THE BULLETIN OF THE SANTA FE INSTITUTE

WINTER 1994 · VOLUME 9 · NUMBER 2

BULLETIN

SFI's FIRST DECADE

For those of us who were present at the creation of the Santa Fe Institute ten years ago, it seems both incredible that the decade has passed so quickly and remarkable that we have met and exceeded our expectations in so short a time.

Our founding group included the late Herbert Anderson, Peter Carruthers, George Cowan, Stirling Colgate, Darragh Nagle, Nicholas Metropolis, Louis Rosen, and myself, all Senior Fellows at Los Alamos National Laboratory. The discussion quickly left the boundaries of the lab. Joining in were mathematician Gian-Carlo Rota from M.I.T., David Pines from the University of Illinois, and chemist Anthony Turkevich from University of Chicago. Murray Gell-Mann heard about the nascent organization and immediately took an active role in its formation. Richard Slanksy, who was a LANL Fellow at the time, joined the group as did John Rubel, a former deputy director for research in the Department of Defense, who provided much practical advice. Alwyn C. Scott, head of the Center for Nonlinear Studies at Los Alamos and later his successor, David Campbell, made contributions to research planning. While everyone was supportive of the notion of a new kind of campus for scientists, of course our visions differed. And, like any good story, each founding member has his variation of how the Institute began. Some are told in this issue of the *Bulletin*.

We set out to create a new kind of research environment, a truly bottom-up culture, and an independent haven for multidisciplinary research. It would have been far easier to create an institute under the umbrella of an existing university or large research laboratory; whether by blind luck, circumstances, or some kind of inchoate wisdom, the founders set SFI out on its own. George Cowan was the first President, and I was the first Vice-Chairman of the Board of Trustees. Mike Simmons, who had spent many years in the management of science at LANL, joined SFI as Vice President in 1986. Having survived the pains of start-up, with very little activity at first and lean finances thereafter, SFI emerged with a degree of independence and a distraction-free operation that made it possible to launch truly new, if not revolutionary, approaches to science. We were fortunate to receive some key support in those early years from the John D. and Catherine T. MacArthur Foundation, from the National Science Foundation, from the Department of Energy, and from Citibank. Those sponsors, who have been with us ever since, gave us great leeway in building a program that didn't fit traditional categories.

Circumstances also favored SFI's tenth anniversary year by placing us in a permanent home. We joke that our early home was a desk drawer in Trustee Art Spiegel's office in Albuquerque, and we progressed to a Santa Fe post office box before we had a place of substance. Our first real home, the small Cristo Rey Convent, set a pattern for research interactions that continues to this day. When we outgrew that space we moved into less agreeable quarters on Old Pecos Trail, always mindful of having traded a kind of intimacy at Cristo Rey for some elbow room. Happily, now that we are ensconced in our home on Hyde Park Road, many of us feel we have recaptured some of the interactive atmosphere of the convent.

Having a home of our own also demonstrates to all that we are here for the long run. Research funding has continued to grow and diversify each year, along with requests for research visits. At the same time we see the multi-disciplinary approach and the growing interest in complexity and adaptation of Santa Fe diffusing to more and more research institutions around the world. We founders hoped that SFI would have an impact; it has, and the impact continues to grow.

Finally, I recently told the SFI family that after four exciting years as President of the Institute, I intend to step down in March 1995 so that I may indulge my own research interests at SFI, something I have reluctantly had to suppress during these years. I am certain that the next President will enjoy a tenure as filled with scientific excitement and impact as mine has been. I want to thank the whole SFI family —scientists, trustees, and a remarkably dedicated staff— for their creativity, support, and friendship during that time. One could not be in finer company.

Edward A. Knapp

SFI's FIRST DECADE

NEWS AND FEATURES

RESEARCH UPDATES

10 FIRST DECADE

Complexity Made Simple: John Holland Gives the First Annual Stanislaw Ulam Memorial Lecture


Photo: Murrae Haynes

In the first of three consecutive public lectures to explain complexity, John Holland turned to his audience and said simply, "Think about New York City."

New York is a large city with no central planning, Holland went on, yet everyone expects to buy food whenever he wants it. Given that the city possesses only a week's supply of food at any time, how does it happen that people always get what they want? And how does the system, with so many individuals ordering goods independently, work as a whole? Compared to twenty years ago, he said, many of the people, government officials, and buildings have changed; yet the city runs along smoothly.

The fact that New York City works "just in time" without central authority is an example of a complex adaptive system. And with this illustration, Holland launched into a detailed description of his research using a delightful array of analogies, metaphors, and commonsense thinking.

"We chose John Holland to give the first Stan Ulam Memorial lectures because he is recognized as the leader in genetic algorithms, a dominant research agenda at the Santa Fe Institute," said Jack Repcheck, a former editor at Addison-Wesley Publishing Company in Reading, Massachusetts and creator of the new lecture series. The idea behind the series is to have a brilliant scientist deliver a series of public talks on a cutting-edge topic, said Mr. Repcheck, who is now at Princeton University Press. Many of the most famous books in science, including *Relativity* by Albert Einstein and *QED* by Richard Feynman, were based on public lectures, he said. Next spring Addison- Wesley will publish a book which expands on Holland's talks, "Complexity Made Simple." The book's tentative title is *Hidden Order: How Adaptation Builds Complexity.*

The lectures were dedicated to the memory of Prof. Stanislaw M. Ulam, a renowned mathematician long associated with Los Alamos National Laboratory who was highly regarded by the SFI scientific community. "He was a thinker's thinker," Mr. Repcheck said. SFI Vice President Mike Simmons said, "The enormous range of Ulam's scientific thought encompassed not only mathematics but also physics, computation, biology, and much else. He would have been very much at home in the present day Santa Fe Institute, which was founded in the year of his death."

In introducing Dr. Holland, SFI economist Brian Arthur described the sixty-five-year-old scientist as "the Enrico Fermi of Computing" and the inventor of genetic algorithms; he is otherwise known as "the man who taught computers how to have sex." Trained at the Massachusetts Institute of Technology, Holland has spent his entire academic career at the University of Michigan in Ann Arbor. Today he holds appointments in two departments —computer science and engineering, and psychology.

Complex adaptive systems come in many shapes and forms, Holland began. Examples include economies, ecosystems, immune systems and nervous systems. Each has the ability to anticipate the future, learn and change in ways that are not well understood. They are diverse and highly innovative. Until more is known about the internal workings of such systems, he said, "we will not be able to solve the AIDS problem, understand trade balances, mimicry among plants and animals, human consciousness, or other phenomena produced by such systems."

"We are in search of a theory," Holland said. "Until we have that, we have no possibility

FIRST 10 DECADE

of understanding complex adaptive systems," which may, he explained, be more complicated than anything tackled in physics thus far.

To probe for a theory, he uses metaphors drawn from all walks of life. "The reason to roam intellectually is that sometimes what is obvious in one complex adaptive system may be less obvious but equally applicable to another system," he said. "At this point, we are making guesses, feeling our way. But most of us think a science is there."

Thus Holland is asking questions such as, How is a business firm like an antibody? How is a tropical tree with ten thousand different insect species like the human brain? How is evolution like learning?

To find answers, he looks for shared properties of complex adaptive systems. A basic feature is that all such systems are composed of adaptive agents —individual actors that give rise to the larger system. Examples include grocery buyers in New York City, antibodies, neurons, business firms, and single ants in an ant colony. Adaptive agents are interactive, and they learn. Moreover, they can be described by simple if/then rules —If such-and-such happens, then the agent carries out a predictable response.

The behavior of agents is a succession of if/then moves, Holland said, and reflects what the system is capable of doing from that point forward. An agent that adapts is able to change its rules in response to experience, a hallmark of complex adaptive systems.

Depending on the system, learning takes place on various time scales. The brain learns in seconds to hours, the immune system in hours to days, a business firm in months to years, a species in days to centuries, and ecosystems in centuries to millennia.

However long it takes, complex adaptive systems concentrate not on becoming the best in their category, but more simply on becoming better at what they do. To get better, he said, they need to discover or create new rules, a process that involves complicated tradeoffs. For example, if a business firm is good at building a product in an established production line, it may thrive. But if the business fails to explore new technologies or innovations, it may fail. On the other hand, if a business spends all its resources exploring new options, it may never knuckle down and build a successful production line, in which case it may fail. The tradeoff is between exploitation and exploration.

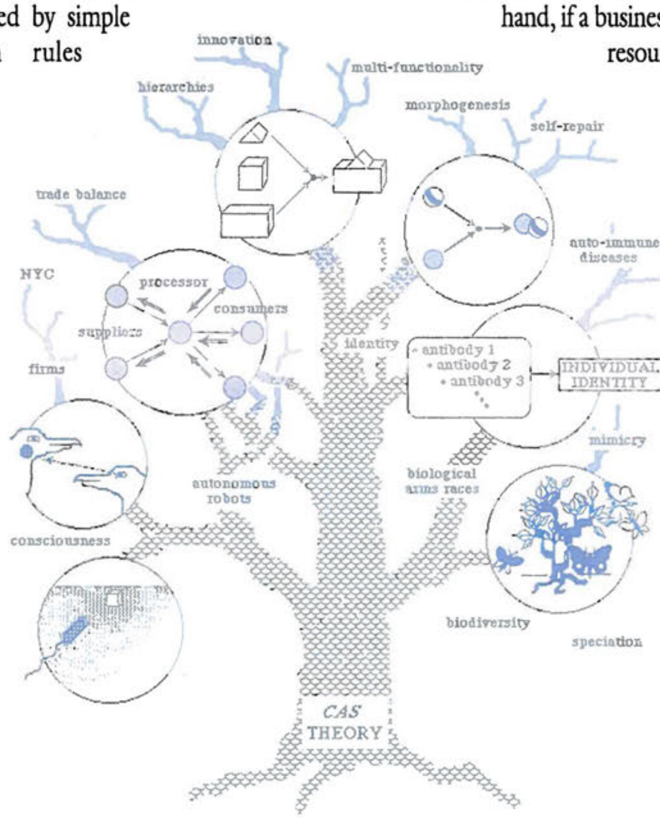

Figure 1. The Role of the Seven Basics in the Study of Complex Adaptive Systems

In this process, some new rules will help the system; but others might hurt it, Holland said. So how does a complex adaptive system go about inserting new rules to replace old rules that are not working well? How does it discover better rules? "You can't tinker at random because complex adaptive systems are so complex," Holland said. "If you tinker, you get garbage. You need plausible new rules."

Complex adaptive systems perform within the framework of collections of rules. While each rule tends to be simple —Holland used the example of a frog turning its head toward a moving object— behavior depends on the summation of many rules, which rules predominate in a situation, and which rules other rules attend to. He emphasized that all the rules in the system are not consistent with each other. For example, it is a good basic rule for frogs to turn their heads in the direction of a moving object if what is moving is an insect. That's food. But if the movement were from something big —like a human foot— the frog should flee, and not turn toward the movement. Layers of if/then rules interact to produce behavior, he said, producing a "performance system" —what the system is capable of at any given moment.

Every time a new rule is added, Holland said, the system needs to determine the rule's place and usefulness in the rule hierarchy. Generally, rules that have worked well in the past will have an advantage over novel or rarely-used rules, he said. To deal with rule hierarchy, complex adaptive systems tend to assign strength to individual rules. The more strength a rule has, the more likely it is to win the competition. Rarely-used rules may even disappear from the system.

Rules gain or lose strength by a complicated set of interactions called credit assignment. If a system makes a mistake, the rule that was followed loses strength. If a rule makes the system stronger or increases its fitness, it gains strength. A deep theoretical question is how to assign credit to a rule that occurs much earlier in a system's behavior, he said. This time Holland used an example involving the game of checkers. Say you make a move early in the game that later sets up a triple jump, he said. How do you get credit assigned back to that earlier rule which later led to the triple jump? This is a difficult problem. Rules used in the past may not get credit for outcomes or behavior that take place in the present.

To further explore rule behavior and evolution, Holland outlined seven basic features of complex adaptive systems.

"Aggregation" refers to the fact that agents get together and act as a group. An economy is an aggregation of firms. An immune system is an aggregation of cells that confer individuality.

"Nonlinearity" refers to the fact that the whole system is more than the sum of its parts. One can't predict the behavior of a complex adaptive system using conventional mathematics based on averages or trend analysis. Complex adaptive systems have trajectories, not end points, and as such they require a new kind of nonlinear mathematics.

"Flows" refer to the fact that resources in any system move from point to point. A fascinating and potentially very important aspect of complex adaptive systems is what Holland terms a multiplier effect —a small input within one part of the system can lead to a large output in

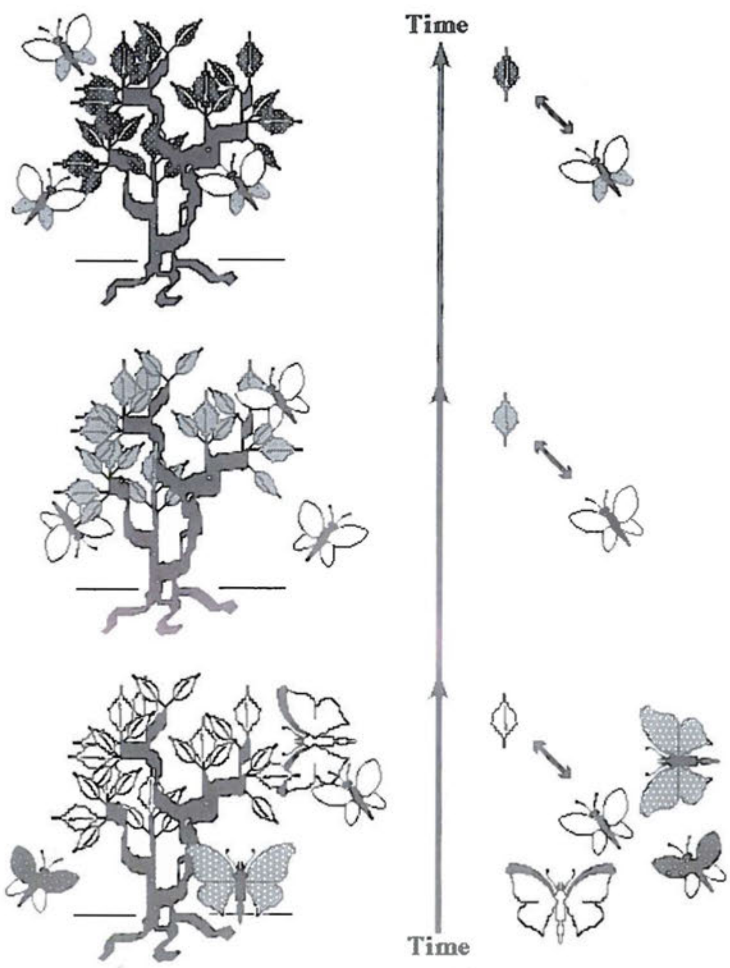

As time passes the plant evolves a succession of biochemicals [ ] that poison the butterfly larva, while the butterfly evolves enzymes [ ] that neutralize these biochemicals.

Figure 2. A Biological Arms Race

GEORGE COWAN

In the spring of 1987 George Cowan wrote, "The world the physical scientist likes most to study is vastly different from the world of our everyday experience. The world of the physical scientist is microscopic in scale. It wants to consist of no more than two particles at a time. The world of everyday experience is, in contrast, incredibly complex. There is no quantitative, general measure of complexity; but there is a hierarchy of complexity." Cowan bridges this contrast. He is a physical scientist who is vastly interested in the world of everyday experience and how science applies.

Cowan is often credited by his colleagues as the stimulus behind the Santa Fe Institute. He was the person who brought the original Los Alamos lunch group together. He served as President of the Institute from its inception until 1991. He continues to be involved in its programs. "Everybody has an original vision of the Institute," Cowan says. Essentially, Cowan's vision has been to apply the quantitative tools of the physical sciences, such as mathematics, to the social sciences. "Subjects that are of interest in the social sciences are typically nonlinear," he says. "So you can't maintain any confidence about the results."

FIRST 10 DECADE

A chemist by training, Cowan came to Los Alamos in 1949 after having worked on the wartime atom bomb project at the Metallurgical Laboratory at the University of Chicago. He is known for a number of things around Los Alamos, including being the founder of Los Alamos National Bank. Although he retired in 1988, he is currently a Senior Fellow Emeritus and Consultant to the labs. He and his wife have lived in the same house, once a government-owned duplex, for nearly all that time. As he talks about the history of the Institute, the house is getting a new roof.

"I think the principal success of the Institute has been in developing the kind of community where physical scientists talk to social scientists and exchange ideas," he says. And then he points out another key ingredient to the Institute's success: "People like to come here. People like to come back."

Cowan hopes to see the work of the Institute move more toward the behavioral sciences. "I think human behavior embraces everything we talk about in the social sciences, including economics," he says. "Human behavior is not a rational computing machine." Cowan sees the same thread of human behavior weaving through the political sciences and history, and he believes that by understanding the processes or "complexities" that lead to the demise of an ancient civilization, we might learn how to impede global warming. He has great ambitions for the Institute. "We can draw up a blueprint for an ideal world that will sustain itself forever," he says. "But the people who are alive one hundred years from now will do what they want. The best that we can hope for is that those people are simply better able to deal with complex problems than we are. I think that maybe the Santa Fe Institute will have something to do with that."

Françoise Ulam Recalls her Husband on the Occasion of the First Ulam Memorial Lecture

It is a great honor for me to be here tonight to celebrate the tenth anniversary of the creation of the Santa Fe Institute, and I also want to express all of my thanks and appreciation to its founder, my good friend George Cowan, and the leaders of the Institute, Ed Knapp and Mike Simmons and all the other people involved who have made this new series of lectures a memorial to my late husband.

For those of you who did not know him, let me just say a word or two about him. In a sense Stan was a sort of one-man Santa Fe Institute because of the diversity of his interests. But this was so long ago that the expression hadn't been coined yet. Were he alive today, he would love the Santa Fe Institute's unstructured informality, for he had very little use for the trappings of bureaucracy or authority. He loved to claim that the only committee he had ever served on was the Wine Tasting Committee of the Junior Fellows at Harvard. At Los Alamos, with the help of Carson Mark, the Theoretical Division leader, he confounded the Lab once by creating and circulating an official interoffice memo that listed numbers from one to one hundred in alphabetical order for quick and easy reference!

