The Bulletin of the

Santa Fe Institute

Fall-Winter, 1992
Volume 7, Number 2

Santa Fe Institute

In This Issue

Research Updates

Reviews

Education

Regular Features

COVER: The cover, "Global CA," depicts the emergent organization of species competing for space in an artificial world and is based on a Connection Machine experiment using the Genetic Algorithm to evolve competitive components of an immune system. The immune system experiment and this image were executed by Ron Hightower, a graduate student of External Faculty member Stephanie Forrest.

The Bulletin of the
Santa Fe Institute
Vol. 7, No. 2 (Fall-Winter, 1992)

L. M. Simmons, Jr., Editor in Chief
G. Richardson, Managing Editor
R. K. Butler-Villa, Designer

The Bulletin of the Santa Fe Institute is published biannually by SFI to keep our friends and supporters informed about the scientific and administrative programs. The Bulletin is free of charge and may be obtained by writing to the Editor at the address below. Letters to the Editor are welcomed.

The Santa Fe Institute (SFI) is a multidisciplinary graduate research and teaching institution formed to nurture research on complex systems and their simpler elements. A private, independent institution, SFI was founded in 1984. Its primary concern is to focus the tools of traditional scientific disciplines and emerging new computer resources on the problems and opportunities that are involved in the multidisciplinary study of complex systems—those fundamental processes that shape almost every aspect of human life. Understanding complex systems is critical to realizing the full potential of science, and may be expected to yield enormous intellectual and practical benefits.

Published by the Santa Fe Institute, 1660 Old Pecos Trail, Suite A, Santa Fe, NM 87501, U.S.A.; Phone 505/984-8800; Fax 505/982-0565.

A Banner Year for the Santa Fe Institute

President's Message

We are now recognized as representing the leading edge in a major new direction in science, the synthesis of complex system behavior from individual interactions between system components—called the science of complexity. Our programs grew by almost 30% in 1992, a reflection of the keen scientific interest in complexity research. Again in 1993 we will expand our quarters to accommodate what we expect to be more healthy growth. We continue to have some difficulty getting support from the traditional science funding sources because of our non-traditional agenda, but support has been growing from individuals, companies, and foundations.

Two popular books about our work were published in the past few months (one of them is reviewed in this issue). More are in preparation. Yet another, Steven Levy's *Artificial Life* published by Pantheon Press, came out during the summer. We are also much more visible in the corporate world, and the establishment of our Business Network for Complex Systems Research has met with gratifying success in its brief existence. And in December a group of SFI scientists traveled to Japan to lead a series of public symposia, co-organized by the University of Tokyo, to introduce our view to Japanese academic and business leaders.

Are there common themes and approaches that run through many kinds of systems composed of many simple, interacting parts? In the summer of 1992, we convened an international meeting devoted solely to probing this question. This "Integrative Themes" meeting better defined the focus or "core" of the studies of SFI. The proceedings of that meeting will be published in 1993 as part of our book series.

In 1992 we officially formed a program in Adaptive Computation. Among other goals, research in this program will develop the computational approaches and tools necessary to describe systems that evolve and adapt as they develop in time. Such systems are central to the SFI agenda. We expect the techniques being developed to be easier to implement, and we expect them to help highlight the common threads that seem to run through many of these studies.

Artificial Life is closely allied to research in Adaptive Computation. In 1992 we began an effort to develop what Christopher Langton calls his "Swarm" model, an attempt to develop a general software structure that can readily accommodate modeling of virtually any kind of complex adaptive system. In 1993, this work may provide a tool that non-computer scientists who are interested in systems behavior can use intuitively and quickly.

Three archaeologists are in residence for our Evolution of Culture program. Though this work is still in early stages, SFI has stimulated innovative collaborative work to consider how the ebb and flow of southwestern prehistoric cultures may have been linked with changes in resources and how the often autonomous cultural centers interacted to define a larger culture linked in adaptive ways. It represents the continuing attempt by SFI to broaden its studies of complexity to include new disciplines and new kinds of thinking.

The Economics and Theoretical Biology programs continue strong at SFI. As in other areas, we are severely resource limited—many more scholars want to take part in our programs than we have the ability to support. However, our message is being widely disseminated in both the academic and business communities, and if a measure of our success is the seriousness with which the science of complexity is taken in these arenas, then we are wildly successful, far beyond what our modest size would suggest.

Edward Knapp

Software Review

SimLife from Maxis: Playing with Virtual Nature

Christopher Langton

SimLife is the newest simulation game from Maxis. In the arena of personal computers, SimLife is by far the most accessible product for experiencing, and experimenting with, Artificial Life. For a product intended to be an educational "Software Toy," it is *awfully* close to being a useful tool for scientific research, and I wouldn't be surprised to see it used as such, although there is certainly room for improvement in this direction.

In the world of games for personal computers, Maxis is in a class by itself. Their series of simulation games—SimCity, SimEarth, SimAnt, and now SimLife—are far more than games. In each of these engaging systems, Maxis provides the user with a fascinating virtual "nature," with its own physics and environment, replete with occupants that "live" their virtual lives within the confines of these artificial realities. The role of the user in these games is not so much as a participant in the action, as is the case with most computer games, but rather as the reigning "God" who designs the universe from the bottom up.

All of the "actual" players in these games are virtual agents who have simple rules dictating their own behaviors, and whose collective interactions lead to the emergence of, often complex and unpredictable, global dynamics. The goal for the user in these games is to attempt to maintain the complex emergent dynamics of the whole system within certain "viable" bounds, and he or she must do so solely by tuning the behaviors of the individual virtual players or by adjusting various aspects of the physical environment.

In short, these games are "flight simulators" for complex systems, and, as is often the case in flight simulators, the user gets an eye-openingly rapid education in the difficulty of managing such systems by "crashing a few" before he or she manages to even get one off the ground successfully, let alone fly reasonably skillfully or make a safe landing.

For example, in SimCity, the first game from Maxis, the user is "The Mayor" of an initially very small town populated by virtual citizens known as "Sims." As Mayor, the user can determine land use via zoning assignments, and has access to funds with which he or she can clear land, build power plants, fire stations, police stations, hospitals, bridges, airports, and harbors for shipping. The more successful the Mayor is in developing an "attractive" living and working climate, the more industry moves in, creating jobs, which attract more Sims, who build more houses and require more services. The tax base goes up, but so do expenses. Traffic becomes a problem, so new roads must be constructed and old ones begin to need repairs. Fire and police stations must be built to keep up with development, or else fires and crime will break out in increasing numbers. If things deteriorate too far, Sims begin to move out, neighborhoods decay and factories close, leading to higher crime in the abandoned neighborhoods, which drives out more Sims. And then there are the natural disasters: planes crash at the airport, ships run into bridges, trains derail and destroy track, and tornadoes pass through. All of these need to be repaired and hospitals need to be adequate to handle the injured. Very soon in the game, there is not enough money in the city treasury to cover all of the expenses. So what do you do as Mayor? Raise taxes, of course! But then Sims begin to leave because the taxes are too high, and so the tax base drops, resulting in less money rather than more.

In SimLife, Maxis has essentially created a flight simulator that gives one a taste of what it would be like to be in the pilot's seat occupied by God. In fact, if God used a computer to create the world and populate it with organisms, his software tools would look a lot like those found in SimLife.

On entering the game, the user, à la "Gen-

Christopher Langton, sometimes called the Father of Artificial Life, is an SFI External Assistant Professor and a member of the Complex Systems Group, Theoretical Division, Los Alamos National Laboratory.

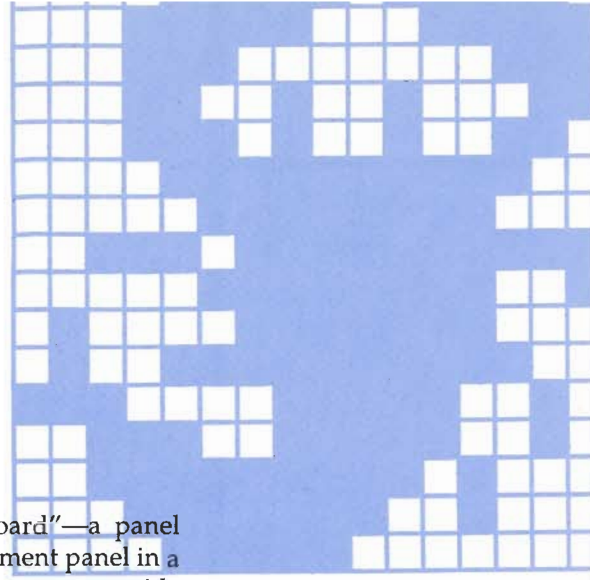

esis," must first define a physical world by creating mountains, lakes, and rivers; establishing a climate, complete with regional and seasonal variation in temperature and moisture; allocating an initial deposit of fertile "topsoil"; and distributing the "Filter-Food" which constitutes the bottom layer of the food chain. As the world is created, the user is able to view his "creation" developing in a "Map" window, which provides a synoptic view of the dynamics of the world throughout the simulation.

Next, following the "Genesis" plot line, the user must create the plants and animals who will "live their lives" in this newly created world. Through the "Dashboard"—a panel that is the analog of the instrument panel in a flight simulator—the user has access to a wide variety of tools for designing plants and animals and for populating the world with the resulting organisms. By clicking a few buttons, the user is presented with a tool that allows him or her to edit the phenotype or even the genotype of any animal or plant, or to construct a new one from scratch.

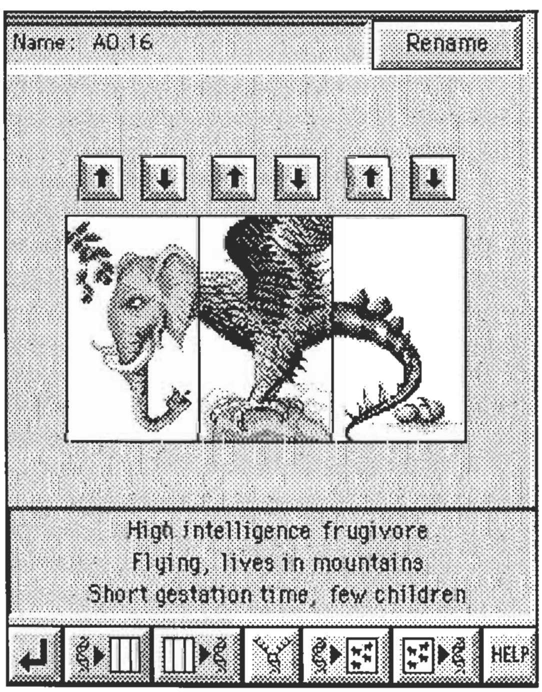

Figure 1. Part of the SimLife control panel for designing organisms. Attributes such as intelligence, mode of locomotion, and gestation time are assigned by selecting from among a set of "flash cards."

The organism design tool is incredibly easy to use, consisting of a series of picture cards that allow one to construct an organism the way police artists reconstruct a face from a collection of eyes, noses, mouths, chins, and hairstyles (Figure 1). For animals, there are three sets of "cards." The first set consists of a variety of heads, which define the "intelligence" (high, medium, or low) and preferred foods (carnivore, herbivore, fructivore, etc.) of the organism. The second set of cards consists of a variety of torsos, defining the mode of locomotion (walking, swimming, or flying) and the preferred habitat (mountains, plains, oceans, etc.). The third set of cards consists of a variety of tails, dictating the mode of reproduction (many offspring, several, or only one) and the average gestation time (long, moderate, or short). The set of cards for plants allows a similar variety of choices. By clicking one's way through these sets of cards, one can easily define somewhere in the neighborhood of 8,000–10,000 different animals and plants.

Screened bitmap image above is from another portion of the menu shown in figure 1.

Beyond this, however, the user can go on to edit the underlying genotype responsible for these phenotypic traits, and many of these underlying genes are actually continuous in definition. Thus, there is an almost infinite variety of distinguishable organisms allowable within a prespecified genotypic space.

Once the user has defined a set of plants and animals to his or her liking, or chosen a few from a large set of predefined organisms, he or she is ready to populate the world with

continued next page

SimLife from Maxis: Playing... (cont'd.)

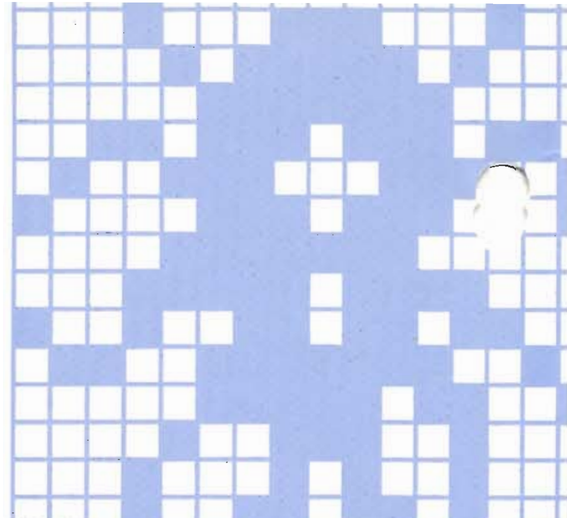

Screened bitmap image above right is from another portion of the menu shown in figure 1.

them. A "Populate" button in the Dashboard makes this a trivial task. It pops up a window that allows one to select a species of animal or plant and deposit a certain number of them wherever the user chooses. One can deposit organisms one at a time or trace out great curving arcs of them using the mouse.

Once the world has been created and populated, there is nothing left to do but to start the simulation going by clicking off the "Pause" button. The simulation starts running; organisms begin to move about, consuming food, mating, giving birth, and dying. The user can pop up a number of different tools for tracking data on the system, including a nicely general time-series graph tool, a window for tracking changes in the gene pool, a food chain analyzer, a historical record of events, a tool for displaying the diversity of species, and so forth.

As days, seasons, and years click past, the user has a great many means for interacting with the simulation directly. One can pick out a particular organism and move it, clone it, view its genotype or its phenotype, or display its internal state variables (health, food store, water store, size, age, etc.). One can even "Smite" a particular organism with a lightning bolt, reducing it to a pile of bleached bones. One can add mountains, alter the climate, or even "part the waters" if one so desires. Mutagens can be distributed around that cause more variation in the genotypes of offspring than was specified originally.

Throughout it all, evolution proceeds, species come and go, ecologies form and crash, and life goes on. That is, it goes on if the user is successful in setting the world upon an open-ended course, which is, of course, the goal of the game.

This goal is not easy to achieve. Establishing a viable ecosystem is not at all straightforward, and many worlds become barren of "life" before one achieves a few long-lived populations.

However, a great deal of familiar ecological behavior is observable along the way, even if it's along the way to extinction! Predator-prey cycles are fairly common. Several times I have observed predator-prey cycles involving phase-shifted oscillations of predator and prey on different "continents," one constantly reinfecting the other if either predator or prey become extinct there. The food webs can become fairly complex, although I have not observed any tendency for food webs to increase in "depth."

One would need to run for weeks on end to observe long-term evolutionary change, which I have not done, but I have achieved populations which I am sure would survive for weeks on end, so long-term evolutionary dynamics is not out of the question.

SimLife does allow the user to log data into a file which can be saved to disk for later analysis

All in all, SimLife is reasonably "wide open" in terms of its potential for capturing a variety of evolutionary and ecological phenomena, although perhaps not at the level of detail or sophistication that would be necessary to adequately capture specific real-world situations. A system much like SimLife that provided the user more access to the underlying methods for determining the physical laws and the behavior of the organisms would save a lot of us from ever having to write a line of code again.

Nonetheless, SimLife allows the construction and study of simple, artificial ecologies, and the surprising complexity and richness of behavior that emerges in even these extremely simple "artificial natures" has already given me a greater appreciation for the real thing—Nature in all her real glory—writ with a capital "N." That glory shines all the more brightly even from the meagre illumination that this simple "Software Toy" is able to shed upon Her.

Complexity, the Emerging Science at the Edge of Order and Chaos

Book Review

Harold Morowitz

Complexity, The Emerging Science at the Edge of Order and Chaos
M. Mitchell Waldrop
(Simon & Schuster, New York, $23.00)

Rapid developments in contemporary science and technology have given rise to a whole new genre of literature, chronicles of emerging concepts. These chronicles have tended to be written by individuals deeply trained in science, often with doctoral degrees. Their role is to be communication links between the scientific innovators and that small, but book-buying segment of the general public who wish for deeper insight into modern ways of understanding the universe.

Among the better known of these authors are Horace Judson (*The Eighth Day of Creation*) and James Gleick (*Chaos*). One of the highly respected of these techno-Boswells is M. Mitchell Waldrop, whose latest work explores a vibrant emergent approach to viewing the world, the men and women who are bringing it about, and the physical and intellectual loci where it is happening. This is a tall order, and Waldrop manages to carry it off with elegance and élan.

The science of complexity is an attempt to break away from the Galilean-Newtonian simplicity that has served much of traditional physics so well and to tackle those real-world problems that do not yield to this neo-Platonic mold. In fact, most of the problems of the modern world, even in physics, have moved away from the earlier model. The new approach is made possible, indeed it is mandated by computers, those piles of hardware elements that have forever altered how mankind views itself and its universe.

