STELLA Modeling as a Tool for Understanding the Dynamics of Earth Systems

David M. Bice
Department of Geosciences
Pennsylvania State University
University Park, PA 16802
dbice@geosc.psu.edu

Abstract

Earth system science represents an important new way of looking at our planet, but it is difficult to help students learn, in an experiential mode, about the complex dynamics of Earth systems. Here, I describe how the computer program STELLA can be used to construct and then experiment with a variety models to illustrate some important concepts of systems dynamics. A very simple model of a bathtub with a faucet and a drain serves to illustrate a wide range of systems concepts including residence time, response time, lag times, and feedback mechanisms. Variations on the bath tub model provide examples that illustrate the problem of model complexity vs. simplicity. A more complicated model of the global carbon cycle is used to demonstrate one means of model validation, testing the model against the historical record of CO_2 buildup in the atmosphere, using as input the historical record of fossil fuel emissions and land use changes.

Introduction

One of the most important recent developments in the Earth sciences has been the move towards looking at the Earth as a large, complex, interconnected system. This field of study, Earth system science, is exciting not just because it is new (in fact, it is not that new — see the appendix of Strahler and Strahler, 1989), but because it deals with relevant issues such as global climate change and also because it is highly integrative. Conversations with undergraduates at Carleton College have revealed, perhaps not surprisingly, that societal relevance and integration of different disciplines are characteristics students desire in their education. For these and other reasons, Earth system science is an attractive new emphasis in many undergraduate Earth science and geology courses. Moreover, the systems approach to understanding whole Earth dynamics was identified as the major objective for the future of the Earth sciences by the National Research Council (NRC, 1993)

But how does one go about teaching modeling in Earth system science? Regardless of how you choose to incorporate Earth system science into geology classes, you face challenges associated with reviewing and/or learning the relevant physics, chemistry, biology, and meteorology that are needed to understand global systems such as the carbon cycle or the global climate system. Other challenges arise when you search for ways to go beyond the passive description of how various Earth systems are constructed and how they operate. Although dynamic systems surround us, we are generally not very good at understanding and predicting their behavior without the aid of some type of model that allows experimentation and observation. In this paper I describe — in the form of several examples — how STELLA can provide an effective means for enabling students to understand and explore the dynamics of Earth systems. This amounts to developing the systems part of Earth system science.

The task of using computer models to explore the dynamics of Earth systems is made much easier through the use of a program called STELLA (High Performance Systems, Inc., http://www.hps-inc.com/), which runs on both Macs and PCs. STELLA is essentially a graphical programming environment that is specifically designed for modeling systems that can be described by a series of differential equations. Users construct graphical representations of systems and then enter values and expressions corresponding to the components of the system. The program uses this information to construct a set of differential equations that are integrated over time using standard numerical techniques. The algorithms used in STELLA are the same used by most programmers to carry out the numerical integration of differential equations. STELLA models seem to be extensively used in ecology and economics (Hannon and Ruth, 1994), but relatively few earth scientists (Mayer, 1990; Moore and Derry, 1995) have made use of this program in teaching and research. This is somewhat surprising given the extensive use of box-models in geochemistry (Garrels et al., 1973, Walker, 1991, Richter and Turekian, 1993) and the obvious possibilities for modeling systems such as the rock cycle, the water cycle, and the carbon cycle, among others.

It is of course possible to understand how a system like the global water cycle works without modeling, but this understanding is largely static and qualitative. Modeling provides a means for deepening this understanding by exploring the dynamics in a quantitative manner. For example, one could look at a typical depiction of the global water cycle and think about the consequences of groundwater withdrawal, irrigation, and reservoir building on the whole system. This thinking cannot progress too far beyond some qualitative conjecture without the aid of a numerical model; but armed with a model, one can understand quantitatively the impacts of these changes on global sea level and others parts of the system.

Introduction to the Modeling Process with STELLA

In order to illustrate some fundamental aspects of modeling with STELLA, we begin with a very simple system — a tub of water with a faucet and drain. Details for constructing the model can be found at http://www.geosc.psu.edu/~dbice/DaveSTELLA/ modeling/ch2.3.html.

