Toward Artificial Life

BY CHRIS LANGTON

Artificial life was an ideal subject for a landmark conference. Here was a new subject lurking in old disciplines, and to see it you had to walk to a crossroads so far out of anyone's current field that it gave a refreshing view of the road ahead. Chris Langton was the ideal organizer. Academically correct (Center for Nonlinear Studies, Los Alamos National Laboratory), eclectically broad-minded (every conceivable approach was represented), and heading it up for the right reason (no one else would stage the conference he had been urging on colleagues for years), Chris was also superbly organized. No less than thirty talks and twenty demonstrations hinging on temperamental equipment, presented in five non-stop days and nights, all without a hitch. [1]

One of the most surprising developments I discovered at this conference was the frequent mention of God. His grand demo running outside looked better by the hour as various artificial life demos struggled inside. I found that camping at night in the crystalline clarity of Bandelier National Monument, New Mexico, a 10-minute ride down from the Oppenheimer Center in Los Alamos, was the perfect counterbalance to the workshop. The hardy, inexhaustible complexity of rustling grass, drifting stars, and hooting owls kept me skeptical of, and impressed with, the fragile life cuddled in the rooms next day. —*Kevin Kelly*

PERHAPS THE MOST intriguing thing about life is that it is a property of the *organization* of matter, rather than a property of matter itself. It is one of those wonderfully mysterious phenomena wherein the whole is more than the sum of its constituent parts: life "emerges" out of the interactions of a great many non-living molecules.

There is no special "vitality" brought to a living system by any of its ingredients. The vitality of living systems depends on the set of *functional relationships* that develop between biomolecules, not on the specific material out of which those biomolecules are constructed. If one could replace the biomolecules of a living system with other entities that engaged in a similar set of functional relationships, the resulting system would exhibit similar vitality. Thus, life is a *process*, one that obeys its own "bio-logic," and as such, should be able to be "lifted out" of the particular physical details of its molecular "wetware."

The rapid increase in our knowledge of the inner workings of living systems, together with the increasingly powerful computational resources at our disposal, will soon give us the capability to create processes obeying very similar "bio-logics" within computers, or in some other medium. When we do this, we will have created "Artificial Life."

Entries In The First Artificial Life 4-H Show

photos and captions by KEVIN KELLY

In September 1987, the first workshop on Artificial Life was held at the Los Alamos National Laboratory. Jointly sponsored by the Center for Nonlinear Studies, the Santa Fe Institute, and Apple Computer Inc., the workshop brought together 160 computer scientists, biologists, physicists, anthropologists, and other assorted -ists, all of whom shared a common interest in the simulation and synthesis of living systems. During five intense days, we saw a wide variety of models of living systems, including mathematical models for the origin of life, self-reproducing automata, computer programs using the mechanisms of Darwinian evolution to produce co-adapted ecosystems, and much more.

Throughout the workshop, there was a growing sense of excitement and camaraderie, perhaps even profound relief, as previously isolated research efforts were opened up to each other for the first time. It quickly became apparent that, despite the isolation, we had all experienced a remarkably similar set of problems, frustrations, successes, doubts, and visions. What follows is a summary of some of the common themes that emerged at the workshop.

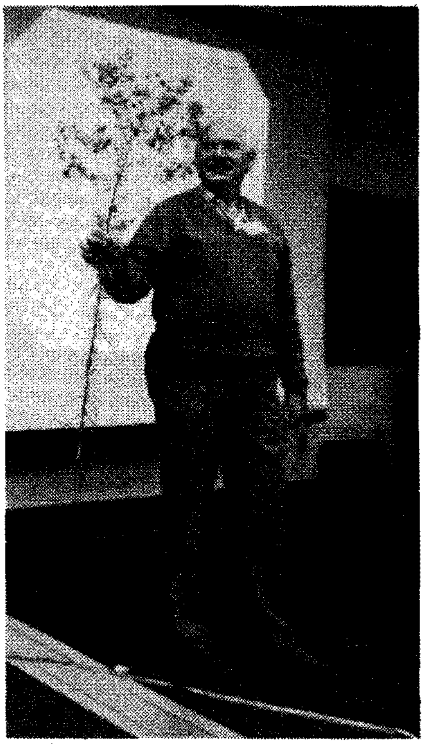

Lessons of the unreal. Dutch mathematician and biologist Aristid Lindenmayer (left) waves a fall aster plant he pulled up from the parking lot perimeter. Lindenmayer is one of the grandfathers of biological mathematics — tracing the mathematical patterns in natural growth. Using computers primed with very simple rules, he has reconstructed the complex growth of wildflowers. He determined that exactly three distinct signals traveling up and down a plant stem will produce nearly all observable budding patterns. Interestingly, although there is an extraordinary visual match between real blossom sequences and artificial ones, there have been no botanical chemical signals discovered yet.