Waldrop chooses to tell his story in terms of the development of the Santa Fe Institute, that wonderous and oracular nucleus of condensation for people interested in "the complexity problem." Thus the book is a series of tales of the savants who have made it happen. They emerge as charismatic, insightful, brilliant, rambunctious, revolutionary, and argumentative. They are clearly individuals who hear beats of different drummers. One of the best features of Waldrop's style is that he presents those individuals as real flesh and blood thinkers, and he does it while minimally intruding on their privacy.

One feature on which the author must be commended is his ability to take difficult abstruse technical material and cast it in a syntax and style that is accessible to a literate readership. He deserves my mythical C. P. Snow award for communicating between the two cultures. Clearly, this is an important book and a masterful piece of writing.

This praise of "Complexity" comes with a caveat or at least a question addressed to this entire genre of literature. When Judson published *The Eighth Day of Creation* in 1979, the revolution in biology was over and he was a chronicler of history. Looking back, Boswell published the biography seven years after Johnson was dead. When Waldrop writes, he is less a historian and more of a reporter. The achievements he writes about have not passed the test of time. We know not what will be of lasting value. Such are the problems of a reporter. But in these days a journalist is also an actor, for support of research is geared to fame, and fame comes in part from inclusion in books of this type. It is a reportorial uncertainty principle, the observer influences the system being observed. This type of writing is in a rapid growth phase. Future historians will have to evaluate its effect on science and scientists, some of whom have a less than humble view of themselves in any case.

Harold Morowitz is an SFI Science Board member and Clarence Robinson Professor of Biology and Natural Philosophy, George Mason University

Book Review

Biology and the Loss of Mechanism, or Why the Physicists Were Wrong

John L. Casti

Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life
Robert Rosen
(Columbia University Press, New York, 1991)

The research agenda for much of modern biology was set by two pivotal events orchestrated primarily by physicists: the publication in 1944 of Erwin Schrödinger's little volume *What is Life?*, and the announcement in 1953 of the double helical structure of the DNA molecule by Francis Crick and James Watson. The Schrödinger volume was a call-to-arms to the community of physicists, arguing that the process of heredity could be understood in purely physico-chemical terms. The Watson-Crick exercise, coupled with the elaboration of the genetic code by Crick and Sydney Brenner a few years later, gave every appearance of vindicating Schrödinger's arguments, offering logical, geometrical, and material paradigms for how the information needed to make a living being is coded, transmitted, and translated from mere matter into what we call life.

The initial euphoria arising from these developments resulted in a massive influx of renegade physicists and near-physicists from the fast-paced world of mathematical mumbo-jumbo and bubble chambers into sleepy backwaters of academia, areas previously inhabited only by the kind of scientist who felt more at home drawing taxonomic charts and collecting butterflies than operating electron microscopes and briefing Pentagon poseurs. If the secrets of the atom could be unlocked, uncovering the secrets of life should be a mere bagatelle. Or so the physicists thought.

But the physicists were wrong.

After nearly half a century of trying to shoehorn the empirical facts about organisms into the mechanistic paradigms of both classical and quantum physics, molecular biologists and biophysicists have yet to be able to delineate what it is *exactly* about organisms that separates them from other types of material objects. The book under review offers a novel, some might say eccentric—or even crazy—answer as to why mechanistic biology is a nonstarter when it comes to providing a scientific account of what's so special about life. To paraphrase another physicist, Wolfgang Pauli, the question confronting us here is whether the theory proposed in this book is crazy enough to be correct.

Studying the Nature of Life

In his approach to studying the nature of life, Rosen follows the dictum of mathematician René Thom, who once remarked that "theoretical biology should be done in mathematical departments." So right from the outset, this is a book guaranteed to alienate just about every biologist, being written more in mathematical symbols and words of philosophy than in chemical formulas. Even worse, it contains not a single account of any laboratory experiment. In short, this is a volume whose natural home is somewhere in that no-(wo)man's land separating the departments of mathematics, philosophy, system theory, and biology. As a result, and despite the author's clear and engaging literary style, it's not easy reading for those who lack at least a nodding acquaintance with the concepts and terminology of these fields. But the ideas the book exposes are of such novelty and broad scientific currency that the reward of making the effort far outstrips the costs. So let me try to briefly summarize the main thrust of the book's argument for the poverty of mechanism as a scientific research strategy, in general, and for studies of life, in particular.

The first pillar upon which Rosen rests his case is the idea of *natural law*. All of theoretical science is based on bringing two inferential structures into congruence: (1) the modes of causal entailment between phenomena in a

John Casti is a long-term visitor at SFI whose research interests include theory and application of metabolism-repair systems, and the structure and dynamics of road traffic networks. Currently he is writing a general reader volume on complex systems to be published by Harper Collins.

So everything about an organism is entailed by something else from within the system itself. The author's general claim then is that the study of organisms—i.e., theoretical biology—is the study of the category of all models of such systems.

natural system N, and (2) the modes of syntactic entailment between the symbols of a formal logical system M, a system that we take to be a model for N. Natural law, then, is a concrete embodiment of the modeling relation that connects these two radically different modes of entailment.

The next step is to revive the Aristotelian notion that there are many different—and inequivalent—types of causality. To be specific, in Aristotle's view of the world causality comes in four flavors: material, efficient, formal, and final causation. And each of these types creates a different way of answering the question: "Why?" Rosen shows how these causal categories give rise to different types of entailments, both in the natural world of N and in the formal world of its model M. In particular, he argues that the entailment structures of physics rely solely upon the first three types of Aristotelian causation, a fact of the greatest significance for those physicists seeking professional salvation in the realms of biology.

Following an illuminating, but rather technical, elaboration on the distinction between analytic and synthetic models, we come to the heart of the book. First of all, the notion of *simulation*. A formalism (read: model) M can be simulated if its inferential structure can be expressed as software to a mathematical machine (read: Turing machine). It turns out that simulability imposes severe restrictions on the inferential structures possible in the model M. But given a natural system N, we necessarily have formalisms M associated with it simply by virtue of the modeling relation imposed by natural law. So we arrive at the idea of a mechanism and a machine: a system is a *mechanism* if and only if all of its models are simulable. And the system is a *machine* if and only if it is a mechanism, such that at least one of its models is already a mathematical machine. In linguistic terms, a mechanism is a system in which syntactics and semantics coincide.

But a modeling relation brings causal entailment in the real world into congruence with the entailment structures of a model. So to assert that the system has a model that is a mathematical machine is saying something very special about the operations of causality. What this boils down to is that certain modes of entailment are just not available in a mechanism. There is always some component f of a mechanism that requires us to jump outside the mechanism itself in order to answer the question: "Why f?"

Finally, we're ready for the book's punch line: *A material system is an organism if and only if it is closed to efficient causation*. So if f is any component of an organism, the question "Why f?" has an answer within the system itself, an answer corresponding to the category of efficient cause of f. This result can be abstractly expressed using a block diagram of sets and maps showing the entailment structure of the system. The diagram for a mechanism has no closed paths of efficient causation; organisms, on the other hand, have diagrams containing **only** such paths. So everything about an organism is entailed by something else from within the system itself. The author's general claim then is that the study of organisms—i.e., theoretical biology—is the study of the category of all models of such systems.

The epistemological consequences of this circle of ideas are profound. First of all, biology is now properly located beyond the scope of mechanism. Moreover, we can finally understand why the question "What is life?" is so refractory when approached from a mechanistic point of view. At the same time, the scope of biology is dramatically enlarged: biology itself becomes identified with the class of material realizations of a certain kind of relational organization. One consequence is that biology becomes divorced from the struc-

continued next page

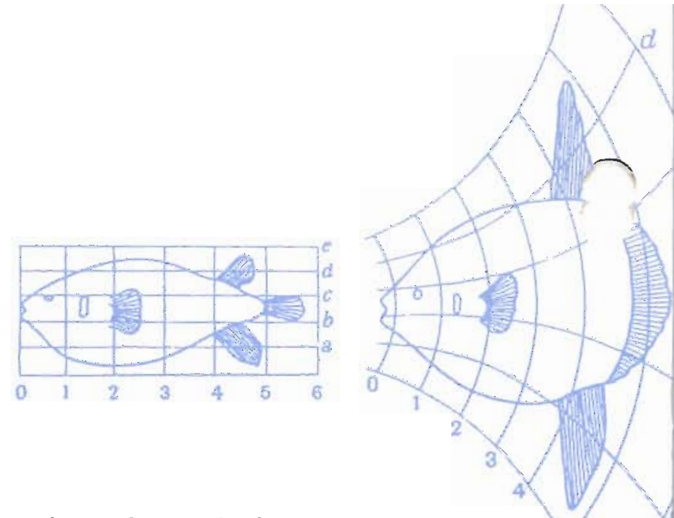

Screened fish image, above right, is from the reviewed book; copyright © 1991 by Columbia University Press. Reprinted by permission.

tural details of any particular kind of realization. From this it follows that the study of biology is a creative process: the fabrication of a material system that possesses the right kind of model. This is nothing less than the creation of a new type of organism.

Seeing the World in a New Way

This cursory summary of the book barely does justice to the depth, originality, and import of the ideas it exposes. But paradoxically, perhaps, it's unlikely that this work is going to generate any major tremors in the world of biology. The overwhelming majority of mainline biologists are likely to find its arguments completely incomprehensible. Even more disheartening, those biologists who do understand the book will probably hate it, simply because it contains the most devastating kind of critique of their research programs. Mathematicians will hate it, too, because it contains no theorems and no new mathematical concepts. And even philosophers are likely to turn up their noses in disdain at Rosen's arguments, since again they run completely counter to almost all conventional wisdoms in the philosophy of biology. So who, besides this reviewer, is going to like this book?

While it's always difficult trying to outguess the tastes of the scientific, as well as the reading, public, I think the real audience for this work consists of that rare breed of intellectual who is not afraid to look beyond the existing paradigms underlying current research in the traditional disciplines, especially biology. This means that the book is for biologists who believe that what's important about living things is not what organisms are made of but the relationship between their parts, applied mathematicians who think that the connection between the real world and its symbolic representations is more important than theorems and proofs, and philosophers of biology who aren't afraid to admit that perhaps their particular brand of epistemological medicine is not uniformly most powerful. It's for such *rara avis* in the world of the intellect that the book was written. In this reviewer's opinion, publication of Rosen's work will come to be seen as a watershed event in our understanding of what it means for an organism to be an "organism" and not a mere "thing." Whether the book's arguments ultimately turn out to be right or wrong is less important than the fact that anyone who understands these arguments will end up seeing the world in a new way. It's hard to ask more of a book than this.

Lest the reader think I'm giving this work a totally clean bill of health, let me close by mentioning two negative factors that, while by no means fatal, certainly detract from an otherwise sterling piece of work. The first is an inordinate number of typographical errors. Fortunately, none of these errors seems especially serious, nor are they difficult to find and correct. Nevertheless, these infelicities contribute to an air of slovenly casualness that takes attention from the main lines of argument.

The second disconcerting aspect is the paucity of references to any of the modern literature in biology and system modeling, other than work by the author himself. Of course, in a book of such originality, it's at least partially defensible to argue that there exists no body of literature from which the ideas have sprung forth. Nevertheless, no work stands alone; every book or research paper has some antecedents.

But these are minor quibbles. The book is a gold mine of ideas on both the philosophy and practice of biological research, setting out a research program that should occupy the attentions of several generations of theoretical biologists and system theorists. It can only be hoped that the scientific community will hear the book's clarion call.

New Books by SFI Research Family

Reality Rules, I & II
John Casti
(John Wiley & Sons, $72.00 set)

In *Reality Rules* John Casti provides a comprehensive and accessible introduction to the theory and practice of mathematical modeling in the physical, biological, and social sciences. Starting out as a simple update of his *Alternative Realities* of 1989, the book became a runaway project reflecting the recent explosive growth in the field of dynamical system theory and finally split into two volumes. It is in large part a new work with more than 300 pages of "new points of contact between the worlds of nature and mathematics."

The core of the book is still how to use the rules of mathematical modeling. In Volume One, "The Fundamentals," Casti provides the essential concepts and results needed by the modeler. In Volume Two, "The Frontier," he introduces the application areas and associated techniques of modeling that complement the ideas of the first volume. A new chapter on computation and complexity treats matters of current concern from Godel's Incompleteness Theorem and its connection to work in artificial intelligence to the problem of NP completeness and the complexity of numerical algorithms. The last chapter addresses the philosophy of modeling with discussions of "good" and "bad" models, the criteria for a good theory, and a look at science, religion, and the nature of belief systems.

American Mathematical Monthly describes the book as "unorthodox, philosophical, and purposeful in offering a highly innovative 'theory of models.'"

Beginnings of Cellular Life, Metabolism Recapitulates Biogenesis
Harold J. Morowitz
(Yale University Press, $27.50)

In this book Harold Morowitz presents a new theory of the origin of life on Earth four billion years ago. He postulates that the first step toward the origin of life was the spontaneous condensation of amphiphilic molecules to form vesicles (or protocells). This hypothesis provides a framework for reexamining the emergence of cellularity. Morowitz further proposes that metabolic processes have not changed for some 3.8 billion years, so we can use a study of modern biochemistry to advance our knowledge about the chemical processes of the earliest photocells. Morowitz views origin-of-life issues from the perspective of certain constructs in the philosophy of science that provide guideposts to formulating and assessing hypotheses. The book presents a unique discussion among origin-of-life books on the relation between science and epistemology on the difficult problem of learning about the very distant past.

"This book is an extraordinatory integration of a lifetime of thought on the subject by an exceptional scholar," says Donald M. Engleman of Yale University. "Morowitz logically, consistently, and persuasively moves to define arguments that will influence thought in a fundamental realm of biology."

The Facts of Life, Science and the Abortion Controversy
Harold J. Morowitz & James S. Trefil
(Oxford University Press, $19.95)

In *The Facts of Life*, Morowitz and Trefil attempt to clarify the issue of when "life" begins for a fetus. They suggest that all life is a continuum from plants to humans, from the earliest form of life billions of years ago to its most sophisticated forms today. But, according to the authors, the crux of the abortion question lies neither in the existence of life in the fertilized ovum, nor in personhood—which only adheres to a being after birth—but in the acquisition of the unique qualities of "humanness."

They argue that the distinctive qualities of

continued next page

"humaneness" lie in our verbal and intellectual abilities, which are located in the cortex of the brain. The cortex clearly begins to acquire its full capabilities in the fetus' 24th week. Perhaps coincidentally, the survival rate of premature infants rises dramatically along with the development of the cortex. From this perspective the authors view Roe's trimester schedule as conforming very closely with what the most reliable scientific evidence suggests. Abortion in the third trimester presents similar moral issues to the withdrawal of life support systems from a comatose patient. In these cases the authors note that the decision "is seen to be too important to be left solely to a single individual without rules and guidelines."

Elizabeth Fox-Genovese calls the book "an admirably clear and concise, if intellectually dense essay" which sheds clear light on a controversial issue.

The Origins of Order: Self-Organization and Selection in Evolution
Stuart A. Kauffman
(Oxford University Press, paperback $35.00)

Kauffman presents a new paradigm for evolutionary biology, one that extends the basic concepts of Darwinian evolution to accommodate recent findings from the fields of biology, physics, chemistry and mathematics. This work focuses on the concept of self-organization: the spontaneous emergence of order widely observed throughout nature. Self-organization plays an important role in the Darwinian process of natural selection and in the primordial emergence of life itself. Yet until now no systematic effort has been made to incorporate the concept into evolutionary theory. The construction requirements which permit complex systems to adapt remain poorly understood, as is the extent to which selection itself can yield systems able to adapt more successfully. The book shows how complex systems can spontaneously exhibit stunning degrees of order, and how this order, in turn, is essential for understanding the emergence and development of life on Earth.

"There are very few people in this world who ever ask the right questions of science, and they are the ones who affect its future most profoundly. Stuart Kauffman is one of these. Read this book," says Philip Anderson, Princeton University.

Analogy-Making as Perception: A Computer Model
Melanie Mitchell
(MIT Press, April, 1993, no price yet)

The centrality and ubiquity of analogy in creative thought has been noted again and again by scientists, artists, and writers, and the understanding and modeling of analogical thought has emerged as one of the most important challenges for cognitive science. In this book, Melanie Mitchell describes Copycat, a computer model of analogy-making, developed by Douglas Hofstadter and herself. This work is based on the premise that analogy-making is fundamentally a high-level perceptual process, in which perception, interacting with concepts, gives rise to "conceptual slippages" which allow an analogy to be made. Copycat is a model of this complex, subconscious interaction between perception and concepts that underlies the creation of analogies.

On one level, the work described here is about analogy-making, but on another level it is about cognition in general, exploring such issues as the nature of concepts and perception, and the way in which highly flexible concepts emerge from a lower-level "subcognitive" substrate. This book is aimed at cognitive scientists, psychologists, philosophers, and artificial-intelligence researchers interested in cognitive modeling, perception, concepts, analogy, and creativity.