1. From Real World to Conceptual Model to Computer Model

The first step in modeling is to define and consider the system as it exists in the real world. This involves identifying the components of the system, the material or entity that is moving through the system, the major processes involved in moving this material or entity, and the other quantities that these processes depend on. This step necessarily involves simplifying the real world because it is obviously impossible to model the full complexity of nature. Drawing a cartoon of the system, as shown in Figure 1, is an important part of this process. When possible, constructing an actual physical model of the system is very helpful (Moore and Derry, 1995) in understanding the dynamics of the system as well as the connection between the physical model and the STELLA model.

The next step in the modeling process involves representing the various inflows and outflows in the form of equations. The two flow processes in this case are represented by the faucet and the drain. The faucet flow is simply an adjustable constant; the drain flow is instead a function of the size of the drain opening and how much water is in the tub, because the weight of the water overlying the drain determines the amount of pressure that is forcing the water through the drain. A more precise description of this dependence is provided by Torricelli's Law, which states that the velocity of water flowing out of a drain is given by:

$$v = \sqrt{2gh} \tag{1}$$

where v is velocity, g is gravity, and h is the water depth. The velocity is then multiplied by the area of the drain opening to give a discharge in volume of water per unit of time.

Figure 2 shows how this system is represented in STELLA using the four building blocks of systems — reservoirs, flows, connectors, and converters (nomenclature of HPS Inc., the authors of STELLA). The connector arrows represent dependence; the depth of water is dependent on the amount of water in the tub and the area of the tub base. The converters represent either constants or variables defined by equations or graphs. Note that the two flows have cloud symbols at the ends away from the tub, indicating that this is an open system, drawing water from some unspecified source, and sending it to an unspecified sink at the other end.

The system described in Figure 2 is essentially a simple differential equation that has the general form of:

$$\frac{dW}{dt} = F \quad kW \qquad (2)$$

Here, W is the volume of water in the tub, F is the faucet flow rate, and k is a constant that incorporates the tub dimensions, the drain dimensions, and gravity (combination of equations shown in Figure 1).

The STELLA model depicted in Figure 2 is complete and ready to run after we specify the initial amount of water in the tub, the inflow rate, the length of time of the simulation, and the increments of time (the time step, or dt of the calculations) over which the program does the calculations. The user can also specify different integration techniques, with the default being Euler's method. The time step (adjustable via a program menu) is critically important, and before any model results can be seriously considered, it is important to try varying the time step to see how the model results change — if the results do not change significantly when you decrease the time step by one-half, then the time step is at an acceptable value; otherwise, the time step must be reduced to smaller values.

2. Using the Model to Illustrate Systems Concepts

There are a number of important concepts of dynamic systems behavior that can be illustrated with this simple bath tub model. These concepts are useful in understanding all dynamic systems, even very complicated ones such as the global climate system.

Steady State

Running the model for 100 seconds (Figure 3), reveals some interesting changes as the system evolves towards a **steady state**, where the amount of water in the tub stays constant. The inflow stays constant over time, but the outflow undergoes an exponential change, increasing until it approaches the same value as the inflow. Initially, the inflow is greater than the outflow, and so the amount (and depth) of the water in the tub increases, thus increasing the drain velocity and the outflow; this continues until the inflow and outflow match, at which point water continues to move through the system, but the amount of water in the tub remains the same.

Residence Time

When the system is in the steady state, we can define another concept — the **residence time**. The residence time is effectively the average length of time that an entity — in this case a water molecule — remains in a reservoir. It is really only meaningful for a reservoir that is at or near a steady state condition. By definition, the residence time is the amount of material in the reservoir divided by either the inflow or the outflow (they are equal when the reservoir is at steady state). If there are multiple inflows or outflows, then we use the sum of the outflows or inflows to determine the residence time. For the water tub system shown here, the residence time is:

$$t_{residence} = \frac{\text{amount in tub}}{\text{outflow}} = \frac{10.16 \text{ m}^3}{1 \text{ m}^3\text{sec}^{-1}} = 10.16 \text{ sec} \qquad (3)$$

If we increase the flow rate, the water moves through the reservoir faster, so the residence time decreases. It is possible then, to calculate any of the above three parameters (residence time, reservoir amount, and inflow or outflow) if the other two are known and if we assume the system is in a steady state. For instance, if we assume that the water cycle is in a steady state and we know much water is in the atmosphere and the global rainfall, then we can calculate the residence time of water in the atmosphere. Residence time is an important concept in problems of pollutants in ground water or surface water reservoirs, and also in understanding the long-term effects of greenhouse gases added to the atmosphere.

