COMPLEX ADAPTIVE SYSTEMS AND FUTURES THINKING: THEORIES, APPLICATIONS, AND METHODS

edited by

Linda Groff and Rima Shaffer

Special Issue *FUTURES RESEARCH QUARTERLY*

This special issue of *Futures Research Quarterly* is on the subject of "Complex Adaptive Systems and Futures Thinking: Theories, Applications, and Methods." The articles cover a wide range of topics from theories and models of complex, adaptive systems (CAS) to applications of complex adaptive systems models and thinking in different areas (including from macro system levels to micro system levels, from interrelated factors driving change of systems in our outer world, the inner world of the psyche and consciousness, and/or their interrelationships, as well as different methods for dealing with complexity of systems and life in the 21st century).

Though systems thinking and futures thinking are separate disciplines, there is a natural overlap between them, which all these articles explore. While some futurists look at change only within a specialized area, most futurists are big picture thinkers, making them also inherently dynamic, interdependent, complex, whole systems thinkers as well. Like complex adaptive systems (CAS) thinking, futurists have always had a model of reality that looks at the interrelationships between different variables, as these interact and change within a whole systems context over time. The overall evolution of different systems over time—including periods of slower change, as well as periods of faster change and evolution, and perhaps of crisis, disruption, and discontinuity, leading to breakdowns of systems, are often then followed by reorganization and breakthroughs to new emerging, larger, more complex system levels.

Rima Shaffer *is a futurist, organization developer and executive coach, Shaffer Synergistics, Inc., Washington, D.C.. She may be contacted at rima1shaffer@verizon.net.*

Linda Groff *is a professor of political science and future studies, California State University, Dominguez Hills, Carson, California. She may be contacted at ljgroff@csudh.edu.*

All of the authors of the articles in this special issue are futurists. Futurists have always looked at factors driving change, as well as crises within systems that propel evolution forward. Writings from wisdom cultures view reality as holistic and interrelated. In addition, wisdom cultures tend to take a long view of reality. Contemporary thinkers also look at phenomena through the lens of complexity and systems thinking, because all aspects of life are interacting on global, environmental, and extra-planetary system levels—not just local and national levels. Governments, trade, science, and economics are viewed on larger system levels today, with more diversity and complexity within them, due to various factors, including globalization of new technologies, major societal changes, an evolution of consciousness, and a number of different crises. Such crises indicate that systems that once worked are no longer working well, implying that solutions require new thinking and a reframing and reorganization of policies on larger, more complex, global and planetary system levels. Policy makers and decision makers are challenged to reframe problems and seek solutions from the perspective of larger system levels.

When systems are viewed as complex and emergent, linear thinking no longer suffices; and solutions may include both technical and materialistic/outer world variables and consciousness perspectives and influences. Each of these special issues includes both macro and micro, outer and inner reality perspectives—all from an evolving systems perspective, so readers can come to their own conclusions on their importance and how they may interrelate.

The articles in this special issue of *Futures Research Quarterly* will (1) look at systems and futures thinking from macro to micro levels, (2) from technical to philosophical perspectives, (3) from outer/materialistic to inner/consciousness worldviews and perspectives, and (4) the interrelationships between these different levels. We hope these articles will generate discussions amongst the authors, and amongst futurists in general, and between the fields of systems and futures thinking–about the changes and challenges, as well as opportunities, confronting humanity and the world today.

Teaching Systems Thinking

by

Peter Bishop

Systems thinking is a fundamental perspective of future studies. Even calling it a "perspective" underestimates its importance. Some even claim that it is *the* paradigm of futures studies. It is at least the lens through which futurists view the world.

Systems thinking embodies some of the principles that lie at the foundation of futures studies:

- Every entity (thing) is a system which consists of parts (subsystems) and which is also a part of larger systems—a "holon" to use Arthur Koestler's term (1968).
- Every system and every part of a system is connected to every other system, at least indirectly.
- Systems and parts of a system interact in ways that can produce surprising and counterintuitive results.
- The tendency to produce unexpected results makes predicting the outcome of systems' interaction difficult, if not impossible.

As a result, it is critical that futurists introduce students and others to these principles if they are to approach the future in a sophisticated and systematic fashion.

Unfortunately, teaching systems thinking is easier said than done. The subject is obvious to those who understand it and opaque to those who don't. Even those who don't *get it* might agree with these principles, and not *see* the world that way. Those who do see the world that way cannot understand why everyone does not. Teaching systems therefore requires communication across a deep paradigmatic boundary in a language that is quite foreign to the listener. That is very hard to do.

Chris Dede, now at Harvard, created the Systems Thinking course at the University of Houston-Clear Lake in 1975. Chris is an

Peter Bishop *associate professor, University of Houston; president, Strategic Foresight and Development, Houston, Texas. He may be contacted at pbishop@uh.edu.*

outstanding educational futurist and brilliant teacher; the course became a tradition. He said that Using Systems Approaches (the name of his course) was the hardest course he ever taught, and he was right. I hope that this reflection might tempt others to travel this journey themselves.

While the principles of systems thinking are embedded in most ancient philosophies, the theory of systems thinking was first articulated in the early 1930s by the biologist, Ludwig von Bertalannfy (1976). Since then, a library of literature has developed around the subject. Other notable contributors were Jay Forrester (*Industrial Dynamics,* 1961), Russell Ackoff (*On Purposeful Systems,* 1972; *Redesigning the Future,* 1974; *Creating the Corporate Future,* 1981), James Grier Miller (*Living Systems,* 1978), Karl Weick (*The Social Psychology of Organizing,* 1979), C. West Churchman (*The Systems Approach,* 1984), Peter Senge (*The Fifth Discipline,* 1990), and now Ken Wilber (*A Theory of Everything*, 2000).

The practical application of systems theory began during World War II in the work of two eminent scientists—Norbert Weiner and John von Neumann. Weiner is credited with articulating the fundamentals of control theory, also called cybernetics, in which negative feedback is applied to changes in a system to keep it within certain limits. The common household thermostat is the most obvious example. Control theory was the basis for the development of much more complicated systems in the Postwar world—from intercontinental ballistic missiles and nuclear submarines to computers and the Internet. Systems engineering has since emerged as a separate discipline with a deep mathematical basis and universal application to all machines.

Jay Forrester, also of MIT, was the first to apply control theory to social systems. Forrester also invented the formal language of causal models (also called influence diagrams) and systems dynamics, which allowed the simulation of first-order differential equations using simple difference equations. Forrester used these tools to describe the development of cities in his 1961 book *Industrial Dynamics.* Dennis and Donella Meadows and Jorgen Rander also used systems dynamics in their famous *Limits to Growth* in 1973. Forrester and his colleagues offered system dynamics to the public in the Apple IIe program called *Dynamo,* which Barry Richmond turned into *Stella* and *iThink* for the MacIntosh and which Ventana Systems turned into *Vensim* for the Windows computers. Today high school students (and probably some elementary students) can simulate quite sophisticated systems using these simple tools. Forrester's tradition became the inspiration for Peter Senge's groundbreaking book *The Fifth Discipline* in 1991 and influenced John Sterman and others at the MIT Systems Dynamics Group. Finally, the Systems Dynamics Society is a well-known and prestigious society of researchers who use these theories and tools today.

