Loop Eigenvalue Elasticity Analysis: Three Case Studies

Christian Erik Kampmann
Copenhagen Business School
Birk Centerpark 40
7400 Herning
Denmark
CEK.LPF@CBS.DK

Rogelio Oliva
Mays Business School
Texas A&M University, 4217
College Station, TX 77843
USA
ROLIVA@TAMU.EDU

May, 2005

Abstract

The paper explores the application of loop eigenvalue elasticity analysis (LEEA) to three models in order to reveal the potential of the method for generating insights about model behavior and to uncover issues in developing the method further. The results indicate that the utility of the method depends upon the character of the model and dynamics involved. In models where the transient behavior is of interest, the method yields insights on par with the pathway participation method, though better tools to link the method to time paths of particular variables is needed. In models involving near-equilibrium oscillation, the LEEA is clearly the most powerful, though more efficient computer programs are needed to handle large-scale models. The analysis reveals a number of technical issues to be resolved relating both to repeated eigenvalues and to interpreting a large number of loop elasticities. In highly non-linear models exhibiting deterministic chaos, LEEA, being based upon linear concepts, does not appear to yield any insight because the eigenvalues may change substantially even when the mode of behavior (expanding oscillations) appears constant. The paper also describes the set of tools and processes that we have developed to date and the design for a web-based toolbox to make the methods readily available to a wider audience in the hope that others will join the efforts to develop analytical methods for interpreting model behavior.

Introduction

There is a long-standing desire in the system dynamics scientific community for tools for analyzing the behavior of large-scale models. The field has relied heavily on intuitive approaches, based on the notions of system archetypes (Senge 1990), simple models linking feedback loops to system behavior (Ford 1999; Richardson, 1995), and a structured approach to model testing and simulation experiments (Barlas 1989; Forrester and Senge 1980).

A more formal approach, developed by Mojtahedzadeh, Richardson and Anderson (2004), uses the so-called pathway participation metric (PPM) to identify the structure that is most influential in affecting the qualitative time path of a given variable. The method works by examining how much the net flow in the state variable under consideration changes relative to the change in the variable itself ($d\dot{x}/dx$), or, equivalently, ($\ddot{x}/\dot{x}$), i.e., this ratio contains information about the slope and the curvature of the variable's behavior. The method proceeds by examining the separate influence on this ratio of the cause-and-effect chain of variables feeding into the behavior of *x*. Specifically, the method assigns a *pathway participation metric* of the relative strength of the influence of the variable *i* on the behavior of variable *k* and then attributes importance to those pathways or feedback loops whose PPM is numerically large and has the same sign as $d\dot{x}/dx$.

The main strength of the approach is that it takes place in the time domain and does not require computing eigenvalues and other highly intensive tasks. There are, however, some limitations to the PPM method. First, the method is not particularly well suited for systems that oscillate. Even though loop gains are frequently constant in oscillating systems (always the case in linear systems), PPM would indicate shifting loop dominance. A second limitation of the current implementation of PPM is that it uses a depth-first search for the single most influential pathway for a variable. This strategy does not capture the situation where more than one structure may contribute significantly to the model behavior and, through the depth-first algorithm, may miss alternative paths that could prove to yield a larger total value of the metric. In other words, since PPM is contingent to the selected variable of interest, the method could lead to "localized" explanations of the variable's behavior – as opposed to a system-wide analysis of the structure.

Another formal approach, which underlies this paper, is the eigenvalue elasticity analysis. The method was first introduced by Forrester (1982) in the context of analyzing stabilization policies in a macro-economic model. The method uses eigenvalues to characterize the behavior of the system, since the behavior of any system at any point in time can be fully described by the eigenvalues and eigenvectors of the linearized system around the current operating point. The importance of a particular link or parameter in the model for a particular behavior mode (eigenvalue) is then measured by how much the eigenvalue changes with a small change in the strength of the link or parameter.

Kampmann (1996) further refined the approach by using graph theory to develop the notion of an independent loop set (ILS), a much smaller set of loops, where the gains of these loops can be *independent* determined and related to the eigenvalues, i.e., one can measure the loop eigenvalue elasticities. Although the LEEA method has been employed in a variety of models and has proven to be a useful tool for gaining deeper insight into model behavior (Gonçalves 2003; Gonçalves, Lerpattarapong, and Hines 2000; Kampmann 1996), it is clear that further development is needed.

One problem with Kampmann's approach, as pointed out by Oliva (2004), is that the elasticity measures depend upon the choice of the independent loop set, which is not unique. Oliva proposed the notion of a *shortest independent loop set* (SILS), consisting entirely of *geodedic* loops where the path between any two variables in the loop is also the shortest path between those variables. Under some conditions, the SILS is unique, and the shorter loops frequently lead to more intuitive interpretations. Furthermore, in a comparison between the PPM method and the SILS method, it appears that the two tend to identify the same structure as being dominant (Oliva and Mojtahedzadeh 2004).

Although these initial results are encouraging, it remains an open question how applicable the eigenvalue approach, which is based on linear systems theory, is to system dynamics models, which are frequently nonlinear. The motivation for the present paper is to explore this issue by applying the LEEA method to three very different models. In choosing the models, we focused on three different types of behavior patterns.

The first model, Industrial Structures, is (essentially) the same as the one used in Mojtahedzadeh et al (2004) to demonstrate the PPM method. The model shows a *single transient behavior* where the growth of industry is limited by a non-renewable resource (water). Thus, as industry grows, water demand increases proportionately, which depletes the water resource base until the point where water shortage affects the growth of industry, changing the behavior to a gradual decline in industry and water. This structure is fairly common, appearing, for instance, in the Urban Dynamics model (Forrester 1969) (from which it is in fact extracted), the Kaibab model (Goodman 1974) (in a variant where the resource is renewable but at a low equilibrium deer population), and the World Dynamics model (Meadows *et al.* 1972). The model is also interesting because it allows for comparisons with the PPM method.

The second model, first proposed by Lorenz (1963), is a classic model exhibiting deterministic chaos. The model cycles around two unstable equilibrium points, tracing out a characteristic butterfly-shaped (strange) attractor. Nonlinearities play a key role in chaotic dynamics, and the model was chosen to explore whether eigenvalue analysis could yield an understanding of the nature of both the behavior phases where the model cycles around a given equilibrium point in expanding oscillations and the phases where it suddenly shifts to cycle around the other equilibrium point.

The third model is a medium-sized model developed by Mass (1975) in which he analyzes the causes of different types of economic cycles. Mass identifies distinct structures responsible for the short-term business cycles (involving inventory and labor adjustment), and longer-term cycles (involving capital investment). The model was chosen both because it is of a fairly large size and because it exhibits (mostly linear) oscillations, an area where the eigenvalue method ought to have the greatest potential. The question is both whether the eigenvalue method for such a large model will yield insights, given the there is bound to be a very large number of potential loops in the system, and whether the results concur with the findings in Mass' analysis.

The paper is structured as follows: In the next section, we present a brief review of the loop eigenvalue elasticity analysis (LEEA) method. The following three sections contain the analysis for the three models. Following this, we describe the set of tools and processes that we have developed to date and the design for a web-based tool to make the analytical tools available to a wider audience. The paper concludes with our thoughts on the strength and weaknesses of LEEA and suggestions for future work.

