

Would-Be-Worlds

John L. Casti

By more-or-less common consensus, Galileo is credited with ushering-in the idea of controlled, repeatable, laboratory experiments for the study of physical systems. And as such experiments are an integral part of the so-called scientific method, it's no exaggeration to say that Galileo's work formed a necessary precondition for Newton's creation of a workable theory of systems composed of interacting particles, a theory that formed the basis for much of modern theoretical science. But Newton's particle systems are what in today's parlance we would term "simple" systems, since for the most part they are formed of either a very small or a very large number of interacting "agents" [i.e., particles] interacting on the basis of purely local information in accordance with rigid, unvarying rules. Complex systems are different.

Typically, complex systems like a stock market or a road-traffic network involve a medium-sized number of agents [traders or drivers] interacting on the basis of limited, partial information. And, most importantly, these agents are intelligent and adaptive. Their behavior is determined by rules, just like that of planets or molecules. But the agents are ready to change their rules in accordance with new information that comes their way, thus continually adapting to their environment so as to prolong their own survival in the system. At present, there exists no decent mathematical theory of such processes. One part of the argument to be made here is that a major stumbling block in the creation of a theory of complex, adaptive systems has been the lack of ability to do the kind of controlled, repeatable experiments that led to theories of simple systems. The second half of our argument is that the micro-simulations, or "would-be worlds," presented at this meeting constitute nothing less than laboratories for carrying out just such experiments. So for the first time in history, we have the experimental tools with which to begin the creation of a bona fide theory of complex, adaptive systems.

Theories, Experiments, and Big Problems

To see the role that micro-simulations will play in the creation of a theoretical framework for complex systems, it's instructive to examine briefly the history of theory construction for several major areas of modern science.

Typically, a theory of something begins its life with what I'll call a "Big Problem." This is some question about the world of nature or humans that cries out for an answer, and that seems approachable by the concepts and tools of its time. Just to get a feel for what such questions are like, here is a rather eclectic list of Big Problems from a few areas of natural and human affairs:

- * Biology: The Structure of DNA — What is the geometrical structure of the DNA molecule, and how does this structure lead to the processes of heredity?
- * Astrophysics: The Expanding Universe — Is the Universe open or closed, i.e., will it continue to expand forever, or will a phase of contraction back to a "Big Crunch" occur?
- * Economics: Equilibrium Prices — In a pure exchange economy, does there exist a set of prices at which all consumers and suppliers are satisfied, i.e., is there a set of prices for goods in the economy at which the supply and demand are in balance?
- * Physics: Stability of the Solar System — Does there exist a finite time in the future at which either there will be a planetary collision, or at which some planet attains a velocity great enough to escape the solar system?

So what we have here are four questions about the real world, each of which arises pretty much from opening our eyes and looking around. And each of these questions has given rise to a theoretical framework within which we can at least ask — if not answer — the question. But these theoretical frameworks, be they the theory of knots for studying the geometry of DNA or the fixed-point theories of economics that tell us about prices, have each come about as the outgrowth of experiments with the system of interest.

For example, it was only by having access to the x-ray crystallographic studies by Rosalind Franklin that James Watson and Francis Crick were able to uncover the double-helix structure of DNA. Similarly, observations by Edwin Hubble using at the Mount Palomar Observatory showed the expansion of the universe, an empirical fact that has led to current theories of dark matter for answering the question of whether or not this expansion will continue indefinitely.

These examples — and the list could be extended almost indefinitely — illustrate the so-called scientific method in action. It consists of four main steps:

observation > theory > hypothesis > experiment

This diagram makes the importance of experimentation evident; in order to test hypotheses suggested by a theory, we must have the ability to perform controlled, repeatable experiments. And this is exactly where the micro-simulations possible using today's computing machines enter into our discussion. In contrast to the more familiar laboratories of the chemist, physicist or biologist, which are devoted to exploring the material structure of simple systems, the computer-as-a-laboratory is a device by which we can probe the informational structure of complex systems. Let me look at this point just a bit further.

Information versus Matter

For the past 300 years or more, science has focused on understanding the material structure of systems. This has been evidenced by the primacy of physics as the science par excellence, with its concern for what things are made of. The most basic fact about science in the 21st century will be the replacement of matter by information. What this means is that the central focus will shift from the material composition of systems — what they are — to their functional characteristics — what they do. The ascendancy of fields like artificial intelligence, cognitive science, and now artificial life are just tips of this iceberg.