A New Direction for SFI Book Series

Ronda K. Butler-Villa

In planning the further development of the SFI publications program, Addison-Wesley and the SFI Editorial Board agreed that it is time for a change. The issue at hand is whether to expand staff to handle all book requests or to limit the books that we can produce while holding staff and costs at the current level.

Of the twenty volumes in our series, thirteen have been proceedings volumes and four lectures volumes. The Editorial Board agreed to move the series emphasis away from conventional proceedings volumes. Future proceedings volume will include only the best papers from the meeting or commissioned chapters of primarily new material; in either case, the volume should be heavily edited to ensure a concise, well-organized book. Another option may be to encourage multiauthored volumes based on key presentations at workshops or collaborative books coauthored by several participants; such volumes may be more suitably marketed in the lecture notes subseries, or in a new subseries. Additionally, the Editorial Board will be actively seeking more lecture notes volumes and monographs authored by members of the SFI family and will select volumes for the series after review of the manuscript or a written proposal.

Every scientist who visits SFI is invited to consider publishing in our series work that is related to SFI programs.

Since book commitments have already been made through 1993, these changes will be more evident in 1994.

For a complete list of our books, please contact the SFI Publications Office or write to Addison-Wesley, Advanced Book Program, One Jacob Way, Reading, MA 01867.

Announcements

Jack Repcheck, previously Economics Editor of Princeton University Press, has joined Addison-Wesley's Advanced Book Program (ABP) as the new Editor-in-Chief, reporting to Vice President David Miller. Under his direction, A-W plans a more aggressive market plan for our series in 1993 than at any time in the past. In addition, A-W plans some changes in the series design while retaining the essence of the current theme, including the SFI logo, but allowing more individual identity.

Recently the highly successful SFI volume *Introduction to the Theory of Neural Computation* by Hertz, Krogh, and Palmer was published in a Japanese language edition.

New Books

The Principles of Organization in Organisms, edited by J. Mittenthal and A. Baskin. Proceedings Volume XIII.

The Double Auction Market: Institutions, Theories, and Evidence, edited by Daniel Friedman and John Rust. Proceedings Volume XIV.

1991 Lectures in Complex Systems, edited by Lynn Nadel and Daniel L. Stein. Lectures Volume IV.

Thinking About Biology, edited by Francisco Varela and Wilfried Stein. Lecture Notes Volume III [available April 1993].

Understanding Complexity in the Prehistoric American Southwest, edited by George Gumerman and Murray Gell-Mann. Proceedings Volume XV [available mid-1993].

To order, contact your local bookstore or, for credit card orders, call the Addison-Wesley order desk at 800-447-2226.

Workshop Review

Artificial Life: New Approaches to Theoretical Biology

This slightly modified report by SFI External Faculty Member Harold Morowitz on the Artificial Life III conference originally appeared in Hospital Practice.

A short two blocks away from the Plaza and the Palace of the Governors in the oldest capitol city in the United States and just a few miles away from some of the most ancient cultural centers on the North American continent seems a somewhat anomalous site for a meeting on ideas that represent revolutionary shifts in the way that mankind conceptualizes the world of human thought. Yet Santa Fe has long had a catalytic effect in bringing together old and new ideas and modes of conceptualization.

If nearby Bandelier National Monument is thought by some to be the navel of mankind, and nearby Los Alamos is envisioned as the birth place of the nuclear age, then the neighboring Santa Fe Institute is emerging as one of the incubators of the sciences of complexity. It is exploring the full range of computer potential for developing radically new approaches to problems ranging from algorithms to zoology—and including a broad domain of social disciplines, where theory has been frustratingly lacking in predictability and rigor.

The magnet that draws me to Santa Fe is Alife III, a conference on artificial life that has attracted 400 scientists from around the world. Artificial life, which suggests and includes robotics, goes far beyond the golem of Prague and the creature of Frankenstein. It attempts to seek general approaches that abstract life from its flesh and blood—even, in some cases, its carbon-centered chemistry. Artificial life has come to encompass a number of machine-centered and abstract chemical approaches to theoretical biology. It has certainly gripped the imaginations of the scientists, engineers, and hackers who are gathered here to exchange ideas.

The first day is given over to the chemists: organic, physical, and biological. The subject matter is diverse, but there are common themes: self-replicating chemical systems and evolving chemical systems. The self-replicating systems range from simple organic molecules to micelles and liposomes. These systems extract molecules from environments and synthesize structures like the starting units. Insofar as the synthesized structures are identical to the original ones, this represents replication; insofar as they are different, this represents the possibility of evolution.

On the fourth and fifth days, there are a number of papers and demonstrations on robotics. Robots are analogous to one or more very specific aspects of living systems. The ability to perform a task robotically may provide insight into how it is done by living systems. And indeed, adaptation and evolution are possible with robots. The robotics and chemical studies share a common feature: Both are hardware embodiments of various biological concepts and approaches. They are tightly constrained by physical and chemical considerations and limited to operating in real time.

Between the chemists and the roboticists, the rest of the meeting takes on a distinctly software approach. Much of this material is the new theoretical biology. One develops a set of formal postulates, allows the computer to work out the consequences of the hypothesis, and compares the output to examples in experiments.

Some of the most dramatic work in this field comes from plant morphogenesis. One of the speakers postulated a set of rules for cell division in developing plants. The output is represented graphically, and one watches the plant grow on the screen. For many plants, programs have been devised so that the representation on the screen is the exact equivalent of a time-lapse series of photographs of the developing plant. At this level, one has a theory of plant morphogenesis. The rules stay fixed, and one generates a repertoire of development depending on initial conditions.

The morphogenetic rules are, of course, the result of genetic, physiologic, and biochemi-

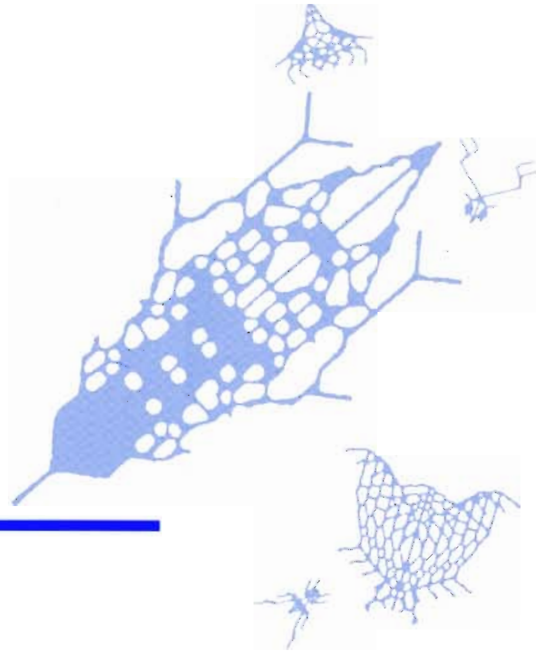

Artificial Life III Invited Speakers

Hardware Synthesis
Rod Brooks, MIT
Maja Mataric, MIT
Mark Tilden, Waterloo
Fred Martin, MIT
Randal Beer, Case Western
Inman Harvey, Sussex
Stevan Harnard, Princeton

Software Synthesis
Kristian Lindgren, Sweden
Leo Buss, Yale
Przemysław Prusinkiewicz, Regina
Thomas Ray, SFI
John Koza, Stanford
Walter Fontana, SFI
Mitchel Resnick, MIT
Karl Sims, TMC
Pattie Maes, MIT

cal processes operating within and between the cells. A more reductionist theory would require that the developmental rules be understood in terms of the underlying interactions (this should ultimately be possible to achieve), but for the moment the developmental programs are impressive all by themselves.

In another type of approach, one models a system with rules of behavior that can change at random (Darwinian) or in response to outputs or environment (Lamarckian). Thus, one can do computer "experiments" on evolution and memory. Here, the computer is not only a calculating machine but also acts as an analogue of biological features that change in different ways over different time domains. The comparison of computer output with observed or experimental features of the biological world is still necessary to complete the epistemological loop.

Quite apart from the detailed presentations, the meeting has an air of excitement. The participants represent a number of countries and an extremely broad range of scientific disciplines and research environments. There is a feeling in the air of anticipation, a feeling that genuinely new approaches to human thought are being explored. On average, the participants are quite young. This is not a field of science that is being transmitted by the elders; rather, it is being developed as the participants move along, trying new approaches.

Since access to computers, workstations, and mainframes is fairly widespread, and since the equipment is so universal, this is a rather egalitarian branch of science, as contrasted with those in which very large setup funds and costly equipment are needed. As a result, many of the participants are somewhat idiosyncratic, and others seem downright flaky. But they all seem to share the excitement over what is taking place.

I don't know whether it is the anoxia of high altitude altering my cerebral function, the multitudinous chills of New Mexican cuisine rumbling in my stomach, or the residuum of margaritas in my nervous system, but I'm feeling very euphoric about artificial life. After years and years of frustration about the progress of theoretical biology, I see new approaches that look very promising.

Theoretical Neurobiology Study Group Convened

Understanding the complexities of the brain demands sophisticated theoretical approaches. Yet today a comprehensive theoretical neurobiology does not exist. This past summer Charles Stevens (Salk Institute) and Michael Stryker (UC San Francisco) convened a working group of experimental and theoretical neurobiologists to try to stimulate some needed intradisciplinary collaboration in this field. Particularly in the biological sciences, this is not as common as might be expected. "The goal of the meeting," says Stevens, "was both to begin an evaluation of the state of theory in neurobiology, and to start identifying research areas in which a combination of theory and experiment might contribute most to our understanding of brain function."

The effort is supported by a major grant from The Pew Charitable Trusts.

Eighteen neuroscientists visited SFI over five weeks in June and July. The group was mixed in at least a couple of ways—about half were senior scientists, half, junior; half were theorists, complemented by an equal number of experimentalists. Everyone, however, works on problems relating to the visual system. "We picked this topic," says Stevens,"because understanding of brain structure and function is most advanced for vision. Also, visual system problems have traditionally been attractive to theorists." Daily talks aimed at uncovering common ground: the theorists presented an extensive summary of theory in visual neurobiology—although the field is far from complete—and the experimentalists in turn described phenomena in need of theoretical descriptions.

Although no existing theory has yet materially changed our understanding of the brain, the group feels that theoretical approaches are on the verge of doing so. They also agree that the theories with the most potential in that regard share a common feature: they relate to some general principle that governs brain structure, function, or development. One general rule, for example, is that the brain uses knowledge about properties of the world (recorded in its circuit organization and in the dynamic neuronal properties of neurons by evolution and by plasticity mechanisms that operate in development) to solve problems for which insufficient information is available from the environment. This principle provides the basis for a theoretical account of retinal receptive fields: information placed in the structure of the retinal circuits by evolution can be used to compensate for defects in the eye's optics, and the computation the eye must do to make this correction explains the retinal field structure. "The fact that theories using general principles are beginning to successfully account for common experimental observations points the way to a growing maturity in neurobiological theory," says Stevens.

Collaborations

Several collaborations immediately resulted from last summer's meeting. Marcus Meister, for example, has developed a unique technology that lets him monitor the activities of many neurons at the same time. His experimental methods are very useful in testing the recently proposed theory of Joseph Atick and Norman Redlich on retinal processing. An experiment that developed out of the Santa Fe meeting is to measure the simultaneous response of many retinal output neurons when the retina is stimulated by natural images. The theory predicts that the output exhibits a pattern of statistically independent activations. The experiment of Marcus will be able to test whether this is the case in actual retinas.

Self-organization of neuronal circuits during development was one of the recurrent themes during the workshop, and Charles Gilbert (Rockefeller U.) and Stevens began a col-

"We are...excited about the unusual opportunity for a working meeting among scientists with very different backgrounds but convergent interests. SFI is an excellent setting because of its tradition of support for new scientific approaches."

laboration that uses detailed information, gathered in Gilbert's laboratory, on the distribution of synapses—the points of contact between nerve cells at which information is exchanged—in the visual system. Their goal is to develop a refined version of current theories of how this self-organization takes place. The detailed data on synapse distribution should permit decisions to be made between several existing ideas about the mechanisms of self-organization.

In another effort, Rama Ranganathan (UC San Diego), Dennis Baylor (Stanford), and Stevens teamed up to construct a theory for the initial steps in phototransduction, the process through which eyes change patterns of light into nerve signals. Ranganathan had new data on phototransduction in eyes from mutant fruit flies that lack a crucial protein known to be involved in the process. The responses of the mutant fly eyes gave clues for how normal transduction works, but a theory was needed to interpret the data. Ranganathan and Baylor have done some experiments this fall that provide data for testing this summer's theory, and the results are being evaluated now.

Goals for 1993

The group reconvenes next summer, led by Stryker and Nancy Kopell (Boston U.). The working title for the 1993 meeting is "Dynamic Control of Stability and Flexibility," focusing on how the central nervous system comes to be organized to be capable both of flexibility and robustness. This theme touches on questions at many different levels of organization and time scale, from rapid signals among neurons to neural development. Questions range from how properties at the level of cellular biophysics affect emergent network behavior to how learning within a network can be done in a stable manner.

One goal for the 1993 working group is to articulate a theoretical framework for addressing these questions. In the past, models of neural networks tended to choose between analytic tractability or depth of detail. The first omits biophysical data that could be crucial to behavior; the latter approaches the overall complexity of the system being modelled. What's needed is a theory that captures what neurons (and assemblies of them) actually do, and which helps to tease out which properties matter to which behavior. "We are both excited about the unusual opportunity for a working meeting among scientists with very different backgrounds but convergent interests," says Kopell. "SFI is an excellent setting because of its tradition of support for new scientific approaches."

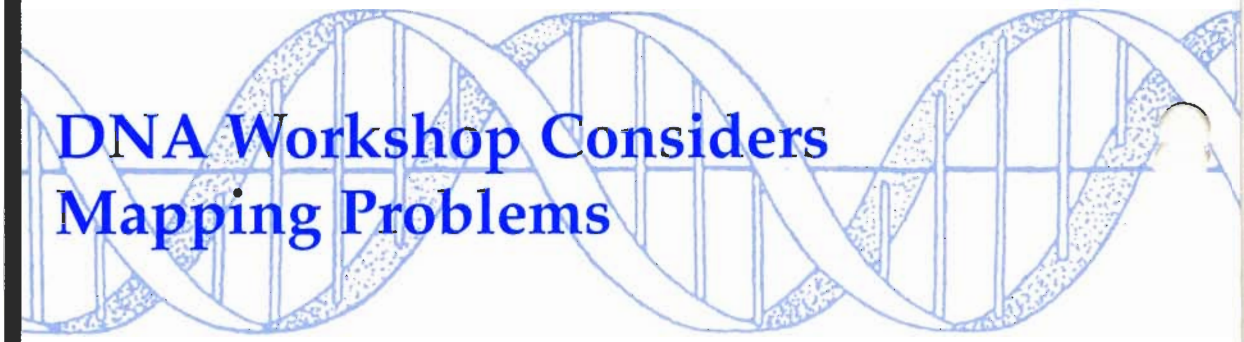

DNA Workshop Considers Mapping Problems

Image is a partial view of the DNA molecular according to Watson and Crick. Figure is redrawn from Cells: Their Structure and Function by E. H. Mercer, copyrighted © 1962 by The American Museum of Natural History.

Earlier this year SFI and the Theoretical Division of Los Alamos National Laboratory (LANL) jointly sponsored an interdisciplinary meeting at the Institute titled "DNA Map Assembly Workshop."

The two dozen workshop participants included physicists, computer scientists, and theoretical and experimental biologists from SFI, the University of New Mexico's medical school and computer science department, and several Laboratory groups including Computer Research, Complex Systems, and the Center for Nonlinear Studies.

The meeting gave participants an opportunity to blend experimental and theoretical approaches to various DNA mapping problems, said workshop organizers Cari Soderlund and Christian Burks of LANL. Talks presented included information on experimental approaches to DNA sequence assembly, DNA physical mapping, and DNA clone hybridization strategies and on theoretical approaches for optimizing the assembly process which is usually so combinatorically complex that approximative solutions such as simulated annealing and genetic algorithms are advantageous.

"The participants had a unique chance to share information about existing data and discuss alternative assembly strategies," Burks said. "As a result of this workshop, we anticipate new collaborations will be established among scientists in the Northern New Mexico research corridor in Albuquerque, Santa Fe, and Los Alamos."

External Faculty To Meet

The central core of SFI's scholarly family is an External Faculty of some 35 people from more than 20 institutions in the U.S. and Europe. SFI's External Faculty members each spend an average of one month a year at the Institute. These visits are often organized around workshops or working groups. Or visits may coincide with the residencies of other colleagues to work on problems of mutual interest.

These small gatherings notwithstanding, the External Faculty has never come together as one group for sustained, serious scientific discussion. To remedy this, SFI will hold an extended scientific meeting of the External Faculty from June 28 through July 10, 1993. This occasion will be an opportunity for the External Faculty to explore in some depth their common scientific interests and to become more deeply involved with the research life of SFI. The meeting will be largely unstructured, leaving ample time for individual discussion.