Loop eigenvalue elasticity analysis (LEEA)

The LEEA method, proposed by Kampmann (1996), seeks to explain system behavior by examining how changing the gains of individual feedback loops in the system affects a particular eigenvalue, λ, or behavior mode – subsequently the two terms will be used interchangeably. The method works on a linearized version of the model around the current operating point, so that the dynamics at any point in time may be approximated by the linear differential equations,

$$\begin{aligned}\dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y} + \mathbf{b},\\ \mathbf{y} &= \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{y} + \mathbf{h}.\end{aligned} \tag{1}$$

As is well known, the behavior of linear systems is completely characterized by the eigenvalues (and eigenvectors) of the system matrix. The measure of influence used in the analysis is the *eigenvalue elasticity*,

$$\varepsilon(c) = \frac{\partial \lambda}{\partial g(c)} \frac{g(c)}{\lambda}, \tag{2}$$

where $g(c)$ is the gain of the feedback loop c. The problem pointed out by Kampmann (1996) is that the number of feedback loops in a system dynamics model is typically very large (in the full version of the Mass model, there is over 5500 loops!), greatly exceeding the degrees of freedom in determining the system structure. In other words, there is a great deal of redundancy in the describing the system by (all) its constituent feedback loops. Instead, one must restrict oneself to a subset of *independent* loops, defined as a maximal set of loops whose gains can be independently determined or changed by an appropriate change in the system matrices. By considering the model as a directed graph, where a link from node x to node y exists whenever x enters into the equation for y (familiar as to system dynamics practitioners as causal loop diagrams), Kampmann showed that if a system has a total of n variables and N links, the number of independent loops in a fully connected system is simply N-n+1, which typically grows linearly with the size of the model n. (A fully connected system is one in which every variable is directly or indirectly influenced by every other variable – in graph theory terms, the graph is strongly connected). If more than one strongly connected component exists, he showed how each component determines a particular subset of eigenvalues independent of the other components, so that each component may be analyzed in isolation. Kampmann also presented an algorithm to construct an independent loop set.

Once an independent loop set is found, one may proceed to find the loop elasticities as follows: If the links (or edges) and the loops in the system are arranged in column vectors, $(e_1,\ldots,e_N)'$ and $(c_1,\ldots,c_{N-n+1})'$ respectively, then one may describe the independent loop set by the $(N-n+1\times n)$ *cycle matrix* $\mathbf{c}_r$, which has a 1 in the i-th row and j-th column iff edge j appears in loop i and 0's everywhere else. If one considers the *edge eigenvalue elasticity,*

$$\varepsilon(e) = \frac{\partial \lambda}{\partial g(e)} \frac{g(e)}{\lambda}, \tag{3}$$

where $g(e)$ is the gain of edge e, one may obtain the *loop eigenvalue elasticity* from the edge elasticities from the set of linear equations $\mathbf{c}_r\mathbf{g} = \mathbf{e}$, where $\mathbf{g}$ and $\mathbf{e}$ are the column vectors of loop and edge elasticities respectively. This system of equations has a unique solution for $\mathbf{g}$ since $\mathbf{c}_r$ has full row rank. Kampmann further showed how the edge elasticities may easily be obtained by linear calculations from the system matrices and eigenvectors.

In summary, the LEEA method consists of first constructing an independent loop set, then finding the system eigenvalues (frequently at different points in time as they change over time in nonlinear systems), and then using the calculated loop elasticities for a particular eigenvalue to assess the relative importance of that loop in that particular behavior mode at that particular point in time. The sign of the elasticity will show whether the loop amplifies or attenuates the mode and changes in loop dominance will show up as changes in the loop elasticities.

If the eigenvalues are complex numbers, this indicates an oscillatory behavior mode. In this case, eigenvalues and elasticities come in complex conjugate pairs, where one may look at various measures of influence, such as the damping ratio, the natural frequency, or the damped frequency (Forrester 1982). In general, the absolute value of the elasticity is a measure of how much the loop affects the behavior mode in question.

Industrial structures model

The model

Figure 1 shows a flow diagram of the simple industrial structures model, originally presented in Mojtahedzadeh et al (2004). The Vensim equations for the model are provided in the appendix. In order to allow analytical differentiation of the equations, the original table functions were replaced by analytical functions. The match in both structure and behavior is very close, giving us confidence that the two models are equivalent.

<u>Insert figure 1 about here</u>

Model behavior

The behavior of the model is shown in Figure 2. Construction of new industry is assumed to be proportional to the stock of industry, and industry depreciates at a constant fractional rate. Initially, when there is ample water available, the fractional investment exceeds the fractional depreciation, and industry grows at a constant exponential rate.

<u>Insert figure 2 about here</u>

However, industry investment is affected by the availability of water, measured by the "normal" water demand compared to actual consumption. Water consumption, in turn, is proportional to the industry stock, modified by an effect of low levels of water resources. As industry stock rises and water resources decline, the resource constraint affects water consumption which in turn affects the growth of industry. Eventually, the effect is sufficiently large to stop the growth. As water is continually depleted, the industry stock starts declining and this process continues for the rest of the simulation as both stocks decay in parallel.

Feedback loops

The model contains six independent feedback loops, as indicated in the flow diagram in Figure 1. In this respect, the model is unusually simple. In the course of the simulation, the loop gains shift, as illustrated in Figure 3. In fact, simply looking at the loop gains in this simple model gives a good impression of the shift in loop dominance.

<u>Insert figure 3 about here</u>

Loop 5, which is positive, is somewhat of an artifact of the formulation for the effect of water shortage on industry growth, which is a function of Water Consumption relative to Water Demand. It would be more correct to combine it with Loop 4 which is negative, indicating the limit on industrial growth for a given level of water consumption. Given the formulation in the model, those two loops exactly cancel each other out and could be eliminated from the model. (This is verified in Figure 3, where the sum of the two loop gains -- line "4+5" -- is exactly 0 at all times.) It can be seen from the following considerations:

$$EWAG = f(WC/WD), \quad WC = WD * EWAC \Rightarrow EWAG = f(EWAC) \tag{4}$$

where the variable names are

EWAG	Effect of water adequacy on growth
WC	Water consumption
WD	Water demand
EWAC	Effect of water availability on consumption

This means that the effect of water shortage on new industry construction is in reality direct function of the effect of water availability on water consumption, a table function of a table function, which is usually not good modeling practice. However, we have chosen to retain the original structure because it highlights the issue of how formulation details in the model may create such "artificial" feedback loops.

Figure 3 shows how around the turning period from t=25-30, the positive loop 1 loses strength while the negative loops 3 and 6 both gain strength. Subsequently, the two latter loops lose some of their strength but the drop in strength of loop 1 is permanent. Notice that the maximum for industrial structures in Figure 2 occurs exactly when the gain of loops 1 and 2 are numerically equal, as would be expected.

PPM analysis

According to the PPM analysis in Oliva and Mojtahedzadeh (2004), the behavior of the industry stock can be divided into four phases, which are also indicated in Figure 2, namely a period of exponential growth (phase 1), leveling growth (phase 2), exponential collapse (phase 3) and exponential decay (phase 4). (In the figure, the phases were identified by considering the sign of the $\ddot{x}/\dot{x}$ ratio). In phase 1, the dominant structure identified by the method is the positive growth loop 1. In phases 2 and 3, the dominant structure is the negative loop 3 controlling water depletion, along with the pathway from Water Consumption to Effect of Water Adequacy on Growth to New Industries to Industrial Structures. In phase 4, the dominant structure is the depreciation loop 2. It is of interest to see whether the same structures are identified as important in the eigenvalue analysis, although the loop elasticities do not lend themselves to direct comparisons with one-way pathways, as is the case in phases 2 and 3.