John von Neumann, Weiner's colleague, is also credited with establishing a different branch of systems theory based on cellular automata (CA). As opposed to cybernetic systems, in which variables are the components, von Neumann's systems consisted of independent agents (the CAs) whose actions depend on the conditions in their immediate environment and on the actions of other CAs close to them. What is now called complexity theory, or agent-based modeling, took longer to develop, since complex systems cannot be modeled using differential equations the way control systems can. They must be simulated in a step-by-step fashion, and the computers required to do any meaningful simulation did not become available until the 1970s. At that time, John Conway invented the famous *Game of Life*, a two-dimensional array of agents operating on very simple rules that produced surprising and beautiful patterns. Stephen Wolfram used a one-dimensional CA to investigate the various states that an agent-based system could take in a famous article in 1982 which he later turned into his book *A New Kind of Science*. The Santa Fe Institute was founded in 1984 to study complex adaptive systems, now that powerful graphical workstations from Sun Microsystems were available. SFI also pioneered the development of network theory, which became staple of many scientific and engineering disciplines.

The abstract (and somewhat arcane) systems theory of the 1950s has come to define our world and to influence the many technologies we have created within it. Earth scientists use systems theory to describe the operation of the inanimate parts of our planet—the oceans, the atmosphere, the land, and the energy that flows among them. Biologists use systems theory to describe living systems—organisms and the ecologies they live in. Psychotherapists use systems theory to describe the interactions among family members or small work groups. Futurists use systems theory to describe larger human systems – communities, organizations, regions, nations and indeed the whole of human society itself. Systems theory, then, is essential for understanding the world and how it might develop and change in the future.

Each course in the University of Houston futures curriculum begins with a course generalization. The generalization is a single statement that embodies the essential learning in that course. It is a vision statement of sorts about what we want the student to learn. The course generalization guides the selection and development of the modules in the course, with each module elaborating and reinforcing the generalization.

The generalizations for many of the courses are obvious and somewhat simplistic, but no generalization is as important as the one for Systems Thinking.

"A SYSTEM'S BEHAVIOR IS A FUNCTION OF ITS STRUCTURE."

Or as Peter Senge put it "Structure influences behavior." (*The Fifth Discipline,* 1990) That simple statement contains the essence of systems thinking, but first some definitions:

System:	a set of parts that interact to produce observable effects (behaviors) outside the system
Behavior:	a change in (or the stability of) an externally observable or measurable unit or quantity associated with (or produced by) the system over time
Structure:	the relationship of the system's parts (subsystems, variables or entities) interacting with each other according to fixed rules

In other words, a system's behavior is a function of the relation and interaction of its parts–its structure. As such, this generalization seems pretty obvious and therefore not too impressive, except for the fact that it is *not* the most common explanation of phenomena in the world. Two other explanations are more commonly advanced for why things (human systems, in particular) behave the way they do: the personal explanation and the external explanation.

The personal explanation claims that systems behave the way they do because of the people in them. According to this theory, people (such as leaders, managers, workers, suppliers, regulators, customers, etc.) account for the system's behavior. Change the people in the system (by retraining, supervising, or replacing them), and you will change the behavior of the system. "If we could only get rid of ..., If the boss would only think..., If the employees would only behave like..., If only ***they*** would do something, then everything would be all right." Systems thinkers claim otherwise; they hold that changing the people in a system rarely changes the behavior of the system.

The U.S. Congress has been around for more than two centuries. Tens of thousands of people have served over that time, yet the institution still seems to behave the same over time. Is it the people? Clearly not. And one could say the same for business, schools, churches, or families. The people in a system cannot explain the behavior of that system when that behavior persists long after those people are gone.

Another popular explanation for a system's behavior is that forces, external to and beyond the control of the system, cause it to behave the way it does. Laws, regulations, the market, the physical world are all used as reasons why the system behaves as it does. That of course does not explain how some systems operating in those same environments seem to behave differently. So some businesses succeed in a heavily regulated environment while others do not. The same can be said of almost any type of environment. Blaming external events for trouble is common, but again systems thinkers do not take that 'easy out' either.

People do make a difference and the environment does influence behavior, but not nearly as much as most believe. The situation is illustrated in the diagram below. While we acknowledge that a system's structure does influence its behavior, we rarely use the structure to explain the behavior because it is "underwater"–invisible and hard to see. People in the empirical West prefer to explain things using tangible evidence (people and events) rather than the apparently ethereal and largely invisible *structure of the system* (whatever that is!). Systems thinking "drains the water from the pond" in order see its structure and allow it to play its proper role in explaining the system's behavior.

A course in systems thinking provides the understanding and the tools to reveal the structure of a system and its effects on the system's behavior. The course achieves this mission by reading what others have said about systems, by reviewing cases of structural explanations of system behavior, and by modeling and simulating systems themselves. The ultimate objective is always to explain a system's behavior in terms of its structure.

FIGURE 1 – CONCEPT OF SYSTEMS THINKING

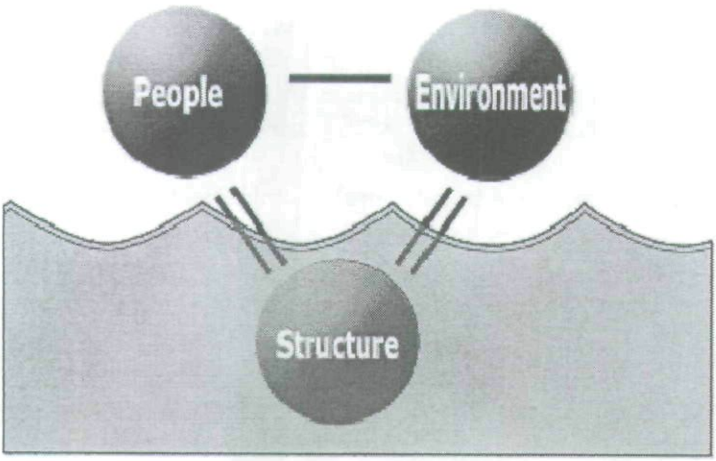

SYSTEMS

The concept of system is so big that it is hard to think of something that is not a system. Some examples of living systems are cells, organs, organisms, ecologies, families, organizations, communities, societies, and even the global society. On the inorganic level, atoms, molecules, crystals, oceans, atmospheres, solar systems, galaxies, machines, circuits, utilities (water, electricity, telephone) and, of course, the Internet are all systems.

Each of these entities has a number of things in common:

1. Each is made of parts.
2. The parts interact with each other.
3. The interaction of the parts produces behavior at an observable level. (The patterned interaction of the parts is the structure.)

Understanding a system and its behavior begins with constructing a model or representation of the system. Models come in various types physical, graphical, mathematical, verbal, and so on. Each has its own use, and most systems can be modeled in many ways. The model focuses on certain aspects of the system to explain the system's behavior. The model is always a simplified representation of the actual system because its simplicity demonstrates how the system operates. An ecological model of a lake would include the species but not the chemical bonds of the water molecules, because those are not required to explain the system behavior.

A system boundary delineates what to include and what not to include. What is left out is the system environment that part of the rest of the universe that interacts with the system and influences its behavior to some extent. In the long run, everything is connected to everything else, so boundaries are arbitrary. The boundary of a system is an analytical concept; it is not part of reality. Rather it is a device created by the analyst to improve understanding.

Establishing boundaries is arbitrary because there is no one way to define a system's boundary. Nevertheless, there are useful boundaries and useless ones. For example, Texarkana is one of the few towns in the United States that has a state boundary (Texas and Arkansas) running through it. That boundary is as arbitrary as any other boundary. It is useful when considering matters of state law and taxes that apply to its citizens. It would be harmful, however, to consider the two parts of the town as separate communities since they act as one system in every other way.