Response Time

A closely related concept is that of the **response time** of a system, which measures how quickly a system recovers and returns to its steady state after some perturbation. This concept can be illustrated by running several simulations in which the starting amount of water in the reservoir is varied and while the inflows and outflows remain unchanged. The results, shown in Figure 4, are somewhat surprising; regardless of how great the initial departure from the ending steady state, this system gets to the steady state at about the same time.

The concept of response time can be more clearly illustrated by an even simpler system, consisting of one reservoir (W) and one outflow (D) defined as a simple fraction (k) of the amount in the reservoir (W) such that at each instant in time,

$$D = kW \tag{4}$$

In this case, the response time is defined as $1/k$ — this turns out to be the time needed for the system to accomplish 63% (1/*e*) of the return to its steady state. Also note that in a very simple system, the residence time and the response time are numerically the same even though they are mathematically and conceptually different (the former reflects a steady state, the latter reflects a transient state). The response time is a very useful concept because if it is known (or hypothesized), we can make some predictions about how quickly the system will respond to a change in a reservoir. Alternatively, if we change one of the inflow or outflow processes, we can predict how that will affect the response time of the system. The concept of response time is important in understanding the future of the global carbon cycle — if we halt the anthropogenic alterations to the carbon cycle, the response time of the system tells us how long it would take for the carbon cycle to return to a more natural state.

A more detailed discussion of the concepts and mathematical formulation of residence time and response time can be found in Rodhe (1992) or through the web site given at the end of this paper.

Another important observation to be made here is that the system evolves to the same steady state in each case, so the steady state of a system (along with the response time) is primarily determined by the nature of its inflows and outflows.

Feedback Mechanisms

This particular system returns to a steady state because it contains a **negative feedback mechanism** in the connection between the drain flow rate and the amount of water in the tub. A negative feedback mechanism is a controlling mechanism, one that tends to counteract some kind of initial imbalance or perturbation. A familiar example of another negative feedback mechanism is a simple thermostat in a home that responds to changes from the steady state, returning the home to a specified temperature. Note that the word negative, as used here, means that this mechanism acts to reverse the change that set it into operation. If our tub is in its steady state, knocking the system out of its steady state by suddenly dumping in more water will cause a response — the drain will increase its flow rate, thus decreasing the amount of water in the tub, bringing it back towards the steady state value. If we instead decrease the amount in the tub, the negative feedback associated with the drain will force the amount of water in the tub to increase until the steady state is returned. The important thing to remember is that negative feedback mechanisms tend to have stabilizing effects on systems.

In contrast, a **positive feedback mechanism** is one that exacerbates some initial change from the steady state, leading to a runaway condition — it acts to promote an enhancement of the initial change. One way to modify the simple bath tub system to create a positive feedback system is to alter the system (Figure 5) so that the inflow depends on how much water is in the tub.

This system has one possible steady state, where the initial amount of water in the tub is 5 m^3; any departure from this value, however slight, leads to runaway behavior and the amount of water in the tub follows an exponential curve of the form:

$$W(t) = \frac{d}{k} + \quad W_0 \quad \frac{d}{k}\, e^{kt}, \qquad (5)$$

where W(t) is the amount of water in the tub at any time, d is the drain constant (1 m^3 /s in our case), k is the faucet rate constant (0.2), W_0 is the initial amount of water in the tub (variable in the 3 model runs shown in Fig. 5), and t is time. A useful thing to remember with exponential growth is that the doubling time can be easily calculated (after some manipulations of the above equation):

$$t_{doubling} = \frac{.693}{k}. \qquad (6)$$

Here, the doubling time works out to be 3.465 seconds.