LEEA

Figure 4 shows a plot of the two system eigenvalues over time. Both are real-valued. Initially, one eigenvalue is positive and equal to the growth rate of industry while the other is zero. The latter indicates a system behavior with no feedback, which is the case for the depletion of water resources, since the effect of water availability is initially inactive. Hence, industry grows independently of the water resources, and the water depletion is driven by the growing industry stock with no feedback from either stock.

<u>Insert figure 4 about here</u>

As the effect of water availability is felt on both water consumption and industry growth, both eigenvalues start to fall. Thus, from around t=22 to t=25, a positive and a negative mode are present. Subsequently, both eigenvalues are negative. The larger eigenvalue levels off slightly above the value of fractional depreciation rate of industry. (There is still some construction even at the end of the run, though it is much smaller than depreciation). The positive mode changes sign to become negative around t=25, which is a few years before the industry stock peaks. This change in sign in eigenvalue coincides closely with the inflection point that defines the transition from phase 1 to phase 2 in the PPM analysis.

Since eigenvalues at times become 0 or close to 0, thus making the elasticity measure (eq. 2) undefined., we adopted *loop influence* – a measure related to the elasticity and defined by

$$\frac{\partial \lambda}{\partial g(c)} g(c) \tag{5}$$

– as the measure for our analysis. A negative influence measure always implies a "stabilizing" influence (positive eigenvalues are reduced in numerical values and negative ones are increased in numerical value), and vice versa. Figures 5 and 6 show the influence measure for the largest and smallest eigenvalue (nos. 1 and 2 in Figure 4), respectively.

<u>Insert figure 5 about here</u>
<u>Insert figure 6 about here</u>

One can define an initial phase in Figure 5 (until around t=20), where the positive eigenvalue is affected by loops 1 and 2. A stronger Loop 1 (growth) increases the growth rate while a stronger Loop 2 (depreciation) decreases the growth rate. The other loops do not play a major role in this mode. Note that the two "artificial" loops 4 and 5 have a smaller equal and opposite influence during this phase—this is *not* the case later in the simulation. The other eigenvalue (cf Figure 6) is zero or almost zero during this time, since the table functions are not in the active range. Consequently, none of the feedback loops can influence this mode either. All of this is as expected and in accordance with the PPM analysis, which also showed the growth loop being the dominant mechanism. The eigenvalue analysis, however, also reveals the role of the depreciation loop in dampening the growth rate.

During the following few years (t=20 to t=25), the table functions resulting from water shortage are starting to become active and the loop gains, eigenvalues and influence measures all change rather quickly, cf. Figures 3-6. Loops 1 and 2 continue to play the dominant role in the growth mode. A new mode (eigenvalue no. 2) now appears, and Figure 6 reveals that loop 3 (which regulates water consumption as water resources fall) plays an important role in this mode. Hence, one may interpret the result as the activation of the first-order control on water resources. Note, however, that the growth mode (eigenvalue no. 1) is not much affected by loop 3 in the eigenvalue analysis. In short, the analysis reveals that loops 1 and 2 primarily regulate the growth mode while loop 3 regulates the relatively rapid adjustment mode. However, the LEEA method does not say anything about how much of these two modes is expressed in the time behavior of a given variable. (This would require attention to the eigenvectors of the system as well).

Notice also that the two "artificial" loops 4 and 5 no longer cancel each other out with respect to the influence on the eigenvalues. This may seem like a paradox, given that the loops are never effectively present in the model (cf. Figure 3 and the discussion above). However, the loop elasticities measure the effect that would happen if one could change the gain of the loop without affecting the others. This is possible in principle, e.g., by adding a new parameter or changing the formulation so that the two loops do not always cancel each other. The fact that the model is formulated in such a way that the two loops always cancel is in some ways immaterial to the method. The results show that one must take great care in interpreting loop elasticities and continually refer back to the model formulations to make sure that the loops are "real".

After time t=25, the largest eigenvalue becomes negative. At the same time, the significance of the growth loop 1 in determining this eigenvalue drops quickly. At the same time, loop 4, which conceptually limits growth via the effect of water adequacy, shows increasing influence on this mode, indicating that were it not for the fact that its gain is neutralized by loop 5, this loop could have a relatively large impact on how rapidly industry declines.

In the subsequent years, loop 2 emerges to become the dominant structure regulating the gradual decay of industry (eigenvalue 1). This is in accordance both with intuition and with the PPM method. The phase transition happens when the influence of loop 2 in eigenvalue no. 1 becomes larger (in absolute terms) than the influence of loop 6.

The other mode (eigenvalue 2) is mostly dominated by loop 3, except for a few years between t=25 and t=30, where loop 6 is also influential. Given that the measures shift around so rapidly during this time, and given that loops 4 and 5 to some extent cancel each other's influence, it is difficult to give a clear account of the dynamics during this period. After time t=30, loop 3 clearly emerges the dominant loop for this mode. The eigenvalue is negative and numerically quite large, indicating a rapid adjustment. An intuitive account of what is going on could be that the level of water reserves is rapidly adjusted to the path of the declining industry stock, via the first-order control loop 3.

In summary, it is clear that, while the eigenvalue analysis is clearly capable of revealing subtleties in how structure affects behavior, the picture can easily become rather complex, particularly in the presence of "artificial" loops that are not in fact active in the model because they cancel each other out. Moreover, as long as the behavior modes (eigenvalues) are not related to the time path of specific variables, it is not easy to make intuitive interpretations of the results. Thus, the analysis clearly points out the need to consider the eigenvectors as well, as they determine how much of a given mode is expressed in the behavior of a given variable.

Lorenz model

The model

The Lorenz model is one of the earliest examples of deterministic chaos. Lorenz was a meteorologist studying weather forecasting and its fundamental limits, i.e., the predictability of the solutions to hydrodynamics equations. His now famous model (Lorenz 1963) can be thought of as a gross simplification fluid motion driven by thermal buoyancy (convection), although his original paper seems to use the model simply as a set of equations "whose solutions afford the simplest example of a deterministic nonperiodic flow of which the writer is aware". The model has three state variables governed by the following equations:

$$\begin{aligned}\dot{X} &= -c(X-Y),\\ \dot{Y} &= aX - Y - XZ,\\ \dot{Z} &= b(XY - Z).\end{aligned} \tag{6}$$

The parameter c depends upon the thermal and viscous properties of the fluid and may attain value anywhere between 0.7 (for an ideal gas) to 1000 for oil. The value used here is the same as that used in Lorenz' original model, c=10. The parameter b is 8/3 and is related to the geometry of the convection cell (the horizontal diameter compared to the height). Finally, the parameter a is the key parameter controlling the nature of the fluctuation and is a measure of the temperature differential between the (hot) bottom of the convection cell and the cool top. For values a<1, the system settles to the equilibrium at 0,0,0 (while the other two equilibria are unstable). For values 1<a<24.5, the equilibrium at the origin is unstable while the other two equilibria are stable. For larger values of a, including the value in the model, a=28, the model exhibits deterministic chaos (see below).