The dance of leaf growth and blossoms opening and fading in ivy-leaved wild lettuce (Mycelis muralis) is governed by "two signals and accumulated delay" in Lindenmayer's color computer graphic display (far left).

A whole meadow of artificial life sprouts on the display screen. The flowers were not "drawn." Seeds of numbers were planted in electronic memory, and their colliding calculations painted the garden patch.

First of all, we saw immediately that the proper way to generate lifelike behavior is from the bottom up, rather than from the top down. The most lifelike behaviors demonstrated were generated by systems that consisted of a set of relatively simple entities, each with its own behavioral repertoire. The behavior of the system as a whole was the result of the aggregate of the local, rule-governed interactions between these simple entities. Nowhere in the system were there rules for the behavior at the global level. The net behavior of the system was entirely emergent, supported on the shoulders, so to speak, of the myriad local, rule-governed interactions.

By contrast, the behaviors that were generated by systems based on top-down specifications tended to be rigid, inflexible, and quite un-lifelike. Top-down systems supply global rules for global behavior. Low-level entities must be moved around to conform to the desired global change of state. These systems must inevitably be very complicated, for they must try to capture, in global-level rules, the results of all of the nonlinear local interactions taking place among the low-level entities. This is not only difficult to do, but probably impossible in the general case. Many results, especially from automata theory and the theory of chaos (chaotic dynamical systems), indicate that our ability to predict the results of nonlinear interactions is limited not only in practice, but also in principle.

If there is an artificial-life equivalent to AI's Turing test, it amounts to the statement "I'll recognize life when I see it." Many of the bottom-up models passed this test, and were met with spontaneous applause at the conference. Few, if any, of the top-down models elicited such a response.

Another common theme to emerge at the workshop was the recognition that it is very easy to underestimate the complexity of environmental interactions. Most models, even the bottom-up ones, provided extremely simple environments with pre-specified responses, and clear-cut boundaries between the environments and the "living" systems they nurtured. Environments were often specified top-down, even when the primary actors in the model were specified bottom-up. In nature, it is often extremely difficult to draw such sharp distinctions between the living-system and its environment, and interactions with the environment are often as complicated as interactions within the living-system.

This became especially apparent in models of evolving systems. Rigid, prespecified, unnatural environments foster rigid, predictable, unlifelike evolutionary progression. Systems adapting within a model where the environment itself is specified only at the low-level, in a bottom-up fashion, have much greater potential for demonstrating genuine evolutionary progression. Thus, it was recognized that the "fitness function," the set of criteria that determines whether an organism is "fit" in its environment, must itself be an emergent property of the system.

A third common theme was an increased appreciation for the behavioral complexity that can be exhibited by even

The power of one gene can be seen in the botanical work of Przemyslaw Prusinkiewicz. Prusinkiewicz, working at the University of Regina in Canada, won the Blue Ribbon prize at the first annual Artificial Life 4-H Show for his colorful garden of artificial flowers grown in a computer. His plants had the individual dignity and distinction you find in real plants — each sample of a species looks similar but individually different. The laws of their growth are complex simplicity. A few principles, governed by a few numbers, develop this complex artificial plant (far left). The same formula, with only one single number accidentally altered late one evening, produced this radically transfigured mutation (left).