If one runs this model with an initial water tub value of less than 5 m^3, you learn another important lesson, which is that in the real world, there are limits to exponential change, regardless of whether that change leads to growth or decline. In this case, the limit is reached when there is no more water left in the tub. In the case of population growth, the limit to exponential growth is reached when the carrying capacity of the ecosystem is approached — then the population growth decelerates and gradually approaches the carrying capacity (with reference to the human population, see Cohen, 1995, and Meadows et al., 1992 for detailed discussions).

Positive feedback mechanisms, like negative feedback mechanisms are not necessarily good or bad. Epidemics and infections have positive feedback mechanisms associated with them, but so does the growth of money in a bank account with compounded interest. The Earth contains a wide variety of both positive feedbacks and negative feedbacks and depending on the conditions, either kind of feedback may dominate. But — and this is very important — the mere fact that we exist, the fact that our planet has water and an atmosphere is compelling evidence to suggest that ultimately, our Earth system is dominated by negative feedback mechanisms (see Kasting, 1989 or Lovelock, 1988 for more discussion of this). However, it is equally important to realize that human time scales are much shorter than the history of the Earth and over periods of time that interest humans, positive feedback mechanisms may be very important; they have the potential to produce dramatic changes.

Lag Time

We next consider a slightly more complex system to illustrate the concept of a **lag time**. In Figure 6, two water tubs have been connected such that the drain from one flows into the adjacent tub. Here, the drains have been simplified greatly — all of the drain parameters in our first model are represented by a rate constants labeled k_1 and k_2; these rate constants get multiplied by the amount in the reservoir at any time to give the volumetric rate of flow out of the drain. Next, we take advantage of a useful feature of STELLA — the ability to define various parameters as graphical functions of other system variables or time. In this case, I want to show how the system responds to a sudden spike in the faucet flow rate, so I first define the faucet rate as being equal to time, then click on a button in the dialog box and a graph appears, enabling you to define the nature of this graphical relationship. The faucet in this case starts out with a value of 1 m^3/sec, which puts the system in a steady state to begin with, then increases to a peak value of 4 m^3/sec, and then quickly returns to 1 m^3/sec again and stays there for the duration of the experiment.

The response of the system is shown in Figure 6. The first tub reaches its peak 1 second behind the peak in the faucet — it lags behind the faucet. The second tub peaks at a lower value and much later than the first — the perturbation is buffered by the first water tub. Note that both tubs return to their steady state after variable amounts of time and that the total area under the two curves is equal, meaning that the same volume of water moved through each reservoir. The pulse of extra water, propagating through the system is very similar to the pulse of water moving down a stream system, with the lag time in this analogy being the time between the peak in the rainfall and the flood peak along a certain reach of the stream. The concept of a lag time is also relevant to systems such as the global carbon cycle — anthropogenic additions of CO_2 into the atmosphere are similar to the faucet here; the climatic response involves a certain lag time. This lag time means that if we halt emissions today, the climate will continue to warm — a fact that many policy-makers and citizens should be aware of.

Model Simplicity vs. Complexity

Models are clearly meant to represent simplified versions of the real world, yet still complex enough capture the essence of the real system. So, what is too simple, and what is simple enough? One way to understand this is by way of three variations on the water tub model, shown in Figure 7. In the simplest version of this system, the faucet and drain flows are simple constants — they do not change over time. The more realistic model uses Torricelli's Law to express the rate of flow out of the drain. The intermediate model simply represents the drain flow as the product of a rate constant times the volume in the tub. From the results (Fig. 7), it is clear that the simplest model does a poor job of representing the real behavior of this system; it drops off at a constant rate and does not achieve a steady state except in the special case where the inflow is set equal to the outflow. This simple system does not have any negative feedback associated with it either. So, this is a good example of a model that is too simple. In contrast, the intermediate model does a remarkably good job of matching the behavior of the more complex model. It is fair to say, then, that the intermediate model is simple enough and yet not too simple — it captures the essence of the more complex model using a more parsimonious mathematical representation. Nevertheless, the simple model is valuable as a starting point in the modeling process; its shortcomings provide suggestions for improvements and added complexity

It is worth noting that even the more complex model shown in Figure 7 is not overly complex. An overly complex model might, for instance, try to represent the friction, viscosity, and 3-D turbulent flow in the water, which would clearly represent overkill in this case.