Model behavior

Figure 7 shows a phase plot of the model behavior in X-Z space. The model has three unstable equilibrium points, indicated in the figure,

$$\begin{aligned}&X = Y = Z = 0,\\ &X = Y = \sqrt{a-1}, z = a-1,\\ &X = Y = -\sqrt{a-1}, z = a-1.\end{aligned} \tag{7}$$

The behavior traces out a characteristic butterfly shaped attractor, where the system initially moves in expanding oscillations around one equilibrium point. When the oscillations become sufficiently large, the system jumps across the Z-axis to start oscillating around the other symmetric equilibrium point. The further away from one equilibrium point the system was when leaving it, the closer it moves to the other equilibrium point as it starts its oscillations around that point. The result is a "strange" attractor (see, e.g., Ott 1993), which never returns to the same point in state space, yet shows characteristic behavior patterns (expanding oscillations) intercepted by seemingly random jumps between the two regions of state space.

<u>Insert figure 7 about here</u>

Feedback loops

The model has six independent feedback loops, three constant-gain first order loops connected to each state variable and three loops with variable gains: $\{\dot{X} \to X \to \dot{Y} \to Y\}$, $\{\dot{Z} \to Z \to \dot{Y} \to Y\}$ and $\{\dot{X} \to X \to \dot{Z} \to Z \to \dot{Y} \to Y\}$. It was our original hope that an eigenvalue analysis would demonstrate how the expanding oscillations were caused by a set of complex conjugate eigenvalues with positive real parts. However, this turns out not to be the case.

LEEA

Figure 8 shows the movements of the three system eigenvalues during the same time period as Figure 7. Although it is not obvious from the figure, an animation of the movements of the eigenvalues reveals that during most of the time when the system is oscillating around one equilibrium point, two of the eigenvalues are indeed complex conjugates, but they move around in the sync with the movement in statespace. (The third eigenvalue is always real and thus moving along the horizontal axis, making it difficult to discern its motion in the figure). Thus, the eigenvalues change a great deal during each oscillation, obtaining both positive and negative real parts (which would generally control the rate of expansion or contraction of the oscillations) and also showing great variation in the imaginary part (which would generally control the frequency of oscillation).

<u>Insert figure 8 about here</u>

It seems clear, therefore, that the nonlinearities in the system play an active role at all times, making eigenvalue analysis invalid. Another way to express it is that the relative strength of the different loops shifts rapidly around during each revolution of the simulation. This is clearly seen in Figure 9, which plots the gains of the three loops that have variable gains against each other (the gain of loop X-Z-Y on the horizontal axis and the gains of loops X-Y and Y-Z on the vertical axis). Clearly, the loop gains shift around substantially with each revolution of the oscillations.

<u>Insert figure 9 about here</u>

It is possible that one might focus on a single cycle of oscillation and look at the shifts in loop dominance etc. over the course of that cycle. Similarly, one might look at the dynamics during a single shift from one region of state space to the other. This, however, is not much different from examining the behavior of the model in state space. It seems that not much additional insight is gained from the loop analysis, and this approach fails to address the characteristic butterfly attractor. Thus, one must conclude that certain nonlinear dynamic phenomena are not readily explained by eigenvalues or feedback loops.

Mass' economic model

Results of the analysis of this model are not available at this time due to unexpected difficulties in implementing the analysis in the software, but will be ready for presentation during the Conference. At that time we also expect to have an updated version of the paper reflecting the results from this analysis.

However, we do have some preliminary issues and results raised. First, it is striking that, even though the model has only 17 state variables, many of which are simple first-order delays, and 94 variables, the total number of feedback loops is 5,548! When one considers that some parameters in the model are switches, which only attain values 0 or 1, the number of loops is reduced somewhat (to 2,578 when all switches are set to 0). The size of the independent loop set is 69, illustrating the substantial reduction in loop set from the requirement of independence.

Second, a technical issue appears when there are repeated eigenvalues (i.e., when the eigenvalues are not distinct), both because the eigenvector space may no longer span the total state space and because it is no longer possible to distinguish modes of behavior that are conceptually different but numerically equal, making interpretation of the results more difficult. The former issue requires using generalized eigenvectors. The latter issue may be addressed by making small parameter changes in time constants etc., so that all eigenvalues become distinct, although we consider this to be somewhat unsatisfactory.

Toolbox

As it currently stands, the toolset to perform the eigenvalue elasticity analysis is limited and the analysis of model behavior is a cumbersome process requiring four separate pieces of software and multiple translations among data formats. Models were originally built, tested and simulated in Vensim (Ventana Systems 2002). We conducted all the eigenvalue analyses using implementations of the routines to identify and analyze feedback loops in Mathematica versions 5.0 and 5.1 (Wolfram 2003), having generated the representation of the Vensim models into Mathematica through a Perl (www.perl.org) script. Vensim simulation output was also imported to Mathematica to calculate the gains, elasticities and eigenvalues. Finally, data were exported from Mathematica into Excel to generate the graphs to facilitate the analysis and reporting of results (see Figure 10 for a detailed flow diagram of the analysis process).

Insert figure 10 about here

Additionally, the need to linearize the model to estimate eigenvalues severely constrains the range of model formulations that can be processed – table functions represent a significant challenge – requiring significant amount of work within Vensim to make system dynamics models amenable to the existing tools.

While we are not certain yet that the proposed method will yield the kind of insights that would warrant incorporation of this toolset into the standard system dynamics packages, we believe that making these tools widely available and as easy to use as possible will encourage practitioners to join the efforts to develop analytical methods for interpreting model behavior. Towards that goal, we are making available the tools developed to date in an integrated website: http://iops.tamu.edu/faculty/roliva/research/sd/leea/. We expect that over the next few months more and more steps of the analysis process will be automated, and that the site, through the functionality of webMathematica[1], will eventually have the ability to perform the analysis in without the need of any specific software beyond a web browser.

Conclusions

The purpose of the paper was to explore the application of the LEEA method to three different models of quite different type. The results show that the utility of the method depends upon the nature of the dynamics. Thus, one must make a distinction between different kinds of nonlinear behavior.

In *single-transient behavior*, as seen in the industrial structures model, the behavior is typically caused by a relatively small number of one-time shifts in loop dominance. The analysis of the Industrial Structures model revealed that the method can reveal subtleties in the dynamics and shifts in loop dominance, but that the picture can easily become quite complex. The conclusions regarding dominant structure was to a large extent in accordance with the findings from the PPM analysis, even though the two methods are very different. The main drawback of the LEEA approach is that it does not relate behavior modes directly to the time paths of given system variables. The next step in developing the LEEA method would therefore be to use eigenvectors to determine how much a given mode is expressed in the behavior of the variable.

In *linear and quasi-linear oscillations*, as seen in the Mass economic model, behavior is dominated by a fairly constant set of loops, involving major negative loops with delays. In the quasi-linear case, cycles of different lengths (or, more generally, cycles and other modes) may interact, as it is seen between the interaction of the business cycle and the long wave in the National Model (Forrester 1989; Forrester, Mass, and Ryan 1976), where business cycles tend to grow larger and more severe during downturns in the long wave. Linear and quasi-linear oscillations is probably the area where the LEEA method will prove most powerful. However, a number of technical issues remain, relating to non-distinct eigenvalues and to the computational burden of large-scale models, not to mention how one may develop tools to sift through a very large number of loop elasticities.