The rule for deciding a system's boundary optimizes two principles: 1) completeness—include all the parts in the model necessary to explain the system behavior, and 2) parsimony—do not include any more parts in the model than are absolutely necessary. The first rule is obvious. If one leaves out an essential part of the system, some of the behavior will not be explained. If one includes too many parts, the model will become too complicated to understand. As someone once said, "Replacing a system that is poorly understood with a model that is poorly understood is no progress."

SYSTEM BEHAVIORS

The central question of systems thinking, "Why does the world act the way it does?" is applied to one system at a time. The world is a complicated place, and we do not understand the half of why things are the way they are. Here are some examples that a class came up with one year in Houston:

- U.S. healthcare system, though the most advanced in the world, does not take care of everyone.
- People don't accept alternative medical treatments despite their proven successes.
- Welfare does not help those who need it the most.
- Although schools spend more money than they used to, students are exhibiting lower skill levels than they used to.
- Slash-and-burn agriculture continues.
- Arabs and Israelis cannot resolve their differences.
- NASA has spent a fortune on organizational consultants, but the culture remains the same.
- Politicians do not fulfill their campaign promises.

Not everyone would agree that all these statements are true. To the extent that they are, they represent a list of curious behaviors of the systems in our world. Systems that are designed to do one thing (health care, education) seem to end up doing something else. As a result, they do their intended mission poorly. Health care is really not taking care of healthy people, but rather treating sick people. It should be called sick care. We build roads, but traffic jams increase. We want security, but end up building 10,000 nuclear missiles. How do such things happen?

Take the experience of dieting. Most people believe that if they eat less, they will lose weight. Why do the people who diet continue to be the heaviest? They should be the lightest. Does anyone understand why this happens?

SYSTEM STRUCTURE

The most common explanation for the fact that heavier people usually don't benefit from dieting is that they lack will power–an explanation rooted in the people themselves. If they would only eat less, then they would lose weight. In fact, some people do eat less, but most don't. Are those that don't eat less therefore to blame for their overweight condition? Most people believe so.

The people themselves, however, have a different explanation. They believe that something outside them forces them to eat, usually identified as stress. That represents the second most popular explanation for a system's behavior–something outside the system is responsible. Businesses blame regulators, regulators blame legislators, legislators blame lobbyists, lobbyists blame regulators. Everyone has some external explanation for their behavior. This explanation is usually not adequate.

The final type of explanation is somewhat more accurate, but still not sophisticated. It is the simple cause or linear explanation. Einstein once said, "To all the complicated problems in the world, there is a simple solution, but it is always wrong." He appreciated how complex and subtle the world is. Simple explanations fail to capture complex reality. So obesity is caused by an eating disorder–nice and simple, but hardly adequate. Corruption is caused by greed; pornography by moral decline; poor educational performance by a lack of family values. All nice and simple, but hardly explanations to count on.

Take the solution of raising taxes to reduce the government deficit. Government deficit is the result of revenue that is less than expenditures. One way to solve the problem of deficits is to raise the tax rate to produce revenue to equal the expenditures–nice, simple straightforward. As many political leaders found out, that solution may not work. They raise the tax rate, and the revenues go down. They raise the tax rate again, and they revenue goes down again! How to understand this system behavior?

Understanding begins by listing the parts of the system that produce the behavior:

Revenues Expenditures Deficit Tax rate	Adjusted gross income (pre-tax) Net income (after-tax) Living expenses Savings
Gross profit (pre-tax) Net earnings (after-tax) Dividends Retained earnings	Investments Productivity Growth

An explanatory model of the system would point out that revenues are produced from two sources: businesses and individuals. Business tax rates apply to gross profits (business revenues less business expenses). The higher the tax rate the lower the net earnings after taxes. With a fixed dividend, the lower the retained earnings means the company has less to invest. The individual sector works the same way. Tax rates apply to adjusted gross income (individual income less deductions). The lower the gross income, the lower the net income and, with fixed living expenses, the lower the savings that would be used to buy stocks and bonds. Therefore, the higher the tax rate, the lower the investments from businesses and individuals. Lower investments lead to lower productivity which in turns leads to lower growth. Lower growth means lower profits for business and lower incomes for individuals resulting in lower revenues for the government. As a result, a higher tax rate leads to lower government revenues–just the opposite that one would expect.

The preceding paragraph is a verbal model of the government revenue system designed to explain the unusual result that higher tax rates may lead to lower revenues. That model would also explain that under certain circumstances, lower tax rates might even lead to higher revenues. That actually happened in the Kennedy administration in 1963. The Reagan administration tried the same thing in 1982, but it did not lead to lower deficits because government expenditures (mostly military spending) increased at the same time. In any case, the verbal model shows how it might happen. Most importantly, the explanation is 1) not due to any person or group of people involved in the system, 2) not due to forces outside the system, and 3) not a simple explanation from just one cause. It is an explanation based on the structure of the system; the interaction of its constituent parts.

THE APPROACH

So if the objective is to learn the course generalization and be able to apply it to explain system behaviors, how do we do that?

The first overriding consideration in designing this course is to distinguish between the two types of system structures–cybernetic and complex. As described above, cybernetic system theories and models are based on control theory; complex system theories and models are based on agents. Cybernetic models are macro, top-down, describing the system as a whole. Complex models are micro, bottom-up describing the actions of individual agents. Each of these paradigms will be described in turn. The approach to learning each paradigm consists of the following elements:

- Instruction: reading, lecture, discussion
- Demonstration: exercises, simulation
- Activity: practice, feedback
- Assessment: tests, products

The first step is, of course, instruction—reading and lecturing on systems theory and the ways to apply it in real situations. Systems thinking is a skill and some instruction is necessary, but the primary strategy is practice and feedback.

CYBERNETIC SYSTEMS

Literature on cybernetic systems theory

The best introduction to systems thinking is contained in two short books by Draper Kauffman titled (cleverly) *Systems I* and *Systems II*. Kauffman's books are deceptively simple. They might seem beneath a university course, but they contain all the important elements of systems theory in an engaging and easily understood manner. Who says that learning can't be fun, too?

The classic text in systems thinking is, of course, Peter Senge's *Fifth Discipline*. Senge not only introduces Forrester's insights about causal modeling, but he provides the rationale for why study systems on the very first page.

From a very early age, we are taught to break apart problems, to fragment the world. This apparently makes complex tasks and subjects more manageable, but we pay a hidden, enormous price. We can no longer see the consequences of our actions; we lose our intrinsic sense of connection a larger whole.

Part of that socialization is a model of how the world works, something cognitive psychologists call a "schema". Futurists point out that we also have schemas for the larger systems in the world—why sales go up or down, why crime occurs in certain neighborhoods, why wars erupt. Some of those schemas are well-supported by scientific evidence, such as the operation of the economy; others are little more than common sense and traditional wisdom.

Not everyone has the same schema or model for the same phenomena. Many schemas are deeply ingrained cultural constructs. These constructs become rote, unconscious, and unquestioned. It is only when we interact with people from different cultures or lifestyles that we realize that the world is made of all kinds of schemas, some apparently quite bizarre.

We also have different schemas or models for how the large systems in the world operate—the physical, biological, and human systems of the planet. For instance, some will disagree on whether nature is there just for human to use as they wish or whether it has independent status and value that must be respected. Schema guide decisions and actions toward nature, such as how people vote, and what teachers teach, what philanthropists donate.