Increased Model Complexity

Natural systems are usually more complicated than our simple water tub model, so it is important to explore some aspects of the dynamics of more complex systems. In the following discussion, we will pay particular attention to how the response time of the system changes with increasing complexity.

To begin with, imagine adding another drain flow to our simple water tub model. If we represent the two drain flows in the manner of the intermediate model of Figure 7, each with its own rate constant, k, we could show that the resulting response time of the model would be $1/(k_1+k_2)$, which is less than the response time with just one drain ($1/k_1$). The implication is that a system with more flows has a shorter response time — it can react to changes faster and moves to its steady state faster. As we'll see, this is not always the case in systems with connected reservoirs.

More complex systems consist of numerous reservoirs, linked together by various flows. In order to understand the effect of these connections, consider the case shown in Figure 8, where two isolated reservoirs with simple drain-like flows are then connected to each other such that the outflow of one reservoir becomes the inflow of the other reservoir, but the flows are mathematically the same and the starting amounts in the two reservoirs are also the same. The result of connecting these reservoirs is to decrease the response time of the system to $1/(k_1+k_2)$, as shown in Figure 8. Another way of saying this is that the connected system is more responsive than the disconnected system.

This last point is often true, but only to a certain point, as illustrated in the next example in which 4 reservoirs with drain-like flows are connected (Figure 9). We can imagine this system in its disconnected state and consider the response times of the individual reservoirs with their draining flows; these response times, given in Figure 9, range from 1.67 to 5 years. If we were to simply extend our analysis of the rate constants, we might hypothesize that for this four-reservoir model, the response time of the whole system would be given by:

$$\tau_{connected} = \frac{1}{k_{12} + k_{23} + k_{34} + k_{41}} = \frac{1}{1.5} = 0.67. \tag{7}$$

When we test this hypothesis by running the model, we find that the real behavior is more complex and the overall response time of the system is difficult to define, but if you focus on reservoir M2, which appears to be the slowest to respond, you can estimate that the response time of the whole system is in the range of 4 to 5 time units. Obviously, our simple mathematical prediction of the response time was wrong; when faced with a more complex system, it is better to define the response time of the whole system by direct observation.

By studying the evolution of this system, we can see that the reservoirs tend to overshoot their eventual steady-state values, delaying their arrival at the steady state. Why do they overshoot? One way of understanding this behavior is to study the apparent response times of each outflow, noting that they are not all the same and they alternate between short and long response times. The M3 resevoir is initially furthest from steady state and it has a shorter reponse time than M3, the next reservoir downstream, with the result that material piles up in M3 faster than it can adjust, moving it far away from its steady state. If we apply a similar analysis to the other pairs of reservoirs, we can better understand the why this system overshoots. Note that in this case, the response time of the whole system is significantly greater than the response times of some of the reservoirs considered in the disconnected state. Recall from our examination of the simple two-reservoir system that this was not the case there — the two reservoir system had a shorter response time than either of the reservoirs considered in the disconnected state. So, in larger systems, the overall response time of the system may be significantly greater than the response times of the separate parts of the system.

Figure 9 also shows the effect of adding more flows to this four-reservoir model — this amounts to increasing the connectedness of the system. Our simple analysis would lead us to believe that making this change will decrease the response time of the system, and in fact when we run this model, we find that such is the case. This further supports the idea that increased connectedness, which amounts to increased complexity, decreases the response time.

Validation and the Significance of Computer Models

The purpose of a model is not to replicate the real world — it is clearly impossible to put the complexity of the real world into a computer. Instead, the goal is usually to understand something about the behavior of a system, including how the system responds to changes. We create and use these models because their real-life versions are so complex, large, and often slow that we cannot generally understand them without some kind of controlled experimentation. But, the obvious simplification of models commonly creates leads to skepticism that ranges from complete rejection of anything the model reveals to a milder form in which the results are accepted as a good possibility for the way the real world behaves. So, how does one develop a sophisticated, nuanced appreciation for the significance of model results?