Attendees to date are W. Brian Arthur, Stanford; Jack Cowan, U. Chicago; Jim Crutchfield, U. California, Berkeley; Rob de Boer, U. Utrecht; Rob Farber, LANL; Doyne Farmer, The Prediction Company; Marcus Feldman, Stanford; Stephanie Forrest, U. New Mexico; John Geanakoplos, Yale; Jonathan Haas, Field Museum of Natural History; Alfred Hubler, U. Illinois; Stuart Kauffman, SFI; Alan Lapedes, LANL; John Miller, Carnegie-Mellon; Richard Palmer, Duk; Alan Perelson, LANL; Jose Scheinkman, U. Chicago; Peter Schuster, Instititute for Molecular Biology, Jena; and Gérard Weisbuch, Ecole Normale Supérieure.

LeBaron to Head Economics Program

Blake LeBaron (U. Wisconsin) will direct the SFI economics program from January through July, 1993. LeBaron has been involved with the Institute's economics program since 1989. He has been active in many of the workshops on time series and financial markets, and was recently in residence during the Fall of 1991. LeBaron's research involves the detection and understanding of nonlinearities in financial time series. Currently he is studying several aspects of financial time series and trader behavior including: the connections between stock trading volume and price movements, the statistical and evolutionary properties of simple price forecasting techniques (technical trading rules), and multivariate properties of stock index movements and lagged price adjustment. In general, LeBaron is interested in understanding what "different learning mechanisms look like" and whether we will be able to recognize them in observed series.

On his leadership role at the Institute, LeBaron notes, "I intend to steer the program in a more empirical direction, using techniques evolved at SFI to explain important economic facts." The research plan for next year, developed by LeBaron in collaboration with the executive committee of the economics board, emphasizes three areas:

- Theoretical approaches to pattern formation and to nonlinearities in the economy. These will combine techniques from probability theory and statistical mechanics to develop methods akin to "interacting particle" techniques in physics and mathematics. The program has done much work along these lines that this project will build on.

- Theory of financial markets. Financial theory is currently shifting from a perfect-rationality, equilibrium approach toward a bounded-rationality, evolutionary approach. The stock market model of Arthur, Holland, Palmer, and Tayler, and the work of Miller and Andreoni will launch new efforts in this direction. This work will rely upon the ideas of how people actually make choices and form market hypotheses, referred to earlier.

- Time-series analysis. The emphasis here will be on developing methods to detect underlying structure and nonlinear mechanisms in financial and other time series.

Plans are also in the works for a full-scale gathering in September, 1993. This major meeting—on the scale and level of the groundbreaking 1987 meeting which brought together economists and physical and biological scientists to forge a new conceptual framework for economics—will provide renewal and new directions for the economics program. The seven- to ten-day workshop will be co-chaired by Kenneth Arrow and Philip Anderson.

"This will be an important year for the economics program," says LeBaron. "I see the venture as being very close to solidifying the nexus from Santa Fe Institute tools to observed economic phenomenon."

Economics Research Program Works to Validate SF Approach

Research at the Santa Fe Institute is an attempt to chart new routes around scientific cul-de-sacs—the frustration that arises from being unable to understand or predict how whole systems work in spite of understanding in sometimes intricate detail how their parts function. It is a frustration shared by many scientists—physicists, biologists, and economists among many. In fact, some of SFI's earliest successes were inspired by the challenge by Citicorp Chairman John Reed to consider an alternative approach to economics, one that might provide better understanding of the world financial market. Reed was reflecting a practitioner's view that, while academic economics had made progress in quantifying the science and developing rigorously mathematical theories of economic behavior, it had done so with the requirement for simplifications and assumptions about behavior that seemed increasingly to diverge from the real world. Indeed, the trend to more rigorously mathematical treatment had tended to pull economics along paths similar to those taken by other disciplines—into finer detail and away from the gross structures.

SFI took up Reed's challenge and, over the past four years, created and began to validate an approach to understanding economic systems that is at once both radical and commonsensical. The environment at SFI has attracted a brash mixture of economists, biologists, and computer scientists who have explored ways to model economic behavior that reflect the limits of human knowledge and the importance of learning—two factors that tend to be "assumed away" in classical economics.

Inspired in part by the emerging "Santa Fe approach," economics is currently undergoing a major long-term shift in its way of looking at the world. Within the discipline there is a growing perception that the standard neoclassical viewpoint, with its emphasis on equilibrium, stasis, linearity, and perfect rationality of agents, has become to some degree "mined out." Many first-rate economists believe that this standard approach needs to be complemented by a viewpoint that emphasizes process, evolution, nonlinearity, and the bounded rationality of agents. Santa Fe has already had a major impact on economics, in that ideas that were not considered "kosher" previously are now routinely invoked by conventional economists at MIT, Stanford, and the University of Chicago.

The program is not just an exercise in academic thinking. Many of the problems that face the U.S. economy—trading internationally in high-tech goods, maneuvering strategically in complicated free-trade negotiations, shifting from a government-directed to a market-directed economy, deregulating industries, advising the emerging Eastern European economies—need theory and methods that include the realities of positive feedbacks, adaptive processes, and limits to human rationality. These are largely missing from conventional textbook theories and, in the past, have led to misguided policies. Santa Fe aims to bring common sense and much improved theories to these policy discussions.

Research Themes of the Program

It is useful to map out some of the broad themes that the program has concerned itself with from the beginning and that are likely to persist over the next several years. These themes consist of a central core and three partially overlapping, peripheral topics.

Core Research

The program envisions the economy as composed of large numbers of interacting agents, mutually adjusting to each other as time passes: in that sense the economy is a massively parallel system that is continually

The program is not just an exercise in academic thinking. Many of the problems that face the U.S. economy...need theory and methods that include the realities of positive feedbacks, adaptive processes, and limits to human rationality. These are largely missing from conventional textbook theories and, in the past, have led to misguided policies.

adapting, continually in process. The agents in this economy—the "interacting particles" of economics—decide their actions consciously, with a view to the possible future actions and reactions of other agents. That is, they formulate strategy and expectations. In doing this they may be faced with complicated and possibly ill-defined problems that are far beyond the scope of normal human intelligence to solve completely. Hence, they are often forced to act inductively: they form internal models; they transfer experience from other, similar problems; they generalize from limited data; and they learn as they go. As this learning and mutual adaptation take place, sometimes new economic structures emerge, and there is a continual formation and reformation of the institutions, behaviors, and technologies that comprise the economy. Some parts of the economy may be "attracted" to an equilibrium; some parts may continually evolve and never settle.

The central task of the program is to develop the methods and to formulate and solve the demonstration problems that will articulate this picture of the economy. Sometimes this calls for mathematical analysis; sometimes, for computer experimentation.

These core themes of interactions among agents, massive parallelism, dynamics, adaptation, and emergence overlap strongly with core concerns in the Institute's research program as a whole and, thus, there is a continuous exchange of ideas with other Institute programs.

Nonlinearities

Nonlinear mechanisms and positive feedbacks are normally present in the economy, but they have been relatively little researched. Today's economies are no longer based on agriculture and bulk manufacturing, but rather on high-technology industries and increasingly sophisticated services. High-technology products—like new pharmaceuticals, Cray computers, or Microsoft Windows—share the property that most of their costs are upfront R&D investment costs. The more market they capture, the lower their per-unit costs. Hence, they show positive feedback or "increasing returns" to market share. (Learning within the service industries shows the same effects.) At Santa Fe, we are beginning to reach a clear understanding of markets that display increasing returns. They tend to be highly unstable at the outset. Hence, a firm or a country that plays in such markets can effectively "tip" them at the outset, and lock in the market to their own advantage. How this should affect U.S. policy toward technology and industry is one of our concerns. How a firm should maneuver in a complicated positive-feedback environment is another.

Pattern Formation

Closely overlapping with the above theme is that of pattern formation—the emergence of new, sometimes unexpected macrostructures from the microinteractions of firms, industries, and financial agencies. Of interest here is the formation of cooperative structures over time and the "selection" of one from many possible potential candidate structures by processes of adaptation. An example would be the evolving pattern of industry in Eastern Europe, or in the new U.S.-Canada-Mexico free-trade environment. How do new macrostructures form? What are the possible candidate patterns of trade and industry? Can pattern formation theories in physics and biology help our understanding in economics? These are some of the questions we are asking in the program.

Empirical Research

At Santa Fe we are not interested in theories built upon imagined, idealized views of human decision-making. Rather we formulate

continued next page

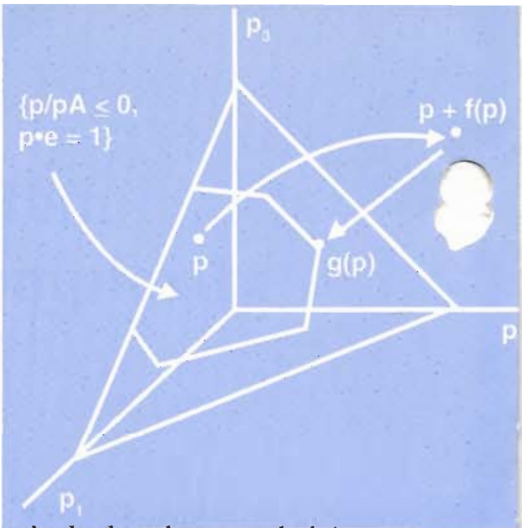

Screened image is redrawn from a continuous function graph that appeared in Timothy Kehoe's chapter in The Economy as an Evolving, Complex System, edited by Anderson, Arrow, and Pines (Addison-Wesley, 1988); reprinted by permission.

the assumptions we use analytically, from rigorous observation of *actual*, human economic behavior. This entails observation and analysis of human behavior in bargaining situations, financial markets, and in experiments in economics laboratories. And it entails the participation of psychologists and others who know how human decision-making *actually* works.

Along the same lines we are also interested in detecting the presence of underlying behavioral patterns in time-series data that are often of limited duration and highly corrupted by noise. Standard theory in economics usually tells us that many series of interest—financial series for example—should show *no* patterns at all: they are merely random walks. We find this not to be the case. One of the most important products of the program, in fact, is methods that detect underlying patterns and that use these to predict. Using the machine-learning techniques of John Holland, John Koza, Doyne Farmer, and other participants, Santa Fe has pioneered approaches that seek out the real behavior in markets and that contrast this with imagined, theoretical behavior.

For these past four years, the economics research program at the Santa Fe Institute has been building an adaptive, complex, evolutionary viewpoint into the central body of economic theory. The long-term objective of this program is to articulate this new viewpoint—and to provide methods, theories, frameworks, and solutions that will help catalyze this change. With the 1993 program in place under the direction of Blake LeBaron, the Institute is poised to lead this effort into the future.

Anticipated 1993 Workshops

SPRING
Adaptive Computation Workshop in Robotics
Melanie Mitchell, University of Michigan

MARCH 7–11
Reinforcement Learning in Robotics
Nils Nilsson, Stanford University
Melanie Mitchell, University of Michigan

APRIL 13-17
Spatio-Temporal Patterns in Nonequilibrium Complex Systems
P.E. Cladis, AT&T Bell Laboratories
P. Palffy-Muhoray, Kent State University

MAY 30-JUNE 25
6th Annual Complex Systems Summer School
Lynn Nadel, University of Arizona
Daniel Stein, University of Arizona

JUNE 11–13
Project 2050: Crude Look at the Whole
Murray Gell-Mann, California Institute of Technology

JUNE 28-JULY 10
Scientific Meeting of the SFI External Faculty

SUMMER
Modeling the Interaction Between Learning and Evolution
Rik Belew, University of California, San Diego
Melanie Mitchell, University of Michigan

JULY/AUGUST
Theoretical Neurobiology Working Group
Nancy Kopell, Boston University
Michael Stryker, University of California Medical School, San Francisco

Competitors to Earn Cash in AZTE Competition

Is "artificial intelligence" superior to human intelligence? In some domains such as chess, computer programs now outperform all but the very best human players. However, in other domains such as speech, handwriting, and other kinds of pattern recognition, computers lag far behind human beings. On Wall Street, computer "program traders" are becoming increasingly common, yet there is substantial controversy over their performance—they have even been blamed as a factor in the October, 1987 stock market crash.

The Arizona Token Exchange (AZTE)—a computerized market in which a fictional commodity called "tokens" are traded—has been created to compare the performance of human and program traders. The purpose of this study, co-sponsored by the University of Arizona's Economic Science Laboratory and the Santa Fe Institute, is to compare the performance of human and program traders to see whether humans can learn to exploit the limitations and idiosyncrasies of computers in repeated interactions.

In each trading session on AZTE, traders are assigned the role of buyer or seller and are given an allocation of tokens. A seller's objective is to sell tokens for as much as possible above the token cost and a buyer's objective is to buy tokens as cheaply as possible below their redemption value.

By ranking the token costs and redemption values, well-defined supply and demand curves can be constructed. The intersection of these curves defines the so-called competitive equilibrium (CE) price and quantity, at which neoclassical economic theory predicts all trading will occur. The complication is that in the AZTE, each trader's token costs and redemption values are private information and differ from trader to trader. Thus, traders in the AZTE face a complex sequential decision problem: how much should they bid or ask for their own tokens, how soon should they place a bid or ask, and under what circumstances should they accept an outstanding bid or ask from some other trader?

Unlike real commodities markets where most traders are humans, in the AZTE, human competitors play against computer programs. The opponent programs will be selected from a field of over 30 different trading strategies including winners of the Santa Fe Institute's Double Auction Tournament held in March, 1990.

To trade on the AZTE, competitors need a UNIX or PC-compatible computer linked to the Internet computer network. The organizers provide the trading interface software that allows a competitor to log on and trade at any time and for as long as he likes (subject to general restrictions). Dollar earnings accumulate in individual accounts until time of cash out.

The software and an ASCII traders' manual (including the application form) are available via anonymous ftp on "fido.econ.arizona.edu," in the azte sub-directory.

For those who don't have access to anonymous ftp, a diskette containing the software and traders manual is available. To cover the costs of a diskette and surface mail send $5.00 to:

Shawn LaMaster
Manager, Economic Science Systems Development
Economic Science Laboratory
McClelland Hall, Room #116
University of Arizona
Tucson, AZ 85719
(602) 621-6218
Internet: lamaster@ziggy.econ.arizona.edu

Research Update

Adaptive Computation at SFI

Computation is an integral part of the study of complex adaptive systems. It plays many roles. First, computers are tools for modeling natural systems, and in many cases, computer simulation of a complex system may be the only way to gain insight into the system's behavior. Second, computer science often borrows ideas from natural complex adaptive systems to create new problem-solving methods; some examples of nature-inspired computational methods are neural networks, genetic algorithms, classifier systems (based on ideas from economic systems), and simulated annealing. Finally, the notion of "computation" is a conceptual framework for scientists thinking about the behavior of complex adaptive systems in nature. For example, one might ask of a complex system, what sort of computation is it doing? How complex is that computation? Such questions lead to a broadening of traditional notions of how and where information processing can take place.

Much of the research at SFI involves computation in one or more of these roles, and the SFI's Adaptive Computation (AC) program was formed to coordinate and unify such research efforts. The idea of a research program in adaptive computation was inspired in part by the Institute's successful economics program, whose purpose is to forge connections between research on economics and on other complex systems, both to discover new ways of thinking about economics and to communicate ideas from economics that might be useful to researchers in other fields. The AC program is the computer science analog of the economics program: its purpose is to make fundamental progress on issues in computer science that are related to complex adaptive systems, and to export the results to researchers in other disciplines. The program, which has existed informally for several years, was officially launched last spring with a Founding Workshop, supported by the Alfred P. Sloan Foundation. Most of the central issues of the field were covered in talks by AC program members and other invitees from the fields of artificial intelligence, mathematical genetics, ecology, economics, psychology, and physics, as well as a number of people in industry.

Melanie Mitchell (U. Michigan) is director of the 1992–93 AC program. She has been a collaborator with SFI workers since 1990, both as a member of the University of Michigan/SFI collaborative research project and the AC program. Mitchell's work focuses on artificial intelligence and machine learning—in particular, genetic algorithms and classifier systems—and on computer modeling of perception and analogy. Mitchell's job as AC resident director is to help set the direction and focus of the program, to oversee its research and fund-raising activities, in addition to conducting her own research. Mitchell wears all of these hats well, and the current program is now in full swing. It includes long-term residential research projects, a set of shorter workshops, and a number of individual visitors.