Part of systems thinking involves surfacing the schemas and mental models that we and others use to understand and explain the world. The behaviors in that world are apparent, but the structures that produced those behaviors are not. So we need a tool, an X-ray machine of sorts, to expose those tacit structures. Once exposed, we can examine them, test them, discuss them, and ultimately come to understand how the world works in a conscious and explicit way not only for ourselves, but in communication and dialogue with others. Once we have revealed the mental models that we and others use, we can compare them and perhaps agree on how the world works or at least understand the different assumptions that each person uses to make sense of the world. One cannot discuss what one cannot say or show. Systems thinking provide the means to identify our deepest assumptions about the world so we can choose which ones we want to use.

DEMONSTRATION OF CYBERNETIC SYSTEMS THEORY

One of the most memorable parts of this course is the participation in simulation that concretely shows that a system's behavior *really* is a function of its structure.

The two most famous simulations are *The Beer Game* and *Fish Banks.*

The Beer Game is written up in Senge's *The Fifth Discipline*. It simulates a four-station supply chain in which retailers, distributors, wholesalers and manufacturers order and receive (or produce) shipments of beer based on their expected demand. Not to give away the plot, but the behavior at every station is almost always shortage followed by a huge oversupply because of the built-in delays in the system. Even when participants have heard or read about *The Beer*

Game, they still exhibit the same behavior! The behavior is a function of the structure, not of the participants or their knowledge.

The Systems Dynamics Society sells the materials for the board game (http://www.albany.edu/cpr/sds/Beer.htm). MIT (http://beergame. Mit .edu/) and MA Systems (www.masystem.com/beergame) offer online versions, and MIT also offers a simulator that plays the game automatically based on input parameters (http://web.mit.edu/jsterman /www/SDG/MFS/simplebeer.html).

Fish Banks is a simulation now distributed through the Sustainability Institute, a successor to the Institute for Policy and Social Science Research at the University of New Hampshire—the same people who produced *Limits to Growth*. The simulation consists of teams fishing in the same water, and produces the same behavior as *Limits*—overshoot and collapse. Even when the participants know about this scenario, the system usually produces the same behavior. In this case, the software is essential since it calculates and keeps track of all the variables in the system (http://www.sustainer.org tools_resources/games.html).

Many other activities and simulations are contained in the *Systems Thinking Playbook* (http://www.sustainer.org/tools_resources /games.html). Nothing is more powerful than demonstrating the power of the course generalization, particularly when the students themselves participate in the system and produce the behavior themselves.

MODELING CYBERNETIC SYSTEMS

Systems thinking is primarily a skill, not just an intellectual pursuit. Our professional program at Houston focuses on honing skills—by constructing models. A model is a representation of reality in some form. All types of models exist, including:

- Physical (scale) models
- Mathematical models (equations)
- Computer models (programs)
- Geographical models (maps)
- Process models (steps)

A model is like the reality, but it is not the reality. The map is not the territory. A model extracts only a limited number of parts of the reality for representation. The model focuses on those parts for better understanding and, in dynamic models, better manipulation in ways that cannot be done with the real system for both practical or ethical reasons.

The systems-thinking course distinguishes four types of models used to articulate the mental models of a system's structure – verbal, formal, simulated, validated.

Verbal models use ordinary language to explain the system's behavior using the system's structure. We really don't need any instruction on how to explain behaviors using language because we do it all the time. Language is highly flexible, but flexibility comes with a price. Language is also ambiguous. Different people can understand different things even when using the exact same words. So language is not a perfect way to articulate a mental model. In fact, there is no perfect way. Different types of models are useful for different purposes.

Formal models solve that problem, to some extent, because they use a formal language to describe the system structure in a precise and unambiguous way. Mathematics is a formal language, and it is used to model most systems in science and engineering. In social systems, however, we need a language that is somewhat more flexible and forgiving, so we turn to Forrester's causal models, also called influence diagrams. Causal models are composed of three types of entities:

- **Variables**—any quantity that can vary
- **Links**—the association of one variable with another
- **Loops**—circular sets of variables and links

Figure 2 shows a simple reinforcing, positive feedback loop that describes wage-based inflation as a function of the structure of the manufacturing system.

Figure 3 shows a simple balancing, negative feedback loop that describes adjustments to the price of gasoline as a function of the structure of the market.

The purpose is to show that a formal language is a way of describing mental models and systems structures more precisely than informal language. Causal models also take the individuals and the events out of the explanation. Any person in these systems is assumed to act in the same way. That is not exactly the case, of course.

Some manufacturers might not increase their wages to meet the cost of living, or they might move their factories overseas to prevent wage increases. Formal models do not ignore the possibility that people and events do influence system behaviors, but they do focus on the system structure as the explanation, since it is so rarely identified as such.

FIGURE 2 - REINFORCING (POSITIVE FEEDBACK) LOOP

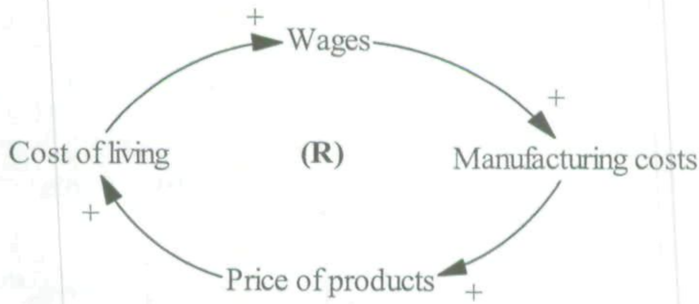

FIGURE 3 - BALANCING (NEGATIVE FEEDBACK) LOOP

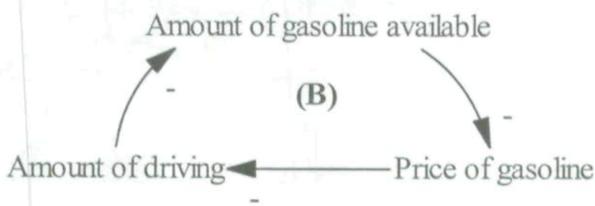

We use Virginia Anderson and Lauren Johnson's *Systems Thinking Basics* as the primary text for teaching causal modeling. Their publisher, Pegasus Communications, is also an excellent source for other materials on causal modeling.

Formal models solve the problem of the ambiguity of language, but they do not directly link the system behavior and its structure. Causal models are pictures, static pictures. We can say, "When A goes up, B goes up," but the picture does not do that itself. The next level of modeling actually produces behaviors as output.

Simulated models produce behaviors using a computer program. Any programming language can be used to simulate a system since they all produce output (values of a variable over time), and most depict those values in graphical form as well. The structure of a system can be modeled using the relationships of variables, and the behavior of the system is the numerical or graphical output of one or more of those variables. The specific target to be explained; is the behavior of a system as manifested in the changes of a variable over time, usually depicted in graphical form. So the model of a system explains why a particular variable acts the way it does, and that action is shown as a graph of the value of that variable over time.

Depicting the behavior of a system as the graph of a variable over time gives one the ability to perform experiments. We first identify the behavior of the system to be explained (in the form of a graph), model the system structure, simulate its operation over time using a computer program, produce the output of the variable to be explained in graphical form, and compare the first graph with the second. If they do not match, we know that we have not modeled the system correctly. If they do match, we have evidence that we *might* have modeled the system correctly.

We do not know that we have modeled the system correctly for sure because many models can produce the same behavior. We know that we have one of them, but only one. We can never be sure that it is *the* one that produced the behavior in the world. That is an assumption, and a pretty good one, barring evidence that another model is better, but it will always remain an assumption. Since the structure of the system is fundamentally unobservable, we can never know for sure that we have the right one. But one or models that produce the targeted behavior is better than none.