A particular form of nonlinear oscillation is the *limit cycle,* such as occurs in the simple long wave model. Kampmann demonstrated (1996) how it was possible to use the LEEA method to gain deeper insight into the behavior of that model by dividing the cycle into phases that were controlled by particular feedback loops, in effect breaking the cycle into short transients.

Finally, there are clearly cases of *chaotic behavior* where the method is fundamentally unsuited to yield insights. The analysis of the Lorenz model showed how the eigenvalues moved rapidly around in sync with the system variables. Using eigenvalues as explanations of system movement requires that they remain relatively constant while the system variables move. The fact that the eigenvalues move around so much in the Lorenz model came as a surprise to us, because we had expected that the gradually expanding oscillations occurring during much of the simulation were the result of a relatively constant pair of complex eigenvalues with positive real parts. In some ways, the LEEA did yield insight into the model, because it clearly showed that the behavior is governed by highly nonlinear dynamics at all times.

[1] http://www.wolfram.com/products/webmathematica/.

References

Barlas Y. 1989. Multiple Tests for Validation of System Dynamics Type of Simulation Models. *European J. of Operations Research* **42**(1): 59-87.

Ford DN. 1999. A behavioral approach to feedback loop dominance analysis. *System Dynamics Review* **15**(1): 3-36.

Forrester JW. 1969. *Urban Dynamics*. Productivity Press: Cambridge MA.

Forrester JW. 1989. The System Dynamics National Model: Macrobehavior from Microstructure. *Proceedings of the Computer-Based Management of Complex Systems: International System Dynamics Conference*, pp.3-12. Stuttgart.

Forrester JW, Mass NJ, Ryan CJ. 1976. The System Dynamics National Model: Understanding Socio-Economic Behavior and Policy Alternatives. *Technological Forecasting and Social Change* **9**(1/2): 51-68.

Forrester JW, Senge PM. 1980. Tests for Building Confidence in System Dynamics Models. *TIMS Studies in Management Science* **14**: 209-228.

Forrester N. 1982. A dynamic synthesis of basic macroeconomic policy: Implications for stabilization policy analysis. PhD Thesis, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA.

Gonçalves P. 2003. Demand bubbles and phantom orders in supply chains. PhD Dissertation, MIT, Sloan School of Management, Cambridge, MA.

Gonçalves P, Lerpattarapong C, Hines JH. 2000. Implementing formal model analysis. *Proceedings of the 2000 Int. System Dynamics Conference*. Bergen, Norway.
Goodman MR. 1974. *Study Notes in System Dynamics*. Productivity Press: Cambridge MA.

Kampmann CE. 1996. Feedback loop gains and system behavior (unpublished manuscript). Summarized in Proceedings of 1996 Int. System Dynamics Conference (p. 260-263). Cambridge MA.

Lorenz E. 1963. Deterministic nonperiod flow. *Journal of Atmospheric Sciences* **20**: 130-141.

Mass NJ. 1975. *Economic Cycles: An Analysis of Underlying Causes*. Productivity Press: Cambridge MA.

Meadows DH, Meadows DL, Randers J, Behrens WW, III. 1972. *The Limits to Growth: A Report for the Club of Rome's Project on the Predicament of Mankind*. Universe Books: New York.

Mojtahedzadeh MT, Andersen D, Richardson GP. 2004. Using *Digest* to implement the pathway participation method for detecting influential system structure. *System Dynamics Review* **20**(1): 1-20.

Oliva R. 2004. Model structure analysis through graph theory: Partition heuristics and feedback structure decomposition. *System Dynamics Review* **20**(4): 313-336.

Oliva R, Mojtahedzadeh MT. 2004. Keep it simple: Dominance assessment of short feedback loops. *Proceedings of the 2004 Int. System Dynamics Conference*. Oxford, UK.

Ott, E. 1993. *Chaos in dynamical systems.* Cambridge Univ. Press: New York.

Richardson GP. 1995. Loop polarity, loop dominance, and the concept of dominant polarity. *System Dynamics Review* **11**(1): 67-88.

Senge PM. 1990. *The fifth discipline*. Doubleday/Currency Press: New York.

Ventana Systems. 2002. *Vensim 5.0 User's Guide*. Ventana Systems: Harvard MA.

Wolfram S. 2003. *The Mathematica Book*. 5th ed. Wolfram Media, Inc.: Champaign, IL.

Appendix

Equations for simple industrial structures model

Industrial Structures = INTEG (new industries-demolitions, 10)
[Units: Structures]
new industries = Industrial Structures * effect of water adequacy on growth * NORMAL_GROWTH
[Units: Structures/Year]
NORMAL_GROWTH=0.12
[Units: Dimensionless/Year]
demolitions = Industrial Structures * DEMOLITION FRACTION
[Units: Structures/Year]
DEMOLITION FRACTION = 0.05
[Units: Dimensionless /Year]

Water Reserves= INTEG (-water consumption, 10000)
[Units: Water]
water consumption= water demand* effect of water availability on consumption
[Units: Water/Year]

Water demand = Industrial Structures*WATER DEMAND PER INDUSTRY
[Units: Water/Year]
WATER DEMAND PER INDUSTRY = 10
[Units: Water/Structure/Year]
REFERENCE WATER RESOURCE COVERAGE = 10
[Units: Year]

effect of water availability on consumption = 1/(1 + EXP(-6.98405* (Water Reserves/water demand/ REFERENCE WATER RESOURCE COVERAGE - 0.4715)))
[Units: Dimensionless]
effect of water adequacy on growth = 1/(1 + EXP(-6.98405* (water consumption/water demand - 0.4715)))
[Units: Dimensionless]

Figures

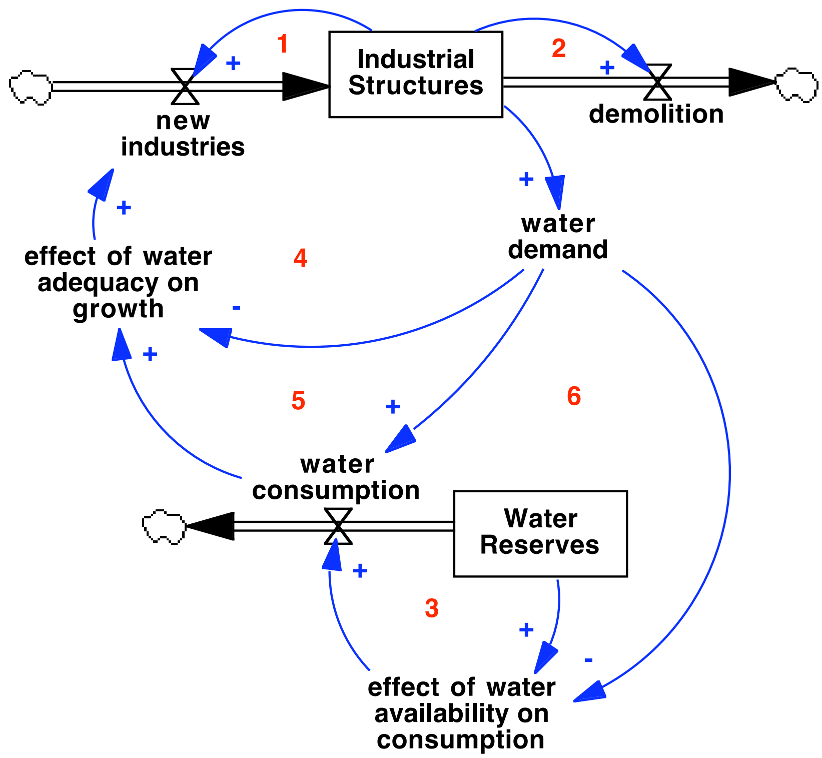

Figure 1
Flow diagram and feedback loops of the simple industrial structures model