Residential Research

Adaptive computation is potentially such a broad area that many of the existing research projects at the SFI could come under its rubric. The Institute has chosen to seek funding for a small subset of those projects, with the hope that this initial core will be expanded later on. Proposals to fund continued work within the central parts of the program are currently under consideration at the National Science Foundation and the Department of Energy. The core projects include research headed by Alan Lapedes on applying machine-learning methods (such as neural networks and genetic algorithms) to the prediction of protein structure; research led by John Holland on the ECHO system, a computational base for studying ecological economic and other phenomena characterized by networks of interacting agents; work directed by Chris Langton on the process gas simulator, a general purpose simulation tool for modeling complex systems; work by John Miller on computational models of adaptive economic agents and the formation of economic markets; Tom Ray's project on the evolution and ecological interaction of digital organisms; work by Stephanie Forrest and Mitchell on the foundations of genetic algorithms; and research by David Wolpert on the foundations of supervised learning methods.

Workshops and Working Groups

The meetings scheduled for Fall, 1992 are:

- Issues in genetic algorithms and classifier systems (September 23–24)
- Approaches to Artificial Intelligence (Nov. 6–9)
- Computation, dynamical systems, and learning (Nov. 16–20)
- Genetic algorithms and neural networks in protein-folding prediction (Dec. 14–15)

Meetings planned for January–August 1993 (so far) are:

- Adaptive computation in robotics
- Adaptive computation and economics (jointly organized with the SFI Economics Research Program)
- Sparse distributed memory
- Genetic algorithms and real genetics
- Modeling the interaction among evolution, learning, and culture

In addition to the small, specific-topic meetings, there are plans for a larger general Adaptive Computation workshop, to take place in late spring or summer 1993. This workshop will bring together many of the current workers in the AC program as well as a number of additional participants. It will serve as a forum for discussion of the scientific advances and questions encountered by the various research projects in the AC program in the past year; as an opportunity to step back and once

continued next page

Jones Developing ECHO in UNIX

Image appears by permission of David Moser.

Terry Jones, a Ph.D. student from Indiana University, has spent the last five months at SFI working closely with Stephanie Forrest and John Holland to develop a UNIX version of John's ECHO model. Terry's dissertation also concerns ECHO. ECHO is a very flexible general-purpose modeling environment in which independent agents interact at many sites in a computer-simulated world. The model allows researchers from many fields to design all the characteristics of these agents, sites, and worlds and to exercise fine-grained control over the simulations they run. Tools like ECHO should enable researchers to gain insights into a variety of complex systems including the immune system, the ecology, and the economy.

One of the first real applications and verifications of the validity of ECHO is likely to be done in conjunction with James Brown in the Biology Department at the University of New Mexico. Brown has wide experience and approximately 25 years of data from ecologies that are perturbed by the removal of species. ECHO is particularly well suited for simulations of this kind. Terry is already reporting exciting ECHO runs that exhibit characteristics of natural ecosystems.

again examine some fundamental questions related to adaptive computation; and as a way in which to gain the participation of outside invitees who can add new ideas and perspectives, and who may eventually become more deeply involved in the program.

Outreach and Fund-Raising Activities

Funding is, of course, a big concern for the program: the various activities of the comprehensive project go beyond what the SFI core budget can provide. Mitchell—along with the advisory committee for the AC program and the SFI staff—devotes a considerable amount of time to fund-raising and outreach activities.

Since the ideas and approaches being explored by the Institute's AC program are at the frontiers of current computer science and computational modeling, it is important to communicate these concepts to a larger audience, including both scientific and industrial communities. "Outreach events" says Mitchell, "help publicize the existence and activities of SFI to people who can potentially be of great help to the Institute, either as donors, or as contacts in other research and applications communities." To this end, Nils Nilsson, Professor of Computer Science at Stanford and member of the AC Program Directorate, and David Liddle, co-founder of Interval Research and SFI Trustee, helped organize an Adaptive Computation Symposium at Stanford in July, 1992. The audience consisted of influential scientists and industrialists from the Bay Area; they spent an afternoon listening to talks by Ed Knapp, John Holland, John Miller, Tom Ray, and Melanie Mitchell about the AC Program and its various activities. More recently, David Liddle arranged an afternoon symposium in Palo Alto for members of the U.S. Venture Partners, a venture capital group which has extensive interest in biotechnology, but which wanted to hear about the "next wave" in computation. The speakers at that event were Mitchell, Marc Feldman, Professor of Biology at Stanford and SFI Science Board member, and John Koza, Professor of Computer Science at Stanford and SFI Trustee.

Mitchell enjoys spreading the word. "Work on adaptive computation at SFI is tremendously exciting," she notes. "In my years in graduate school at Michigan, professors such as John Holland and Art Burks infused students with the notion that computer science is rightly the study of information processing, not only in computers, but throughout all of nature. Computational research at SFI captures that remarkable spirit."

Holland Named Henry Russel Lecturer

John Holland, AC network member and Professor of Computer Science and Engineering and Professor of Psychology at the University of Michigan, has been named the University's 1993 Henry Russel Lecturer. The annual lectureship is the University's highest honor given to a senior faculty member. Holland will deliver the Russel Lecture in March, 1993.

Holland, who received the MacArthur Foundation Fellowship last June, is an expert on cognition and artificial intelligence. He studies complex adaptive systems—constantly evolving, changing, and interacting mechanisms that control everything from how people learn, to global economics and politics.

Mitchell, Crutchfield, and Hraber Evolve CA Rules

It has long been known that extraordinarily simple machines can be assembled to make more complicated devices. In association, these simple machines can be used to do interesting and useful things. These devices manifest themselves in many ways, from those like the wheel and axle and lever which make physical labor a little less demanding, to tiny electronic components that modify a current to a given degree. Within the microcosm of the computer, there exists such a simple device as the cellular automaton. In its purest form, such an automaton consists of "cells" which can either be "on" or "off" at any given moment, and owes whether it will be on or off at the next moment solely to the current on/off states of itself and two or more neighboring cells. The possibilities can be summarized in a single table of rules corresponding to all of the possible states for a cell and its "neighbors."

What is so intriguing about the cellular automaton is that a collection of cells, when linked to form a long lattice, can behave in interesting ways. Many people at the Santa Fe Institute are currently engaged in studies of this emergent complex behavior, including Melanie Mitchell, Resident Director of the Adaptive Computation Program, and her assistant, Graduate Fellow Peter Hraber. Together they are working with cellular automata rules, particularly their emergent ability to perform logical functions, or calculations. It has been theorized in past work that some of the most interesting things which a CA rule can do occur under conditions in which the rule is on the threshold between chaotic behavior and intermediate-term periodicity. There have been many attempts to devise means for quantifying this transition. In examining the relation between a CA rule's dynamical class and its ability to perform useful computations, Melanie and Peter hope to understand salient features among the various indices of complexity, as well as address the issue of whether this regime of useful computation does indeed belong to the transition between intermediate-term periodicity and chaotic behavior.

At the same time, Melanie and Peter are implementing a genetic algorithm to find those CA rules that perform interesting computations. The genetic algorithm is a programming paradigm inspired by evolutionary dynamics and the notion that an organism is a unique solution to a set of problems. This approach is exciting because it is an example of a case in which a general set of programming rules can be used to develop more specific "programs" to perform any number of user-defined computational tasks. Melanie and Peter are currently experimenting with various attributes of the genetic algorithm that affect its ability to effectively "navigate" through the vast computational space of all possible CA rules, and the capacity, in general, for the genetic algorithm to search this space in comparison to other exploratory algorithms.

In the long term their work may lead to a clearer conception of the mechanisms of computation in general, including the notion that other dynamical systems, such as those physical and biological, may have dynamical attributes in common with those of computational systems, and may be conceived as thriving on the same threshold between chaotic dynamical behavior and the more static intermediate-term periodic dynamics. The work also potentially offers insights into the implementation of self-programming paradigms, particularly in light of the inclinations toward massively parallel computation, a task for which CA are particularly well suited.

Research Update

Swarm Simulation System Focuses on Roots of Complexity

The "Swarm" Simulation System is a generalized programming framework for simulating and studying the complex behaviors that arise in systems composed of many components. This project, part of SFI's Alife program, will create a fundamental tool for complex systems research. The project, led by SFI External Faculty member Chris Langton, builds on earlier work by him and others but began new directions this fall with support from the Carol O'Donnell Foundation.

A Common Architecture

Most of the complex systems that occur in nature share a common architecture, which we refer to as a "Swarm." The term refers to a large collection of simple "agents" interacting with each other. The classic example of a Swarm is a swarm of bees, whether in flight or inhabiting a nest. The notion of a Swarm can be readily extended to other systems with a similar architecture. Thus, an ant colony is a Swarm whose individual agents are ants. A flock of birds is a Swarm whose agents are birds. Traffic is a Swarm of cars, and a crowd is a Swarm of people. Extending the concept even further, a gas is a Swarm of molecules, an immune system is a Swarm of cells, and an economy is a Swarm of economic agents.

The operating requirement of a Swarm is that it shows collective behavior (of many types, not just motion). For example, a Swarm of bees behaves in very complex ways, even when all of the bees are sitting in the nest. Likewise, the individual agents that constitute an economy do not necessarily move in space, but the dynamics of the transactions they effect among themselves are quite complex.

The essential structure of a Swarm is a collection of relatively autonomous entities *with no central organization*. In a Swarm, each of the thousands of individuals makes its own behavioral choices, based upon its own evaluation of the local environment and upon its communication with other nearby individuals. There is no central authority (no "drill sergeant," leader, or choreographer) directing the behavior of the Swarm.

What makes such Swarms scientifically important is that, although the individuals have limited intelligence and exhibit simple behavior, their collective behavior can appear to be highly intelligent (and complex). Thus, Swarms, by illuminating the point at which simplicity becomes intelligence, allow us to focus directly on the fundamental roots of complexity and intelligence.

Beneficiaries of Swarm

Nearly three dozen simulation projects, at SFI and elsewhere, have been already identified as possible beneficiaries of the Swarm programming framework. These include such applications as self-assembly at the molecular level (relating to origins of life, synthesis of drugs, actions of antibodies), building an "artificial" cell or economy or stock market, understanding and applying to physical systems evolution and coevolution, simulating long-term global change, building "artificial" insects as a route to autonomous, adaptive robots, transportation system design, and geopolitical stability. The goal of the project is to produce a "proof of principle" that other organizations can carry to higher levels of sophistication and to specific applications. It will be designed to be applicable to many computer architectures, allowing it to run on a single computer, over a heterogeneous network of workstations, or on a massively parallel computer. It is expected that the project will take one to two years to complete. Active full-time effort began this October with the arrival of David Hiebeler, a programmer from Thinking Machines Corporation (TMC), who

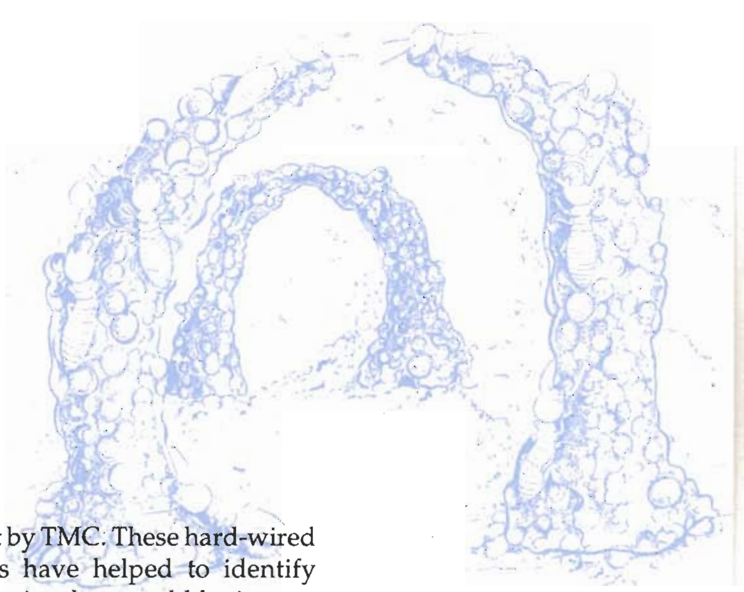

Image from The Insect Societies by Edward O. Wilson, copyrighted © 1971 by the President and Fellows of Harvard College. Reprinted by permission. Illustration by Turid Hölldobler; based on Grassé (Experientia 15(10) (1959): 365) and Chauvin (Animal Societies from the Bee to the Gorilla, Hill and Wang, 1968).

is in residence at SFI working full-time under the direction of Langton.

The general form of the Swarm Simulation System (SSS) will be as a distributed, concurrent, object-oriented programming language combined with a set of window-based tools for the specification and management of Swarm simulations. Users will be able to specify the behaviors of the individual "agents" (bees, ants, stock traders, antibodies, etc.) making up the Swarm in any of several familiar programming languages, such as C, C++, LISP, or Fortran. They will be able to specify the properties and behaviors of the "physical environment" in which the Swarm is embedded. Users will be able to specify the initial state and distribution of the Swarm in that environment. They will be able to "instrument" the simulation with various predefined data-collecting and data-graphing objects, or they may define their own. Finally, users will be able to run the simulation interactively, altering the state, arrangement, or even the rules of the individual agents or the environment as they wish during the simulation.

The ideal implementation of a Swarm system would be as a parallel object-oriented programming language running on a massively parallel computer. In practice, the implementation of such a simulation system will go beyond the implementation of a programming language. It will also involve the implementation of the operating system needed to manage the execution of, and communications between, many thousands of individuals.

SFI's Role

SFI has done a good deal of work on identifying the computational issues involved in a Swarm system and has implemented several specific (not general-purpose) simulations on a massively parallel 64,000-processor Connection Machine, built by TMC. These hard-wired Swarm prototypes have helped to identify many of the properties that would be important in a general-purpose version. Ultimately, Swarm will lead to a new programming environment consisting of an integrated set of tools for creating, running, and analyzing distributed, multiagent systems that would capture many of the features of the LISP programming environments developed for the pursuit of artificial intelligence. In fact, it is quite likely that many of the artificial intelligence programming tools could be incorporated into a Swarm programming environment with little alteration to their basic functionality.

It is important to note that most of the current wave of business and government interest in "massively parallel computing" is headed in a very different direction. The mainstream work on parallel computation involves extending traditional serial computing into the parallel domain. Such approaches devote a good deal of effort to eliminating exactly what Swarm is based on: the nonlinear interactions that represent the way real-world systems behave.

Although SFI plans to develop a good, working prototype of a Swarm Simulation System, the implementation of a full production version is beyond the Institute's present capabilities. SFI's course of action will be to construct a "proof of principle" system, while at the same time recruiting other research groups involved in similar implementation efforts to join SFI in a collective, multi-institution effort to produce a full production version of an SSS. SFI's international recognition as a pioneer in studies of complex adaptive systems, and its undisputed leadership in the rapidly growing field of Artificial Life, makes SFI the natural focus for these efforts. We expect a rapid convergence of researchers from other organizations around the Swarm project.

Research Update

Prehistoric Villages as Complex Adaptive Systems

Tim Kohler

If you were living near Dolores, Colorado, about A.D. 800, or near Los Alamos, New Mexico, around A.D. 1250, chances are very good that you lived in a small hamlet with no more than a few kin as next-door neighbors. Fifty years later, in both areas, your children would be living in villages several times that size. Archaeologists refer to this "phase transition" as aggregation. Why villages form out of dispersed habitations here in the Southwest, and in Neolithic and peasant societies all over the world, has become a central problem in archaeology.

Not that there aren't some obvious factors that do seem to explain aggregation, at least in some proximate sense, in some times and places. Anyone who has traveled through southern France, for example, has seen the walled villages that first formed on hilltops for defense from Medieval bandits.

But in other times and places, villages seem to form in periods of local peace and prosperity. Given the apparent disadvantages of village life (including increased distance to fields, potential for conflict among neighbors, and threats to community health), why bother? Microeconomic theory and the theory of non-zero-sum games—two of the simplest tools in the arsenal of methods for getting into complex adaptive social systems—are providing some surprising insights into the process as it took place in the Mesa Verde region of southwestern Colorado between A.D. 900 and 1300.

Tree Rings, Corn, and Cooperation

Southwest Colorado north of Mesa Verde National Park—still the "Pinto Bean Capital of the World"—contains some of the densest remains of prehistoric habitations in the American Southwest. The Anasazi occupants of these sites depended on rainfall agriculture for most of their caloric intake. But the summer precipitation needed in this area for corn to mature is highly variable and often localized. When populations were low, farmers with agricultural shortfalls could probably make up their deficits by foraging for wild resources. But another way to buffer such risks—more workable when population levels were high—might be to have many households farming different areas pool their produce. This thought experiment leads to a simple model for explaining aggregation and the "public architecture" (such as towers and great kivas) that goes with it: villages ought to emerge out of dispersed settlements in times and places where dry-farming production is highly variable over time and space. And social complexity within such villages ought to increase as managers who can assist with this shuffling of production across households become worth their overhead.

Carla Van West, a recent doctoral student at WSU, developed a database that made it possible to test this idea. Van West produced computerized maps which retrodicted potential maize production across a 1500 km^2 portion of Southwest Colorado, for every year between A.D. 900 and 1300, at a spatial scale of 4 ha. For our purposes here, all that we need to know about this feat is that it was possible because information is embedded in ancient tree rings about prehistoric precipitation. This, plus the twentieth-century weather records, agricultural records, soil surveys, and tree-ring records in this area make it possible to establish regular relations for various soils between precipitation and maize production for both the modern and the prehistoric case.