But to create scientific theories of the functional/informational structure of a system requires employment of a totally different type of laboratory than one filled with retorts, test tubes or bunsen burners. Rather than these labs and their equipment designed to probe the material structure of objects, we now require laboratories that allow us to study the way components of systems are connected, what happens when we add/subtract connections, and in general, experiment with how individual agents interact to create emergent, global behavioral patterns.

Not only are these "information labs" different from their "matter labs" counterparts. There is a further distinction to be made even within the class of information labs. Just as even the most well-equipped chemistry lab will help not one bit in examining the material structure of, say, a frog or a proton, a would-be world designed to explore traders in a financial market will shed little, if any, light on molecular evolution. So let me conclude this short discussion by considering some would-be worlds, each having its own characteristic set of questions that it's designed to address.

Would-Be Worlds

In the past few years, a number of electronic worlds have been created by researchers associated with the Santa Fe Institute to study the properties of complex, adaptive systems. Let me cite just three such worlds here as prototypical examples of the kind of information laboratory we have been discussing.

* *Tierra* — This world, created by naturalist Tom Ray [1], is populated by binary strings that serve as electronic surrogates for genetic material. As time unfolds, these strings compete with each other for resources, with which they create copies of themselves. New strings are also created by computational counterparts of the real-world processes of mutation and crossover. Over the course of time, the world of *Tierra* displays many of the features associated with evolutionary processes seen in the natural world, and hence can be used as a way of experimenting with such processes — without having to wait millions of years to bring the experiment to a conclusion. But it's important to keep in mind that *Tierra* is not designed to mimic any particular real-world biological process; rather, it is a laboratory within which to study neo-Darwinian evolution, in general.

* *TRANSIMS* — For the past three years, a team of researchers at the Los Alamos National Laboratory headed by Chris Barrett has built an electronic counterpart of the city of Albuquerque, New Mexico inside their computers. The purpose of this world, which is called *TRANSIMS*, is to provide a testbed for studying the flow of road traffic in an urban area of nearly half a million people. In contrast to *Tierra*, *TRANSIMS* is explicitly designed to mirror the real world of Albuquerque as faithfully as possible, or at least to mirror those aspects of the city that are relevant for road-traffic flow. Thus, the simulation contains the entire road traffic network from freeways to back alleys, together with information about where people live and work, as well as demographic information about incomes, children, type of cars and so forth. So here we have a would-be world whose goal is indeed to duplicate as closely as possible a specific real-world situation.

* *Sugarscape* — Somewhere in between *Tierra* and *TRANSIMS* is the would-be world called *Sugarscape*, which was created by Joshua Epstein and Rob Axtell of The Brookings Institution in Washington, DC. This world [2] is designed as a tool by which to study processes of cultural and economic evolution. On the one hand, the assumptions about how individuals behave and the spectrum of possible actions at their disposal is a vast simplification of the possibilities open to real people as they go through everyday life. On the other hand, *Sugarscape* makes fairly realistic assumptions about the things that motivate people to act in the way they do, as well as about how they go about trying to attain their goals. What is of considerable interest is the rich variety of behaviors that emerge from simple rules for individual action, and the uncanny resemblance these emergent behaviors have to what's actually seen in real life.

The main point of bringing up *Tierra*, *TRANSIMS*, and *Sugarscape* is to emphasize two points:

* We need different types of would-be worlds to study different sorts of questions, and

* each of these worlds has the capability of serving as a laboratory within which to test hypotheses about the phenomena they can represent. And, of course, it is this latter property that encourages the view that such computational universes will play the same role for the creation of theories of complex systems that chemistry labs and particle accelerators have

played in the creation of scientific theories of simple systems. For a fuller account of the technical, philosophical and theoretical problems surrounding the construction and use of these silicon worlds, see the author's volume [3] which will appear in the fall of 1996.

[1] Ray, T., *An Approach to the Synthesis of Life,*" in *Artificial Life - II*, C. Langton et al, eds. Addison-Wesley, 1991, pp. 371—408

[2] Epstein, J. and R. Axtell, *Growing Societies*, MIT Press, 1996

[3] Casti, J., *Would-Be Worlds*, John Wiley & Sons, 1996