Computer models such as the ones shown in this paper are nothing more than a set of differential equations that are integrated over time. The significance of the model results are therefore dependent on the nature of the equations. Just as there is a range in the quality of equations from highly abstract to highly realistic, there is a range in the significance of the model results. If the equations were pulled out of thin air, then the results are not significant relative to any real world system (but they are nevertheless meaningful in a purely mathematical sense). If the equations are designed to express the general relationships of a real world system (e.g., the intermediate model of Fig.7), the model results might be only qualitatively meaningful. If the equations are designed and tested such that they mimic the important processes of the real world system, then the model results may be quantitatively significant, and the model might have some important predictive capabilities. A model such as this last type could be tested against known histories of the real world system; models that pass these tests are validated, and their results can be considered to be reliable, at least within a certain range of conditions.

The task of model validation is simple in some cases. For instance, the water tub model is easily tested by constructing a real version of the model (see Moore and Derry, 1995) and collecting some real data to compare with the computer model results. If this succeeds, then our model is validated and we could say that it captures the essence of the system; if not, then we must examine the model carefully, make revisions, and then re-test. Torricelli's Law, used to construct the more realistic version of the water tub model in Figure 7, has been tested and generally gives quite reliable results.

But for larger systems, such as the global carbon cycle (Figure 11), the process of model validation is a bit trickier — you clearly can't create a lab-version of such a vast, global-scale model. Instead, we have to rely on some kind of natural experiment in which we have some knowledge of an imposed change and some record of the model's performance or state over time. Unwittingly, humans have been conducting such an experiment and the state of the carbon cycle is partly available in the form of instrumental measurements since the late 1950's and ice core records of atmospheric CO_2 concentrations before then (Figure 12).

Figure 11 shows a STELLA diagram of a simple global carbon cycle that can be tested by adding in the anthropogenic carbon additions given in Table1. This model incorporates the best estimates of reservoir sizes and mathematical representations of the major processes, including a fairly detailed carbonate chemistry scheme and a detailed photosynthesis scheme. Details for constructing the model and downloadable versions can be found at http://www.geosc.psu.edu/~dbice/DaveSTELLA/Carbon/c_cycle_models.htm#top. The anthropogenic effects are added by making three changes to the system. The fossil fuel emissions, varying as a graphical function of time, are added straight to the atmosphere. The land-use changes represent carbon transfer from burning trees (75% of the total at each point in time) and through accelerating soil respiration (the remaining 25% of the total at each point in time). These transfers are represented by adding two new flows that go from the Land Biota and Soil reservoirs to the Atmosphere reservoir (not shown in Figure 11). As can be seen in Figure 12, the model results match the observed record of atmospheric CO_2 fairly closely, and so to a first approximation, this model passes a validation test. This means that models such as this — very similar to the one used by the IPCC in their exploration of the effects of various future emissions scenarios (IPCC, 1995) generate results that are quantitatively meaningful. This successful validation does not make this carbon cycle a perfect predictor of the future performance of the real global carbon cycle, but it probably does give a pretty good estimate of what will happen, so long as there is not some complex non-linear aspect of the system that kicks in once we stray too far from the states experienced in the past hundred years. In other words, the initial validation of the model earns it a status of cautious respectability.

Conclusions

Computer modeling offers the possibility of developing a deep understanding of the dynamics of Earth systems. While numerous computer programs could be used to this end, STELLA is particularly well-suited, especially for beginning students because of the ease of use relative to other programming languages. Simple models are easily created, modified, run, and observed, providing a superb tool for experiential learning. Experimenting with these modes, making changes and trying to predict the response, followed by comparing predictions with actual performance leads to the development of a fairly sophisticated, almost intuitive sense of system dynamics.

A surprisingly complete range of systems dynamics concepts can be illustrated through modifications of a very simple system that represents flow of water in and out of a tub. Although these models are far removed from the complexity of real Earth systems, they are simple to create, modify, and understand; in fact, more complex systems are often so challenging to understand that simple concepts may be obscured.

At one level, experimentation with computer models might have as its goal the development of general, abstract systems concepts. In this case, there is no need to consider the relation between the model results and the real world — the model is significant in and of itself, making no claims to represent the behavior of a real world system. But inevitably, students of Earth system science will want to explore systems with greater relevance to the real world, and this raises the question of whether or not the model captures the essence of the real world system. Various model validation approaches can help answer this question; for instance, the global carbon cycle model shown here did a fairly good job of matching the system response to the anthropogenic alterations over the past 100 years, and therefore seems to warrant a level of cautious acceptance.