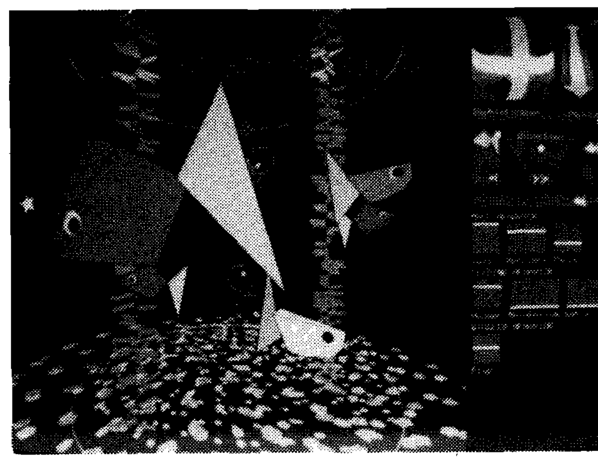

Learning how to school. Peter Broadwell of the Media Lab's Vivarium project had a story to tell about the fishes in his "Fishbowl." He designed the two different-colored fishes in his computer aquarium to swim round and round in an invisible glass bowl. The fishes would eat others of a different color, grow larger, mate to produce offspring of the same color, and die after a certain duration of time. He could alter the rates by tweaking the parameters on the side of the screen. Usually the aquarium would stabilize to a half dozen adult fishes, as shown here. Once, at a computer graphics show, he set the machine up as a visual soother in a room where computer artists were resting. During the evening when he was gone, they fiddled with the parameters and left it on overnight. The next morning he came in to see unanticipated evolution: sixty very tiny fish, all of one species, crammed into the bowl like sardines. They were swimming round in circles as a school, a behavior he had never designed into the system.

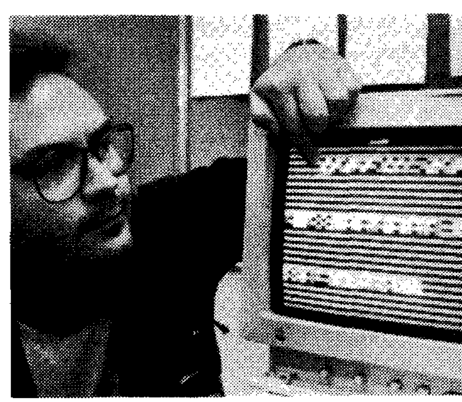

Software parasites, like this worm, are one of the earliest forms of artificial life. James Hauser points to a worm of his creation which prowls through the memory core of an Apple IIe. The worm is actually a software program coded to zip through all the sections of the computer's chips while producing a visual record of its journey. It appears on the screen as a snake of multi-colored segments. It will go round and round endlessly, until the power is killed.

Hauser and his partner Bill Buckley decided to see what would happen if you let two worms loose into one computer — and the worms could "eat" each other. That was the first battle in an ongoing championship called "Core Wars." The object is to write a simple worm program that can replicate itself faster than the other worm program can eat it. The one alive at the end wins. Some of the winning programs have a chromosome consisting of a mere four lines of code. Longer genes can't execute as fast as short ones, so they tend to get weeded out. Nicknames of current parasites like Dwarf, Locust, Mice and Imp indicate the sneakiness an organism needs to survive in Core Wars. [2]

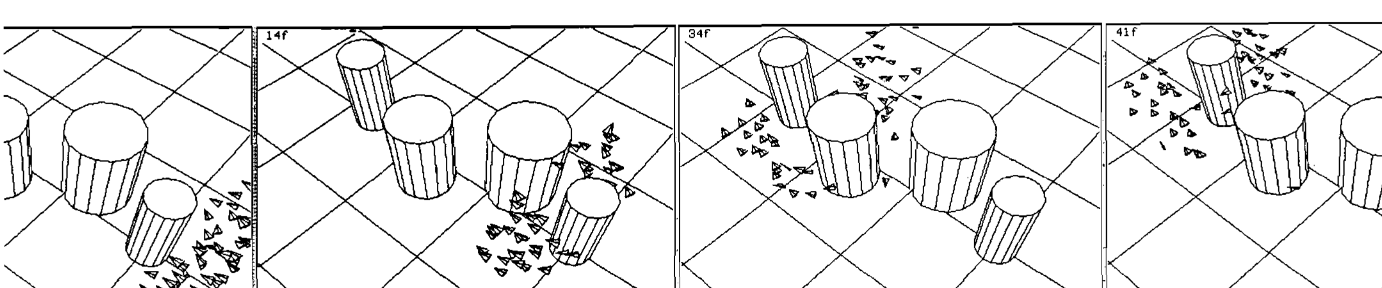

the simplest machinery. During the transition from the industrial era to the computer era, our notion of a machine has changed radically. We have come to believe that the essence of a mechanism — the "ghost in the machine" — the "thing" that is responsible for its dynamic behavior, is not a thing at all, but an abstract control structure, or program. Furthermore, we recognize that the essential features of this control structure can be captured within an abstract set of rules — a formal specification — without regard to the material out of which the machine is constructed. We have learned to separate the logical form of a machine from the material of its construction, and have found that "machineness" is a property of the former, not of the latter.