When he was promoted to Group Leader, he delighted in the fact that he was group leader of one, namely himself, for he was the only member of his group! So you see, Stan was a very playful man, and he never considered thinking to be work but rather play — like playing with ideas or inventing games; and he also took great delight in playing with words. The quite clever title of this series, "Complexity Made Simple," would please him very much, I think, for it's the kind of paradox he liked.

another part of the system. Such lever points can produce tremendous changes and may be applicable to social ills such as violent crime in cities, Holland notes.

"Diversity" is a hallmark of complex adaptive systems, which are always changing in response to new information flowing toward them. Nature provides countless examples of how animals, insects, and plants compete in endless arms races. For example, a caterpillar begins feeding on a plant which, to defend itself, learns how to synthesize a chemical that poisons the caterpillar; but then new generations of the caterpillar learn how to digest the chemical, and so the competition continues, each time producing more diversity in nature. Such competitions are often called Red Queen games, referring to the incident in *Alice in Wonderland* when the Red Queen said that you must run as fast as you can to stay in the same place. Mimicry arises from these battlefronts, Holland said. A well-known example is the Monarch butterfly, which feeds on milkweed, ingesting a substance that is disgusting to birds; Monarchs fly freely because they are not tasty to birds. The Viceroy butterfly mimics the Monarch's bright coloration and pattern and thus flies freely even though Viceroys are tasty to birds. But as with all mimics, Viceroys remain lower in number relative to Monarchs since, if birds discovered that a significant portion of the butterflies were tasty, they would eat them more often. Mimicry is found in all complex adaptive systems, including economies and the immune system.

Holland uses the term "Tag" to name emblems that identify individual agents. Examples include political banners, the surface molecules that antigens present to certain immune cells, and the singular spots beneath a seagull's eye which signal its identity to other seagulls. The natural world employs tags to break symmetries, Holland said, which can be described by another analogy. Imagine a billiard table full of cue balls. Since all are ivory colored, it would be difficult to keep track of any individual ball moving around on the table. But if one ball were painted red —breaking its symmetry with the others— its position could be followed easily. Complex adaptive systems employ this tactic to identify individuals, he said.

All complex adaptive systems use "building blocks" to evolve and change, Holland said. And new inventions and innovations usually involve assembling known building blocks in novel ways. For example, internal combustion engines employ Venturi's perfume spraying device (carburetor), Volta's sparking device (spark plug), the pistons of a mine pump, the gear wheels of a water mill, and so on. The history of science and technology is replete with examples of new objects assembled from available building blocks, he said, including automobiles, jet engines, and computers.

Finally, complex adaptive systems can build "internal models" of their world. They can run these models in virtual mode—sometimes called "lookahead"— to anticipate the future. Even something as simple as a bacterium anticipates the future when it swims up a chemical gradient. This ability to react on the basis of anticipations confounds the use of trends and polls to understand complex adaptive systems.

The lectures, held at the Greer Garson auditorium at the College of Santa Fe on June 6, 7 and 8 drew over four hundred people each night. Holland seemed surprised at the turnout, but it was the audience who went away impressed by the sheer breadth of the information presented. "John Holland has set the precedent with his broad, deep grasp of the subject of complexity and with his entertaining style for what should become a historical series of presentations" said SFI President Ed Knapp. "We look forward to living up to this legacy in future Ulam lectures."

Sandra Blakeslee is a journalist living in the Santa Fe area.

PROFILE

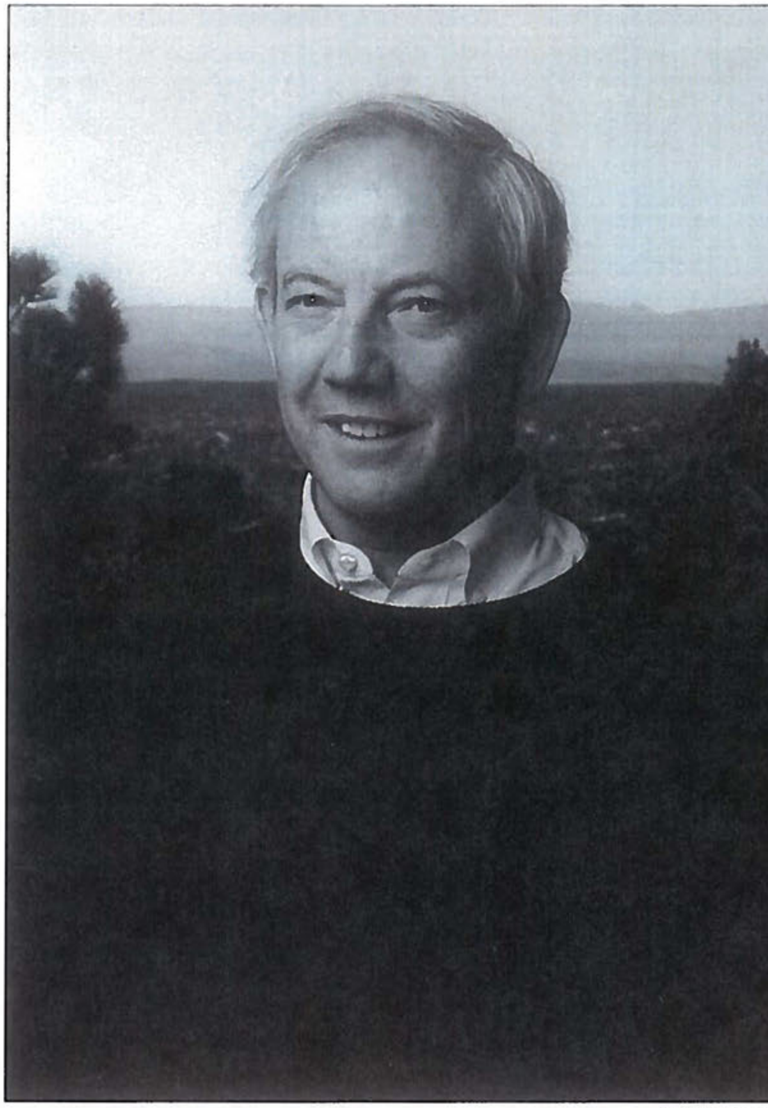

Photos: Dan Barsotti

W. BRIAN ARTHUR

Janet Stites

"The writer Peter Matthiessen once said, 'The secret of well-being is simplicity.' True. Yet the secret of evolution is the continual emergence of complexity. Simplicity brings a spareness, a grit; it cuts the fat. Yet complexity makes organisms like us possible in the first place. Complexity is indeed a marvel when it evolves naturally and delivers powerful performance. But when we seek it as an end or allow it to go unchecked, it merely hampers. It is then that we need to discover the new modes, the bold strokes, that bring fresh simplicity to our organizations, our technology, our government, our lives."

—from Brian Arthur, "Why Do Things Become More Complex?" —an essay in *Scientific American*, May 1993.

Economist W. Brian Arthur's office at the Santa Fe Institute's new digs in the hills north of Santa Fe is across the driveway from the main mansion, in the chauffeur's quarters. He prefers this isolated workplace because he is trying to work out a new theory of financial decision-making based on pattern recognition and inductive reasoning, and he needs the peace. "I'll come over when I want to talk," he says, which is another good reason for the distance. (Everyone knows Arthur likes to talk.)

Arthur has been puttering around the Institute since 1987 when he was invited by economist and Nobel Laureate Kenneth Arrow to participate in the Institute's first economics meeting, a meeting which set in motion a number of events resulting in what is now the Institute's pivotal economics program. As an External Professor (from Stanford University), Arthur has been instrumental in almost every aspect of the program's evolution. Currently, he is the Citibank Professor at the Institute.

His background is the definition of nonlinear. He was born in Belfast, Northern Ireland, did his undergraduate work in electrical engineering at Queen's University in Belfast, started graduate work in operations research in England, and, after a time, traded the wet and cold of England for the cold and cold of Ann Arbor. At University of Michigan he found the environment he needed to pursue his intellectual interests, but not his personal interests. He arranged to finish his work at University of California at Berkeley, which was considerably closer to the sea. While transferring to Berkeley, he was offered the opportunity to work as a management consultant for McKinsey & Company in Germany, so he took a break from his studies. It was during that time he became fascinated with economics. He eventually finished his Ph.D. in operations research and stayed on at Berkeley as a postdoc in the economics department. In still another departure from his education, he worked as a demographer for the Population Council in New York.

"I saw a lot of the Third World," he says, "India, Bangladesh, Kuwait, Syria." He wound up in 1982 in Stanford with the Morrison chair in economics and population studies. At the time, he was Stanford's youngest endowed chair holder.

What explains this Irish engineer turning demographer then economist ending up in Santa Fe? What explains his coming to have an office attached to a garage? To borrow from Newton, what makes Arthur tick? The answer is enormous curiosity.

"I've always been fascinated by economic development, and I wondered why some economies developed and some did not. I wondered what it would be like to have an economics that took *process* seriously. I wondered about the mechanisms by which technologies evolve, and how that might affect economies. I began to realize that much of our theorizing in economics is rigorous, but based on dubious assumptions; and I wondered what economics would look like if it were based on realistic assumptions."

In the Santa Fe Institute he has found people with similar concerns —and similar curiosity. "I feel at home here," he says.

The assumptions Arthur has taken to task are the very tenets of contemporary economics, the four commandments, as it were: 1) the assumption that the economy is based on diminishing returns, 2) the assumption that all the action that is interesting happens at equilibrium, 3) the assumption that there is a fixed number of goods and services, 4) the assumption that we can regard people as infinitely rational. Arthur compares the economy to a speeded-up evolutionary system where everything is adapting to everything else. "What's truly important in the economy is not equilibrium, but adaptation," he says. "The challenge for economists is to understand how adaptation works."

For modern economies, the tenets are simply outdated, according to Arthur. He explains that economists developed the conventional economic ideas one hundred or so years ago during an age of agriculture and bulk-manufactured goods. At the time, the economy relied heavily upon resources —fertile soil, coal, and mineral ores— that were scarce. Expansion of production of any one product, say wheat or steel, could easily run into resource limitations, thus producing diminishing returns and, eventually, equilibrium in the market.

"Diminishing returns means the more you do something, the harder or less interesting it gets," he says. "The more coal you mine the harder it is, eventually, to reach good coal. So coal becomes more expensive, and it shares the energy market with hydroelectric." Such behavior brings equilibrium, and it allows for a certain predictability in the economy.

This picture no longer reflects the reality of a high-tech economy, an economy now largely driven by knowledge. As Arthur sees it, traditional goods like coal, coke and iron are largely "congealed resources." High-tech products such as software and sophisticated aircraft are largely "congealed knowledge." "Whether you're dealing with pharmaceuticals, telecommunications, or computer software, you're dealing with huge up-front R & D costs. So these get cheaper per unit as production volume increases," he says. "The first floppy disk to go out the door for Microsoft Windows cost the company $50 million. The second ten dollars —to copy the disk, print the manual, and whatnot." Such behavior produces *increasing* not diminishing returns.

One consequence of increasing returns is that the high-tech products that get ahead often stay ahead, whether they are the best technology on the market or not. "What happens in technologies is that we appear to sort of groove out paths we go down, almost like ruts, that are very hard to get out of," Arthur says. "There are positive feedbacks, meaning that the more market share you get, the more advantage you have." He points to the competition between VHS and Beta video players as an example, noting that at some point VHS monopolized the market, although it's widely acknowledged to be inferior to Beta. "I can't take up Beta no matter how good it is if everyone else— and Blockbuster, too—has VHS." Arthur refers to such a situation as lock-in. Much earlier, and perhaps more important, was the race between gas and steam to power cars. There had been one hundred years' worth of experience with steam, and it required very few moving parts," he says. "Had steam been developed, engineers tell me it might have been just as good or better than what we wound up with."

What gives a high-tech product the advantage may have more to do with chance than technology, according to Arthur. "Small events can sway the situation," he says, "whims and who got what at the start." Ransom Olds, of subsequent Oldsmobile fame, went to a car race in 1895 where a gasoline car won, so he switched from making steam cars to gas cars. What's more, research on the steam car went by the wayside when one of the Stanley brothers, designers of the Stanley Steamer, was killed in a steam-car accident. In the world according to Arthur, given different historical events China rather than Japan could be a capitalist giant, and computers could run on totally different principles —parallel rather than sequential processing, say. (The early ENIAC computer, for example, was a parallel processor.)

Long before his association with the Institute, Arthur found the seed for his theories in the sciences, particularly molecular biology and physics. "I believe that all the sciences are intertwined," he says. "And so ideas tend to travel together and permeate many fields at once." He was especially influenced by notions of nonlinear dynamics posited by Herman Haken and Nobel Laureates Jacques Monod and Ilya Prigogine. "When I read their work in 1979 I was taken by the notion that small events — fluctuations— could become magnified by nonlinear effects and have a huge influence on things," Arthur says, "the idea that some slight chance event, some small deviation, could lead through these mechanisms to a larger deviation. I realized immediately that the counterpart in economics was increasing returns."

Arthur contends that economics is at a pivotal point in its development, following, albeit a little late, the path of science. "The history of science in the 20th century has moved away from a deterministic, mechanistic view into something more sophisticated and less simple," he says. "It's suffered a loss of innocence, you could say. This is true of physics, mathematics, chemistry, biology, immunology. It's true of philosophy. The sciences more and more are seeing the world as process-dependent, organic, indeterminate, and always evolving." Economics, he contends, has always had an affinity with ideas in evolution. "Darwin got his key idea of competition for survival by reading the economist Thomas Malthus. And there's been a thread, after Darwin came along, where economists have been

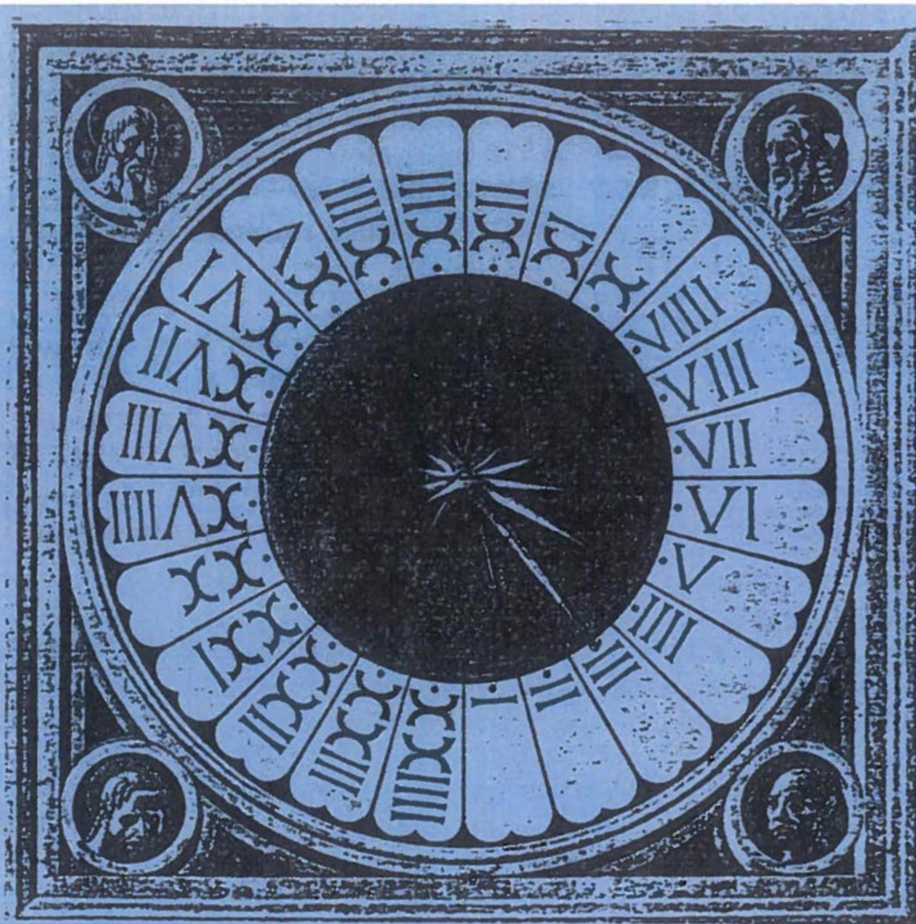

Asked at a 1982 lecture to give an example of a technology in history that would have been as good as the one that became locked in, Arthur replied, "Clocks. For all I know clocks could as easily have gone 'anti-clockwise' as the convention we take for granted."

Three weeks later a vacationing colleague sent him this picture of the 1443 Uccello clock in Florence cathedral. The back of the postcard read simply, "Congratulations." Twenty-four hour faces and counterclockwise direction persisted until about 1550 when they were crowded out by the standard we have today.

looking at firms and consumers almost the way you look at an ecology," he says. In this thinking, posited by Alfred Marshall as early as 1890, economic firms or products have to compete against each other in an environment they created themselves.

The problem with this, Arthur explains, is that the Victorian economists had neither the mathematics nor modern ideas of evolution to back up their instincts. "What's happening in economics is that things are beginning to change and we're coming back into this Darwinian way of looking at things," Arthur says. "We're moving steadily from looking at statics to dynamics, from looking at equilibrium to looking at process. There's a growing movement in economics in this direction. We are filling out a part of the picture that has long been missing."

Along these same lines is the notion that the idea of equilibrium should be replaced with the idea of "transience." Arthur uses game theory to demonstrate the difference between an economy based on equilibrium and an economy based on transience. He explains that tic-tack-toe is a simple game where strategies become evident —and therefore static— after only a few games. Chess, on the other hand, is much more complicated. Players are always developing new strategies. "A grand master today could beat a grand master from one hundred years ago," Arthur says. "Not because he's smarter, but because chess masters continually push out the envelope of new strategies." With this shift in thinking comes an emphasis on adaptation, learning, process and evolution — on transience rather than equilibrium.

Arthur contends there is a movement in science to look more holistically at individual entities as they interact and form patterns. "Reductionism narrows something down to a single mechanism and then puts it under a powerful metaphorical microscope," he says. "With reductionism, we lose sight of how these individual mechanisms could work in concert. Looking at how elements together produce pattern or structure is a major change in science that, in my opinion, is not going to go away," he says. He attributes this sea-change in science to the computer workstation, which allows scientists to re-create systems within the computer, analyze pattern formation, and explore process. Sitting not far from his own NeXT Workstation, he says, "Imagine what Darwin could have done if he had a computer and let little organisms evolve inside it. Imagine Henri Poincaré, the mathematician, doing dynamics on a Connection Machine."