Jay Forrester developed another formal language, called stock-flow or systems dynamics, for simulating systems. Stock-flow models contain three types of variables:

- **Stocks**—variables that retain their value over time. They are like tanks that hold water.
- **Flows**—variables that adjust the value of stocks, either increasing (inflows) or decreasing (outflows) them. They are like the faucets and drains connected to the tank.
- **Auxiliaries**—variables that hold parameters or perform calculations during the simulation.

Figure 4 contains a classic stock-flow model of population change (absent immigration).

FIGURE 4 - STOCK-FLOW MODEL

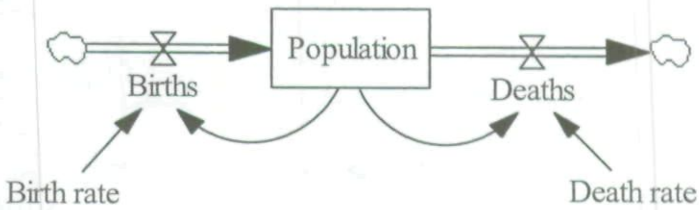

In this model, the number of individuals in the Population is the stock; it persists over time. Individuals enter the population by birth and leave the population by death (the flows). The rates of those flows are held in the birth and death rates (the auxiliaries). The actual number of births and deaths in any time period is the size of the Population times the respective rate.

This model can exhibit three different behaviors, depending on the relative size of the birth and death rates. The Population is stable (constant) when the rates are equal; the Population increases when the birth rate is higher than the death rate, and it declines when the birth rate is lower. Figure 5 shows the graph of Population increase from 1,000 to about 1,800 when the birth rate is 40 per 1000 and the death rate is 10 per 1000, as exists in many developing countries.

The purpose is not to teach systems dynamics or stock-flow models but to show that simulated models are useful in understanding systems thinking. We can verbally state how a systems structure explains a behavior using ordinary language and we can draw that structure using a causal model. However, there is no substitute for actually producing that behavior with a modeling program and comparing the output to the expectation. That is the real test of systems thinking.

At the same time, modeling is no easy task. Aside from getting the structure correct, it also involves finding the right formula for the equations and the right value for the parameters contained in the auxiliaries to produce a behavior that looks like the system's behavior in the world. So some knowledge of how variables in different equations behave and a lot of fiddling with parameters is necessary to get the behavior one wants. The reason for introducing simulation into an introductory course in systems thinking is 1) to demonstrate how simulated models work, 2) to examine the assumptions necessary in modeling, and 3) to show how the structure (the model) actually produces the behavior (the graph).

FIGURE 5 - POPULATION INCREASE

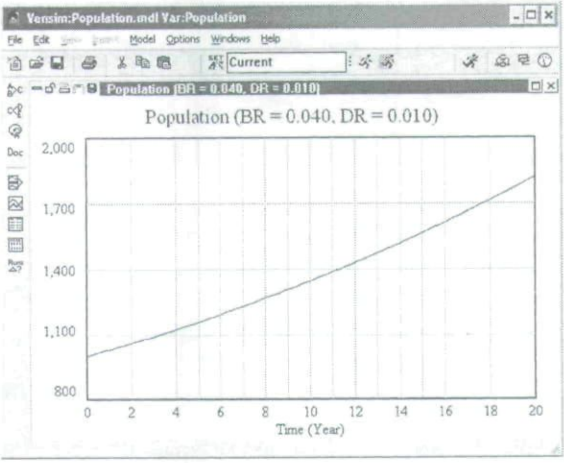

Numerous resources exist to learn systems dynamics. The best discursive introduction is probably Michael Radzicki's *Introduction to System Dynamics,* produced for the Department of Energy (http://www.systemdynamics.org/DL-IntroSysDyn/index.html). Jay Forrester's group has also produced a set of excellent tutorials called *The Road Map,* available at http://sysdyn.clexchange.org/road-maps/rm-toc.html. The definitive text for systems dynamics is probably John Sterman's *Business Dynamics (2000)*, but it is expensive.

Forrester originally programmed his stock-flow models on a computer program called Dynamo (for Dynamic Models). Barry Richmond, founder of High Performance Systems (now isee), developed *Stella*, a modeling program for the Apple Macintosh (http://www.iseesystems.com/softwares/Education/StellaSoftware.aspx). *Stella* also runs on Windows, but most use *Vensim* from Ventana Systems (http://www.vensim.com/download.html) because it is free for educators and students.

The purpose of simulation is to produce the *shape* of the system behavior, not the actual values. While real values are the output of the model, they are not necessarily the values that the variable would have in the world. Shapes are usually enough to understand and explain the behavior of a system. For prediction, we need to know, not only the shape, but also the actual values of those variables. For that, we turn to the final level of system modeling.

Validated or calibrated models produce not just the shape of the behavior, but also the values themselves. These models are "validated" because they are fitted to some historical time series to be sure that the structure, the parameters and the initial conditions of the model are correct before extrapolating the model into the future. Validated models go well beyond an introductory course in systems thinking. They are used extensively in physical science (such as modeling the effects of CO_2 and the other greenhouse bases in the global atmosphere) and economics (such as forecasting the growth of the economy over the next year).

The most famous validated systems model was called World3 in *Limits to Growth (1973)*. Published just months before the OPEC oil embargo, the model predicted long-term scenarios of overshoot and collapse for the world's economy. The original and the two subsequent revisions (*Beyond the Limits* and *The Limits to Growth: The 30-year Update)* make fascinating reading, but students in this course can get the essence from a small pamphlet entitled *A Synopsis: The Limits to Growth* (http://www.sustainer.org/tools_resources /games.html).

COMPLEX ADAPTIVE SYSTEMS

Complex adaptive systems (CAS), the term now used for von Neumann's approach to system structure, are based on cellular automata and independent agents. CAS was in its infancy in the 1970s when the UH-Clear Lake course was established. It took the development of more powerful computers before any meaningful agent-based models could be simulated. Even today, the materials, the demonstrations and the tools available to most people are many years behind what they are in cybernetic systems. CAS is basically where cybernetic systems modeling was in the 1970s—before *Stella/Vensim, The Fifth Discipline,* and *The Road Map*.

Nevertheless, a reduced treatment of CAS was introduced to the Houston systems course in the late 1990s. Today, about 20% of the course is devoted to CAS, because it is essential to understanding that a system's behavior is a function of its structure.

INSTRUCTION ON CAS THEORY

The first objective of this part of the course is to clear up the confusion surrounding recently-developed terms associated with the notion of complex adaptive systems. Coincidentally, all of them begin with "C"—chaos, catastrophe, criticality, and complexity. And regrettably, all have connotations in ordinary language that have little or no relation to their actual meaning in systems thinking. As a result, they are often thought to be other than they are.

Chaos is the first and most widely used term associated with CAS. It often appears with complexity, as in "chaos and complexity," just like "ham and eggs" or "peanut butter and jelly." It is similar to complexity since 1) it does begin with "C," 2) Chaos theory was devised after World War II, and 3) it is a type of system behavior that is unpredictable in the medium-term. But that is where the similarity ends.

Chaos is one of three types of behaviors that a system can exhibit, the first three of which are:

- **Fixed**–a static equilibrium state (e.g., the bottom of the ocean)
- **Periodic**–oscillations between two or more fixed states (e.g., the ocean tides)
- **Chaotic**–movement from one state to another, but never returning to any previous state (e.g., the surf crashing on rocks)

Chaotic phenomena were first identified by Henri Poincare in trying to explain the orbit of Neptune. Though considered the "Father of chaos theory," Poincare never did explain that behavior because it was chaotic.