Once we have separated the "form" from the "matter" of machines, it becomes relatively simple to give formal specifications for a wide variety of machines that we would probably never commit to hardware, and to experiment with their dynamics. When we do this, it becomes apparent that *extremely* complicated behavior can arise in "machines" governed by extremely simple rules. One of the surprising results of recent work in nonlinear dynamics is that complex behavior need not have complex roots: even mechanisms that are governed by very simple *deterministic* rules can generate behavior that is extremely complicated and difficult — even impossible — to predict.

Thus, rather than degrading life by reducing it in rank to the equivalent of the machines of our everyday experience — such as toasters, dishwashers, and automobiles — we have increased our appreciation of what a machine can be to the point that we now believe that behavior as complex as life itself is achievable by machines.

Finally, there was the sobering realization that, on the scale of geological time, we are at the close of a major stage in the history of the evolution of life on Earth, and at the beginning of another.

With the discovery of the structure of DNA and the interpretation of the genetic code, a feedback loop stretching from molecules to men and back again has finally closed. In biological terms, a human being is the physical result (phenotype) of the interpretation of its genetic information (genotype) in the context of a specific environment. The process of biological evolution throughout the last 3.5 billion years has, in us, yielded a genotype that codes for a phenotype capable of manipulating its own genotype directly: copying it, altering it — or replacing it altogether in the case of artificial life.

There remain many, many issues that must be addressed in the pursuit of Artificial Life. By the middle of this century, mankind had acquired the power to extinguish life on Earth. By the end of the century, he will be able to create it. Of the two, it is hard to say which places the larger burden of responsibility on our shoulders. The future effects of changes we make now are, in principle, unpredictable — we cannot foresee all of the possible consequences of the kinds of alterations we are now capable of inflicting on the fabric of inheritance. Yet if we make changes, we are responsible for the consequences. How can we justify our manipulations? How can we take it upon ourselves to create life, even within the artificial domain of computers, and then snuff it out again by halting the program or pulling the plug? What right to existence does a physical process acquire when it is a "living process," whatever the medium in which it occurs? Why should these rights accrue only to processes with one particular material constitution and not another? Whether or not "correct" answers exist to such questions,

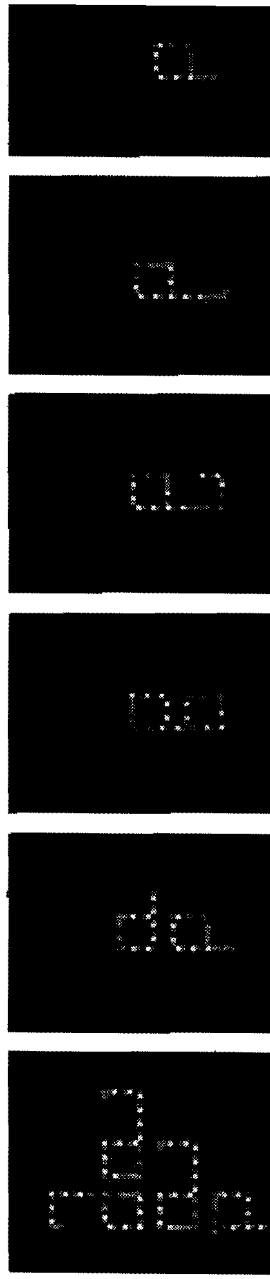

A machine no bigger than its information. It looks like a duplicating "Q," but Chris Langton, the creator, says it is the smallest self-reproducing manmade structure. Self-reproducing devices were conceived and outlined by Von Neumann a generation ago. A representation of Von Neumann's Universal Machine would take up a grid several hundred thousand units wide, still bigger than anybody's computer screen. Langton's universal machine runs in an environment of a hundred units or so. Only eight signals govern his device, versus dozens for Von Neumann's.