Arthur's ideas weren't readily accepted by the establishment of economics, which makes it even more significant that his introduction to the Santa Fe Institute came from Arrow. Arrow, who was (and is) also at Stanford, had had occasional conversations with Arthur and had followed his work on increasing returns. He immediately saw the parallels between Arthur's work and the work physicists such as Philip Anderson were doing with randomness. "I thought his work was an insightful departure from conventional ways of thinking," Arrow says now. At the time Arrow was helping to put together the economics meeting, a conference that was going to bring together economists, biologists, and physicists. "It was my job to find the economists," Arrow says. Arrow attributes some of the resistance to Arthur's theories to the simple fact that he wasn't a member of the club. "He didn't belong to the Guild," Arrow says. "His Ph.D. was in operations research and his experience in demography. He was coming in from the outside."

Arthur finds these days that economics is swinging in his direction. "SFI is a place where leading-edge economists can come and talk about adaptation, evolution and learning in the economy. We put them together with their counterparts in physics, computer science, biology. Nobody's carping about standard economics. We're just saying 'Hey, let's play in a different sandbox.'" Playing in Arthur's sandbox just now is a team that is building an artificial stock market inside a computer. (See the *Bulletin*, Fall 1993). Included are University of Michigan computer scientist John Holland, who developed the genetic algorithm, University of Wisconsin economist Blake LeBaron, Duke physicist and neural net expert Richard Palmer, and Arthur. "One way to do economics is to study data from the real economy," Arthur says. "Another way is to do pencil and paper theory. Another way, we began to realize here, is to actually create little economies inside the computer and study them." Arthur came to this conclusion in 1987 after talking to Holland about his work using artificial agents to learn complex tasks. "We didn't want to create a bunch of equations and solve them on machines," he says. "We wanted to create entities so that each investor is its own little computer program—smart little critters that can learn and adapt."

The Santa Fe virtual stock market is made up of one hundred or so agents who can buy or sell stock or place money in the bank. The agents, which have the capability to learn, begin trading with zero intelligence. They learn and adapt by spotting patterns in the prices that result from their trading. The researchers can change the number of traders, the information the traders get, and the trading rules, to analyze their behavior. In the stock prices that result, bubbles and crashes emerge; sometimes investors get scared out of the market; the market develops a "psychology." In one experiment Arthur calls the Jurassic Park experiment, the team takes the three smartest traders out of the market and freezes them. "We let the program run for a day or two then inject the traders back in," Arthur says, "so that they are like the Velociraptors and T-rexes." Arthur explains that the revived traders generally do poorly in the new market because it has evolved.

The stock market is only one of a number of projects Arthur is working on. "I am fascinated by cognition and decision making these days. We assume, for easy analysis in standard economics, that people are infinitely smart, and this works fine for many purposes. But in reality, as humans, we are no great shakes at deductive logic. What we are incredible at, instead, is pattern recognition. So I am interested in figuring out how economics — finance theory in particular—would look under realistic human rationality. Under pattern-recognizing, inductive, honest-to-god, subjective human reasoning, I've learned a lot from the psychologists and adaptive-computation folks at Santa Fe."

Once desperate for someone to validate his theories, Arthur is now besieged by journalists, analysts, and business people alike. Some he meets at his modest outpost at the Institute; others he meets on their own turf. He is fond of telling the story of a visit in June 1991 to the office of then-Senator Albert Gore. After a brief description of his theories to Gore and his staff, the Senator asked for an example of "increasing returns." "I thought a bit," Arthur says, "and said, 'the presidential primaries.' They all sat bolt upright immediately."

Of late, Arthur has found the business community open to change. "Business executives tend to be impatient with standard economics," he says. "They say it doesn't apply much to what they do."

One businessman with a different view is Henry Lichstein, Vice President at Citicorp, whose Chairman, John Reed, has been a consistent financial supporter of SFI's economics program. Lichstein has known Arthur since 1989 when he came to Santa Fe for the second economics meeting. "Arthur is a particularly creative economist," he says. "The kinds of things he's done are the kinds of things economics has left undone." Lichstein finds Arthur's work particularly valuable in explaining economic development and the process of innovation, an area he considers weak in classical economics. "Arthur brings insight into a problem by allowing us to think in different ways," Lichstein says.

What Arthur does have to say to the business community at large is this: "The name of the game in business is no longer optimization. The challenge is much less to perfect an unchanging operation these days than to figure out what markets to be in, what products to launch, what strategies to bet on. So it's not about optimization but about adaptation." As for his grander quest to change the axioms of contemporary economics, Arthur is a patient man. He believes in evolution. He believes in process. "If you watch a photograph developed in a darkroom, the paper is blank at first," he says. "But then you see shapes start to emerge; contrast starts to increase; details appear, and eventually you can see what the subject is. This is very much like what it would be to figure out an economics based on process." The picture is already clear to Arthur and, like a father brandishing photos of his children, he can't wait to show it to you.

Janet Stites is a free-lance writer whose work appears in Omni *and other magazines. Stites also profiles SFI's founders in this issue.*

Further reading: W. Brian Arthur: "Positive Feedbacks in the Economy," *Scientific American*, Feb 1990, and *Increasing Returns and Path-Dependence in the Economy*, University of Michigan Press, Ann Arbor, 1994. M. Mitchell Waldrop, *Complexity: The Emerging Science at the Edge of Order and Chaos*, Simon and Schuster, New York, 1992.

Building New Models for Change in Organizations

Leading organizational theorists met at SFI in August and began to carve out a set of building blocks for a basic "organizational simulation library." The creation of this generic set is significant because it will put existing organizational models on a uniform footing by implementing them within a common framework. Such a framework — despite some recognized imperfections— could bring order to the field of organizational modeling and be used for teaching students who are new to modeling what the major functional elements of a model are, for guiding authors and reviewers within the literature, and for commercial use.

"Computer modeling of adaptive organizations is a field just taking shape in academic and industrial circles," says program organizer Michael Cohen, "and this could help crystallize it. Software companies, university labs, and consulting firms are beginning to produce simulation models for industrial clients at a rapidly-increasing pace; but they need better, more uniform tools for modeling the tasks actors do, the structures of organizations, and the way that organizational structure shapes the action of its members."

RESEARCH UPDATES

The Echo and Swarm computational systems—general frameworks for simulating multi-actor, complex systems currently being developed at SFI— could provide especially easy-to-use environments for such a catalog of organizational modeling elements.

The participants at the August workshop began by dismantling and analyzing — "reverse engineering," as it were— five classic models in the field. Computer-based modeling of organizations has a long history, going back before 1960; but the area has grown much more slowly than similar modeling of individual cognitive processes. (An exception is operations research modeling in areas like inventory and queuing, but that style of work was not considered at this meeting.) The models chosen were widely-respected examples that span thirty years. Each model was presented by someone who was not its author, although in most cases the author was present and so could comment before a general discussion. The aim was to create a common ground of detailed understanding that could support the subsequent inquiries of the group.

With these examples fresh in mind, they tackled the problem of crafting a framework that lays out the principal building blocks required to construct a useful computer-based model of an organization. The group began by clustering elements under four headings: agents who compose the organization; relations among the agents (such as authority or communication); tasks that the organization is to perform; and resources required in the performance of those tasks. Then they refined the definitions of the major clusters, suggesting subdivisions and developing examples from the existing literature. The group intends to continue working on this scaffolding through informal collaboration and further meetings.

On a parallel track, the group discussed the potential content of a jointly-developed graduate curriculum in this field and how an emerging framework might affect it. They exchanged and analyzed syllabi of existing courses, finding many similar elements; and they also talked about some of the common problems with current courses and how to address them. For example, classic models are not easily available to students in source code; existing models run on an unwieldy variety of computational platforms; and it is difficult to compare the results of new modeling exercises to prior results. Possible solutions might be an Internet server that provides access to source code of classic models, production of video lectures about major modeling issues to be distributed nationally, reimplementing classic models on a common platform, and setting standards within the literature for making code available.

The last suggestion may be put into play immediately. *The Journal of Computational and Mathematical Organization Theory* published by Kluwer will begin production soon. The group was asked by co-editor Kathleen Carley to provide advice on strategies to foster the volume and quality of work available for publication.

On the basis of the August meeting, the group envisions that in the near future public-domain libraries of task, resource, actor, and relational objects can be used in conjunction with platforms like Swarm, Echo, and other systems to rapidly develop new and more comparable models. "There is a real and exciting possibility for developing shared resources that will ease the time needed to learn to build models, the effort required to actually build them, and the problems of comparing results," says Cohen. "And we plan to pursue this aggressively at SFI."

REPORT FROM THE FRONT: ALIFE IV

RESEARCH UPDATES

The fourth Artificial Life Conference was held in July at MIT in Cambridge, Mass. After having organized the three previous conferences through the SFI, it was quite a pleasure to attend an Alife conference as a participant, with someone else taking care of the organizational responsibilities! This time, the conference was organized by Rodney Brooks (Assistant Director of MIT's Artificial Intelligence Laboratory) and Pattie Maes (of MIT's Media Laboratory).

Below, I report on several of the papers presented at Alife IV. (These are contained in the Proceedings Volume, *Artificial Life* —available from MIT Press.) From my point of view, these papers stood out above the rest in overall quality and/or significance. As a consequence, I am publishing these papers together as a special issue of the journal *Artificial Life* (Vol I, No. 4).

The first two research projects, by Karl Sims and by Dimitri Terzopoulos *et al.*, achieve what I consider to be significant new levels of success in the physical modeling of life. Both research projects involve the construction of artificial worlds with "real" physics, in which artificial organisms with "real" physical morphologies behave and interact with one another. They conquer new ground in an area that might be called "virtual robotics."

The paper by Terzopoulos *et al.* reports on current progress in their project to construct a complete artificial marine world, replete with seaweed, plankton, and fully-functional artificial fish. The long-term goal of their research is to develop "a computational theory that can potentially account for the interplay of physics, locomotion, perception, behavior, and learning in higher animals." The artificial fish themselves are quite literally software robots. Each fish consists of "23 nodal-point masses and 91 springs. The spring arrangement maintains the structural stability of the body while allowing it to flex. Twelve of the springs running the length of the body also serve as muscles." (See Figure 1). Swimming involves the flexing of the muscles, which deforms the body which pushes against the water using simulated hydrodynamics. Over time the fish learn how to control their muscles to effect swimming. The visual images perceived by the fish are run through a relatively simple network of goal and intention generators, which results in the production of signals to the muscles. Thus, the fish react in a fairly low-level way to the stimuli in their environment, swimming toward food or away from predators. The physical modeling and graphic rendering is so sophisticated that the results resemble underwater video documentaries of real fish. (In fact, one sequence shown at the workshop was entitled, "The Undersea World of Jack Cousto.") This is a long-term project just getting under way, but it has already demonstrated the capability of its software to serve as a significant virtual laboratory for implementing and testing models of the physiology and "psychology" of real fish.

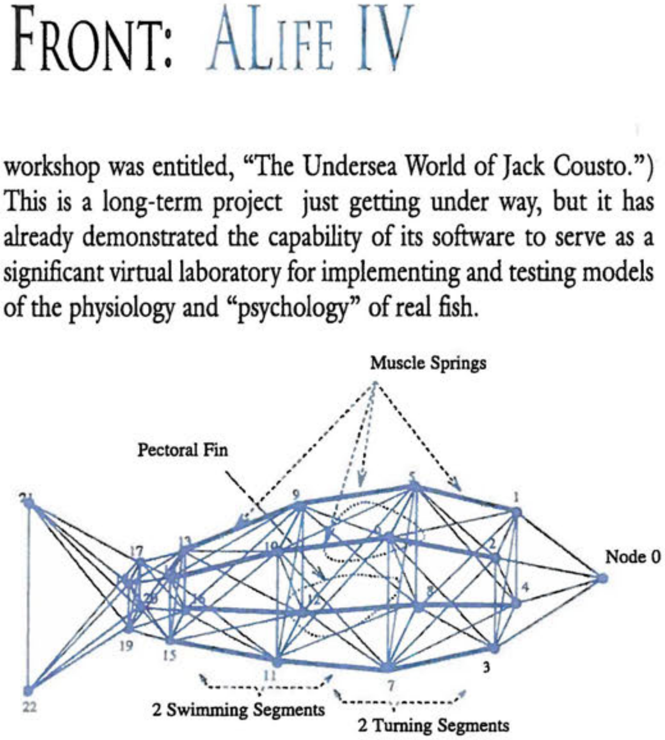

Figure 1

The work of Karl Sims was, for me, the most significant presented at the workshop, partly for what he has accomplished so far, but more importantly for the potential his modeling technology holds for the study of the co-evolution of physical structure and behavior in open-ended ecosystems. Karl constructs 3-dimensional artificial worlds incorporating Newtonian physics, gravity, surface friction, fluid-dynamics, and so forth. He populates these worlds with artificial organisms whose physical morphology is determined by an elegant genetic representation, which also codes for a nervous system to respond to stimuli and control the joints connecting the body segments. (See Figure 2).

Karl starts out with a few simple random genotypes, and then applies selective pressures for the performance of specific tasks such as swimming, walking, jumping, or competing in various pair-wise contests. The results are remarkable. Evolutionary runs lead to the development of complex morphological and neural structures, resulting in organisms with truly astonishingly realistic physical adaptations to the tasks imposed on them and to one other's capabilities. Figure 3 illustrates some of the creatures that have been evolved in the context of a competition for possession of a cube. The physical simulation is realistic enough so that accidental consequences of the physics can be capitalized on by evolution. For instance, in wrapping an arm around the cube in order to gain possession of it, the arm may accidentally hit the

Figure 2

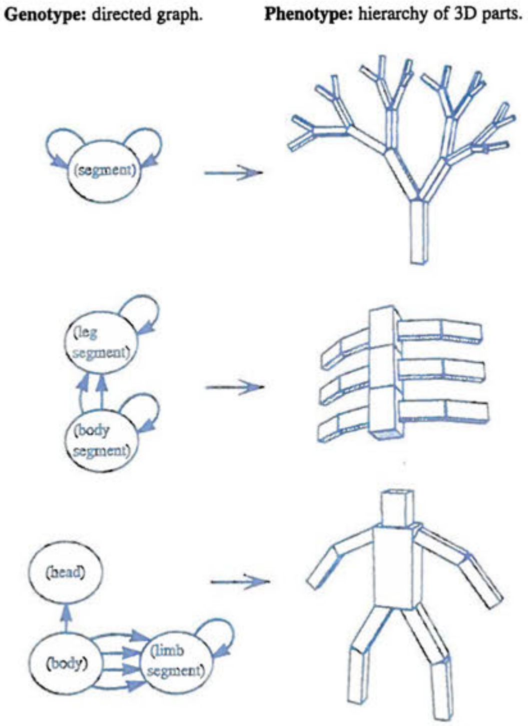

Figure 3

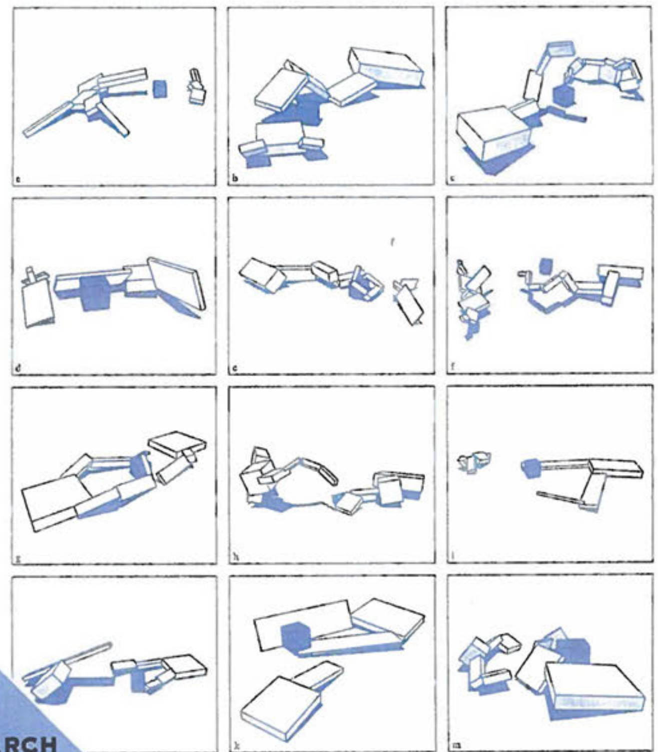

opponent, pushing it away from the cube. This adds to the selective advantage of the organism, and so this accidental side effect can be distinguished and then selected.

The capability of this system to model organisms that can interact with one another, vying for resources according to their physical capabilities, promises to provide an extremely useful platform for the study of open-ended morphological evolution. With such a system, one might introduce a primitive ecology consisting of primary food producers (like plants) that convert sunlight into energy, herbivores that feed on the plants, and carnivores that feed on the herbivores. Physical evolution could then be allowed to take its course, with plants learning to shade out competitors, and to protect themselves from herbivores; with herbivores learning how to get around the defenses of the plants and to avoid the carnivores; and with the carnivores learning to get around the defenses of the herbivores and outwit each other in the competition for resources. The demonstration of the capability to implement this kind of evolution involving natural selection working among varieties of physical capabilities is what makes Karl Sims' work so exciting.