When we compared her results with the relatively well-known archaeological record for this area, our first conclusion was that villages were *not* forming in times of economic stress (which would, presumably, also have been times when raiding was common). This weakens the defensive hypothesis for village formation. Instead, these villages were form-

Tim Kohler is an Associate Professor of Anthropology at Washington State University–Pullman, and a Member of the Santa Fe Institute.

> The village is one of humanity's oldest and most enduring social institutions, emerging almost universally with or after the development of agriculture in societies around the world. But how and exactly why villages form remains a mystery.

ing when production was variable across time and space, as suggested by the thought experiment, but only in the context of periods when production was also generally high. We explain this result in terms of microeconomic theory.

Figure 1 summarizes what we hypothesize to be the relation between the value of a certain level of production to a household (on the y axis) and the absolute level of that production (on the x axis). When year-to-year production per household is generally high (as in Figure 1(a)), households maximize the value of that production by sharing it with other households. However, when production is generally low, as in Figure 1(b), households maximize the average value they realize from their production by hoarding it.

By extending this same logic, it can be shown that the conditions which most favor food sharing among households are spatially patchy and temporally variable production in the context of generally high yields. This is also a good characterization of the periods in which villages form, which makes sense if we make the additional assumption that one important reason for villages to form in this society is to facilitate and formalize food sharing among households. On the other hand, the conditions which least favor food sharing among households are generally low yields coupled with spatial patchiness and high temporal variability. In fact, in our study area, villages break up into their smaller constituent kin units during such periods. It is also during these unfavorable periods that we see the most evidence for such antisocial activities as burning of structures, interpersonal violence, and cannibalism.

From Analysis to Simulation

Now that we know something about the conditions under which villages form in this

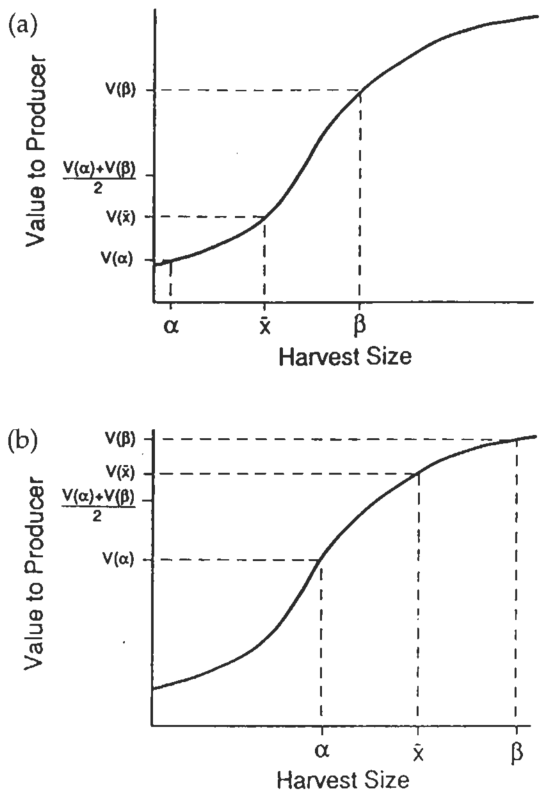

Figure 1. (a) Good-year economics. Units of production on the x axis; units of value to the producer of a certain level of harvest on the y axis. The variability notated by α and β could represent different plots in the same year, or different years in a series of generally good years. $V(\bar{x})$ represents the average value realized after pooling; $(V(\alpha)+V(\beta))/2$ represents the average value realized by not pooling. (b) Bad-year economics. In contrast to good years, in bad years $V(\bar{x})$(the expected value of pooling) is less than $(V(\alpha)+V(\beta))/2$ (the expected value of not pooling).

continued next page

Prehistoric Villages as Complex Adaptive... (cont'd.)

Image shows partial map of Sand Canyon Pueblo: an example of a mid-1200s town in SW Colorado. Courtesy Crow Canyon Archaeological Center, Cortez.

area, we will be able to move from a reductionist/analysis mode to a simulation mode, more in the spirit of other Santa Fe Institute inquiry. For example, this system could be modeled as an iterated n-person game, with payoffs structured according to Axelrod's famous Prisoner's Dilemma. Because we have shown that village formation is a behavior which can emerge from households simply acting in their own best interests, it should also be possible to simulate village formation as a "stigmergic process" through Swarm Intelligence approaches, used profitably for studying how social insects (for example) can achieve a collective task without any central organizer or blueprint. Computational approaches to understanding the development of complex adaptive systems are relatively new in any field, but especially novel in archaeology. SFI provides an atmosphere conducive to the development of such interdisciplinary ventures that offer great potential for generating new insights into past cultural systems.

Diehl Joins Prehistory Program

Michael Diehl (SUNY Buffalo) is at the Institute through next May as a Predoctoral Fellow in SFI's Cultural Evolution/Southwest Prehistory program. His work resonates with Kohler's on prehistoric village formation.

Some human societies rely on subsistence resources that are spatially patchy, temporally unpredictable, or easily overexploited. Based upon limited ethnographic data, anthropologists have suggested that in such situations competitive accumulation of goods to enhance personal prestige or wealth is socially suppressed. Status competition destabilizes social relations among people who are highly interdependent, and can lead to the destruction of the resource base through overconsumption. Is social suppression of status competition typical of groups who rely on unpredictable or easily overused resources?

Diehl's work compares attributes of the subsistence economies of the Early Pithouse (A.D. 200–600) and Late Pithouse (A.D. 600–1000) periods of the Mogollon region of the prehistoric Southwest. He is also looking for evidence of extreme differences in household possession of items that may have been used to mark status or prestige. By assessing changes in parameters of subsistence and wealth, he can determine whether unequal accumulation of prestige-enhancing items became more acceptable as prehistoric Mogollones increasingly relied on a more stable and productive resource base.

This research is important for anthropologists who attempt to model the emergence of formalized leadership positions, because many scholars have observed that chiefs, religious leaders, formal ruling class members, or less formally empowered "big men" get and build power in part by ostentatious display of rare goods, and by competitive accumulation and consumption. Diehl's work will contribute to the discussion by assessing whether the emergence of privileged leaders is an inherent tendency of human societies.

Cowan Chairs Conference on Common Themes

SFI's recent conference *Common Themes in Complex Adaptive Systems* set itself no small task. One of its goals was to push forward the Institute's quest to explicitly define the fundamental concepts concerning the nature and operation of self-organizing and adaptive complex systems. The second aim was to define a major set of questions concerning these systems to guide both near- and long-term research in this growing field.

Meeting chairman George Cowan opened the program by asking the participants to help define the common patterns, if any, that shape the remarkable behavior of the various complex adaptive systems of interest to SFI. More than two dozen members of the SFI family presented papers on a wide variety of topics illustrating various levels in the hierarchy of complexity. Subjects ranged from cellular automata, which are not complex but share some of the properties of complex systems; through the origin of life and of proteins; the functioning of RNA and viruses; developmental complexity and evolution, including the evolution of individuality and of the mammalian brain; ecological systems; cognition and human learning; concluding with the social sciences, particularly economics but also touching on a broader set of questions dealing with global sustainability and human behavior.

Mirroring its subject, the workshop was an exercise in self-adaptation: talk bounced back and forth from broad, fundamental questions like the definitions and measures of complexity and the differing views of the nature of learning to discussion of specific systems and reality checks. People continually revisited the early general discussions, refining original ideas in light of criticism by the other general theorists, and especially in the context of what they heard from people doing work in the specific fields. In this way the meeting encouraged thinking and talking about complex systems in terms of pragmatic usefulness—not just in light of universal theories or potential experiments.

Participants talked about mathematical and computational methods for describing the behavior of complex systems—more or less approximately. These tools include neural nets, genetic algorithms, nonlinear dynamical methods for dealing with chaos, and the notion of self-organized criticality leading to "avalanches" at all scales.

Did the workshop meet its goals? "I'm very excited about the feeling that there is a growing synthesis," notes Science Board Co-chair Murray Gell-Mann. "It's premature to say there exists a rigidly defined science of 'complexity,'" cautions Co-chair David Pines. But "we have come up with a whole set of candidate metaphors for complex adaptive systems that look very promising for a broad range of systems." This material will be available in *Integrative Themes*, the proceedings volume for this meeting due out from SFI/Addison Wesley by late summer.

Complex Systems in the Schools

This election year there has been much concern about the scientific illiteracy of American youth. Beginning last Spring SFI took a step beyond the rhetoric to initiate its own community-based secondary school program. The pilot project is designed to introduce the concepts of complex systems to local high school students. It consists of a series of in-school lectures by the Institute's research staff and visitors each followed by a discussion session for a more limited number of students. About sixty students from all six of the city's public and private high schools are taking part.

The project is headed by Suzy Pines, who brings to the work extensive experience organizing school-based programs. She is assisted by administrative coordinator Margaret Alexander. Support is provided by a seed grant from the Pinewood Foundation.

The program has three goals: to extend complex systems topics to high school students; to encourage more minorities and women to pursue a career in science; and to introduce scientific information and values to the nonscientist student. In-school lecture topics for the spring term included Chris Langton on artificial life, Alan Perelson on immunology and AIDS, Erica Jen on cellular automata, and Geoffrey West on scaling phenomena. Students gathered for presentations on roughly a monthly basis: a 90-minute introductory talk on a topic was featured each Monday; this was followed up later in the week by an informal discussion at SFI—often including hands-on computer simulations—which probed that topic in more detail. The program for the 1992–1993 academic year covers subjects like prehistoric archeaology, origins of life, neural biology, adaptive computation, and astrophysics—all considered within a complex systems context.

From the start an integral part of the Institute's mission has been to provide educational opportunities in the sciences of complexity, and SFI has done so at the postdoctoral, graduate, and even the undergraduate level. But extending its educational mission to the local high school community offers some particular challenges. First, there is the task of presenting material at the right level and in a provocative format. For most of the speakers, used to communicating in scientific shorthand with their colleagues, this is no easy task. "It was very hard but also very rewarding to prepare and deliver a talk to kids who have grown up on Nintendo and personal computers with buttons that automatically generate fractal landscapes," say Jen. "I wanted most of all to communicate to them the value of mathematics in formulating, and in providing solid answers to, the basic questions of complex systems and to do so without putting half the audience to sleep. The surprising part to me was their receptivity to some of the novel ideas of complex systems and their perceptiveness in asking right on (and to me humbling) questions like 'How do you know when a model is right?' and 'What does it mean to be random?'"

In addition, younger students in particular want to know how such basic research is relevant to today's world. They also want to be introduced not only to a topic, but to a researcher's persona, style, and history—including his route to science. The organizers agree that offering new perspectives about career choices is an important part of the program.

The Institute has formed an active partnership with Santa Fe's educational community to craft a program to best meet these student needs. Using their own criteria, local teachers select students from their school to attend any number of the talks. Teachers act as co-hosts for the talks and they meet beforehand with the scientists so that the material is presented in a way appropriate to intelligent and curious high schoolers. Students and teachers alike

"The surprising part to me was their receptivity to some of the novel ideas of complex systems and their perceptiveness in asking right on (and to me humbling) questions like 'How do you know when a model is right?' and 'What does it mean to be random?'" –Erica Jen, Los Alamos National Laboratory

take an active role in the evaluation process.

Even in its early stages, the project is a success. "We are off to an exciting start," says Pines. "Santa Fe students and teachers are equal partners as they stretch to understand how the familiar walls around disciplines may be breached. SFI scientists are providing a window through which we can all look to get a different perspective about science at the frontier—maybe the science of the 21st century."

On the one hand, students are fascinated by their first look at complexity sciences. One student assessed his experience with the project as "fascinating. The speakers were succinct and eloquent." The excitement spills over into the classroom and, according to one teacher, "sparks debates and discussions both in the classroom and in my office. The kids are really *thinking* and I appreciate the opportunity to offer that to them." On another front there is the beginning of a network of community science teachers connecting local educators together and to SFI. The Institute, for instance, works with this local network directing teachers to sophisticated software packages not readily available on the educational market.

Said Mike Simmons, SFI's Vice President for Academic Affairs, "The Institute is an international research center, but we're also very much a part of the Santa Fe community. We intend to continue exploring ways in which SFI can contribute to the local community and the state." During the next year, the project will consider several options for building more "hands-on" computational activities into the program such as Saturday-morning computer simulation workshops or the possible establishment of a computer network between SFI and community high schools.

1993 Computational Science Workshop

The High-Performance Computing Research Center (HPCRC) and the Computational Testbed for Industry at Los Alamos National Laboratory invite scientists and engineers to participate in the summer session of the 1993 Computational Science Workshop.

Participants will have access to high-performance workstations and innovative computer architectures including leading technologies from Thinking Machines Corporation; Cray Research, Inc.; Motorola; IBM; Silicon Graphics, Inc.; and Intel. Workshop attendees are offered three progressive levels of participation—a one-week overview, a three-week overview and tutorial, or a twelve-week extended workshop—all beginning June 14, 1993. The application deadline is February 1, 1993.

For a complete description of the Workshop and application materials, please contact:

Marilyn Foster or David G. Simmons
Computational Science Workshop
Los Alamos National Laboratory, MS M986
Los Alamos, NM 87545
Email: CSI93@c3.lanl.gov
Voice: (505)665-5636; Fax: (505)665-5638

The Computational Science Workshop is a DOE High-Performance Computing and Communications Initiative.

Complex Systems Summer School, A Retrospective

Daniel L. Stein

The Sixth Annual Complex Systems Summer School runs May 30–June 25, 1993 in Santa Fe, New Mexico.

For information contact:
Andi Sutherland
Santa Fe Institute
1660 Old Pecos Trail
Suite A
Santa Fe, NM 87501
phone: (505)989-7372
fax: (505)982-0565
e-mail: summerschool@santafe.edu.

Application deadline is February 1, 1993.

The Santa Fe Institute has now hosted five Summer Schools on Complex Systems. Held during June for a period of four weeks, each Summer School brings together over sixty students and about two dozen faculty. Because both the students and faculty overlap among Summer Schools is small, the five-year-old program has now brought over three hundred graduate students and postdocs to Santa Fe and introduced them, in many cases for the first time, to different aspects of the sciences of complexity. Moreover, well over a hundred people doing active research in areas of mutual interest with the SFI have become acquainted (and in some case, affiliated) with the Institute. The Proceedings, which are published in book form within a year of each school, have spread the gospel much further and contribute to both the evolving definition and the direction of the field of complexity.

The original, and still primary, idea of the School is to provide a concrete manifestation of the notion that a new science of Complex Systems has emerged. While many of the topics can be found individually in some university courses or other summer schools, there is no other school which tries to bring together such an outlandishly varied assortment of problems, approaches, and subjects under a unifying idea. Even more ambitiously (and certainly uniquely), we bring together students from a wide-ranging array of scientific disciplines and attempt to find common grounds and languages. The difficulties and risks are not unlike those facing the builders of the Tower of Babel; the most surprising and unexpected result emerging from the schools is how well this has worked. But more on all that later.

The beginnings were not easy. There were no precedents. Nobody knew precisely what areas such a school should cover, or who should come. The funding agencies, having an institutional aversion to both interdisciplinary and fundamentally new research, were initially less than enthusiastic. To prove the scientific and administrative feasibility of such an undertaking seemed overwhelmingly difficult during that first year. However, thanks to the strong support of both the internal SFI administration and many of the external affiliated scientists, it did happen. More, it has succeeded beyond anything we dared hope in those early days.

Part of our success is due to the fact that the Summer Schools are a national effort, supported by a number of institutions, national labs, and funding agencies. Since the beginning the program has had strong support not only from SFI but also from CNLS at Los Alamos; the Universities of Arizona, Michigan, Illinois, Texas and New Mexico; and Sandia National Laboratories. We've also had financial support from the University of California system, Brandeis, Stanford, Columbia, Yale, Princeton, and the Universities of Pennsylvania and Florida. The Alfred P. Sloan Foundation and Research Corporation provided seed funding, and every year NSF, DOE, NIMH, and ONR provide significant funding.

Putting together a Summer School on Complex Systems presents an organizer with a set of problems and challenges that are simply absent when the subject is well defined and/or fits within a single scientific discipline. Should the themes and subjects be narrowly focused and strongly unified or broad-based and loosely knit? Should we aim for a definition of complexity? Can we take large risks and gamble on subjects or speakers whose connection to complexity is a matter of (sometimes strong) disagreement within the SFI community? Should we repeat bread-and-butter subjects year to year? Should we first assemble a list of topics and then choose speakers, or simply look for good speakers doing interesting things?