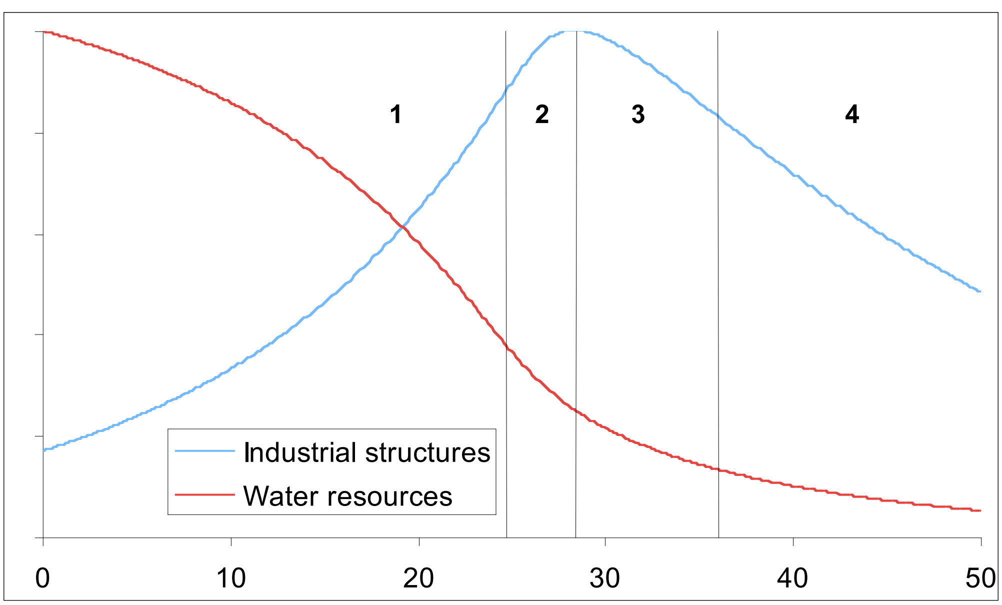

Figure 2
Behavior of the simple industrial structures model.
Numbers indicate phases identified by PPM method.