If the message held in a "Q" can generate another "Q" whose message is "Q," then you have a Universal Machine — a thing whose self-contained information will make more of itself. Langton's "Q" begins as a square patch which sends out signals to make adjoining squares. Secondary signals are sent out which interact with previously sent signals to determine where and whether another square should be built. The signals are ingeniously designed by Langton to keep extending the machine. He was able to pack into a loop of that tiny size the information that will create a loop of similar size, a trick the earliest specks of living matter somehow managed.

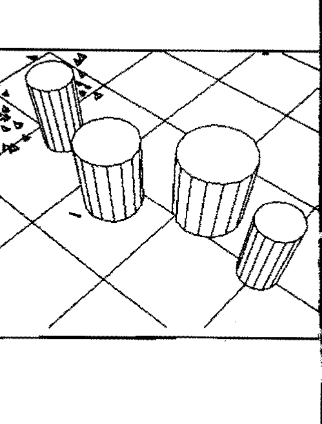

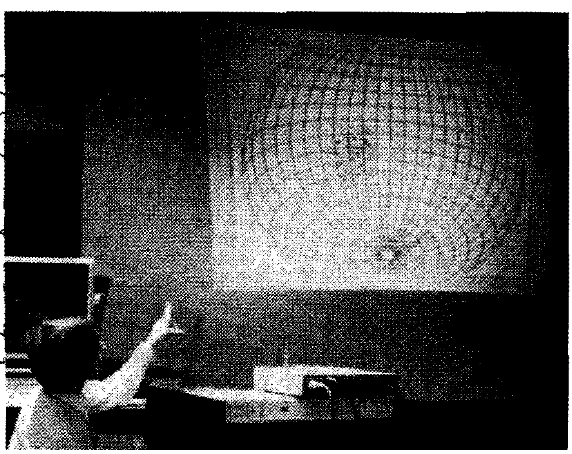

Emergent behavior was THE keyword at the conference. Craig Reynolds (below right) of Symbolics, Inc., a high-powered graphic computer developer, points out the flocking behavior of winged creatures (called boids) in a film sketch for the color video *Breaking the Ice*. The black-and-white line drawings in the clips are later rendered in color and in volume for the final version.

The flight of individual boids is not pre-calculated. Each boid is set flying with only a few instructions: look out for obstacles and don't bump into your neighbor, but don't stray too far away either. Everything else that happens is "emergent" — not pre-planned, not fixed, and not expected. The boids fly as a flock on a preordained route, yet each boid can do what it wants, and does. In one trial episode (left), a flock of boids divides to fly around a pillar. One boid conks into the pillar, flutters momentarily, then straggles behind. Nobody ever plotted that.