For the most part, our strategy has been not

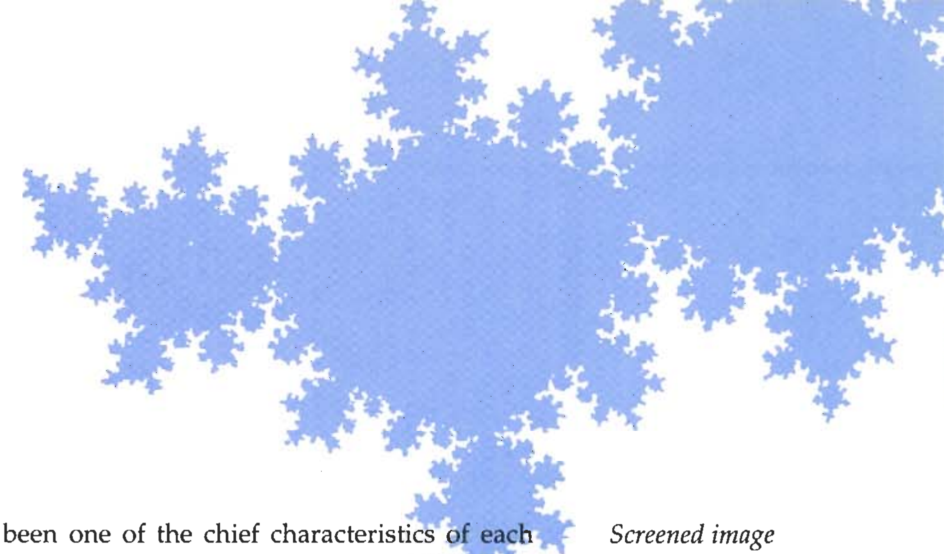

Screened image above is a filled Julia set from 1989 Lectures in Complex Systems, edited by Erica Jen; copyrighted © 1990 by Addison-Wesley Publishing Company. Reprinted by permission.

to have lectures that attempt to define complexity (though we've had a few), but rather to get the best people doing research in different well-defined areas to talk about their subjects and their work. That is, we try to present individual topics within their own frameworks and contexts, with their concepts, histories, and phenomenologies described with the necessary attention to detail that a heavily involved researcher can offer. Only after a system is described, and to whatever extent possible understood, do the lecturers attempt (and invite the students) to see connections to other systems or the overall field of complexity. We do take risks here, which sometimes pay off and sometimes don't. But without these risks, there is little excitement, and excitement has been one of the chief characteristics of each summer school. The students almost universally state that they feel they're part of some pathbreaking endeavor, whose destination is not yet in sight.

> "Rarely have I so enjoyed a scientific experience, and rarely have I learned so many new concepts in a month. As I slowly reimmerse myself in my thesis work, it is not difficult to see that the Summer School has wrought several significant changes in my thinking. I have many fresh ideas for my thesis, and will be using a genetic algorithm to solve a problem that has proven difficult for months. There are vast expanses of literature in nonlinear mathematics and chaos that are now much more accessible to me. Perhaps the most important changes are in my career plans, as I am seriously considering some of the fields to which I was introduced in the Summer School."
>
> Bennett Levitan,
> University of Pennsylvania,
> June 1992

The schools are open-ended; we don't know one year what topics will appear in the next year's school. We do try to maintain a balance among the various disciplines which are represented within the field of complexity; it shifts somewhat from year to year, but over the years most fields receive exposure. In some sense, the Summer Schools are a cumulative enterprise; their integrated history (as recorded by the proceedings volumes) provides a record of the path and new directions of the field of complexity itself.

We try to keep our eyes open both to interesting new subjects and to the best people within a promising field. It's not enough, though, to be a leading researcher in a given area; we try hard to find those people who are best able to communicate ideas to a diverse audience. This is a difficult undertaking, and we sometimes find that a person with a reputation as an excellent lecturer within his/her own field has difficulty communicating across disciplinary boundaries; because this is a new experience for most people, it's hard to anticipate when this will happen. Fortunately, while it does happen, it's relatively infrequent; and just as there is no one who has ever pleased everybody, there is also no one who has pleased absolutely nobody.

In choosing speakers, we also try to achieve a balance between the older "gurus" and the younger up-and-coming researchers. It seems that no matter how friendly and approachable a faculty member is, some students are naturally intimidated by the presence of a Great One (fortunately, most students have no such qualms). Part of the job of Director is to take some students by the hand and help break the ice with a faculty member in whom the

continued next page

students might naturally be interested (after which things usually proceed of their own accord). Moreover, I've seen that it's important for the students to see at close hand someone close to their own age, at a not much more advanced stage of their career, already doing important and interesting things that have achieved recognition. This helps many students become encouraged and excited about their own careers within the field of complexity.

The actual selection of topics and speakers begins as early as the previous year's summer school—many of the desired lecturers are often in high demand elsewhere, so beginning early is crucial. Certain themes repeatedly appear in successive summer schools, although often in different guises; these include chaos, pattern formation (ranging from convective instabilities in fluids to morphogenesis), computational neuroscience, quenched disorder, adaptive computation, cellular automata, evolution, stochastic processes, and many more. Because these general areas have a rich array of manifestations and applications, repetition is easily avoided. Other topics appear less frequently, or only once, so that each summer school maintains continuity with previous ones but retains its own distinct flavor.

While the structure of the school varies somewhat from year to year, the basic format is fairly constant. We typically have two lectures each morning, and frequently one seminar in the afternoon. In addition, early on the students (with a little help from the resident Director) organize their own seminar series. By the end of the Summer, a large fraction of the participants have given talks on their own research, or on a topic they find of interest. The introduction of the student lecture series has been extremely popular.

Indeed, the most gratifying (and amazing) aspect of each Summer School is how well the students interact, self-organize, exchange ideas, and form their own intact community. How this will occur is always completely unpredictable, but it always does. The resident Director nudges this along to the extent possible, by initiating some early organization of research groups, selecting out student coordinators for various tasks (such as organizing the student lecture series), and encouraging interaction between students (and also students with faculty), but the wisest course is to interfere as little as possible. The extreme enthusiasm which the students display as the summer school progresses is more due to this interaction than any other aspect of the School. Many report themselves to be amazed at how much they have in common with people from

"The interdisciplinary nature of the school really gives it a unique flavor and importance. Interdisciplinary studies are essential in today's industrial research environment—we are asked to do so many things that cross the boundaries between disciplines and involve the interactions of more than one discipline. Interdisciplinary work is hard to do: language barriers have to be crossed and we have to be careful about reinventing the wheel. The format of the school and the attitude of the people at the Santa Fe Institute go a long way to avoid the pitfalls of interdisciplinary work and I expect great things to come from the Institute and the students that pass through the school. You are doing important and highly relevant work."

Derek Smith,
Integrated Systems Laboratory,
July 1992

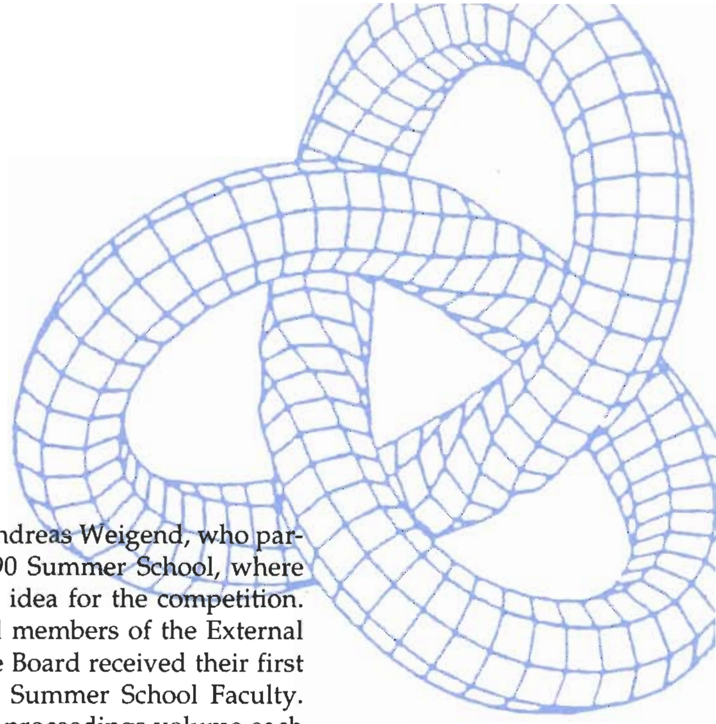

Dan Stein

entirely different disciplines, and consider the opportunity to live and work with very smart and interesting people of many scientific backgrounds as the most exciting (and unique) feature of the School. Many also report that this is one of the most profound impressions they're left with afterwards.

This interaction takes on different forms each year, and provides each Summer School with its own character. I feel that it is by far the most important part of the School—more than anything, it sets us apart from other (single-discipline) Summer Schools. In some of our Schools, the atmosphere has taken on an almost Woodstock-like atmosphere (in spirit, not explicit form!).

The benefits of the School for SFI are far-reaching and often intangible; above all else, it introduces a broad spectrum of some of the brightest graduate students and postdocs to the field of complexity, and also to the Santa Fe Institute. We foster strong interaction with SFI. Students are encouraged to attend SFI seminars, workshops, and public lectures, and many do; every year some Summer School faculty are chosen from the SFI postdocs, External Faculty, or Science Board members; and certain sessions are set aside for the students to learn more about research activities at the SFI. This last is an innovation begun at the 1991 School.

A number of our alumni have gone on to become formally affiliated with SFI, and have conducted research there. One direct outgrowth of the Summer School was the Time Series Analysis Competition, initiated by Neil Gershenfeld and Andreas Weigend, who participated in the 1990 Summer School, where they originated the idea for the competition. In addition, several members of the External Faculty and Science Board received their first exposure to SFI as Summer School Faculty. And, of course, the proceedings volume each year not only gives a sense of where the field has been but also where it is going.

And what about the Summer School itself? Our funding from external agencies, universities, and laboratories is now as stable as anything can be in these unstable times. We have produced a very wide network of converts who are spreading the word about complexity and the SFI. Thanks to the suggestions and help from many people affiliated with the SFI, we have a list of excellent potential speakers which won't be exhausted for several years, even if we restrict ourselves solely to this list (which we won't). We continue to make changes (hopefully improvements), identifying problem areas and incorporating suggestions from former students and from other interested parties.

But perhaps most of all, if forced to put my finger on what is responsible for the sense of excitement that pervades each school and affects all those involved with it, that is somehow unique to this enterprise, it's that sense of being present at the creation.

Image is a tubeplot, courtesy of Art Winfree whose chapter will appear in 1992 Lectures in Complex Systems, edited by L. Nadel and D. L. Stein. Preprinted by permission of the publisher.

Postdoc Report

Korber Joins Project ARIEL

Bette Korber from the Theoretical Biology and Biophysics Group at Los Alamos National Lab has joined the SFI research staff on a half-time basis to work on Project ARIEL, a program supported by the Pediatric AIDS Foundation and the Magic Johnson Foundation. Project ARIEL brings together the expertise of several of the top AIDS laboratories in the country to explore many aspects of mother-infant transmission of HIV-1, the virus that causes AIDS.

Not all pregnant HIV-1 infected women transmit the virus to their offspring; epidemiological studies indicate that there is a transmission rate of 15–30%. Many elements of this transmission are still a mystery. Why does transmission occur in some, but not all, women? When transmission occurs, is it predominantly happening early during pregnancy, across the placenta, or is it happening during birth? Does the mother's immune system play a role in preventing transmission? Do viruses that are transmitted to babies have common characteristics? Understanding the answers to these questions may give insight into methods that could help reduce the risk of transmission.

What is unique about Project ARIEL is that several laboratories will work together, using the same mother-infant blood samples. Different tests will be conducted in different laboratories, each with different expertise (immunology, viral biology, viral sequencing, etc.). The effectiveness of the mother's immune response against the virus obtained from her own blood samples, as well as virus isolated from her baby, will be assessed. The viral load in the mothers will be estimated, and the biological properties of the viral isolates will be characterized. Because HIV-1 is highly variable—it varies even within a person—the evolution of the virus in the mothers will be studied throughout gestation and will also be followed in infected infants after birth.

Bette Korber

(Korber's related work at Los Alamos National Laboratory is to study the extent of variation of HIV. The specific goals are to contribute to the understanding of viral evolution on a global scale, and to design optimal HIV vaccine components.)

All of the ARIEL data—along with patient profiles describing the health status of mother-infant pairs—will be brought together by Korber at the Santa Fe Institute. Her residency is sponsored by SFI External Faculty member Alan Lapedes. She'll look for meaningful patterns in the DNA sequence data, and correlations between the different kinds of data sets. "The hope is that the understanding that develops through the course of this study will point the way to designing effective means of intervention," says Korber.

SFI Postdoc Positions

Candidates for postdoctoral positions should submit a curriculum vitae, list of publications, and statement of research interests, and arrange for three letters of recommendation.

Send applications to: Postdoctoral Committee, Santa Fe Institute, 1660 Old Pecos Trail, Suite A, Santa Fe, NM 87501. Applications or inquiries may also be sent by electronic mail to: postdoc@sfi.santafe.edu

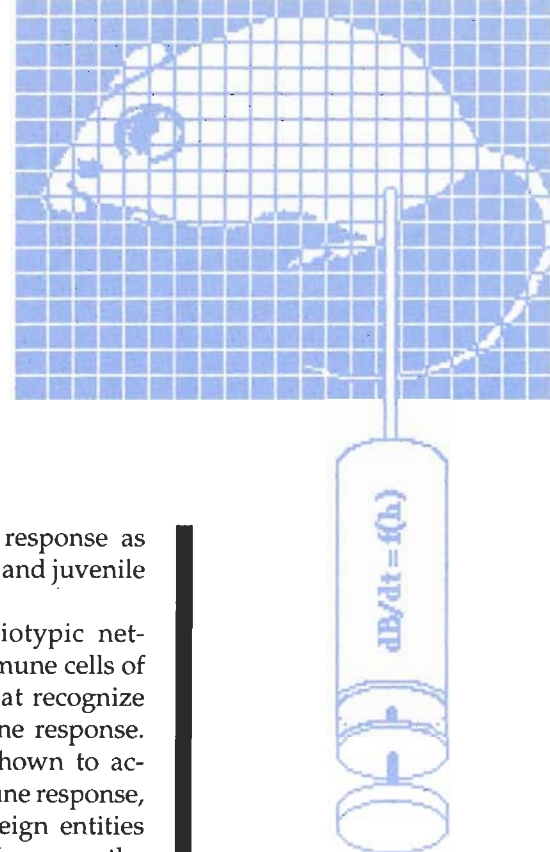

Neumann Named SFI Sullivan Scholar

Avidan Neumann of the Weizman Institute became the Sullivan Scholar at SFI this fall. He will conduct research in theoretical immunology, concentrating on autoimmune diseases such as systemic lupus erythematosus (SLE) and multiple sclerosis (MS). His work in collaboration with SFI External Professor Alan Perelson is supported by a two-year grant from the Joseph P. Sullivan and Jeanne M. Sullivan Foundation.

The immune system is responsible for our day-to-day defense against different viruses and bacteria that invade our body. Several mechanisms serve to respond to and eliminate these invaders and, moreover, to memorize them so as to obtain a more efficient response in the future. It is, of course, essential for the immune system to be able to discriminate between self cells and foreign invaders. Failure to do so triggers an autoimmune response as seen in diseases such as MS, SLE, and juvenile diabetes.

Avidan Neumann with his son Ellan.

Jerne postulated that the idiotypic network—a network connecting immune cells of different shapes, or idiotypes, that recognize each other—regulates the immune response. Idiotypic networks have been shown to account for the memory of the immune response, either in vaccination against foreign entities or tolerance toward self cells. However, the body's discrimination between self and nonself is not yet fully explained, and many aspects of the immune response as modeled by the idiotypic network are still problematic.

For example, experimental results recently obtained for SLE disease show that the idiotypic network has a functional role in the development of this autoimmune disease. Using the idiotypic network model, Neumann is studying the theoretical basis for the development of SLE and other related diseases. In a related effort he is working on constructing schemes for experiments that will give better data on the structure of the immune network, and shed light on the functional role of the idiotypic network in normal immune response.

Postdoc Report

Macready Explores Complexity in Adaptive Systems

Screened image is a K=1 Boolean net in the ordered regime.

Titles of papers in the new science of complexity often include the phrase "complex adaptive systems," and yet the relationship between complexity and adaptability remains an open question—one which is the focus of postdoc Bill Macready's research. This fall Macready, the recipient of a Canadian National Sciences and Engineering Research Council Postdoctoral Fellowship, joined the Artificial Life program lead by External Faculty member Chris Langton.

The first step toward characterizing the relationship between complexity and adaptability is to find an appropriate measure of complexity. We do not think of orderly patterns as being complex because we can describe them with very little information. Disordered or random patterns we might think of as being complex, but actually their statistical properties can also be described with limited information. It is only the patterns at the "phase transition" between order and disorder that are hard to describe, and it is these we call complex.

Macready is exploring the adaptability inherent in systems that are "living at the phase transition," as well as how adaptation can drive a system to this complex region. One definition of adaptability is the ability of a system to respond to a variety of constraints with minimal changes in its structure. Using such a definition, it is not entirely clear that adaptable systems are necessarily complex ones. If we ask a system to adapt to a particular set of immutable constraints, we do not necessarily expect the system to become complex. In fact, we might expect the system to lower its complexity and solve the constraints in the simplest possible way. Thus we might expect the complexity of a system will only increase if we ask the system to adapt to a changing environment.