Another domain of artificial life research concerns the application of evolutionary techniques to engineering problems. It is becoming apparent that evolutionary search procedures can produce solutions to complex engineering problems that rival or even improve on the best efforts of human designers. The work of Johnson, Mas, and Darrel in evolving routines for visual perception is a case in point. The problem of identifying features in images is a well-known hard-problem in Artificial Intelligence. Much of the work in this area has involved the hand-coding of routines to pick out certain features in an image. However, these hand-coded solutions are often quite time consuming and error prone; and they hardly ever generalize. Thus, Johnson *et al.* decided to use the Genetic Programming paradigm introduced by Koza to attempt to evolve feature recognition algorithms. The specific feature recognition task they chose was that of finding the hands of a person in a bitmapped visual image. (See Figure 4). Although the problem domain is a relatively simple one, they were able to evolve routines that consistently beat their hand-coded attempts. The important point about evolving algorithms to solve engineering problems is that human programmers tend to write routines that "make sense" to them, whereas evolution can put together algorithms that people would never write. At first glance, evolved algorithms seem to do things for no apparent reason, or to do things that they undo later. However, on closer analysis, it becomes apparent that the algorithms are exploiting extremely subtle side-effects of

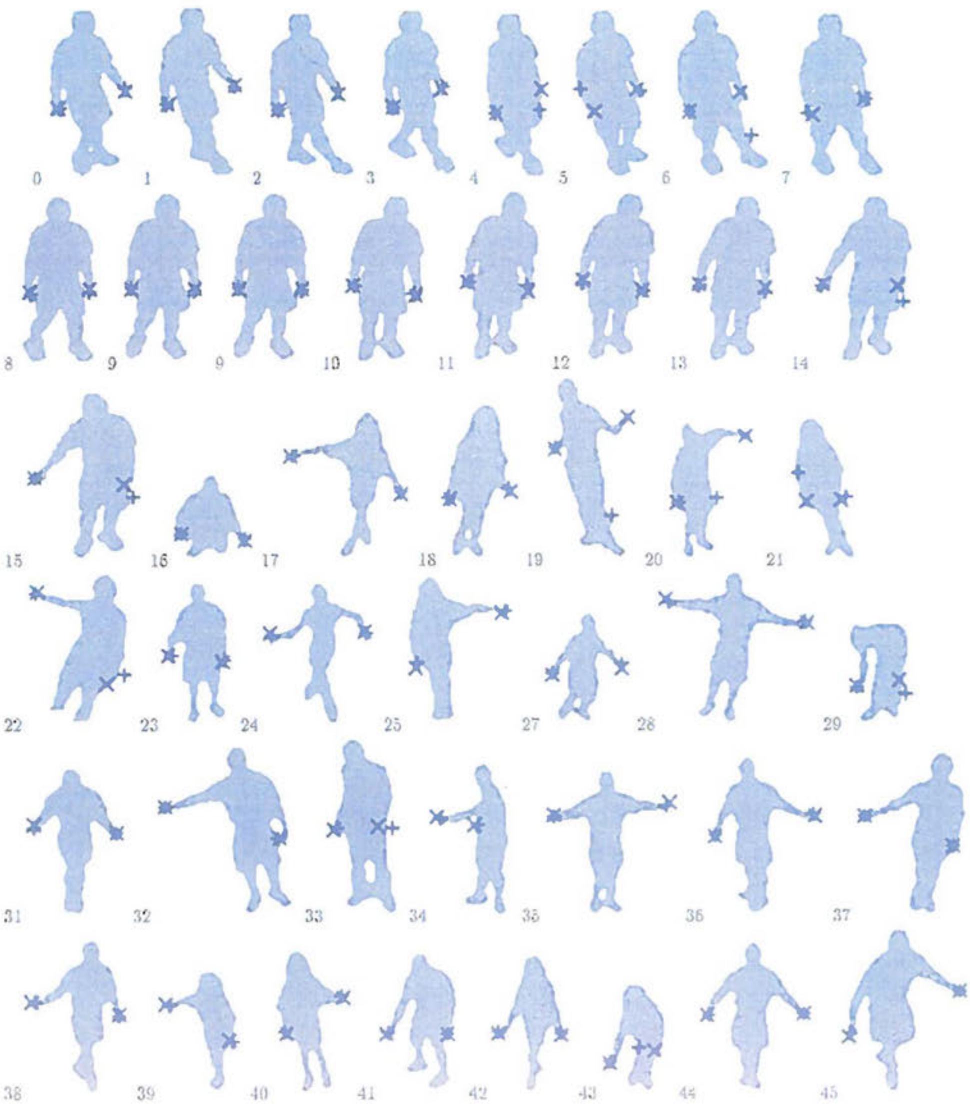

Figure 4

instructions rather than the primary effect of the instruction that we tend to think of when we write code. Evolution has no "conception" of what an instruction's primary function is, however, and will merely follow up on whatever works. The routines evolved by Johnson *et al.* have this character.

The development of evolutionary approaches to engineering seems completely orthogonal to previous engineering approaches because it is quite likely that we will not understand why some solutions work. This is potentially dangerous since, if we can't understand a solution, we might not be able to guarantee that it will not fail catastrophically in some cases. However, the other side of the coin is that evolution can be applied to the testing of evolved solutions as well, evolving test-suites that attempt to break the evolved solutions. Danny Hillis points out that such evolutionary testing will also try cases that we would not think of, and may discover cases for which the solutions break down ("Co-Evolving Parasites Improve Simulated Evolution as an Optimization Procedure" in *Artificial Life*, 1992). By coupling the evolution of the solutions with the evolution of the test-suites, Hillis suggests that the solutions arrived at will be more reliable than those which have only been tested on the kinds of standardized test-suites currently employed in engineering, and the field of solution validation by evolution is as promising as the field of solution discovery by evolution.

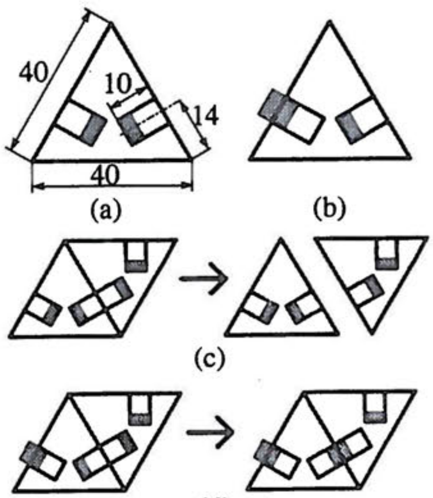

Figure 5

A primary goal of Artificial Life is, of course, to help us understand real biological systems and processes. The best work in this area comes from collaborations between biologists and computer modelers. The paper by Toquenaga, Kajitan and Hoshino examining group foraging and colonial formation in avian species is a good example of such collaboration. The authors set about to test a theory due to Horn, which predicts that colonial organization has an advantage over territorial organization in environments with patchy resource distribution. They evolve foraging strategies for birds in environments with either uniform or patchy resource distributions. They find that when resources are clumped, the artificial birds form colonies to collectively find and exploit the resources, whereas when resources are evenly distributed, colony formation is never observed, with individual birds exploiting the territories in the neighborhood of their own nests. Models such as these allow researchers to test hypotheses about the benefits of collective action on the part of organisms that are difficult or impossible to investigate analytically or in the field.

The paper by Hosokawa, Shimoyama, and Miura is a nice example of Artificial Life in hardware. They investigate the dynamics of self-assembling systems using a very clever physical device as an analog of a molecule that can bind to copies of itself. (See Figure 5.) The equations governing the dynamics of self-assembly and the production of various end-products in this artificial system are found to correspond to those of real chemical kinetics. This work is being carried out in the Department of Mechano-Informatics at the University of Tokyo, which I reported on in my article on Artificial Life in Japan in the last issue of the *Bulletin*.

Finally, the paper by Chris Adami of Caltech takes a major step in the theoretical understanding of the evolutionary process. Adami investigates the statistical mechanics of self-replicating bit strings, both experimentally and theoretically. He finds that evolution in such systems proceeds via stable periods "punctuated" by rapid evolutionary successions, à la the punctuated equilibria proposed by Gould and Eldridge. However, he also finds that these punctuations are distributed fractally in a manner suggestive of self-organized criticality (Per Bak, C. Tang and K. Wiesenfeld in *Physics Review Letters*, No 59, 1987). Adami also associates an entropy measure with the diversity of the population and investigates the changes in this entropy measure over evolutionary time. Punctuation, in which a new strategy emerges and takes over the population, results in a significant decrease in this entropy measure as the diversity of the population is drastically reduced. Thereafter, the entropy gradually recovers as the new species diversifies. Many people have attempted, with little success, to apply principles of statistical mechanics to the evolutionary process over the years. Adami's work should form the basis for significant improvements of the success of such efforts.

RESEARCH UPDATES

Stepping back from this work, my overall impression of Alife IV was very positive. Although there were not so many surprises as in the previous Alife conferences, it is clear that work in the field is beginning to mature, with more and more of the work taking on the characteristics of any other scientific or engineering discipline.

The next Alife meeting in this series will be held in Kyoto, Japan in 1996. The Third European Conference on Artificial Life will be held in Granada, Spain in the summer of 1995.

For further information on events, publications and software in the field of Artificial Life, check our World-Wide-Web pages on Artificial Life at <http://alife.santafe.edu/> or subscribe to the journal *Artificial Life*, published by the MIT Press.

SFI Professor Christopher Langton is often called "The Father of Artificial Life."

SWARM: THE NEXT GENERATION

Over the past summer, the Swarm team —Roger Burkhart, Howard Gutowitz, David Hiebeler, and Nelson Minar, working under the leadership of Chris Langton— focused on a major redesign of Swarm, based on the experiences gained from the system prototype.

Swarm is a general-purpose simulation package for the investigation of concurrent, distributed systems in which a large number of autonomous agents interact with one another and with a dynamically-changing environment. The package provides utilities for designing, implementing, running, and analyzing these multi-agent systems. The primary goal of the Swarm simulation system is to save researchers from having to deal with all of the computer-science issues involved in the implementation of concurrent, distributed artificial worlds. Swarm provides a wide spectrum of "generic" artificial worlds populated with "generic" agents, a large library of design and analysis tools, and a kernel to drive the simulation. These artificial worlds can vary greatly in their properties, from 2-D spatial worlds in which mobile agents move about "physically," to graphs representing telecommunication networks through which sessile agents trade messages and commodities. Whatever the specific "physical" characteristics of the universe of discourse, Swarm's uniform framework frees researchers to concentrate on their specific system of interest, to directly compare scientific results with other users of Swarm, and to eliminate wasteful duplication of common simulation functions from model to model.

Swarm, though currently written in pure C, is object-oriented in style: Everything in Swarm is an object. Objects communicate with other objects by sending messages. The simulation is driven by a special object, the Object List Manager (OLM), which maintains a list of the active object population and sends "step" messages when it is time for objects to update themselves.

All inhabitants of the artificial world (bugs, economic agents, molecules, etc.) are objects. In addition, the environment (or space) the objects live in is itself an abstract object. A space object defines the geometry of a space and manages a set of spatial variables associated with it by the user. The space object defines methods for actions which depend only on spatial geometry, such as the movement of an agent, and does so in such a way that user-defined agents can take immediate advantage of these methods. A space object also makes available to the user a set of general functions operating on spatial variables.

The interface to Swarm currently consists of a collection of tools for analysis and visualization. Data analysis objects (such as objects that compute averages) exist within the Swarm world; these objects then talk to user-interface objects to present X-window displays of data, as well as storage to files for batch-time data collection. Swarm provides the interface; the researcher spends his or her time developing the simulation. To date, Swarm has been used in an agent-based study of the economics of environmental pollution (by G. Weisbuch, H. Gutowitz, and G. Guillemette Duchateau-Nguyen), a study of spatial influences in algorithmic chemistry (by W. Fontana, L. Buss, and N. Minar) and a four-species ecological simulation (by D. Hiebeler).

As is, Swarm is primarily a prototype that was developed to explore the issues involved in implementing such a simulation tool. Although it has been released internally within the Santa Fe Institute in order to get feedback from real users, it will not be released externally, but rather will be replaced by a new, built-from-scratch Swarm that incorporates everything learned from the prototype but also adds new features.

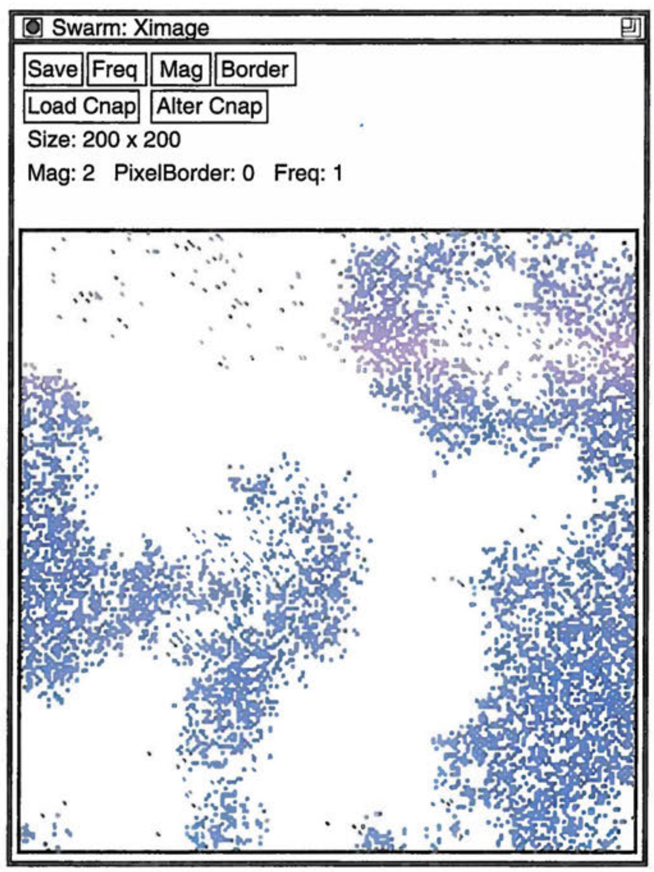

In order to better express the object-oriented design of Swarm simulations, Swarm is being rewritten within a more formal, object-oriented framework. The language chosen is Objective C, a useful compromise between the ubiquity and efficiency of C and the flexibility provided by a run-time, object-oriented environment. Design elements being incorporated in the new version of Swarm include:

- Making Swarm recursive, supporting a hierarchy of functional levels. (Thus, an agent in a Swarm at one level of the hierarchy may be a whole Swarm itself, consisting of multiple agents interacting with each other one level down in the hierarchy.)
- Allowing agents to associate and dissociate into new agents or Swarms so as to permit emergent structural and functional entities.
- Providing a flexible update scheme, allowing a wide variety of synchronous and asynchronous, time-stepped and event-driven simulations.
- Supplying general-purpose, efficient algorithms for common simulation needs, such as finding nearest neighbors in spaces of various geometries.
- Supplying a basic collection of evolution and learning algorithms.
- Providing a convenient, powerful user-interface, building on the Tk/Tcl library, to allow both graphical interaction and batchmode processing.
- Supplying a "Physics Toolkit" which allows users to simply specify the kinds of space and time operating in their artificial world. In particular, the toolkit gives the user 0, 1, 2, ... -dimensional geometries which he can plug and play, as well as choices for continuous or discrete-time updating in each of several sub-phases of a simulation.
- Developing a general-purpose scheme for querying data from sub-collections of agents. Making this query-language flexible enough so that agents themselves can emit queries, and build sub-collections of other agents.
- Supplying more facilities for data I/O: instrumentation of the simulation.
- Satisfying "Thearling's Law," which states that any simulation package should include twice as many analysis objects as simulation objects.
- Simplifying, as much as possible, the user's task of developing new agents and spaces to run under Swarm as well as the task of developing new analysis objects to be added to the user library.
- Documenting the system thoroughly, including a design document, a user's manual, and a Swarm kernel hacker's manual.

Improvements along the lines sketched above will help make Swarm sufficiently complete and stable so that it can be released for general use in the ALife and Complexity communities. This release should bring a many-fold increase in the user community directed at improvement and extension of Swarm. It is hoped that users will write additional modules of general applicability and reusability.

The core Swarm team will continue to concentrate on basic architecture and common needs and will test the Swarm design across a wide spectrum of potential applications. Long-term goals include making Swarm portable to both small environments like Windows and the Macintosh as well as taking advantage of parallel resources like a CM-5 or a network of workstations. In addition, higher-level programmer interfaces (such as *Logo) are being considered to further simplify the Swarm user's work. An internal design document is currently being developed, which should be available shortly. A releasable alpha-version of the new Swarm is intended for early 1995 as well. A full Swarm Version-1.0 should be ready for release by the end of the summer of 1995. Swarm releases will be made available to the public free of charge. Swarm will initially assume a Unix workstation with X-Windows, followed up with Macintosh and PC versions. (At present, Swarm runs on PCs running Linux.)

For further information about Swarm, send requests to <swarm-request@santafe.edu>, or refer to the Swarm WWW page at <www.santafe.edu> under Projects.