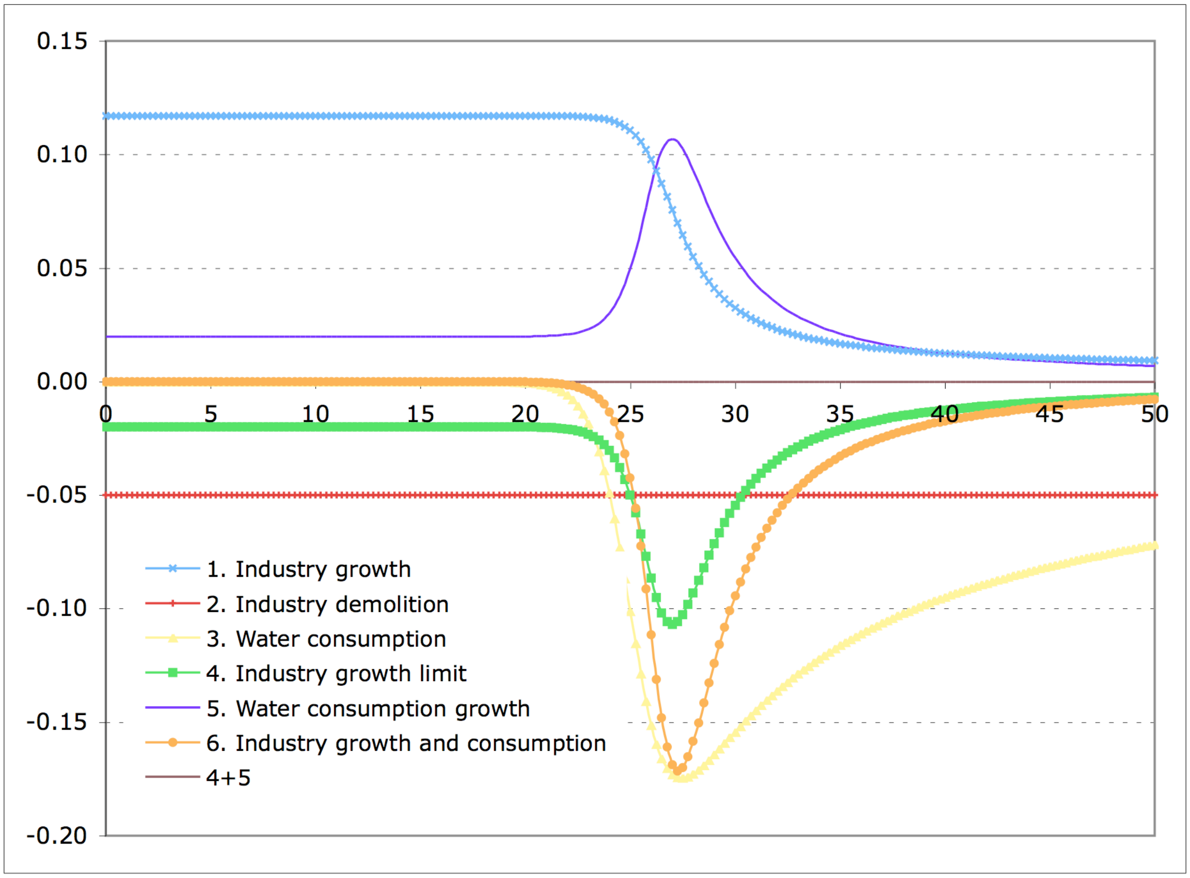

Figure 3
Feedback loop gains, industrial structures model

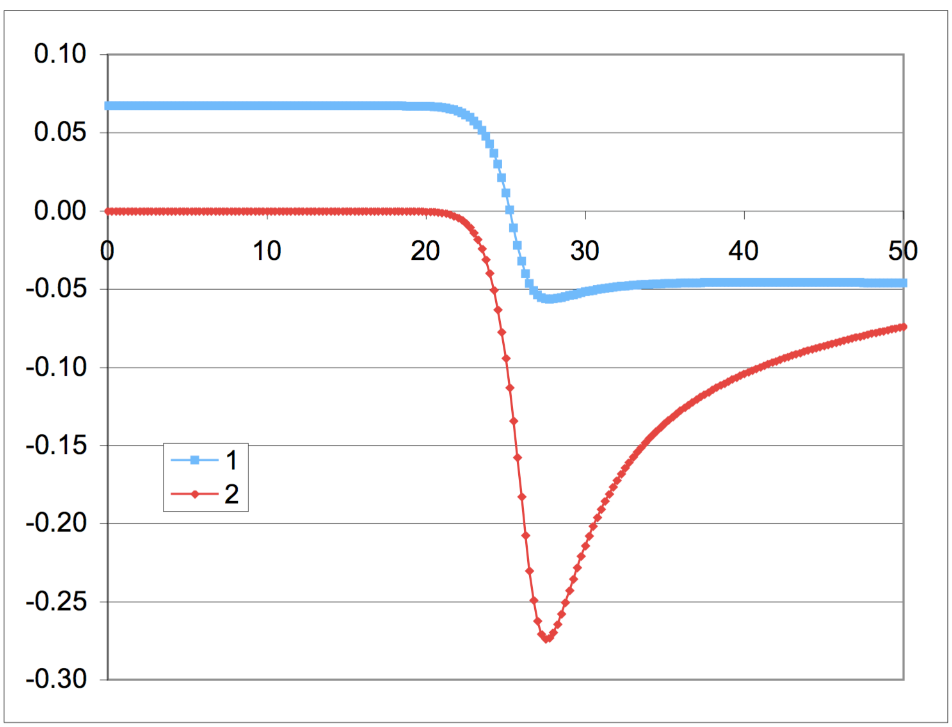

Figure 4
Eigenvalues in the industrial structures model

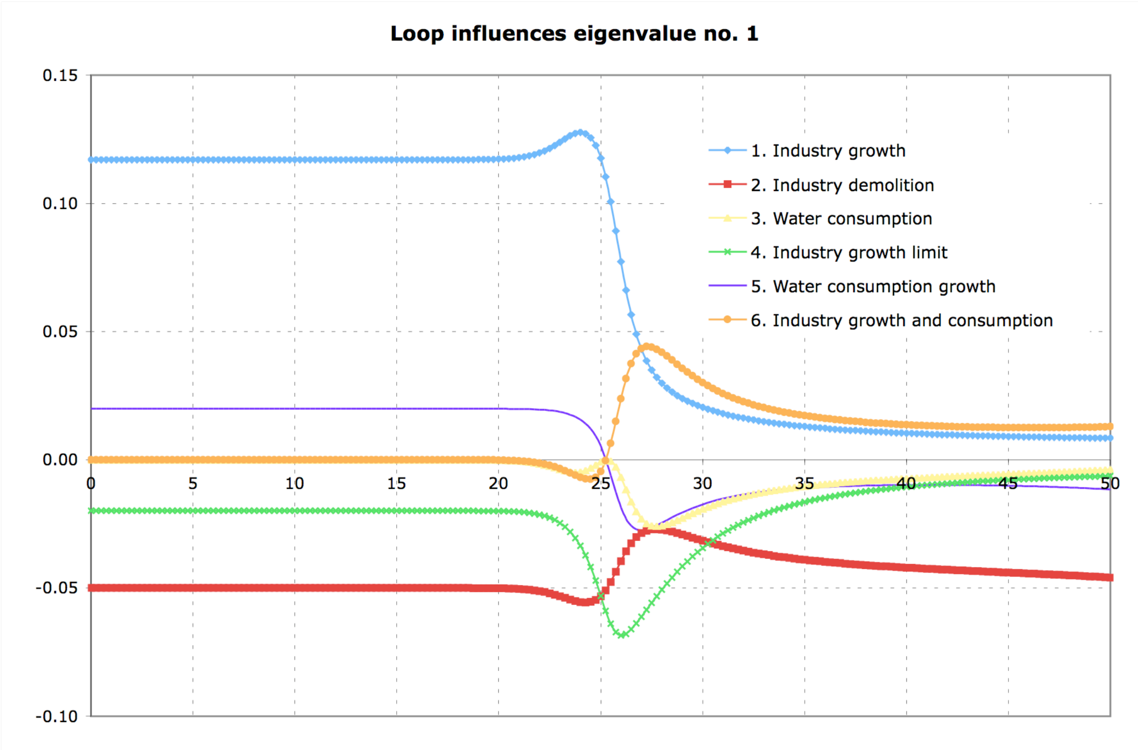

Figure 5

Loop influences, $\frac{\partial \lambda}{\partial g} g$, for the largest eigenvalue (no. 1), industrial structures model

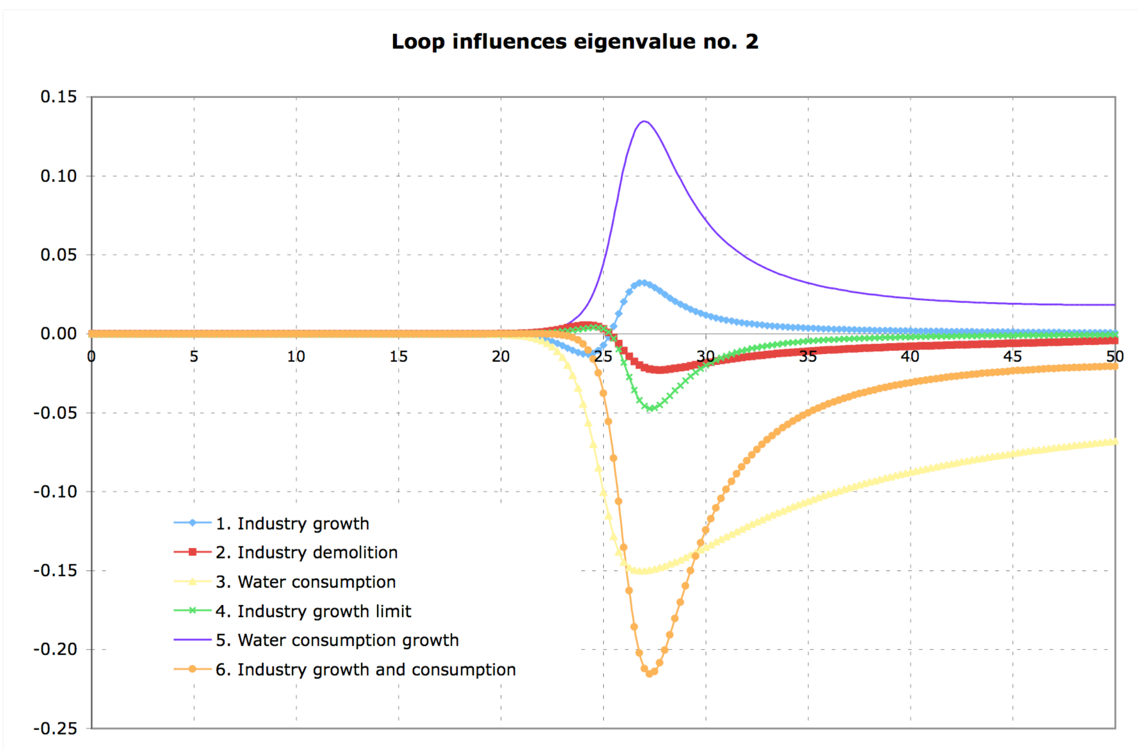

Figure 6

Loop influences, $\frac{\partial \lambda}{\partial g} g$, for the smallest eigenvalue (no. 2), industrial structures model

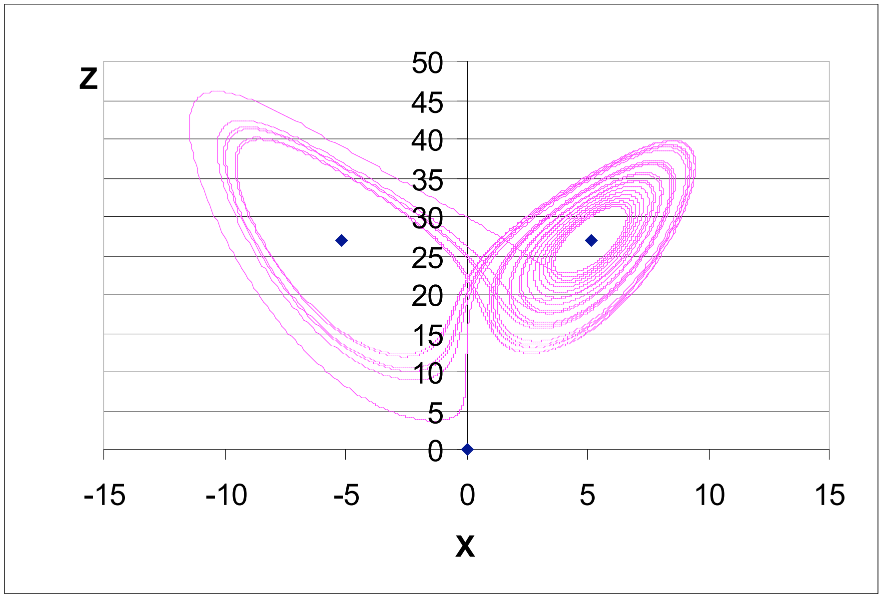

Figure 7
Phase plot of the behavior of the Lorenz model (X vs. Z).