It is Macready's aim to prove, by direct experimentation, independent of any particular model, whether or not complexity is an inherent characteristic of adaptive dynamical systems. Specific instances of dynamical systems are ubiquitous in nature, and dynamical models of natural systems include cellular automata, neural networks, and boolean network models of molecular evolution. A generic dynamical system is obtained by referring only to the state transition graphs without attaching any meaning to the states themselves.

These generic systems are then asked to evolve in order to solve a constantly changing task (again defined without reference to a particular model). If complexity is a necessary feature of adaptive systems, then we expect to evolve systems in the complex or "phase transition" region between order and disorder. However, if these expectations are not borne out, it will call into question a fundamental assumption about adaptive systems.

In short, is there an unwritten law in nature which states that adaptive systems are necessarily complex? Stay tuned.

Theiler Investigates Chaotic EEG

James Theiler, formerly with the Complex Systems Group at Los Alamos National Lab, is at SFI as a Postdoctoral Fellow through 1993. The National Institute of Mental Health has awarded SFI External Professor Doyne Farmer and Theiler a grant to investigate allegations that the electro-encephalogram (EEG) is chaotic, and in general to explore the use of nonlinear time series methods in the characterization of the dynamical behavior of the nervous system.

Neurons are highly nonlinear. It remains unclear, however, exactly what role nonlinearity plays in their information processing capabilities, and how this is expressed in the behavior of systems of neurons.

Chaos theory offers an exciting new possibility for explaining apparently random behavior in the nervous system. It has been suggested as the underlying cause of randomness in several different neural systems. However, in most cases the evidence for chaos remains inconclusive, in large part because the data analysis is based on techniques that are notoriously unreliable, such as currently popular algorithms for computing fractal dimension.

On the other hand, there are many mathematical models of the nervous system where chaos occurs. Since these are well-defined mathematical models, it is possible to analyze them and determine unambiguously the existence and properties of any underlying chaos. In general, though, in neuroscience there is a large gap between theoretical models and the real nervous system; there are assumptions and unknown parameters in the models that leave their relevance to real neurophysical phenomena uncertain.

One such system is the electro-encephalogram (EEG), which has become a widely used tool for the monitoring of electrical brain activity, and whose potential for diagnosis is still being explored. Over the last decade, there have been many published claims that various EEG time series exhibit evidence of low-dimensional chaos. If true, the implications are striking. However, the interpretation of this evidence has been controversial. One problem is that the algorithms for characterizing chaotic time series were originally developed in the context of large, relatively noise-free data sets, whereas EEG time series are often short and noisy. The algorithms can in principle distinguish chaos from noise, but in practice they often fail.

James Theiler

While low-dimensional chaos may be exhibited in free-running oscillations, a more common situation in neuroscience is for a response to depend on a stimulus. Theiler and his colleagues have developed methods for analyzing the possible nonlinear dependence of the response to the stimulus. The method distinguishes between chaotic and nonchaotic responses. They have recently acquired a considerable data base of evoked-response EEG time series, and intend to apply these methods to determine if there is a nonlinear relationship. "Our ultimate purpose," says Theiler, "is to discover any underlying deterministic structure that may currently lie hidden in apparently random neural phenomena."

Research Visitors May–December, 1992

Harald Altmanspacher, Max Planck Institute, Garching
Charles Anderson, Jet Propulsion Laboratory
Brian Arthur, Stanford University
Joseph Atick, Institute for Advanced Study
Per Bak, Brookhaven National Laboratory
Wyeth Baer, California Institute of Technology
Nils Baas, University of Trondheim
Denis Baylor, Stanford University Medical School
Aviv Bergman, Interval Research Corporation
William Bialek, NEC Research Institute
Carlton Caves, University of New Mexico
Françoise Chatelin, University of Paris IX/ Thomson–CSF
Eric Chopin, Ecole Normale Supérieure de Lyons
Hirsh Cohen, Sloan Foundation
Jim Crutchfield, University of California, Berkeley
Rob de Boer, University of Utrecht
Gary de Young, University of California, Davis
Michael Diehl, State University of New York at Buffalo
Murray Gell-Mann, California Institute of Technology
Charles Gilbert, Rockefeller University
Brian Goodwin, Open University
Jim Hanson, University of California, Berkeley
Peter Heilbrun, University of Utah
David Hiebeler, Thinking Machines Corporation
Stefan Helmreich, Stanford University
Brant Hinrichs, University of Illinois, Urbana-Champaign
Ivo Hofacker, University of Vienna
John Holland, University of Michigan
Alfred Hubler, University of Illinois
Terry Jones, University of Indiana
Stuart Kauffman, University of Pennsylvania
Joel Keizer, University of California, Davis
Christof Koch, California Institute of Technology
Tim Kohler, Washington State University
Blake LeBaron, University of Wisconsin
Michael Lewicki, California Institute of Technology
Ralph Lewis, Dartmouth University
Zhaoping Li, Institute for Advanced Study
André Longtin, University of Ottawa
Carlo Lucheroni, Universita Delgi Studi di Perugia
David Mackay, Cambridge University
Bill Macready, Canada
Günter Mahler, University of Stüttgart
Gottfried Mayer-Kress, University of Illinois, Urbana-Champaign
Marcus Meister, Harvard University
Birgit Merté, Technical University of Munich
John Miller, Carnegie-Mellon University
Kenneth Miller, California Institute of Technology
Melanie Mitchell, University of Michigan
Harold Morowitz, George Mason University
Avidan Neumann, Weizmann Institute
Steve Omohundro, University of California, Berkeley
Richard Palmer, Duke University
Klaus Pawelzik, University of Frankfurt
David Pines, University of Illinois, Urbana-Champaign
Dan Pirone, University of Washington
Jordan Pollock, Ohio State University
Rama Ranganathan, University of California, San Diego
Tom Ray, University of Delaware
R. Clay Ried, Rockefeller University
Neantro Saavedra-Rivano, University of Tsukuba
H.G. Schuster, University of Kiel
Peter Schuster, Institut für Molekulare Biotechnologie, Jena
Terrence Sejnowsi, The Salk Institute
Martin Shubik, Yale University
Joshua Smith, Cambridge University
Matthew Sobel, State University of New York at Stony Brook
Peter Stadler, University of Vienna
Chris Stephenson, IBM Research, Yorktown Heights
Charles Stevens, The Salk Institute
Michael Stryker, University of California Medical School
William Sudderth, University of Minnesota
Walter Tackett, Hughes Aircraft/University of Southern California
Kurt Thearling, Thinking Machines Corporation
Guy Theraulaz, CNRS, Marseille
David Van Essen, California Institute of Technology
Udi Zohary, Stanford University Medical School

Board News

SFI welcomes three new members to its Board of Trustees.

David E. Liddle is President and CEO of Interval Research Corporation. Liddle has spent his professional career focusing on client/server architecture and on user interfaces designed for business professionals. In March of this year Liddle, with Paul Allen, co-founded IRC to perform research and advanced development in the areas of information systems, communications, and computer science. Liddle was founder, president, and CEO of Metaphor Computer Systems, Inc. and president and CEO of Patriot Partners, a joint venture between IBM and Metaphor. He later became vice president, New Systems Business Development, Personal Systems, at IBM Corporation, following IBM's acquisition of Metaphor in 1991.

Prior to founding Metaphor in 1982, Liddle spent ten years at Xerox Corporation in Palo Alto, initially as a research scientist, and ultimately as vice president and general manager of the Office Systems Division. He was responsible for the planning, development, and launch of Xerox's pioneering client/server architecture, including the Ethernet LAN and the Xerox Star—the first commercial implementation of the Graphic User Interface.

Daniel C. Lynch is Founder and President of Interop Company, specialists in computer networking education. Prior to founding Interop in the mid-1980s, Lynch managed the computer facility at Information Sciences Institute (ISI) in Los Angeles and, before that, had a similar role at the Artificial Intelligence Center at Stanford Research Institute in Menlo Park.

Lynch is a member of the ACM, IEEE, and IAB. *Communications Week* named him one of the most promising industry leaders of the 1990s and honored him as one of the top 25 Visionaries of 1991. "My excitement for the Santa Fe Institute comes from my belief that complexity science has a wide range of real-world applications," says Lynch. "SFI is creating the conceptual models for the next wave of human understanding."

Robert R. Maxfield was a co-founder of ROLM Corporation, a manufacturer of computer and telecommunications equipment, and served as Executive Vice President and Director until its acquisition by IBM. He is President of the Maxfield Foundation; a Consulting Professor of Electrical Engineering and Engineering–Economic Systems at Stanford University; and serves on the Board of Directors of Software Publishing Corporation and Echelon Corporation.

"I am very impressed by the quality of the people involved with SFI," says Maxfield. "I believe that the theories they are developing for complex systems, coupled with advances in computer technology, will yield valuable insights for some of the most important problems of our time."

David E. Liddle

Daniel C. Lynch

Robert R. Maxfield

1992 Contributors

Academic Institutions
George Mason University
National Bureau of Economic Research, Inc.
NEC Research Institute, Inc.
St. Johns College
Society for Mathematical Biology
University of Arizona
University of Florida Foundation
University of Illinois
University of New Mexico

Corporations
Adager Corporation
Apple Computer
Batterymarch Financial Management
Campbell, Carr, Berge & Sheridan, P.A.
Citibank, N.A.
Computerland of Santa Fe
Deere & Company
Digital Equipment Corporation
Electronic Films
Harcourt Brace Jovanovich, Ltd.
Interval Research Corporation
Maxis
Microsoft Corporation
Ox-Bow Press
Pergamon Press Ltd.
Pittsburgh National Bank (Alex C. Walker Foundation)
Presidio Management Group
Public Affairs Television
Rodey, Dickason, Sloan, Akin & Robb, P.A.
Santa Fe Audio Visual
Shuttlejack
Silicon Graphics
Sun Microsystems, Inc.
Sunwest Bank of Santa Fe
Trinity Capital Corporation
Volvo North American
Westwind Travel

Foundations
Fondazione Eni Enrico Mattei
Richard Lounsbery Foundation
John D. and Catherine T. MacArthur Foundation
Maxfield Foundation
Carol O'Donnell Foundation
Pediatric Aids Fdn.-The Ariel Project
Pew Charitable Trusts
Pinewood Foundation
Alfred P. Sloan Foundation

Government Agencies
Los Alamos National Laboratory
The NATO International Science Exchange Programs
National Institutes of Health
National Institute of Mental Health
National Science Foundation
U.S. Department of Energy
U.S. Forest Service
U.S. Office of Naval Research

Individuals
Bruce A. & Nancy S. Abell, Santa Fe, NM
Ruth & Robert McCormick Adams, Washington, DC
Lewis E. Agnew, Jr. & Margaret Agnew, Los Alamos, NM
Margaret Alexander, Santa Fe, NM
Jean Anderson, Los Alamos, NM
Philip W. & Joyce G. Anderson, Hopewell, NJ
Anonymous (in recognition of Marcia Southwick Gell-Mann)
Jarratt Applewhite, Santa Fe, NM
Marcella D. Austin, Española, NM
Sam Ballen, Santa Fe, NM
Armand & Celeste Bartos, New York, NY
George I. & Virginia L. Bell, Los Alamos, NM
Timothy K. Bliss, Santa Barbara, CA
Erich & Renee Bloch, Washington, DC
Stewart B. Brand, Sausalito, CA
Richard Brandt, Westport, CT (in honor of Murray Gell-Mann)
Per Bro, Santa Fe, NM
Joanne B. Bross, McLean, VA
Ashley S. Campbell, Santa Fe, NM
Mr. & Mrs. J. W. Campbell, Pampa, TX (in memory of Bruce Walters)
John L. Casti, Santa Fe, NM
Albert M. & Molly W. Clogston, Tesuque, NM
Stirling A. & Rosemary Colgate, Los Alamos, NM
Thomas J. Cook, Los Alamos, NM
Jacqueline B. Cotsen, Los Angeles, CA
George A. and Helen Cowan, Los Alamos, NM
Robert W. Craig, Keystone, CO
D. Ronald Daniel, New York, NY
Nicholas & Margaret DeWolf, Aspen, CO
Carl Djerassi, San Francisco, CA
James B. Downey, Santa Fe, NM
John Downing, Leucadia, CA
Esther Dyson, New York, NY
Carole Ely, Santa Fe, NM
J. Doyne Farmer, Santa Fe, NM
Marcus W. & Shirley Feldman, Stanford, CA
Wolfgang & Irmgard Fikentscher, Starnberg, W. Germany
J. D. & Sharon Geist, Albuquerque, NM
Murray Gell-Mann, Pasadena, CA
Garrett Gruener, Berkeley, CA

Donald C. & Margaret J. Hagerman, Los Alamos, NM
Michael J. & Ann Tukey Harrison, East Lansing, MI
Lawrence S. Huntington, New York, NY
Mitch Kapor, Cambridge, MA
Stuart & Elizabeth Kauffman, Santa Fe, NM
William E. Keller, Santa Fe, NM
George A. Keyworth, II, Santa Fe, NM
Richard H. & Kathryn V. Kimball, San Francisco, CA
Edward A. & Jean Knapp, Santa Fe, NM
Jean-Claude Kommer, London, U.K.
Nancy Kopell, Cambridge, MA
George A. Kozmetsky, Austin, TX
Donald Robert Kummer, St. Louis, MO
Simon A. & Carole L. Levin, Princeton, NJ
David E. Liddle, Portola Valley, CA
Daniel C. & Karen D. Lynch, Mountain View, CA
Donald L. Maggin, New York, NY
Robert R. Maxfield, Saratoga, CA
Christine Malina Maxwell, Berkeley, CA
Charles Miller & Beth Jtwros, Houston, TX
Richard D. Moore, Warren, VT
Harold J. Morowitz, Fairfax, VA
Darragh E. & Avery Nagle, Los Alamos, NM
Ann Richards Nitze, Washington, DC
Barbara Norris, Aspen, CO
David Padwa, Santa Fe, NM
Jim L. & Dorothea M. Pelkey, Santa Fe, NM
Alan & Janet Perelson, Santa Fe, NM
Thomas F. Pick, Chicago, IL
David & Suzy Pines, Urbana, IL
John G. & Kimiko Powers, Carbondale, CO
Leland S. & Vivian B. Prussia, Piedmont, CA
David Z. Robinson, New York, NY
Ray D. Sena, Santa Fe, NM
Luigi Sertorio, Brussels, Belgium
Martin & Julia Shubik, New Haven, CT
Leonard M. & Margaret L. Simmons, Santa Fe, NM
Arthur & Libby Spiegel, Albuquerque, NM
J. I. Staley, Wichita Falls, TX
Daniel L. & Bernadette L. Stein, Tucson, AZ
Sidney N. and Marcia M. Stone, Albuquerque, NM
Asis Kumar Sarkar, Philadelphia, PA
Joseph P. & Jeanne M. Sullivan, Chicago, IL
Martin & Gail Tierney, Santa Fe, NM
Charles Tyng, Santa Fe, NM
Philip F. Voegelin, Santa Fe, NM
William Wagner, Los Angeles, CA
Thomas W. Weisel, San Francisco, CA
Bob Wickham, Santa Fe, NM
Henry & Mary C. Wider, Albuquerque, NM
Susan Wider, Albuquerque, NM

SFI Business Network

In the summer of 1992 SFI inaugurated its Business Network for Complex Systems Research as a mechanism for companies to affiliate with SFI and to help support research. In return for an annual fee of $25,000, a company may send people to attend workshops and symposiums; it routinely receives SFI Working Papers, books, and other publications; and it has opportunities to meet with and explore problems of mutual interest with the hundreds of scientists who are associated with SFI in the course of a year.

The Network was established in response to the growing interest in SFI research by companies. This has been especially evident in two areas: finance and computation. SFI's work in economics and in time-series data attracts the attention of banks and investment managers. And the work in computation has obvious appeal for computer companies—both hardware and software. Companies have been especially attracted by the opportunity to stay abreast of novel and wide-ranging research that is virtually unbounded by constraints of disciplines or departments. For them the association with SFI offers a window on a unique realm of research—that of complex adaptive systems. It also gives them a chance to become part of the many ongoing research networks that are centered at SFI.

Early members of the Business Network include Citicorp, Digital Equipment Corporation, Interval Research, John Deere Company, Batterymarch Financial Corporation, and Maxis. SFI's goal is to bring the membership to 20 companies by the end of 1993.

Anyone wanting further information can contact Bruce Abell or Mike Simmons at SFI.

SFI Officers & Staff

SANTA FE INSTITUTE
1660 Old Pecos Trail, Suite A
Santa Fe, NM 87501

Address Correction Requested

NON-PROFIT
ORGANIZATION
U.S. POSTAGE
PAID
PERMIT NO. 415
Santa Fe, NM