RESEARCH UPDATES

SFI's FIRST DECADE

84

MAY Articles of incorporation for the Institute are approved under the name of "Rio Grande Institute"
OCTOBER The Institute receives a non-profit designation from the IRS
OCTOBER First interim Board meeting. Murray Gell-Mann elected Chairman, George Cowan elected President. Throughout first 2-1/2 years Helene Slansky provides administrative support for the organization. The Albuquerque office of Board Treasurer Arthur Spiegel provides financial record-keeping and oversight

85

(Annual budget, $83,000—$8.3x10^{-4}$ "units")
MARCH First meeting of the Board of Trustees. Murray Gell-Mann elected Chairman, Ed Knapp, Vice Chairman of the Board. George Cowan continues as President
NOVEMBER The Institute's name changed to the Santa Fe Institute
NOVEMBER First SFI sponsored meeting, Superstring Workshop, chaired by Murray Gell-Mann and Richard Slansky. Hailed by participants as ground-breaking, no written records of the meeting exist

86

(Annual budget, $97,000)
MARCH David Pines elected Chairman of the Board of Trustees. Peter Carruthers, John Rubel and Mike Simmons elected as Vice-Presidents
MARCH First meeting of the Board of Advisors and Planning Committee—the nascent Science Board—chaired by Murray Gell-Mann
JUNE First issue of the SFI *Bulletin* produced (on a xerox machine) by Simmons, with Carruthers, Pines and Simmons as editors
JULY First SFI community "lecture" (and perhaps the longest)—David Pines, Philip Anderson, Jack Campbell, Jack Cowan, Marc Feldman, Murray Gell-Mann, and Larry Smarr on "Understanding Complexity: An Introduction to SFI"
JULY Founding Workshop on Complex Adaptive Systems, funded by Alfred P. Sloan Foundation
AUGUST SFI establishes first office, one room at First Interstate Plaza in Santa Fe. Ron Zee hired as Executive Director
AUGUST Bob Adams, Phil Anderson, George Cowan, David Pines, John Holland and others meet at Rancho Encantado with John Reed and Citicorp staff to discuss the global economy, setting the stage for SFI's program in economics
AUGUST First national press coverage of SFI in *Business Week*
SEPTEMBER Ginger Richardson joins staff as Program Coordinator

87

(Annual budget, $571,000)
MARCH Ed Knapp elected Chairman of the Board of Trustees; Murray Gell-Mann and David Pines, Co-Chairs of the Science Board
FEBRUARY SFI offices move to former convent, 1120 Canyon Road
MARCH SFI Library established with gift of the Stanislaw Ulam Collection by Françoise Ulam; additional gifts from Nicholas Metropolis and Herbert Anderson
MARCH Addison-Wesley Publishing and SFI sign contract to establish the Studies in the Sciences of Complexity series
APRIL Ronda Butler-Villa joins staff as Publications Coordinator.
JUNE Mike Simmons, on sabbatical from LANL, hired as part-time Vice President for Academic Affairs
JUNE First SFI workshop on Theoretical Immunology, chaired by George Bell and Alan Perelson
JULY First summer school, Matrix of Biological Knowledge, headed by Harold Morowitz. With program funds SFI acquires 6 IBM PC computers. Ultimately these are recalled by NIH, allegedly for use in Operation Desert Storm
SEPTEMBER Founding workshop on Evolutionary Paths of the Global Economy funded by Russell Sage Foundation and Citicorp, chaired by Philip Anderson and Kenneth Arrow
SEPTEMBER First Artificial Life conference held in Los Alamos, headed by Chris Langton. About 125 people attend
SEPTEMBER First workshop on Computational Approaches to Evolutionary Biology, chaired by Marc Feldman. This meeting focused on mathematical modeling of biological phenomena fundamental to SFI's eventual program in adaptive computation
SEPTEMBER Andi Sutherland joins SFI staff as Office Manager
SEPTEMBER First Visiting Fellows—Philip Anderson, Kenneth Arrow, W. Brian Arthur, John Holland and Stuart Kauffman—in residence at SFI
SEPTEMBER Community lecture by John Holland on complex adaptive systems draws 50 people
NOVEMBER First Wall Street Journal article on SFI research
DECEMBER First printed issue of the SFI *Bulletin*

88

(Annual budget, $1,174,500)
JANUARY First volume in SFI Studies in the Sciences of Complexity series, *Emerging Syntheses in Science* edited by David Pines, is published by Addison-Wesley
JANUARY First NSF award for a "Broad Research Program in the Sciences of Complexity" at SFI
JANUARY First graduate student, Jasmina Arifovic, from University of Chicago in residence at SFI as part of the Economics program
MARCH J. Burchenal Ault hired as SFI's first Director of Development
APRIL First DOE award for a "Broad Research Program in the Sciences of Complexity" at SFI
APRIL First Sun workstations arrive (3 of them). Local network formed at the Institute and SFI hooks up to the world through Internet
JUNE First award from John D. and Catherine T. MacArthur Foundation for a "Broad Research Program in the Sciences of Complexity" at SFI
JUNE Marcella Austin joins SFI staff as Financial Assistant
JUNE Stephen Pope hired as SFI's first system manager, his official title being "Computer Expert"
JUNE W. Brian Arthur, first long-term residential visitor, comes to SFI to head Economics program for the next fifteen months
JUNE First Complex Systems Summer School headed by Dan Stein takes place at St. John's College. Ninety students apply for fifty slots
JUNE SFI colloquium series established
JULY Della Ulibarri comes to SFI as Research/Publications Secretary
SEPTEMBER John Miller, first SFI Postdoctoral Fellow, takes up residence at the Institute to work in the economics program
NOVEMBER First workshop on issues of sustainability, "The Elements of Global Security," chaired by George Cowan

89

(Annual budget, $1,545,000)

FEBRUARY SFI Working Paper series initiated

MARCH Robert O. Anderson named Chairman of the Board of Trustees

MARCH Susan Wider joins the SFI staff as Director of Development

MARCH Initial 24 External Faculty members of SFI are appointed

MARCH First SFI workshop on Applied Molecular Evolution, chaired by Stuart Kauffman and Alan Perelson

MIDYEAR Alan Lapedes gathers a SFI group to work on pattern recognition in biological sequences and to explore computational approaches to genetic data analysis.

MAY First Complexity, Entropy and Physics of Information meeting, chaired by Wojciech Zurek

JUNE Margaret Alexander joins SFI staff as Archivist/Librarian

FALL First workshops at SFI focusing explicitly on man, environment and society: Theoretical Ecology, Prehistoric Southwestern Archaeology, Public Policy Studies, Human Cognition and Emotion, Evolution of Human Languages

90

(Annual budget, $1,807,500)

JANUARY First undergraduate intern, Julie Pullen from Macalester College, in residence at SFI

FEBRUARY Artificial Life II conference held in Santa Fe. Three hundred people attend

SPRING First formal SFI Postdoctoral Fellowship in Complex Systems Studies established

AUGUST Patrisia Brunello hired as SFI's receptionist

SEPTEMBER George Cowan and Mike Simmons travel to Washington to brief Senator Al Gore on SFI's research

SEPTEMBER First SFI workshop on Nonlinear Modeling and Forecasting, chaired by Martin Casdagli, Stephen Eubank and J. Doyne Farmer

NOVEMBER First formal institutional collaboration formed, the University of Michigan/SFI joint research program

DECEMBER Tenth book published in SFI series *Complexity, Entropy and the Physics of Information*, edited by W. H. Zurek

91

(Annual budget, $2,312,750)

MARCH George Cowan steps down as first SFI President. Edward A. Knapp succeeds him

MAY Kimberly Bodelson joins SFI staff as Executive Assistant

JULY SFI moves to interim quarters at 1660 Old Pecos Trail

JULY Barbara Hodges joins SFI staff working with the Development Office staff

JULY Deborah Smith joins the Program Office staff at SFI, specializing in researcher support

SEPTEMBER SFI extends its educational programs to the Santa Fe secondary schools in a project coordinated by Suzy Pines, supported by the Pinewood Foundation

SEPTEMBER External Faculty members J. Doyne Farmer and Norman Packard form Prediction Company, a business using time-series techniques for financial analysis

OCTOBER Bruce Abell joins SFI Staff as Vice President for Finance and Operations

NOVEMBER Scott Yelich hired as System Manager for SFI computer environment

NOVEMBER Jean Farrar hired as Financial Assistant

92

(Annual budget, $2,843,700)

JANUARY First Complex Systems Winter School, co-sponsored with the University of Arizona, held in Tucson, with Peter Carruthers as Director

MARCH James Pelkey elected Chairman of the SFI Board of Trustees

MARCH Founding Workshop in Adaptive Computation, chaired by John Holland and John Miller, funded by the Alfred P. Sloan Foundation

MARCH Michael Cohen and David Lane co-chair first SFI meeting focusing on Adaptive Processes and Organizations

MAY Melanie Mitchell is named Adaptive Computation program director

MAY Andi Sutherland named an Amiga de Santa Fe by the Santa Fe Chamber of Commerce in recognition for SFI's contribution to city tourism

JUNE Business Network for Complex Systems Research formed

JUNE Artificial Life III held in Santa Fe. Nearly 500 people attend

JUNE Chuck Stevens and Michael Stryker convene SFI's first Theoretical Neurobiology working group supported by The Pew Charitable Trusts

JULY "Refounding" meeting—Integrative Workshop: Common Principles of Complex Systems—chaired by George Cowan, held in Santa Fe

MIDYEAR Two trade books published about SFI—*Complexity: The Emerging Science at the Edge of Order and Chaos* by M. Mitchell Waldrop and *Complexity: Life at the Edge of Chaos* by Roger Lewin

OCTOBER First SFI conference on Auditory Display takes place, chaired by Gregory Kramer

OCTOBER Twentieth book published in SFI series, *The Principles of Organization in Organisms*, edited by J. E. Mittenthal and A.B. Baskin

93

(Annual budget, $3,292,500)

JANUARY Stuart Kauffman appointed SFI Professor

MARCH David Pines retires as Co-Chair of the Science Board and is replaced by John Holland

MARCH Susan Ballati joins SFI staff as Director of Development

APRIL SFI receives unanimous approval from Santa Fe City Zoning Commission to relocate to what will be its permanent campus at 1399 Hyde Park Road; decision is appealed to City Council. Same night SFI public lecture by Oliver Sacks draws 750 people

MAY George Cowan and Bela Julesz host first SFI workshop on Cortical Plasticity

JUNE Tradition of afternoon tea at SFI established

JUNE Seventh annual Complex Systems Summer School. Nearly 200 students apply for 50 slots

JUNE Bruce Abell presents SFI case and receives unanimous approval from Santa Fe City Council for SFI to locate to its permanent campus

AUGUST Marita Prandoni hired as SFI receptionist

OCTOBER Murray Gell-Mann appointed SFI Professor

OCTOBER Diane Lams hired as Research Assistant to Murray Gell-Mann

FALL First issue of new journal *Artificial Life* appears, edited by Chris Langton

94

(Annual budget, $3,900,000)

JANUARY W. Brian Arthur returns to SFI as Citibank Professor

MARCH David Liddle elected Chairman of the Board of Trustees; Robert Maxfield, Vice-Chairman

MARCH SFI signs contract with J. Wiley & Sons to establish the independent journal *Complexity: The Journal of Complex Adaptive Systems*, with first issue to appear in early 1995

MARCH SFI Postdoctoral Fellow Cris Moore elected to Santa Fe City Council

APRIL At his request Vice President Al Gore visits SFI for briefings on Institute research

APRIL Christopher Langton joins SFI research staff on a full-time basis as Director of the Swarm simulation program

APRIL SFI receives its first grant from the Advanced Research Projects Agency (ARPA) for work on the Swarm simulation system

JUNE John Holland talks about "Complexity Made Simple" in the first Stanislaw Ulam Memorial Lecture series. More than 500 people attend

JULY SFI moves to its permanent campus at 1399 Hyde Park Road

AUGUST Diversity Biotechnology Consortium formed by Stuart Kauffman, Alan Lapedes, Peter Schuster and others

OCTOBER Joleen Rocque-Frank hired as Research Secretary

OCTOBER Marylee Thompson hired as Publications Secretary

NOVEMBER Ed Knapp announces his retirement from SFI. Executive Committee of the Board of Trustees begins search for new SFI president

DECEMBER SFI acquires its first recreational facility, a ping pong table

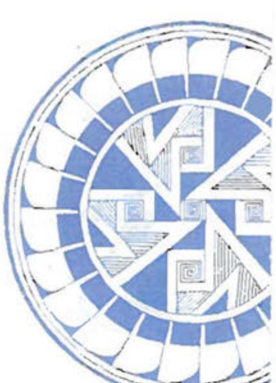

Diversity Biotechnology Consortium Searches for "Irrational" Cures

In a move marked by measured excitement and cautious optimism, the Santa Fe Institute, Los Alamos National Laboratory (LANL), Duke University and four companies have recently fostered the establishment of a non-profit organization dedicated specifically to the advancement of the burgeoning field of molecular diversity —the creation of trillions of entirely new molecules not already present in the biosphere today and the selection, cataloging, and manipulation of those that may serve useful purposes for us.

In August of this year, SFI biologist Stuart Kauffman and LANL and SFI physicist Alan Lapedes joined Mario Geysen, Director of Biotechnology at Glaxo, one of the largest pharmaceutical companies in the world, and six other top researchers from corporate, academic, and government arenas to found the Diversity Biotechnology Consortium. Their aim is to integrate theory with experiments in the rapidly expanding field of chemical and biological diversity. "To achieve this goal," says Geysen, "we are focusing very diverse resources and expertise at a scale that is much larger than any single institution or company could have provided."

RESEARCH UPDATES

The founders were among those who gathered at the Institute in April of this year for a conference entitled "Searching Sequence Space." That title refers to the theoretical and experimental exploration of sequences of amino acids, the building blocks of all living things, which researchers are now generating in phenomenal numbers in laboratories across the United States and elsewhere in the world. The researchers who attended concurred that they were actually poised at center field in "a whole new science," which many refer to as "the second generation of biotechnology."

The "first generation" of biotechnology involved the cloning of existing proteins and genes. The "second generation" involves the random generation of millions of diverse molecules, peptides, proteins, DNA and RNA, and then screening them for utility. "The medical importance," says Kauffman, with a gesture that sweeps the landscape of northern New Mexico, "is immense."

It doesn't take much math to see that only a fraction of possible molecules now exist or have, for that matter, ever been generated throughout evolutionary time. The number of possible proteins with 100 amino acids, for instance, is approximately 10^{130}, or 10^{70} times as many as could ever have been generated in the biosphere and 10^{117} times as many as the estimated ten trillion (10^{13}) in existence today. (By comparison, the number of hydrogen atoms in the universe is estimated to be 10^{60}.) "For this reason," explains Jane Richardson, "we need both theoretical frameworks that give us guidance about where and how we should look and efficient experimental techniques for generating and selecting among working 'libraries' that have billions or trillions of molecules."

For nearly a decade now, Geysen, Huse, Kauffman, Kay, Schuster, Stammer —See box— and others who have pioneered the sequence-library field have suggested that the science could enable us to design small molecules and even proteins to serve specific purposes. Since the chemical synthesis of more than 10 billion peptides by Geysen et al. a decade ago, experiment has begun to validate some of the most basic concepts in this rapidly-expanding field. Geysen, Huse, Kay, Kauffman, Keene and many others have developed libraries of billions of molecules which they have generated randomly, or "irrationally," and then screened, or "selected," for certain properties that suit their needs, like keys fit locks. David and Jane Richardson, on the other hand, have been designing proteins specifically to fold in specific ways, and testing them laboriously, one at a time, throughout this time. They are now embracing molecular-diversity techniques to permit bigger jumps into the "useful unknown." Furthermore, Jane Richardson, William Stemmer, and others have begun to modify desirable molecules to improve their function, their "fitness" to perform certain tasks, such as folding into a stable configuration or binding to the surface of a virus or cancer cell. Most recently, new laboratory techniques have resulted in an explosion of diversity in the test tube, in some cases rivaling the molecular diversity we witness today as the culmination of all evolution. Novel, massively-parallel screening techniques (developed in Göttingen and Jena, Germany) are based on single molecule detection and allow testing of millions of tiny probes in a short time, thus providing the tool for selection of the best molecular candidates for predefined purposes. At the same time, David Richardson, Alan Lapedes, Peter Schuster and others have proven to be masterful at using

DIVERSITY BIOTECHNOLOGY CONSORTIUM DIRECTORS

Mario Geysen
Director of Biotechnology
Glaxo, Inc.

William Huse
Founder and Chief Scientific Officer
Ixsys, Inc.

Stuart Kauffman
Santa Fe Institute

Brian Kay
Department of Biology
University of North Carolina

Jack Keene
Chairman of Microbiology
Duke University

Alan Lapedes
Complex Systems Group
Los Alamos National Laboratory &
External Faculty
Santa Fe Institute

David Richardson
Department of Biophysics
Duke University

Jane Richardson
Department of Biophysics
Duke University

John Rodwell
Vice President and Chief Scientist
Cytogen Corporation

Peter Schuster
University of Vienna &
Institute for Theoretical Chemistry

William Stemmer
Scientist
Affymax Research Institute

today's computing power to recognize patterns in the diversity we are generating daily. Because of this convergence, researchers have reason to believe that we are on the verge of comprehending natural mechanisms for generating diversity and natural patterns for selecting molecules.

"The word is evolution," says John Rodwell, Vice President and Chief Scientist at Cytogen Corporation. "We're looking at very efficient ways to harness in a very short time the forces that have operated throughout evolutionary history." William Huse, Chief Scientific Officer at Ixsys, Inc., agrees. "The study of diversity is not a 'stop-gap effort' that will be used only until a new, more-rational approach supersedes it. I believe the study of diversity is now and will always be the primary source of new information for drug design." The procedures, patterns and insights that pertain to this comparatively simple level of biological systems are almost inevitably instructive for a surprising range of disciplines, including immunology, neurobiology, electrical engineering, economics, and computer and social sciences.

While the refereed literature is the standard way of distributing information in the sciences, this traditional forum doesn't foster the intense collaboration and prompt coordination that the second-generation biotechnologists feel is called for at this pivotal time. The refereed literature tends to entrench a pattern of competitive science-by-individuals with separate territories, "science as a lot of small businesses," says Kauffman. For those who actually wish to put their heads together to jointly formulate a language for conceptualizing the territory, to hone in on key questions, and to develop effective protocols —and for those who foresee more business than any one researcher or company could possibly handle— the round table is the preferred forum.

Why have they come together at this particular time in our history? "Biology," Stemmer observes, "is generally viewed as a very experimental science because it is notoriously unpredictable. Unlike theoretical physics, theoretical biology has been much maligned in principle; it's extremely complex, and it's often unreadable.... If we can come up with simple, practically-applicable procedures —if we can illustrate the theory with just one practical example that works, that leads to a product— then we will have come a long way toward establishing the credibility of theory in this field."

The consortium is especially valuable as a means of agreeing upon and executing the less glamorous, systematic survey work. "The experimentalists, the theoreticians, and the pharmaceutical companies all want the results of that fundamental work," says Jane Richardson, "but no one has the resources, motivation, or persistence to do it alone. The fact that all these people are looking over our shoulders, eager to use the results, is motivating us to do the more laborious, statistically-correct experiments above and beyond the quick ones to answer our own immediate questions."