The practical application of chaos theory was developed by Edward Lorenz, a meteorologist, in 1963. Lorenz was running a weather simulation that he had run before, but this time he interrupted the simulation and restarted it using the last numbers on the printout. He noticed, to his surprise, that the simulation produced entirely different results after the first few time periods compared to the first run. He thought he had entered one of the numbers incorrectly, but he had not. It turned out that he had re-entered the numbers using the first six digits that the computer was printing out, but the computer was actually calculating the numbers using ten digits internally. So the numbers on the restarted run were too small by less than 0.0001%; yet that incredibly small difference produced a significant difference in a relatively short time.

Prior to this discovery, there were thought to be only two types of systems – deterministic and stochastic. First developed by Galileo, Kepler, and Renaissance scientists and later perfected by Newton, deterministic systems acted according to fixed laws, expressed as mathematical equations. They could be used to predict the future state of the system within a fairly narrow range, leading Enlightenment philosophers to believe that we could know the future. Before that, however, some French mathematicians identified probability theory in the study of a game of chance. Stochastic systems, as they came to be called, are systems whose values are independent of each other. They form a distribution of possible outcomes, each with its own probability, but no one outcome could be predicted from the previous data or from the overall distribution. So deterministic systems were predictable; stochastic systems were not.

Lorenz discovered a third type of behavior, a deterministic system (a computer program) that was unpredictable due to its "sensitivity to initial conditions." In other words, the system is sensitive to the incredibly small difference in the initial conditions. And those differences rapidly build up to create large differences in output.

Given the same initial conditions in a computer simulation, the system will behave exactly the same way for as long as you run the simulation. In the real world, however, it is impossible to measure the initial conditions with infinite precision. There is always some measurement "error," some difference between the measure and the reality. It is that difference that builds up to produce a measurably different behavior after a short time.

Chaos behavior is often confused with stochastic behavior because they are both unpredictable. People think that chaotic behavior is disordered and random, when things get out of control, when nothing makes sense. "All chaos breaks out!" Chaos is not disordered or random; it is deterministic. One can predict the very next state with mathematical precision. One could even predict all future states if one knew the initial conditions exactly, but that is not possible. Those quite minor differences in the initial conditions produce measurable differences after a short while.

And, unlike stochastic systems, no system is inherently chaotic.

The weather is the best example of a system that displays chaotic behavior. Predicting the weather from one hour to the next is not very hard, more difficult for the next day, and just about impossible for the next week or two. Just three well-known equations describe the behavior of a weather system using only three well-understood variables—temperature, pressure, and humidity. Weather in the world is chaotic (deterministic but unpredictable), but the "weather" in a building could be stable or oscillating. There are no inherently chaotic systems; there are only systems that have the potential of exhibiting chaotic behavior.

These three types of system behaviors (fixed, periodic and chaotic) can be produced in the same system depending on the choice of parameters. Stephen Langton at SFI depicted these states in his "football" image.

FIGURE 7 - THREE TYPES OF SYSTEMS BEHAVIOR

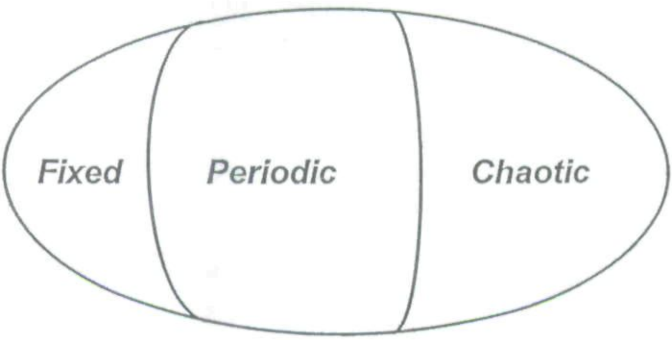

Certain human systems are thought to have chaotic behaviors although we do not have the equations to describe them. Markets of all types, especially stock and commodity markets, are thought to be chaotic.

The occurrence of chaos (in the mathematical sense) is an important part of systems thinking because it gives us reason to distrust predictions of future system behavior. Some of those predictions might come about, but we cannot tell which ones. If human systems are predominantly chaotic, then the results of intervening in those systems are inherently unpredictable. That does not mean that we should not act on those systems. Rather it means that when we do act, we should do so with caution and prudence lest we produce harmful effects that we did not expect or intend.

Stephen Wolfram's one-dimensional CAS models also produced a fourth type of behavior including interesting, even engaging, patterns that lasted for only a short time. They are not mathematically equivalent to the first three because they are only reproducible in CAS simulations. He labeled these behaviors "complex." Complex behaviors lie in a shadowy region between the periodic and the chaotic. Chris Langton from SFI called that region the "edge of chaos," another unfortunate, Madison Avenue label. That region, however, does contain some unique properties, most importantly a balance between order and disorder—enough order to keep the system together, and enough disorder to allow change and adaptation. For that reason, most believe that that behavior describes living systems, including social systems, very well..

FIGURE 8 - FOUR TYPES OF SYSTEMS BEHAVIOR

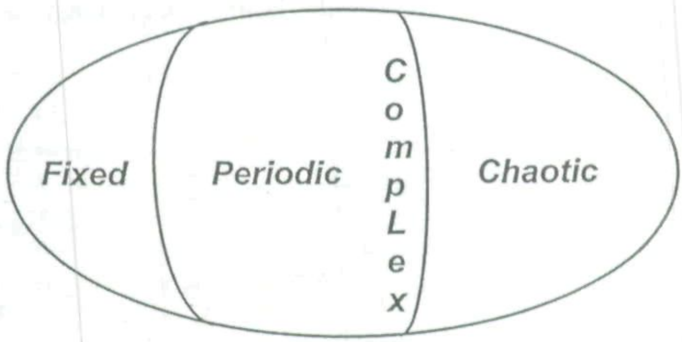

Before complexity, however, the star of the show, we have to stop by two other "C" words—catastrophe and criticality, which describe a different type of behavior from the ones considered so far.

Catastrophe and criticality are behaviors that shift suddenly from one stable state to another. Continuous behavior is smooth; it does not jump; all the points lie along a line. Discontinuities exist, however, in mathematics and in nature, and catastrophe theory and criticality describe those behaviors.

A simple example of discontinuity is a bottle that is stable sitting on its bottom. One can even push the top gently to one side and the bottle will return to an upright position, as long as it is not pushed too far. That range of variation in the vertical orientation of the bottle is called a basin of attraction. The image is that of a marble rolling around on a surface consisting of a number of bowls or depressions. If we tilt the surface, the marble rolls around in its bowl and returns to the bottom when we stop. But if we tilt the surface too much, the marble leaves the first bowl by going over a ridge and enters another in which it will stay. That is a discontinuous change.

Catastrophe theory was developed by Rene Thom to describe certain types of discontinuous change. The mathematics is quite complicated and the applications quite narrow, so few people actually learn and use the theory today. Wikipedia actually has a good description of Thom's catastrophe theory (http://en.wikipedia.org/wiki/Catastrophe_theory).

Criticality, on the other hand, is a common way of describing discontinuous behavior. The image here is "the straw that broke the camel's back." One piece of straw cannot do that, but when added one piece at a time, sooner or later the camel's back will fail, due to the addition of one piece of straw. The more common analogy is adding sand to a sand pile, one grain at a time. A sand pile is a cone whose sides form an angle that depends on the sand's viscosity (stickiness). Adding one grain of sand at a time allows the pile to grow beyond its natural angle, but only for a while. Sooner or later, one more grain will cause the pile to collapse in a little avalanche and return to the natural angle.