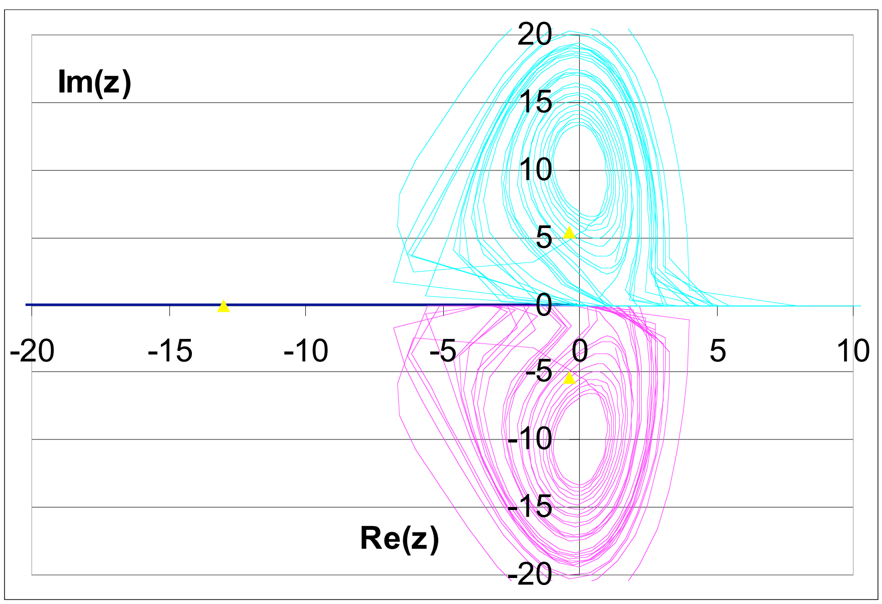

Figure 8
Movements of eigenvalues in the Lorenz model

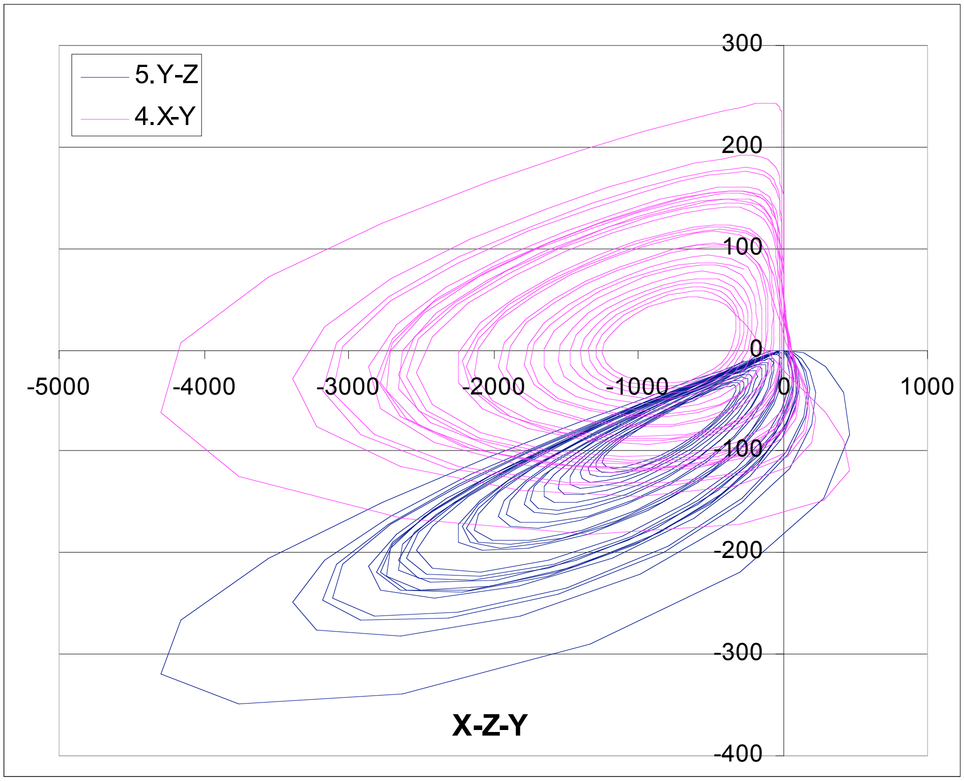

Figure 9
Movements of loop gains in the Lorenz model

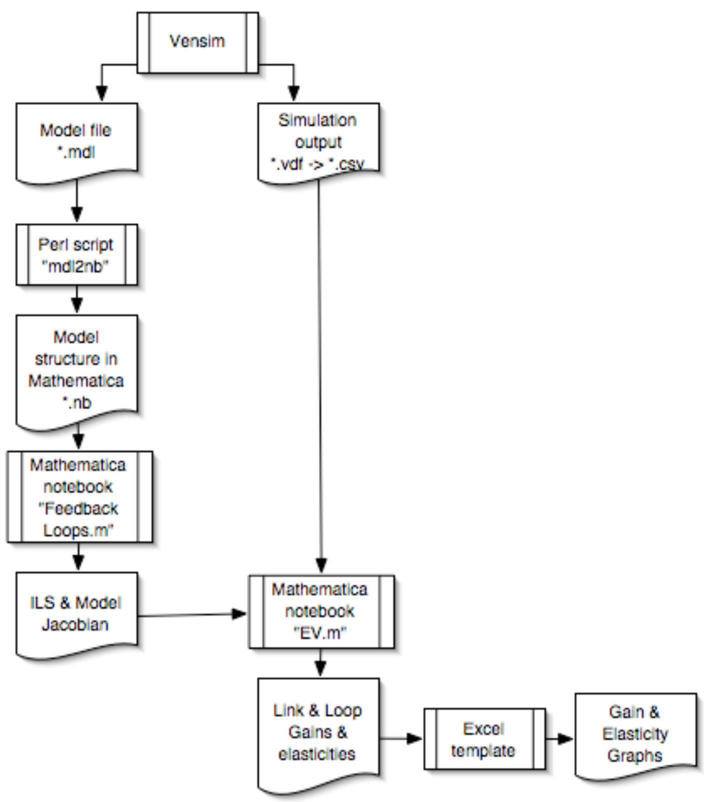

Figure 10
Process flow for Loop Eigenvalue Elasticity Analysis