Acknowledgements

This work represents the result of an on-going collaboration in learning with the many outstanding students at Carleton College. In addition, I thank my colleagues in the Geology Department, especially Cathy Manduca, for their insight and encouragement.

Much of this work was made possible through a grant from Carleton College.

References Cited

Andres, R.J., Marland, G., Boden, T., and Bischof, S., 2000, Carbon dioxide emissions from fossil fuel consumption and cement manufacture, 1751-1991, and an estimate of their isotopic composition and latitudinal distribution, *in*, Wigley, T.M.L., and Schimel, D.S. (eds.), The Carbon Cycle, Cambridge University Press, Cambridge, UK, 292 p.

Broecker, W.S., and Peng, H.-S., 1993, Greenhouse Puzzles, New York, Eldigio Press, 251 p.

Cohen, J.E., 1995, How many people can the Earth support?, New York, Norton, 532 p.

Garrels, R. M., McKenzie, F. T., and Hunt, C., 1973, Chemical Cycles and the Global Environment, W. Kaufmann, Los Altos, CA, 206 p.

Gifford, R.M., 1993, Implications of CO2 effects on vegetation for the global carbon budget, *in*, Heimann, M., (ed.), The Global Carbon Cycle, NATO ASI Series, v. 115, p. 159-200.

Hannon, B., and Ruth, M., 1994, Dynamic Modeling, Springer-Verlag, New York, 248 p.

Houghton, R.A., 2000, Emissions of carbon from land-use change, *in*, Wigley, T.M.L., and Schimel, D.S. (eds.), The Carbon Cycle, Cambridge University Press, Cambridge, UK, 292 p.

IPCC, 1995, Climate Change 1994, Cambridge, Cambridge University Press, 339 p.

Kasting, J., 1989, Long-term stability of the Earth's climate: Global and Planetary Change, v. 1, p. 83-97.

Kwon, O.-Y., and Schnoor, J.L., 1994, Simple global carbon model: the atmosphere-terrestrial biosphere - ocean interaction, Global Biogeochemical Cycles, v. 8, p. 295-305.

Lovelock, J., 1988, The Ages of Gaia, New York, Norton, 325 p.

Mayer, L., 1990, Introduction to quantitative geomorphology; an exercise manual, Prentice Hall, Englewood Cliffs, NJ, p. 380

Meadows, D.H., Meadows, D.L., and Randers, J., 1992, Beyond the Limits, Chelsea Green Publishing, Post Mills, Vermont, 300 p.

Moore, A. and Derry, L., 1995, Understanding Natural Systems through simple dynamical systems modeling: Journal of Geological Education, v. 43, p. 152-157.

National Research Council, 1993, Solid-Earth Sciences and Society, National Academy Press, Washington, D.C., 46 p.

Richter, F., and Turekian, K., 1993, Simple models for the geochemical response of the oceans to climatic and tectonic forcing: Earth and Planetary Science Letters, v. 119, p. 121-131.

Rodhe, H., Modeling Biogeochemical Cycles, in, Butcher, S.S., Charlson, R.J., Orians, G.H., and Wolfe, G. V., (eds.): Global Biogeochemical Cycles, San Diego, Academic Press, 379 p.

Strahler, A.N., and Strahler, A.H., 1989, Physical Geography, 4th Ed., New York, John Wiley and Sons, 562 p.

Walker, J.C.G., 1991, Numerical Adventures with Geochemical Cycles, Oxford, Oxford University Press, 192 p.

TABLE 1: HISTORY OF ANTHROPOGENIC CARBON ADDITIONS TO THE ATMOSPHERE

Year	**Fossil Fuel**	**Land-Use**
	Gt C/yr	**Gt C/yr**
0	0.365	0.6
10	0.548	0.6
20	0.837	0.65
30	1.030	0.65
40	1.046	0.7
50	1.190	0.7
60	1.620	0.8
70	2.543	1.1
80	4.006	1.3
90	5.172	1.25
100	5.941	1.5

Sources: Andres et al., (2000) and Houghton (2000). Year 0 corresponds to calendar year 1890.