While neither of these models is worth covering in-depth in a course in systems thinking, it is worth mentioning because not all system behavior is continuous. Tipping points do exist, after which the system behavior changes dramatically. Examples of discontinuous change abound in physics, chemistry, biology, and in all of the social sciences:

- **Anthropology**–societal collapse
- **Psychology**–conversion
- **Economics**–asset bubbles bursting
- **Political Science**—revolution
- **Sociology**–white flight

The best book on criticality is Per Bak's *How Nature Works.* Bak and his coauthors introduced the concept in a 1988 article "Self-organized criticality" in *Physical Review.*

All of these terms are examples of a broader category of behaviors called non-linear dynamics. A system is linear when its output (behavior) is proportional to its input. The classic linear equation is y = kx, a straight line on a graph. One application is the relation of the force pulling on a spring to the distance the spring travels. "k" is the spring constant—larger for looser springs, smaller for tighter ones. The point is that doubling the force will double the distance; halving the force will halve the distance. The output is proportional to the input. It describes a linear system.

A nonlinear system occurs when the output is not proportional to the input. Technically, any curved line is nonlinear. So compound interest, which grows exponentially, is not linear because one year's interest late in the series returns more than one year's interest earlier in the series.

The importance of recognizing nonlinear behavior in systems thinking is that we are often surprised at nonlinear behavior, even though we can calculate the future of many of those systems exactly. Linear behavior seems somehow built-in and easy to imagine. When asked to draw a trend, most people will draw a line—equal amounts of change in equal time periods. On the other hand, exponential increase, diminishing returns, oscillation, and overshoot and collapse all seem harder to imagine and therefore more surprising when they do occur. And discontinuous change, the fundamental shift from one state to another, seems even harder.

It is more strange that nonlinear behavior is hard to imagine and expect because some would say that all change is nonlinear. In other words, change does not happen in a linear way. That point was made by Story Musgrave, a famous NASA astronaut in the Shuttle era, when he said that all the straight lines he could see on the Earth from space were man-made—contrails, ship wakes, roads, pipelines. Even the famous border between Israel and the Sinai desert is a straight line—green to the East and brown to the West. So with change. All systems behaviors are nonlinear. Getting used to that fact is one of the most important skills in systems thinking.

A **complex** system is one that consists of agents acting independently according to often simple rules based only on information from their local environment. Given that definition, complex systems are quite different from the cybernetic systems in classical systems thinking. The complex perspective takes the ground-level view of the individual agent; the cybernetic perspective takes the global view of the whole system.

TABLE 1 - CYBERNETIC AND COMPLEX SYSTEMS

Cybernetic	Complex
Macro behavior	Micro behavior
Top down	Bottom up
Rational and intelligible	un intelligible, unpredictable
Direct causal relations	No direct causality
Direct feedback	Reciprocal feedback
Explanation and prediction	Explanation but not prediction
Possibility of control	Surprising, creative, innovative
Model of mechanical systems	Model of living ecologies

At the same time, global patterns do emerge from local interactions. These patterns are called emergent because they emerge from the untold number of interactions that agents have with each other. There is no master control, no blueprint, and no overall rule book. Each agent acts according to its own rule book, yet order and pattern emerge nevertheless.

The clearest examples are biological organisms, which are fundamentally complex systems. Each cell is an agent acting on information in its local environment. Some cells, like axons, are long, so they transmit electrical impulses for relatively long distances, but all the inputs and the outputs, even of axons, are just local to that cell. Some organs send information to distant cells by releasing hormones or enzymes, but the distant cell only receives that information in its local environment. We think of our bodies as machines, designed and organized for life. But we can think of them just as readily as a colony of agents cooperating to perform that same function. The latter seems even more miraculous than the former.

Order arises even though there is no overall blueprint and no master control. In *The Ghost in the Machine (1968),* Arthur Koestler noted how wondrous it was that every person in Manhattan ate everyday even though the system that delivered that food (to all the homes, stories, restaurants, carts, etc.) was not planned or designed by anyone. It was an emergent property of the millions of interactions that constituted the food system of that city.

Most, though not all, complex systems exhibit emergence. And the emergent patterns cannot be explained or predicted from knowledge of the agents and their rules. Future emergent patterns are unpredictable, they may even be creative, generating new patterns that persist over time. The development of consciousness, the appearance of different species, and even life itself was an unpredictable emergent pattern based on the interaction of independent agents. Emergence is another reason to be humble and cautious when trying to understand, much less predict, the future of complex adaptive systems. They can easily surprise us.

The text usually used to investigate agent-based systems is *Harnessing Complexity* by Robert Axelrod and Michael Cohen. But a number of other excellent books on this subject are also available. Two histories of the development of complexity science are Roger Lewin's *Complexity: Life at the Edge of Chaos* and Mitchell Waldrop's *Complexity: The Emerging Science at the Edge of Order and Chaos*. They cover the same ground, but both have their own interesting stories and anecdotes about the characters that developed this field. And Stephen Levy's *Artificial Life* is another excellent treatment of the development of this field. John Holland is probably the best known theoretician of complex adaptive systems, genetic algorithms and artificial life so any of his books are always excellent, including his three relatively non-technical introductions—*Adaptation in Natural and Artificial Systems, Emergence,* and *Hidden Order*.

DEMONSTRATION OF CHAOS AND COMPLEXITY

The demonstrations of chaotic and complex behaviors are best done with simple computer programs that show these behaviors quite dramatically.

For chaotic behavior, the most complete set of computer simulations is from Rudy Rucker and is called The Chaos Game http://www.cs.sjsu.edu/faculty/rucker/chaos.htm. It runs a number of chaos and fractal routines that are quite amazing.

The Chaos Game with the magnets is also an interesting visual representation of chaotic behavior. Another even more dramatic example is the Waterwheel Lab, produced by Fritz Gasmann at the Paul Scherrer Institute in Switzerland http://people.web.psi.ch/gassmann/waterwheel/WaterwheelLab.html. It's an animation of the chaotic behavior that results from a constant supply of water to a waterwheel.

Jos Thijssen, a professor of computational physics at Delft University of Technology in the Netherlands, provides a simulation of self-organized criticality at http://www.tn.tudelft.nl/tn/People/Staff /Thijssen/sandexpl.html.

And finally, many have provided simulations of complex adaptive systems themselves, the most famous being John Conway's *Game of Life*. The *Game of Life* is a two-dimensional grid of cells each of which can assume two states—on or off—in successive generations. A cell turns on if three of its eight neighboring cells are on, and they stay on if two or three of its eight neighbors are on. Otherwise, it turns off. Simple rules, but complex patterns emerge. Some of those patterns and a list of the more popular programs can be found at http://en.wikipedia.org/wiki/Conway's_Game_of_Life. And Mirek Wojtowicz has assembled an amazing gallery of all types of cellular automata at Mirek's Celebration (http://www.Mirekwcom /ca /index.html).

Hundreds of programs demonstrate CAS behaviors. Two long lists are at Major Complex Systems Software from the Swarm Development Group http://oasis-edu.com/Oasis/synergie/accueil/soft .htm and the Artificial Life Section of the DMOZ Open Directory Project http://www.dmoz.org/Computers/Artificial_Life/. Some of my favorites are Boids by Craig Reynolds http://www.red3d.com /cwr/boids/ and Microants by Stephen Wright (http://www. calresco.org/sos/mants21.zip). Stephen Prata's *Artificial Life Playhouse* can be purchased second hand http://www.alibris.com/search/books/ It contains a number of genetic algorithms, including WordEvol http://www.jmu.edu/geologyevolutionarysytems/programs/wordevolexp.pdf.