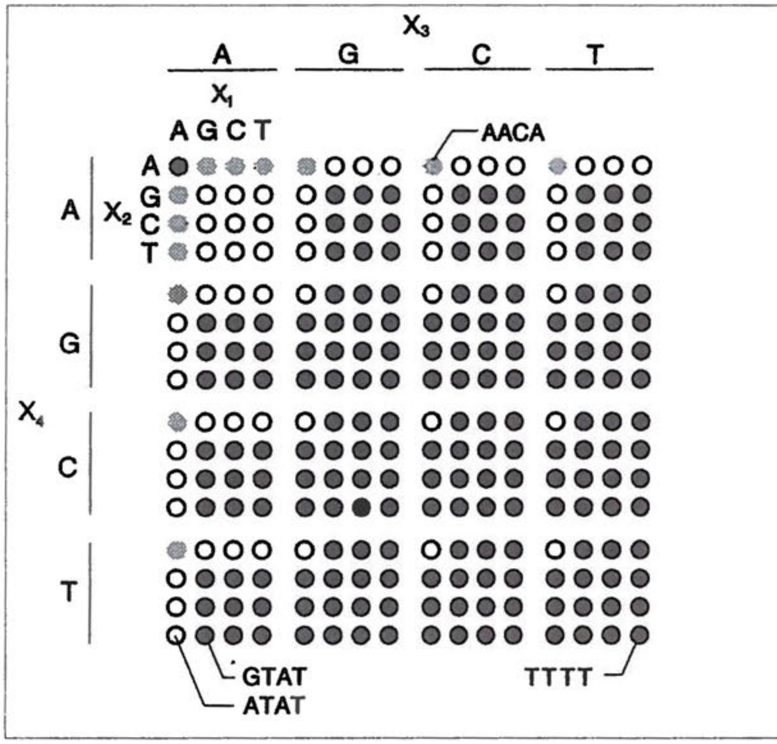

A sequence space is an ordered set containing all different arrangements of atoms or small molecules that can form a large molecule. Much of the Consortium's work focuses on finding new and efficient methods to search through huge sequence spaces of molecules for those molecules near-optimal in some property. This figure shows the sequence space of 4-base long, single-stranded DNA sequences. X_1 - X_4 represent the positions of the four bases, $[X_1 X_2 X_3 X_4]$. Each position contains one of the four DNA bases, A, G, C or T. Neighbor relationships between the molecules can be depicted in this space. For example, consider AAAA (upper-left dot, shown in dark grey) as a wild type sequence. All one-position mutants (e.g. AAGA)are shown as light grey dots, and all two-position mutants (e.g. TATA) are shown as empty circles. Two-dimensional sequence space representations can be used for any length and type of polymer. One way in which they can be used is to plot the fitnesses of all sequences (e.g. their affinities for a receptor) as heights in a third-dimension. This creates a fitness landscape over the sequence space in which good molecules are represented as peaks in the landscape. Figure and planar representation developed by Bennett Levitan, a post-doctoral fellow at SFI.

"There is finally a critical mass of people doing the science," says Kauffman. "At the same time, some aspects of the science are at a pre-competitive stage. That's where we can collaborate. The vast value of this work, after all, lies in the molecules we come up with rather than in the production procedures." Lapedes, too, hopes to see "the broad base of the technology remain public." The companies, in the words of John Rodwell, "recognize the opportunity to accelerate the drug-discovery process tremendously." He is certain, for instance, that "the practical application of the Diversity Consortium's theories is going to give us a whole new generation of pharmaceutical compounds for cancer therapy. It's the patients who are motivating me," he says.

The challenge that the founders have accepted by forming the Consortium is to grapple with the intellectual-property issues and advance the concepts and procedures so that, ultimately, as Kauffman says, "there's more pie to share."

"The synergism," reasons Stemmer, "far outweighs any competitive disadvantages." The Consortium's non-profit status does not in any way foreclose the opportunities for participants to spin off collaborative, for-profit ventures that might pursue certain avenues of research or product development. "In fact," observes Kauffman, "this is precisely what we hope to catalyze."

NICHOLAS METROPOLIS

At seventy-nine, physicist Nicholas Metropolis still goes daily to his office in the Theoretical Division of Los Alamos National Lab. Adjacent to the division is the Lab's center for supercomputing; behind it is the Center for Nonlinear Studies. The geometry seems apropos, for Metropolis, who has been both a strong influence on the evolution of computers at the Lab, and, in turn, computer simulation, was one of the initial visionaries of the Santa Fe Institute. It's not surprising that Metropolis, as one of the original Senior Fellows who founded SFI, thought computing should be a vital part of the Institute's program. While on paper computing itself isn't one of the Institute's mandates, it certainly has emerged, in the form, for example, of Christopher Langton's SWARM model and John Holland's ECHO, as a driving force.

A number of the initial conversations about the Institute were held in Metropolis' office. He remembers them as amorphous. "We thought Santa Fe would be a natural place for the institute," he says, "separated from the laboratory." He believes the institute has turned a corner and moved in a direction of greater generality, varied scope. Although the Institute hasn't had a large influence on his own work, he tries to go to some of the lectures and is proud to have been part of the Institute's inception.

FIRST 10 DECADE

Establishment of the Consortium in August was funded by a generous grant from Glaxo, Inc. Duke University, the University of North Carolina, and the Institute for Molecular Biotechnology (IMB) will be the centers for experimentation; IMB, LANL, and SFI will be the centers for computation. The pharmaceutical companies will advise the group on the most pressing questions which must be addressed in order to advance the applied science and join in the development of experiments and theory.

The Board of Directors of the Consortium is currently discussing exactly which experiments to pursue, how to allocate funds to coordinate theory with experiments, how to reach out to other companies and individuals, how to grow sensibly, and whether to produce a journal. By the first of the year, the Consortium aims to submit at least one grant proposal to DOE, which has administered the Human Genome Project since the early 1970s and the HIV Database since the beginning of this decade; DOE has already demonstrated strong support and encouragement for the current initiative.

"This collaboration is a spontaneous event," says Keene, chairman of Microbiology at Duke and head of an effort to found a center for Molecular Diversity there. "It's very exciting, very experimental, but we can't predict what will come of it. Maybe together we can get this science into the labs; we can get people using these technologies to find answers to important biological questions about growth and development, disease, how cells communicate, how gene expression is controlled, and much more."

This is not the first time that scientists have formed a consortium around an intellectual scientific concept, "but it is, to the best of our knowledge, the first in this area of biotechnology," says Kauffman, who currently serves as the Consortium's president. Rodwell adds, "We have the right people at the right time with a sense of trust among the individuals, a sense of enthusiasm in the field and an opportunity to advance and establish a brand-new science as expeditiously as possible."

"If, in a few years, the Consortium has grown," says Kauffman, "and if we've attracted financial support, done good science, published it, and enlarged the scope of the science as a whole, then we'll know we've succeeded."

Marty Peale is a technical writer and editor in Santa Fe.

RECONSIDERING THE "LONG TWILIGHT STRUGGLE"

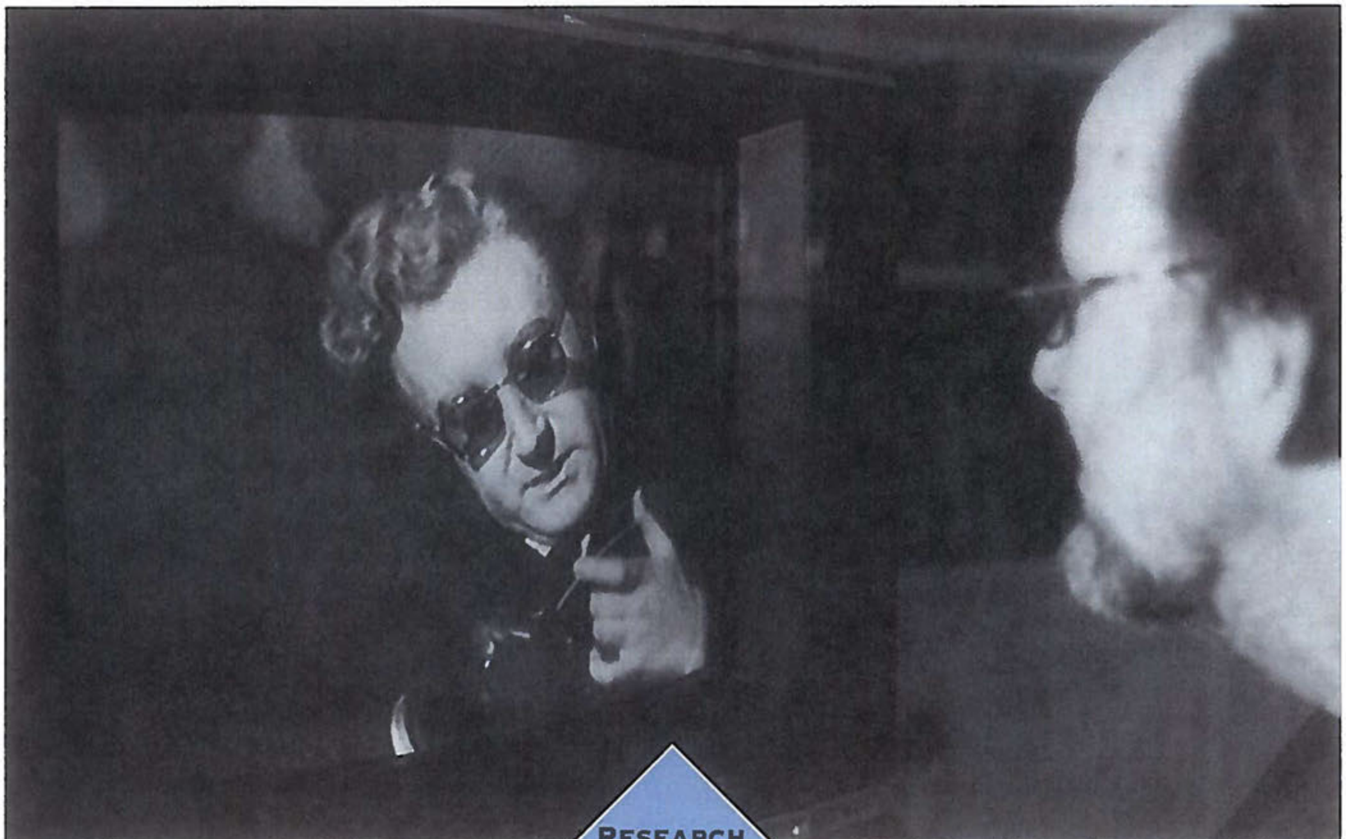

Looking back at Dr. Strangelove

RESEARCH UPDATES

Photo: Dan Barsotti

The SFI workshop,"Cold War Science and Technology," supported by the National Science Foundation, gathered economists, historians, sociologists, political scientists, and physical scientists to review existing scholarship on Cold War Science and to consider new research opportunities, now that the era is over. An updated "case-study" of this period is important in part because the Cold War offers a real-world example of cultural change — a phenomenon of particular interest to SFI's research agenda.

World War II and what JFK called the "long twilight struggle" that followed profoundly changed the course of science and technology in both the United States — a democracy— and the Soviet Union — a totalitarian state. In the United States, such institutions as the national weapons laboratories, military "think tanks," and defense contractors attracted a remarkable amount of the nation's intellectual and financial resources and were sustained for nearly half a century—almost a quarter of U.S. history. Many of these developments were draped in a veil of secrecy — contrary both to the tenets of democratic society and the traditional openness of science itself. American universities were transformed by the heavy infusion of research monies from the military, raising fundamental issues about the role of the state in higher education, scientific development, and industrial policy. For the Soviet Union, the Cold War spanned over half of that state's history and commanded a very large fraction of its scientific, technical, and economic resources. There, too, the Cold War reordered policies for economic development and scientific and technical research. In Europe, the Cold War manifested itself in ways that differed significantly from nation to nation within the more democratic NATO countries as well as the Warsaw Pact

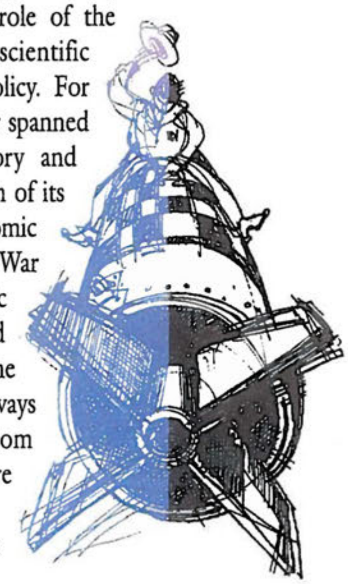

countries, and especially across the East-West divide. The Third World became contested territory for the competing ideologies of the Cold War, and both sides poured resources into the area to secure the allegiances of those nations.

The long duration of this competition meant that three generations of citizens on both sides of the "Iron Curtain" lived lives that were profoundly, yet often imperceptibly, conditioned by the Cold War environment. The collapse of the Soviet Union and the Eastern Bloc offered little more than a symbolic end to the Cold War; after all, many institutions, structures, patterns of thoughts, habits, cultural assumptions, and mentalities of the Cold War period remain.

This meeting focused on identifying researchable questions concerning science and technology in democracies and non-democracies during the Cold War and how science policy will likely change in the post-Cold War era. Although scholars from a variety of disciplines have conducted research on various topics surrounding this issue, it is only because of the end of the Cold War that it at last becomes possible to examine fully the interaction of science, technology, democracy, and national security and welfare during this era. It is now feasible to assess more objectively how the end of the Cold War could change the interaction among these factors. Of particular interest is how science and technology during the Cold War —and the policies associated with its pursuit— challenged, transformed, and strengthened democratic ideals at home and abroad.

Participants focused their attention on six areas: evolution of science and technology policy and research funding; production of knowledge during the Cold War; institutions of the Cold War; economic impact of the Cold War; comparative and international dimensions of the Cold War; and the Cold War and culture. Each section was introduced at the workshop by a scholar whose work was closely related to the area. These experts offered assessments of the state-of-the-art research and suggested important questions that should be part of an updated agenda. Small working groups then reflected on these presentations, formulated conclusions about what new knowledge should be pursued in each area, and posed specific, researchable questions.

Many of these questions involve cultural components —factors like the interaction of resources and technology; information and communication; and policy perspectives and social priorities — that are notoriously difficult to quantify and hence have often been neglected. Here, SFI might be productively involved. Within its organization and sustainability initiatives, for instance, SFI modelers are exploring how to build better scenarios incorporating political, cultural, and ideological systems. Their insights might be applied to this project. In turn, the Cold War experience itself offers a specific case study in which to explore evolutionary modeling from a social science perspective.

In the meantime the workshop findings will be incorporated into a report by program chair David Hounshell, Luce Professor of Technology and Social Change at Carnegie Mellon University, and his students David Jardini, Hugh Gorman, and Daniel Holbrook. This may become part of a National Science Foundation announcement of a new research initiative on *Science, Technology, and Democracy in the Cold War and After: Research Opportunities in the Transformation and Consolidation of Democracies.*

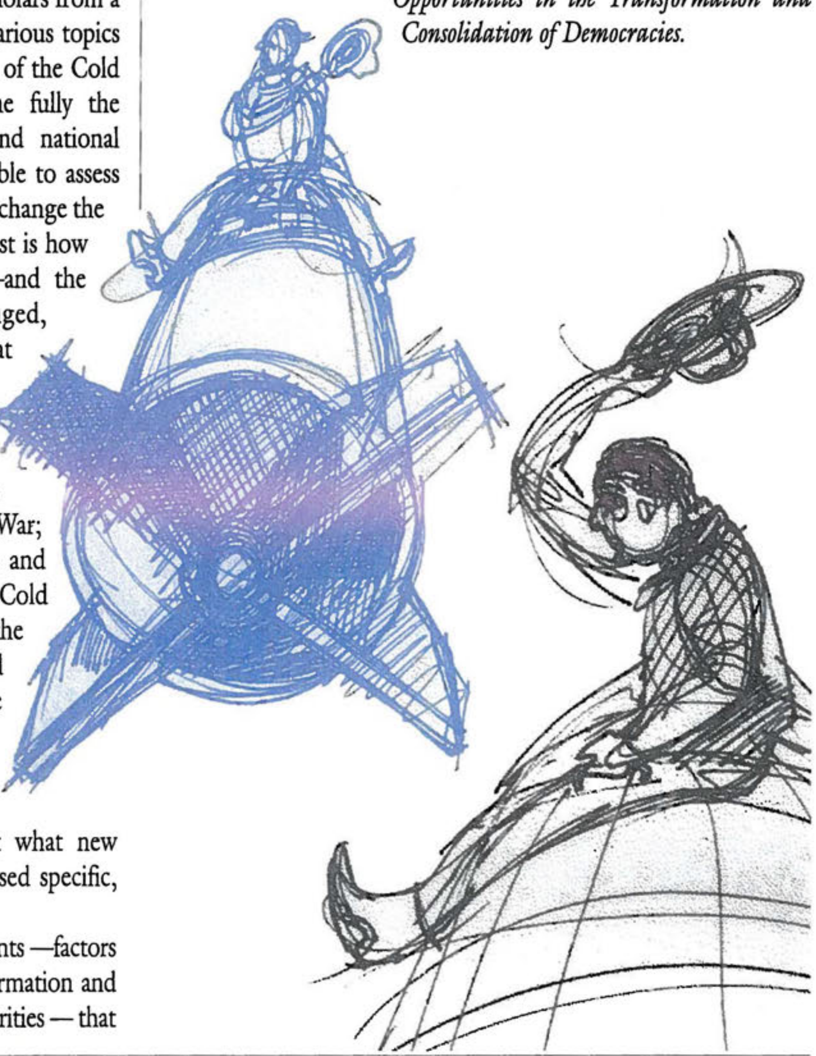

Deciphering the Complexities of the Brain

Understanding the complexities of the brain demands sophisticated approaches grounded in intradisciplinary collaboration between experimentalists and theoreticians. Surprisingly, such interactions are not as common as might be expected. With support from the Alfred P. Sloan Foundation and the Pew Charitable Trusts, SFI's Theoretical Neurobiology study group was formed in 1992 to spur these interactions. Last summer, Francis Crick, Cristof Koch, and Charles Stevens spent time at SFI as part of this effort.

Francis Crick

Francis Crick has commented that at the beginning of this century there were three "Big Questions": the nature of matter, the nature of life, and the nature of mind. With the tremendous advances in physics beginning in the first half of this century, and in biology in the second half —in no small part due to Crick's work, with James Watson, on the double helix—, we now believe we know at least the form of the answer to the first two of these questions. Certainly, many fundamental questions remain in both arenas; but most of the questions are well formed, and we have at least a rough idea of what the answers look like.

