to calculate the probability that at least one was stable. No attempt at this was presented in (May, 1972), and hence his conclusions should be rejected. His calculations underestimate, potentially by a massive factor, the probability of stability in systems "which in general may obey some quite nonlinear set of first-order differential equations".

The ambiguous circuits proposed by Thomas and colleagues in their analysis of multistationarity have been used above to explain how a plurality of equilibria can be generated by nonmonotonic functions. But even simple monotonic functions such as a sigmoid can generate alternating stable and unstable equilibria. A number of different simple nonlinear functions were analysed, to demonstrate just how easy it is to breach the assumptions upon which May was relying.

Does this matter?

Daisyworld models, particularly as the number of variables increase, are just one example of a complex nonlinear system where one would expect May's analysis to be relevant. These demonstrate typical properties of many families of complex nonlinear systems: if one treats the slower variables as parameters and the faster variables as thermal noise, then the remaining variables at intermediate timescales will settle down to a metastable equilibrium (that may be disturbed at a 'tipping point' when a 'parameter' shifts enough). In ecology it used to be a common view that ecosystems developed through succession towards a single equilibrium state or 'climax' (Clements, 1916); but nowadays ecologists are more open to the possible of multiple possible equilibria in an ecosystem.

Our intuitions based on understanding simple linear systems can all too easily lead us into error when considering complex nonlinear ones, with multiple overlapping circuits of feedback. This appears to be the root of the problem here. We are not aware of any previous exposure of these flaws in May (1972); indeed the author is still citing it without qualification (Haldane and May, 2011) in the context of 'banking ecosystems' where clearly there are nonlinearities. It took nearly 30 years for the basic programming errors in G&A to be pointed out in print, and 40 years is too long for these further significant flaws to remain unchallenged.

Acknowledgments

I thank Nathaniel Virgo and David Waxman for discussions, and for checking that my code was working properly when it failed to replicate G&A. NV found the prior reference to this anomaly (Solow et al., 1999). Anonymous reviewers' comments stimulated the expansion and clarification of some sections; remaining errors and obscurities are the writer's.

References

Clements, F. E. (1916). *Plant succession: an analysis of the development of vegetation*. Carnegie Institution of Washington.

Dyke, J., McDonald-Gibson, J., Di Paolo, E. and Harvey, I. (2007). Increasing Complexity can increase Stability in a Self-regulating Ecosystem. In Almeida e Costa, F., Rocha, L. M., Costa, E., Harvey, I. and Coutinho, A, editors, *Eur. Conf. Artificial Life ECAL 2007,* pages 123–132. LNCS 4648, Springer.

Dyke, J. (2010). The Daisystat: A model to explore multidimensional homeostasis. In Fellermann, H., Dörr, M., Hanczyc, M., Laursen, L., Maurer, S., Merkle, D., Monnard, P-A., Stoy K. and Rasmussen S. (Eds.): *Artificial Life XII*. MIT Press.

Gardner, M. R. and Ashby, W. R. (1970). Connectance of Large Dynamic (Cybernetic) Systems: Critical Values for Stability. *Nature*, 228:784.

Haldane, A. G. and May, R. M. (2011). Systemic risk in banking ecosystems. *Nature*, 469:351–355.

Harvey, I. (2004). Homeostasis and Rein Control: From Daisyworld to Active Perception. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T. and Watson, R. A., editors, *Artificial Life IX*, pages 309–314. MIT Press, Cambridge, MA.

Kaufman, M. and Thomas, R. (2002). Emergence of complex behaviour from simple circuit structures. *C. R. Biologies*, 326:205–214.

Kirchner, J. W. (2002). The Gaia hypothesis : fact, theory and wishful thinking. *Climatic Change*, 52:391-408.

Lovelock, J. E. (1972). Gaia as seen through the Atmosphere. *Atmospheric Environment*, 6(8):579–580.

May, R. M. (1972). Will a Large Complex System be Stable? *Nature*, 238:413–414.

May, R. M. (1973). *Stability and Complexity in Model Ecosystems*. Princeton University Press, Princeton, New Jersey.

McDonald-Gibson, J., Dyke, J. G., Di Paolo, E. A. and Harvey, I. R., (2008). Environmental Regulation can arise Under Minimal Assumptions, *J. Th. Biol.*, 251(4):653–666.

Sagan, C. and Mullen, G. (1972). Earth and Mars : Evolution of Atmospheres and Surface Temperatures, *Science*, 177(4043):52–56.

Solow, A. R., Costello, C. and Beet, A. (1999). On an Early Result on Stability and Complexity. *The American Naturalist*, 154(5):587–588.

Thomas, R. and D'Ari, R. (1990). *Biological Feedback*. CRC Press, Boca Raton, Florida, USA.

Thomas, R. L. and Kaufman, R. (2001a). Multistationarity, the basis of cell differentiation and memory. I. Structural conditions of multistationarity and other nontrivial behavior. C*haos*, 11(1):170–179.

Thomas, R. L. and Kaufman, R. (2001b). Multistationarity, the basis of cell differentiation and memory. II. Logical analysis of regulatory networks in terms of feedback circuits. C*haos*, 11(1):180–195.

Watson, A. J. and Lovelock, J. E. (1983). Biological Homeostasis of the Global Environment : the Parable of Daisyworld. *Tellus*, 35B:286–289.