MODELING CAS

Modeling programs for CAS have also existed for a long time. They are called event modeling programs because they program a series of events, like cars arriving at an intersection or products moving down a manufacturing line. The most highly developed agent-based modeling language for teaching systems thinking is NetLogo from the Center for Connected Learning (CCL) at Northwestern University http://ccl.northwestern.edu/netlogo/. NetLogo, like StarLogo offered previously by MIT http://education.mit. edu/starlogo, is a modeling language based on Logo, a programming language developed by Seymour Papert in the 1960s. (Papert played the same role in the development of agent-based modeling that Forrester played in cybernetic modeling.) Logo is language that controls a "turtle" on the screen that can move and draw lines. It is a rich and exciting programming environment.

StarLogo and NetLogo use the turtle concept, but rather than the program controlling one turtle, it controls many—each turtle being an agent in the simulation. Rather than programming the agents and their environment, MIT and Northwestern offer ready-to-use simulations that illustrate most of the important system behaviors and structures that one would like to investigate in a course like this. One can run some of these simulations right from a browser http://ccl.northwestern.edu/netlogo/models/ or download the NetLogo program and associated files http://ccl.northwestern.edu/netlogo/download.shtml and run them locally.

The CCL also has developed two variations of agent-based modeling, called Participatory Simulations (http://ccl.northwestern.edu/ps/) and Integrated Simulation and Modeling Environment (http://ccl.northwestern.edu/isme/) respectively. Both are server-based applications running the HubNet version of NetLogo (http://ccl.northwestern.edu/netlogo/hubnet.html).

Participatory Simulations allow students to interact with each other and with computer controlled agents using computers or TI graphing calculators. One of the simulations lets students control the traffic lights in a city grid to see how they can increase the flow of traffic in the grid.

The Integrated Simulation and Modeling Environment is another project that uses the HubNet application. The project's premise is very much the same as this course—that there are two paradigms of systems modeling today, cybernetic (or what they call aggregate) and agent-based.

These two forms of reasoning are very powerful ways of making sense of complexity in the world—yet, the communities who practice them and the literature describing them are largely separate and distinct. The aggregate and agent-based modeling tools themselves are deployed by different communities—each community focused on its tool and attendant form of reasoning. We believe that at both the cognitive level and the tool level the time has come for a synthesis of these two approaches. Accordingly, we explore how the two forms of reasoning complement each other in making sense of complexity and change—"Overview and Rationale," Integrated Simulation and Modeling Environment, The Center for Connected Learning, Northwestern University (http://ccl.northwestern.edu/isme /purpose.html)

Perhaps someday we will be able to teach systems thinking in an integrated manner.

CONCLUSION

This article has described systems thinking as taught at the University of Houston. As noted at the outset, the course generalization is the heart of this course.

A SYSTEM'S BEHAVIOR IS A FUNCTION OF ITS STRUCTURE

We explored the meaning of those terms (system, behavior, and structure), described the behavior in the form of graphs of key variables over time and modeled the structure using the cybernetic and CAS paradigms. The course teaches systems *thinking* with demonstrations and practice, as well as instruction to hone students systems thinking skills.

The major tenets include:

- Every thing is a system consisting of parts that is itself part of larger systems.
- Every system and every part is connected to every other system, at least indirectly.
- Systems and parts of a system interact in ways that can produce surprising and counter-intuitive results.
- The tendency to produce unexpected results makes predicting the outcome of systems' interaction difficult, if not impossible.

And once you see the world that way, you cannot see it any other way. The process of acquiring a systems perspective is irreversible. Once done, it's that way forever.

BIBLIOGRAPHY

Ackoff, R. *Redesigning the Future* (Hoboken, NJ: John Wiley, 1974).

Ackoff, R. *Creating the Corporate Future: Plan or Be Planned For* (Hoboken, NJ: John Wiley, 1974).

Ackoff, R., Emery, F., and Ruben, B. *On Purposeful Systems: An*

Interdisciplinary Analysis of Individual and Social Behavior as a System of Purposeful Events (Piscataway, NJ: AldineTransaction 1972).

Anderson, V., and Johnson, L. *Systems Thinking Basics: From Concepts to Causal Loops* (Walthum, MA: Pegasus Communications, 1997).

Axelrod, R., & Cohen, M. D. *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. (New York: Basic Books, 2001).

Bak. P. *How Nature Works: The Science of Self-Organized Criticality* (St. Emeryville, CA: Springer-Verlag Telos, 1999).

Bak, P., Tang, C., & Wiesenfeld, K. (1988). "Self-organized criticality," *Physical Review A*, Vol 38, Issue 1, pp 364-374.

Churchman, C. W. *The Systems Approach* (New York: Dell, 1984).

Forrester, Jay W. *Urban Dynamics* (Cambridge: MIT Press, 1970).

Holland. J. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (Cambridge: MIT Press, 1992).

Holland. J. *Hidden Order: How Adaptation Builds Complexity* (New York: Basic Books, 1996).

Holland. J. *Emergence: From Chaos To Order* (New York: Basic Books, 1999).

Kauffman, D. L. *Systems One: An Introduction to Systems Thinking* (Future Systems, 1980).

Kauffman, D. L. *Systems Two: An Introduction to Systems Thinking* (Future Systems, 1980).

Koestler, A. *The Ghost in the Machine* (New York: Macmillan, 1968).

Levy. S. *Artificial Life: A Report from the Frontier Where Computers Meet Biology*. (New York: Vintage, 1993).

Lewin, R. *Complexity: Life at the Edge of Chaos*. (University of Chicago Press, 2000).

Meadows, D. H., Meadows, D. L., Randers, J. & Behrens, W. *The Limits to Growth* (Signet, (1973).

Meadows. D. H., Meadows, D. L., & Randers, J. *Beyond the Limits: Confronting Global Collapse, Envisioning a Sustainable Future* (White River Jct., VM: Chelsea Green, 1993).

Meadows. D. H., Randers, J., & Meadows, D. L. *The Limits to Growth: The 30-year Update* (White River Jct, VM: Chelsea Green, 2004).

Miller, J. G. *Living Systems* (New York: McGraw-Hill, 1978).

Radzicki, M. *System Dynamics Tutorial* Department of Energy. http://www.systemdynamics.org/DL-IntroSysDyn/index.html

Senge, P. *The Fifth Discipline* (New York: Doubleday, 1990).

Sterman. J. D. *Business Dynamics: Systems Thinking and Modeling for a Complex World* (New York: McGraw-Hill/Irwin, 2000).

Von Bertalanffy, L. *General System Theory: Foundations, Development, Applications* (Revised edition). (New York: George Braziller, 1976).

Waldrop, M. M. *Complexity: The Emerging Science at the Edge of Order and Chaos* (New York: Simon Schuster, 1992).

Weick, K. E. *The Social Psychology of Organizing* (Saddle River, NK: Addison-Wesley, 1979).

Wilber, K. *A Theory of Everything: an Integral Vision for Business, Politics, Science, and Spirituality* (Boston, MA: Shambhala, 2000).

Wolfram, S. "Cellular Automata as Models of Complexity," *Nature,* (1984), Volume 311, Issue 5985, pp. 419-424.

Wolfram, S. *A New Kind of Science* (Champaign, IL: Wolfram Media, 2002).