Research Updates

The answers have, of course, completely redefined the questions. Today when we ask, "What is life?," we are no longer asking whether there is some vital essence or humor imbued with the stuff of life. We understand that life is the result of a complex interaction of proteins, nucleic acids, lipids and a few other biochemical building blocks which, when combined in certain very special ways, behave in a qualitatively different way. We might reasonably ask, "Is a virus alive?" But we would recognize that this is purely a matter of semantics; the science would end with a description of viruses, their protein coats and their reproductive habits.

Yet as we near the end of the century, the nature of mind, of consciousness, remains as elusive as ever. It is not even clear exactly what the question is, or how to formulate it. What is mind? Where is consciousness? What is special about this three pounds of brain tissue, apparently no more different from the liver than the liver is from the lung? Can a machine be conscious? In spite of efforts by philosophers, psychologists and computer scientists, we know hardly more today than one hundred years

ago; our ideas today may well turn out to be the equivalent of humors circulating through the ether. The irony is that while this question is really at the core of neuroscience —-and is ultimately what draws many to the field—it is rarely broached in polite company by respectable neuroscientists. Each new generation is told that it is just too soon to go for the Big Question and that we must limit ourselves to the smaller questions that we can surely answer.

For more than five years, Francis Crick and Christof Koch have challenged this timidity. Crick, who, with James Watson, published the double-helical structure of DNA in 1952 (for which they were awarded the Nobel Prize in 1962) has been at the Salk Institute since its inception and for much of that time has focused on neurobiology. Koch is a Computational Neuroscientist at the California Institute of Technology. They collaborate throughout the year, but during the few weeks they spend at the Santa Fe Institute they can focus one-hundred percent of their attention on consciousness. The interdisciplinary atmosphere at the Institute provides the ideal environment for taking the bold leaps required in this work.

Crick and Koch have sought to construct an experimentally testable theory of consciousness. As a first step, they have focused on visual awareness, developing a framework for understanding just what goes on when we see. They hope within this framework that we can ultimately explain why, for example, our visual experience of driving on the freeway is so different from looking at a Rembrandt. On the freeway, our "mind" is barely aware of our visual experience, yet our "brain" is using the visual information to navigate through traffic. Viewing the Rembrandt, our mind is consumed with the image before us while our brain, it might be argued, isn't using the information for anything much at all.

What is the difference between these two cases? (To those who wonder what the difference is between "consciousness" and "awareness," Crick says, "I use the terms awareness and consciousness more or less interchangeably.... I must confess that in conversation I find I say 'consciousness' when I want to startle people and 'awareness' when I am trying not to.")

Their theory rests on the suggestion that there is some neuronal correlate of visual awareness. The advantages of focusing on the visual system are at least two-fold: First, a tremendous amount is known about its organization —from the molecular to the network level. Second, one of the main animal models of vision, the macaque monkey, is close enough to humans so that we can have a reasonable expectation that these monkeys actually possess visual awareness, and, thus, have a similar underlying mechanism.

Some neuronal subpopulation or group of neurons, Crick and Koch propose, behaves in a characteristic way when the animal is visually aware; this behavior would be a marker for visual awareness. This simple notion places their hypothesis squarely in the domain of experimental neuroscience, which is where they would like it to be. It establishes a research agenda: Find the special subpopulation of neurons, or their special mode of behavior that instantiates visual awareness.

The idea that the correlate will be observed in a definable subpopulation, but not in all neurons, is central. The visual system could be loosely defined to include much of the neocortex, the thalamus, the hippocampus and more —in short, a sizable fraction of the entire brain. If at any moment even one percent of neurons displays higher than background activity, then over a billion neurons could be involved. With this many active cells, "we would never be able to distinguish any particular event out of this vast sea of active nerve cells," says Crick. They posit, therefore, that "the majority of neurons will be involved in doing computations while only a much smaller number will express the results of these computations."

RESEARCH UPDATES

STIRLING COLGATE

As a member of the original Los Alamos lunch group, physicist Stirling Colgate was another Santa Fe Institute early convert and he served as a founding board member.

Colgate's interest in SFI was heightened by his desire to upgrade the general level of education in the state. He had been President of the New Mexico Institute of Mining and Technology from 1965 to 1974, and remained engrossed in the promotion of higher education. His interest as a physicist range from nuclear physics to epidemiology, and so the idea of an interdisciplinary entity certainly seemed appropriate to Colgate. He continues to attend seminars and contribute financially to the Institute.

He is, however, disconcerted with how the process of fund raising affects the Institute. "The process of doing science becomes a public curiosity that benefits raising money," Colgate says, "As a consequence it is hard to do science while on display." Still, Colgate has high hopes for SFI. "I hope we will understand the origin of life," he says, "the origin of the diversity of life on the planet, the origin of the human species and the selection mechanism that creates its evolving complexity and diversity."

FIRST 10 DECADE

One might hope to find a single cortical area in which activity corresponds to "what we see." But this does not appear to be the case. Rather, the visual information appears to be distributed across different cortical areas, each performing a different kind of visual processing. Since we are aware of the visual scene as a unified whole, there must be some way in which these different streams of processing are joined; how this is done is called the "binding problem." Crick and Koch have ideas about what kinds of activity to look for, and where in the brain to look for it. The activity signature they propose is the oscillatory and semisynchronous firing of a collection of neurons. This suggestion is motivated by the observation in the cat visual cortex of 40Hz oscillations. The oscillations are synchronized in distant (10 mm) regions of the cortex when the stimulus is a single object presented in the preferred orientation but are not synchronized when the stimulus is presented as two distinct objects. Firing synchrony, then, would be the mechanism of binding. As for where in the brain to look, since the neurons bound together by oscillations span many different areas, they posit that the "where" must be determined not by the location but by some other criterion. From knowledge of the basic microcircuitry of the cortex, they offer as a very simple first guess —the "large pyramidal cells in layer 5 that fire in bursts and project outside the cortical system"; but they comment that though "it would be marvelous if this were true, the answer is unlikely to be as simple as that."

This approach avoids, by its nature, many of the pitfalls of previous attempts to explain awareness. For example, our first intuition about visual awareness often leads us to posit a homunculus, a little replica of ourself somewhere in our brain that watches, as though on a screen, the images arriving on our retina (or, in more sophisticated versions, in our higher cortical areas). The gaze of the homunculus then provides an excellent correlation with visual attention: we are aware of whatever the homunculus is looking at. Of course, this has bought us nothing at all since we are confronted with infinite regression: How is the homunculus aware? Although in its simple form the homunculus is easy enough to avoid, it nevertheless often sneaks subtly through even the most vigilant defenses of other classes of theory. In the neuronal reductionist approach, there can be no homunculus since the explanation is entirely in terms of well-defined, physical entities (neurons).

Of course, this current approach is just a first step. After all, how could knowing the activity of even all the neurons in my brain tell you anything about what I experience when I see a Rembrandt? This question of qualia is the same one

RESEARCH UPDATES

MURRAY GELL-MANN

Particle physicist Murray Gell-Mann did nothing less than discover the "quarks," the subatomic building blocks that make up protons and neutrons. In 1969 he won a Nobel Prize for his earlier work in the theory of elementary particles, and although the field in which he made his mark is somewhat specialized, it's his interest in culture as a whole that attracted him to the idea of the Santa Fe Institute. He became involved with the original "lunch group" while working as a consultant for Los Alamos National Lab and very quickly became a primary intellectual force behind the establishment of the Institute.

He is one of the few founding members to take advantage of the Institute itself to delve further into the subject of complexity. In fact the book he published earlier this year, *The Quark and the Jaguar: Adventures in the Simple and the Complex* (W.H. Freeman & Co., New York), deals at length with the issues of the "simple and the complex" and makes clear his affection for the work done at the Santa Fe Institute. Last year he retired from the California Institute of Technology, where he was the Robert Andrews Millikan Professor of Theoretical Physics; he is now one of the few long-term faculty at the Santa Fe Institute.

Gell-Mann is a hard man to please. He has strong convictions about how things should be done, so it's a tribute to the Institute as a whole that Gell-Mann has decided to take up full-time residency, even though things are not exactly how he used to envision them. It was Gell-Mann who wanted the Institute to be a larger and broader organization. That his original vision was never realized seems inconsequential to him now. He regards it as only natural that the idea of the Institute has evolved, along with its research agenda.

"In the beginning, we couldn't see clearly what sorts of emerging scientific syntheses we should seek," he says. "Now I think we can." Gell-Mann believes the intellectual mandate of the Institute should be to "understand the various meanings of simplicity and complexity and how they are related, to learn about the ways in which complexity emerges from simplicity, and especially to study complex adaptive systems —those that learn or adapt or evolve as living things do." Like a number of the founding members of SFI, he would like to see a broader scope of study that would encompass more of the social and behavioral sciences, as well as some fields of the humanities. He is particularly concerned with psychology, anthropology, and history.

He is also worried about the physical space. "The Institute now has wonderful places for people to interact," he says and shortly follows with a call for more capital. "We need to build more small offices where the researchers can go off and think."

FIRST 10 DECADE

EDWARD KNAPP

Edward Knapp is one of the unfortunate scientists at the Institute whose job isn't to do science, but to get the funding and create an environment where others can do science. "It's very exciting to see people get involved in new ways of thinking," he says. "In some ways, I wish I were doing it." However, he points out that his background is in experimental, not theoretical physics. "There are no soldering irons at the Institute," he says.

Knapp is a founding member and was the first Vice-Chairman of the Board of the Santa Fe Institute. He has been President since 1991. He was founding leader of the Accelerator Technology Division of Los Alamos National Lab, Director of the National Science Foundation, and President of the Universities Research Association. His background qualifies him to make statements such as, "If we're successful here, the whole direction of science might change in the next five to ten years."

As SFI's current president, his aim is to provide a culture where people can interact with one another and learn to think in new, untried ways. For Knapp, the Institute works because the problems worked on are not dictated by the administration, but conceived by the researchers themselves. Through his eyes, the collaboration is indicative of the Institute as a legacy from Los Alamos and the founding members of SFI. "The way that Los Alamos differed from universities was that research was team-based," he says. "That was very effective for certain problems." According to Knapp, it appeared to the founding members that there were a number of problems not being looked at in the university system because they required the team approach. "One of the things that we talked about was that we wanted to make an organization where there were no departmental walls," he says.

Knapp looks forward to the day when the Institute has the funding to actively recruit researchers. "I would like to see us broaden our family of people so that, in turn, we have a broader cross-section of science," he says. "Get people who are leaders in fields which maybe we don't touch very much now." Knapp would like to see the SFI research family work on problems that are good, basic science, yet impact people's lives and society. He believes the Institute is in a perfect position to do just that. "The way we've evolved, we're filling a niche which is unique," he says. "I don't think there is anybody else close to doing what we do."

FIRST 10 DECADE

philosophers usually pose as, "How do I know that when I see red it is the same as what you see when you see red?" These are the really hard questions which Crick and Koch are just beginning to address. Perhaps, if their research agenda pays off, this question will be answered in the same sense as whether a virus is alive.

RESEARCH UPDATES

CHUCK STEVENS

Chuck Stevens of the Salk Institute would like to know how Nature wires up the nervous system. It's not an easy question. An analog problem exists in the design of silicon chips, that is, the precise layout of circuit elements to allow efficient wiring. On modern, complex chips, a substantial fraction of the total landscape is devoted to wires. But while even the most advanced chips (like the Pentium) today have only a few million circuit elements, the human cortex has over 10^{10} elements (neurons). And since a typical pyramidal neuron in the cortex receives more than 10^4 axonal connections from other neurons, efficient wiring is critical.

For example, if each neuron were equally likely to be connected to neurons at any distance, then the $(10^{10})(10^4) = 10^{14}$ axons, each having a volume of at least 10^{-8} ml/cm, and each traveling on the average 10 cm (about the radius of the cortex), would yield about $(10^{15}$ cm$)(10^{-8}$ ml/cm$)$ $= 10^7$ ml—1000 kg! Clearly, if we want to be able to fit our brain in our head, we have to be careful about how we connect our neurons. The circuit designer working in wetware would find that this system differs in at least two important respects from silicon hardware. First, in the nervous system, wiring is set down in three dimensions rather than in just two, as on a chip. On the one hand this is an advantage since it means that not too much care need be taken to keep wires from crossing. At a radius of $< 1\mu m$, the axons can be thought of as one-dimensional structures in three-dimensional space. On the other hand, since the dendritic trees that are the axon's targets are themselves nearly one-dimensional, making sure that the axons ever make contact with their target can be tricky.

The second difference is that a chip works only if precisely the right components are connected —every element on my Pentium is wired the same as on yours. Indeed some simple biological systems seem to adopt the same strategy; in the invertebrate *C. elegans*, the entire wiring pattern —the connectivity of every neuron to every other neuron— is specified in the genome, and as a result the nervous system of every *C. elegans* is identical to that of

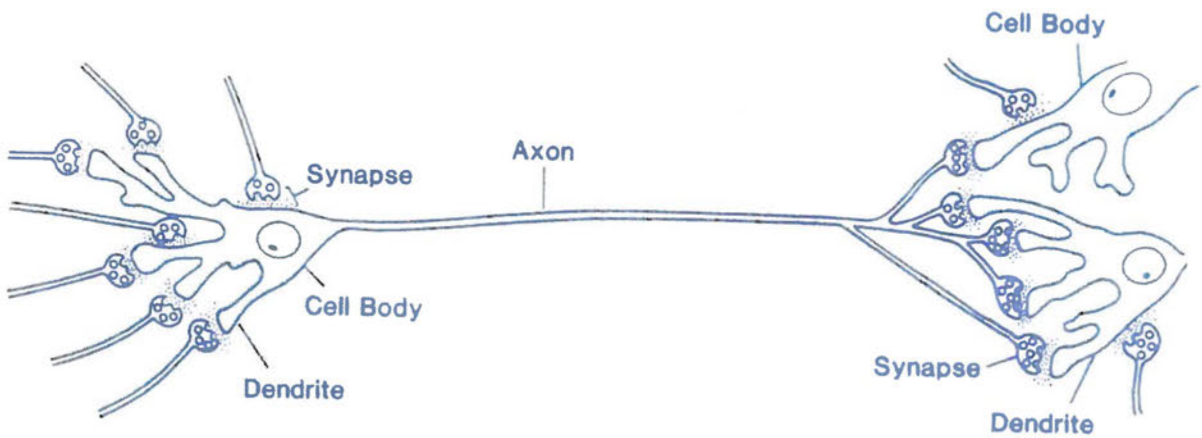

every other. However, this strategy does not scale well to a brain as complex as ours. There are just too many neurons and not enough DNA. In fact, since our entire genome has only about 6×10^9 nucleotides, there are fewer nucleotides in our genome than there are neurons in our brain!

RESEARCH UPDATES

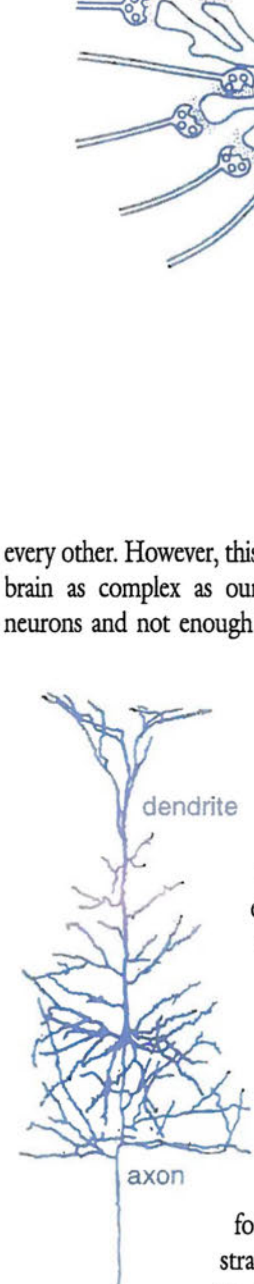

Unlike *C. elegans*, there simply isn't enough space in our genome to specify the entire connectivity pattern explicitly. Since the genome can't spell out the neuron-by-neuron connectivity, it must specify wiring indirectly, through a set of (presumably local) strategies that allow axons and dendrites to interact. These strategies would determine how axons meander in their search for dendritic targets, and whether they form a synapse when they find one. The strategies would have to be simple enough to fit comfortably in the genome and would have to address the issue of how nearly-one-dimensional structures can find each other in three-dimensional space.

Stevens thinks the fractal structure of dendritic trees may provide a clue to Nature's wiring strategy. Dendritic trees are composed of thin processes, some tapering to less than 0.5 µm, extending out from the cell body as far as 1 mm. Axons course through this field of dendrites. Both axons and dendrites behave almost like one-dimensional objects. Stevens observes that if a branched structure like a dendritic tree is fractal, then it can "fill" a higher dimensional space. The degree to which it completely fills the space is given by its fractal dimension. Any "traditional" one-dimensional object, like a line or the perimeter of a circle, has a dimension of one; and any simple, two-dimensional object has a dimension of two. But a highly branched objects can have a fractional dimension between one and two. If a dendritic tree has a high fractal dimension, then it can behave more like a space-filling, two-dimensional structure and can reach out and "sample" axons that meander through.

Stevens is developing experimental tests of this theory. The first step is to measure the fractal dimension of dendritic trees. For some neuronal types such as retinal ganglion cells, this has been estimated at about 1.7; but the fractal dimension of cortical neurons remains to be determined. Ultimately though, the utility of this approach depends on the extent to which the interactions between real neuronal processes obey such simple dynamics.

phenomenon—say ant foraging—has one learned anything which can be brought to bear on other models of the same phenomenon?

Many-To-One

The answer to this question is complicated by the fact that many different sets of agent local rules might produce the same kind of global output. That is, the mapping from micro-rules to macro-structures may be many-to-one. In such circumstances, one may encounter quite different models, all having similar behavior—all producing caricatures of ant foraging, say. When this happens one may be skeptical that anything cumulative has been achieved, although some intuition about the multiplicity of this mapping has probably been produced. However, if some of the models perform at a higher level than others then one has learned something cumulative. For instance, say that two competing ant foraging models have the agents following pheromone gradients, but one model supplements the agents with simple memory. If both models produce interesting caricatures of overall colony behavior in the presence of multiple food sources but only the model with agent memory gets the distribution of agents around the sources qualitatively right, then one reasonably feels that something has been learned about the importance of memory in such models.

Let a Thousand Artificial Flowers Bloom!

To-date it is perhaps fair to say that most work with agent-based models has taken place at the lower levels, mostly due to the novelty of the modeling approach. However, particular models often display elements of several levels. For example, in our own Sugarscape model we have devoted a significant fraction of the overall development effort to creating specialized analytical routines for the purpose of assessing the performance of our model at Levels 1 and 2. We have found it useful to not only plot the distribution of wealth and note that it is highly skewed, as in real societies (Level 1 agreement), but to also compute Lorenz curves and Gini coefficients in search of quantitative (Level 2) agreement with empirical data. Analysis routines comprise somewhere between 1/3 to 1/2 of the entire code, a ratio that has stayed more or less constant as additional behavioral modes have been added to the model. Most recently, in thinking about Level 3 performance— looking at individual agents — we have concluded that nothing short of a full real-time database engine, capable of performing cross-sectional and longitudinal analysis of the agent population, will suffice. The necessity of tying the simulation system to such a database will certainly drive the amount of code dedicated to analysis to well beyond 50% of the total.

Attempts to build models meeting the "caricature" standard could powerfully advance certain fields. For example, since Thucydides, the study of international relations has centered on the macrophenomena of arms races, alliances, and wars. It would be healthy for political scientists to try to "grow" these phenomena. Those purporting to know *why* the international system looks as it does might attempt to specify the *rules* they think the *agents* (states) are executing, put them on a computer, and see if those agents and rules in fact generate a world that looks more or less recognizable. It is at this lowest, "caricature," level that agent-based models are usefully thought of as software laboratories, in which alternative rule systems can be quickly and qualitatively studied.

However, it is also true that in many scientific fields there are probably few problems remaining for which significant progress can be made with simulations which are not developed past the level of "caricature." There is simply too much research which has already come; too much is already known. Yet at the next higher level, Level 1, there may be many problems which could profit from agent-based modeling. For example, an artificial ecology model in which food webs emerge having the same qualitative structure (connectivity, hierarchy) as real food webs could help ecologists to understand, say, how such structures evolve in response to environmental fluctuations. In some fields there may be very few problems left for which Level 1 models would provide important new information. Economics is a highly quantitative discipline and any agent-based artificial economy must be equipped with tools for quantitatively analyzing its performance at Level 2 or above.

Appropriate objectives will, of course, vary with the state of a field, available hardware and software, modeling resources, and so on.

DARRAGH NAGLE

Physicist Darragh Nagle lives in a modest home in the quiet neighborhood in the Los Alamos suburb of White Rock. Outside, children are walking home from the nearby school, kicking leaves and dangling books. Inside, Nagle is talking about the genesis of the Santa Fe Institute, astrophysics, and Los Alamos. Darragh Nagle is another one of the Senior Fellows who conceived of SFI and currently serves on the Board of Trustees. He is delighted with the way the Institute is going.

Nagle recounts the first few meetings of the Senior Fellows. "The question was, 'What new initiative could we as the Senior Fellows take that would lead to a new science?'" he says. "In the initial discussions people were thinking supercomputers because at that time parallel computers were a hot idea. The Laboratory had not yet embarked on this." Nagle explains that most of the Senior Fellows had been around the lab enough to remember the days when there was a lot of discretionary money and people could undertake special projects. In fact, at Los Alamos Nagle had co-invented, along with Santa Fe Institute President Edward Knapp, the side-coupled accelerator.

Nagle believes the Institute has evolved in ways the Fellows hadn't predicted, but he remains close to the idea of studying emerging synthesis. "The Institute extends its influence by not being a large organization," he says. "I think it should foster similar organizations in other countries, building a worldwide network of Santa Fe Institutes." His own work reaches even farther as he oversees the installation of a telescope to look at gamma rays from distant stars. Being semi-retired doesn't seem to daunt him. Of the telescope he says smiling, "I have told them to put in a wheelchair ramp."

10 FIRST DECADE

Software Problems and Solutions Unique to Agent-Based Models

Now, when we speak of "understanding our creations," we must include our computer programs themselves: "Is my program doing what I want it to?" Although one can rarely have more than a probabilistic answer to this crucial question, there are both important pitfalls and powerful diagnostic tools unique to agent-based simulation, which we need to appreciate if we are to have confidence in our results.

Software "bugs" can have special characteristics in agent-based models. We often focus on emergent macro-structures in our models. However, certain pathologies in the agent population due to "bugs" may not be revealed by the macro-structures. For example, imagine that some agents have their internal data over-written occasionally due to an array indexing error in some (seemingly) unrelated procedure. An agent whose wealth is actually 100 might have her wealth field over-written to 0. Since 0 is not an impossible value it goes unnoticed in building up, say, a wealth histogram. To make matters worse, imagine that memory management by the operating system is moving the agents around in memory so that the agent being modified by the "bug" is different each time the "buggy" procedure is called. Such software problems are difficult to discover due precisely to the highly distributed nature of agent-based models. Indeed, the "robustness" of macro-structures to perturbations in individual agent performance— "graceful degradation," to borrow a phrase from neural networks—is often a property of agent-based models and exacerbates the problem of detecting "bugs." By contrast, it is also possible that agent-based models can display sensitive dependence on agent initial conditions. For example, in our Sugarscape model we have studied one set of runs in which particular agents were removed at the start of each run, yielding very different societal evolutions. Some agents matter a lot—they are critical or keystone agents—while others are not so important. When a model is known to have such sensitive dependence then one should be particularly wary of the possible existence of "bugs," since these could produce very large changes in the output of the model. Interestingly, the agent-based modeling approach offers novel ways to systematically address questions of software validity.

Data Gathering Agents!

By mimicking the way human societies gather data about themselves—namely, by sending specialized data-gathering agents out to observe regular (non-data-gathering) agents—it is possible to uncover software bugs. Such agents can be part of the regular agent population, executing the same rules the regular agents execute, but having their rule systems supplemented by data-gathering activities. They track-down bugs by finding and reporting anomalous events—agents with unusual or rapidly-changing internal states, "forbidden" behavioral modes, catatonia, and so on. These agents cannot guarantee software validity but, if they are given enough clock cycles, they can go a long way toward increasing a programmer's confidence that his or her program is operating correctly. Incidentally, the use of data-gathering agents may allow us to systematically study how the presence of observers distorts that which is under observation, an enduring problem for field anthropologists and other social scientists!

Desirability of Analytically-Tractable Special Cases

Modelers in the agent-based tradition are often most interested in the transient, non-equilibrium, non-stationary behavior of their creations. However, it is clearly very desirable, given the complexity of

Richard Slansky

At the time of the initial discussions that led to the founding of the Santa Fe Institute, physicist Richard Slanksy was not a Senior Fellow, but no one held that against him. "They knew I was interested in interdisciplinary problems," he says, "so I was invited to sit in." It was fortuitous that someone thought to ask because he was instrumental in getting SFI going and now, a decade later, Slansky is the Director of the Theoretical Division ("T" Division, as it's called by insiders) at Los Alamos National Laboratory.

There has always been overlap between the Theoretical Division, which, among other things, serves as a parent for the Center for Nonlinear Studies, and the Santa Fe Institute due to the interdisciplinary nature of both institutions. Slansky traces this overlap to the beginning. "Our main goal was to try to have a place where working between disciplines was not only acceptable but encouraged," Slansky says. "That's what the Lab had been doing for some time, but there were a number of subjects that we felt the Lab wasn't covering as well as it might." Chief among those subjects were issues having to do with complexity, or how complex behavior emerged from simple systems. "One reason we got on that theme is because there really wasn't a place where that work was being done," Slanksy says.

Slanksy helped organize the founding workshops of the Institute and currently serves on the Science Board. However, his duties as Director of the Theoretical Division don't allow him to spend much time at SFI. Still, he's strongly involved in an indirect way. "We have a number of people in the Theoretical Division who are very involved with the Santa Fe Institute," he says. "These people interact at the Institute and interact at the Lab. I think that broadening of scientific interaction is really important. The most important thing I'm doing right now is to make sure that they feel comfortable spending a certain amount of time at the Santa Fe Institute."

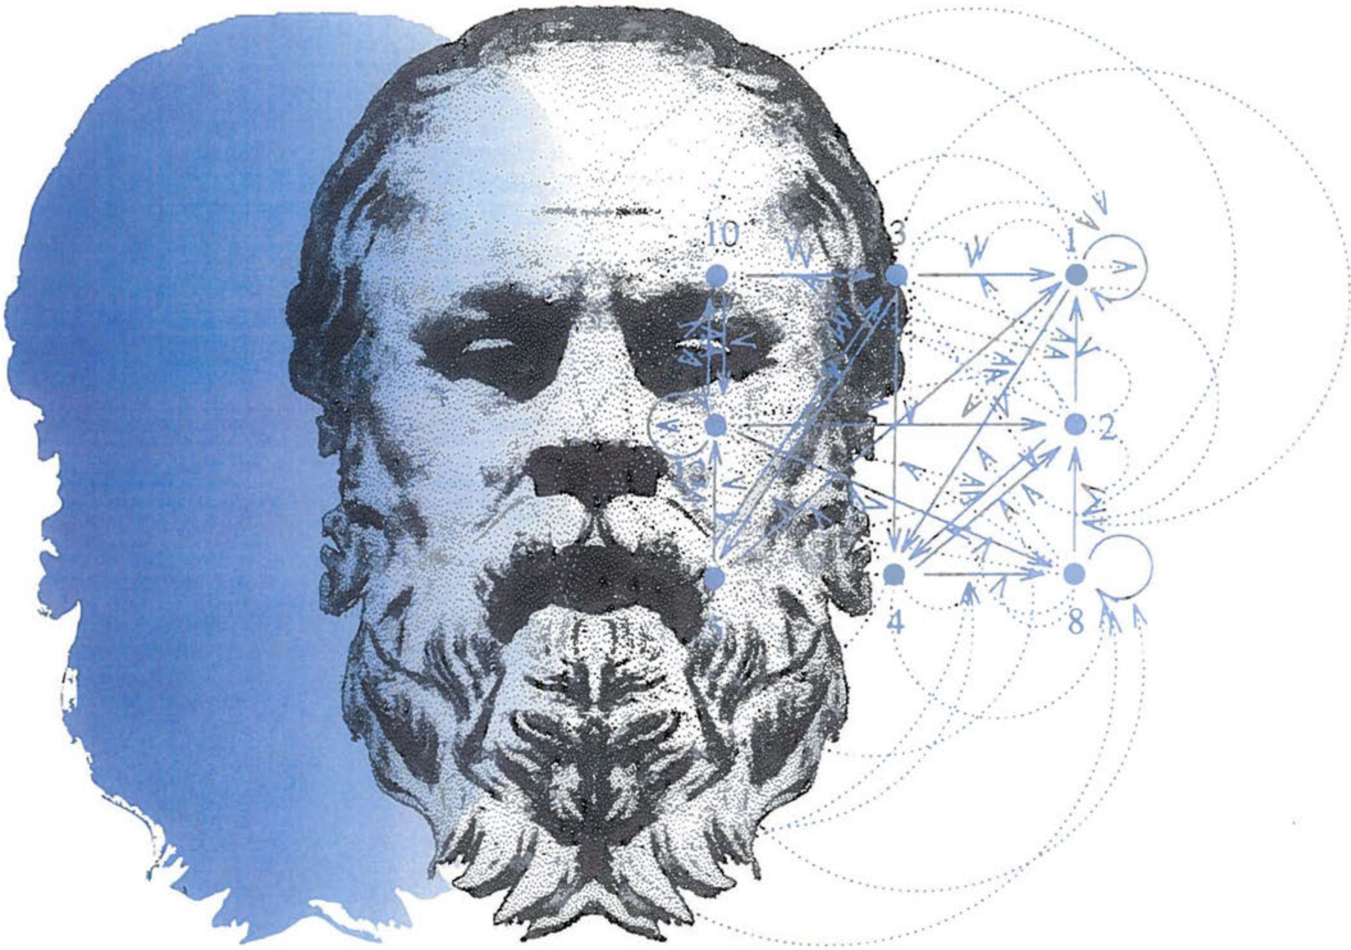

agent-based software, to be able to construct special cases of the model which are analytically tractable. For example, imagine an artificial economy model in which agents are in all respects neoclassical, except that preferences change due to cultural transmission. And suppose prices are observed to vary erratically. To study whether this nonequilibrium price behavior is caused by the changing preferences one could merely "turn off" the cultural transmission behavioral mode, yielding a model with fixed preferences and—dare we say it?—equilibrium behavior, for which there are known analytical results in the literature. It may also be possible to develop new analytical results. Ultimately, as has occurred in other sciences, interesting behavior may inspire the development of entirely new formal methods for analyzing agent-based models. After all, for confidence that we "understand our creations," there's nothing like outright proofs.

Conclusion

In summary, there are many ways to go about understanding our complex agent-based creations. Our own Sugarscape work certainly reflects the high variance of the field as a whole: different areas of the research are at different levels. We have some caricatures, some qualitative agreements with macroscopic observables, some quantitative agreements, a few glimmers of hope that certain of the rules our agents follow might actually be close to real human behaviors, and a body of formal mathematical work applicable to some special cases. But, we have tried to move beyond mere simulation and have begun to develop tools for examining the specimens we grow. For, as Socrates would surely have said, "The unexamined ALife is not worth living."

SFI's horizons are as wide as those of the nearby peaks —and its intellectual fresh air as refreshing.

— Nils Nilsson, *Stanford University*

Newcomers to the field of complexity often have unrealistic expectations. The SFI provides both a sanity check and a set of realistic expectations for my applications.

— Bill Fulkerson, *Deere & Company*

From the people I have met at SFI, the combinations of disciplines, and the free-ranging conversations, I have taken away a different attitude towards my own work. My approach to the study of evolution has evolved at ten times the rate that would have been possible without my participation in SFI. SFI stands for the Scientific Fun Institute.

—Marcus Feldman, *Stanford University*

As an Earth scientist, I know of many significant questions about our planet. But overarching them all is one single question—easy to state, hard to answer, and almost never even asked: "Why is the Earth organized in complex ways?" It is a Santa Fe kind of question. So when I can find time, I come to the Santa Fe Institute to walk among the red, rocky, complex landscape where the Sangre de Cristo Mountains meet the Rio Grande Rift, to talk with friends for whom this is a natural question, and to seek to understand this greatest of geological mysteries.

—Walter Alvarez, *University of California at Berkeley*

Science, Santa Fe-style, with its emphasis on decentralized, transdisciplinary investigations and bottom-up computer simulations of complex, adaptive systems, offers the best possible preview I can envision of what scientists will be doing —and how they'll be doing it—in the 21st century.

In my pre-SFI life, I thought that the existence of an institution devoted exclusively to the furtherance of genuine research on complex systems was about as likely to happen as the attainment of world peace —merely a nice dream. But SFI has proved that sometimes dreams do come true.

—John Casti, Editor, *Complexity*, Santa Fe Institute

FIRST 10 DECADE

For me, SFI is a wonderful, stimulating environment that is ideal for creative work. Nothing gets in the way of putting everything into what you are doing.

— Chuck Stevens, *Salk Institute*

WINTER 1994
Volume 9 • Number 2

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

The Santa Fe Institute is a private, independent, multidisciplinary research and education center, founded in 1984. Since its founding SFI has devoted itself to creating a new kind of scientific research community, pursuing emerging synthesis in science. Operating as a small, visiting institution, SFI seeks to catalyze new collaborative, multidisciplinary research, to break down the barriers between the traditional disciplines, to spread its ideas and methodologies to other institutions, and to encourage the practical applications of its results.

Published by the Santa Fe Institute
1399 Hyde Park Road
Santa Fe, New Mexico 87501
U.S.A.
phone: (505) 984-8800
fax: (505) 982-0565

Illustrations for "Deciphering the Complexities of the Brain" courtesy Francis Crick.
Illustrations for "Alife IV" courtesy of MIT Press.

L.M. Simmons, Jr., *Editor-in-Chief*
G. Richardson, *Managing Editor*
P. McFarlin, *Designer*

Some may consider the research done at Santa Fe rather esoteric and therefore not directly applicable to their lives. But we've used many of the ideas developed there to introduce millions of kids (and adults) to complex systems through our computer games.

— Will Wright, *MAXIS*

Simply... it's exciting! The Santa Fe Institute addresses the emerging, emergent mysteries of the 21st century: economics and markets, game theory, biotech, artificial life, neurobiology, immunology, the evolution of natural language . . . . And it's a model of what it studies: how a simple set of interacting elements — or scientists— can interact to form something exciting, unpredictable and beautiful as it grows and flowers.

—Ester Dyson, *EDventure Holdings Inc.*

SFI is unique because of its focus on interdisciplinary cooperation. Insights in one body of knowledge can advance progress in others. Learning something about complex systems at SFI has convinced me that law is also a complex system which can be analyzed and studied in ways similar to complex systems in other fields.

—Wallace R. Baker, *Baker & McKenzie*

In the past decade, SFI has formed a catalytic core of innovative thought leaders who collectively have transformed traditional, linear ways we think about ourselves and the world. One early down-to-earth application is TeleSim, a management training-oriented "flight simulator" that uses adaptive agent technology to illustrate increasing return dynamics of the information and communication industry. But watch out! To paraphrase Cervantes, "Thou hast seen nothing yet."

—Win Farrell, *Principal, Coopers & Lybrand Consulting*

During the "Evolution and Organization of Society in the Prehistoric Southwest," one of the woman participants who was attending with her new husband said, "I have never been so stimulated in my entire life." Another participant noted in a loud stage whisper that was really quite amazing since she had only been married three weeks.

—George Gumerman, *Santa Fe Institute*

The Santa Fe concepts of complexity and chaos have upset the thinking of many economists, including myself. The mysterious dynamics of the economy at least resonates with the many parallels with physics, biology, and chemistry.

— Ken Arrow, *Stanford University*

FIRST 10 DECADE

SFI is more than a research institute; it is a state of mind that takes over on the road from Albuquerque. From then on the cerebral cortex is kept oscillating by dazzling displays of intellectual give and take. What makes it work is an extraordinarily dedicated staff willing to put up with a choir of tenors. I am most grateful for the opportunity to have added my cacophony to the song.

—Harold Morowitz, *George Mason University*

Santa Fe Institute
1399 Hyde Park Road
Santa Fe, New Mexico 87501
U.S.A.