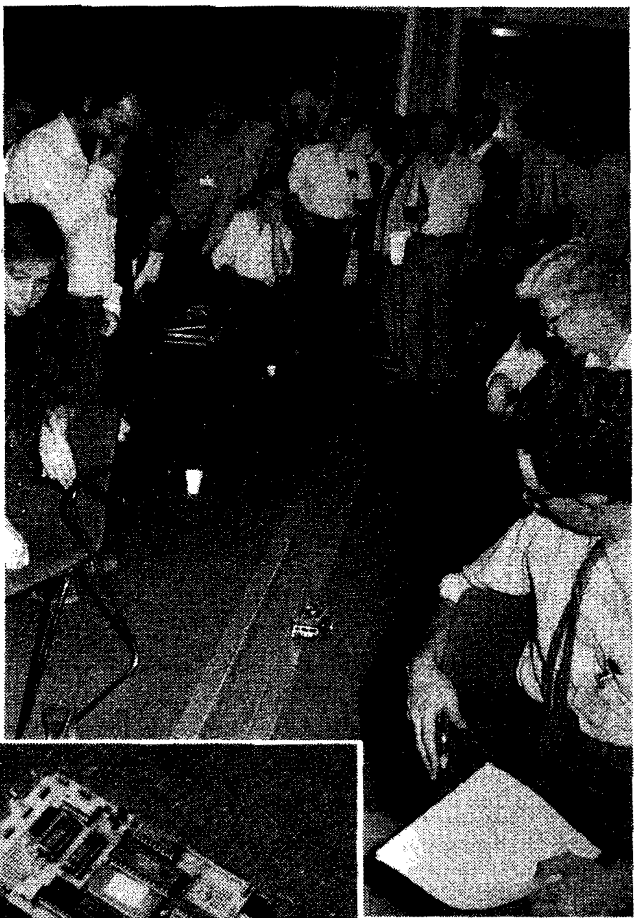

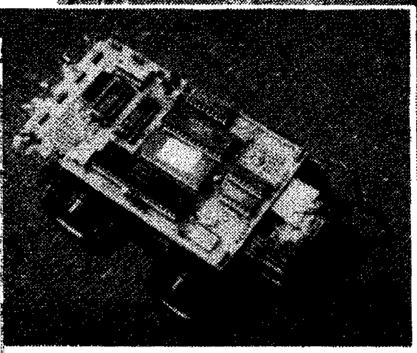

Light mice (left), with a myopic attraction to a flashlight. Eight tiny photo sensors, off-the-shelf parts for a toy dune buggy, and a homemade microchip make up the critters. The mice scamper around a room and run toward the light source. This seemingly simple behavior is astoundingly difficult to program; it takes a small on-board computer to figure it. Counting the computer, the creature has the brain equivalent of an earthworm.

The mice were built by John Wharton as a part-time hack. Their conceptual ancestors were the thought experiments of Valentino Braitenberg at MIT, who imagined an ecology of various species of wheeled vehicles reacting to each other. These simple machines would steer by the direction of light and shadows they cast upon each other, thus functioning as a mechanical environmental selection. His ideas, which arise from neurophysiology, are superbly outlined in a thin, influential book called *Vehicles: Experiments in Synthetic Psychology*. [3]

The Holy Grail of desktop genetic engineering is in Richard Dawkins's (author of *The Selfish Gene* and *The Extended Phenotype*) addictive software program, "The Blind Watchmaker." This Macintosh program breeds creatures by asexual genesis. It produces offspring with slight to severe changes from the original. You select which of those offspring you prefer, and let the program breed it again. In a couple of generations you have a critter you could have never imagined. The mutation rate can be adjusted, as well as 15 other genes which control the image, such as height, scale, segmentation, and branching. Echoing nature, the genes can be set with gradients, or turned off and on by other genes.

You can start with a tiny stick and begin breeding that, or as Dawkins put it, "you can put the program on genetic drift, and when you see a nice one, you can go for a little breed." Human intervention is allowed by genetic-engineering mode; you alter the image on the screen by manipulating it with an icon of a hypodermic needle. Genealogy of your work is easy to look up. You draw out the pedigree from the fossil record in chart form.

None of the forms found in Dawkins's albums were preconceived. Each one was a surprise. "I'm looking for a system that is pregnant with evolution," he said. He spoke of the maniacal drive to explore this world which he had created — awake late at night, nervous with anticipation, as he would sleeplessly breed creatures till morning. Among his trophies is this page (bottom left): a collection of animals vaguely resembling those of the Echinoderm phylum (sea urchins, etc.). Another page displays insectoids. He calls his inhabitants biomorphs, and their domain Biomorph Land. Buried deep in a remote corner of the Land, Dawkins discovered a tiny jewel figure, an image of the Holy Grail. Its genetic formula is "lost." Dawkins has offered a prize of $1,000 to the first person who can dictate the biomorph gene code that will exactly match the bit-map picture of the Holy Grail. [4]

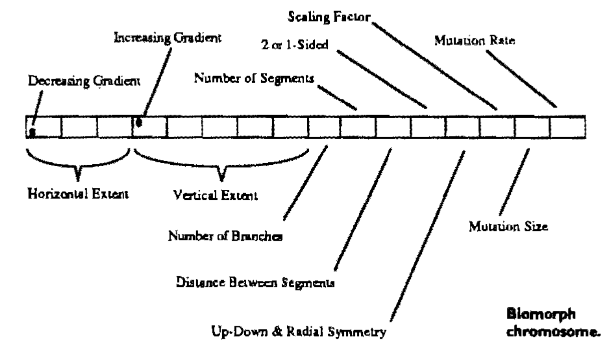

Biomorph chromosome.

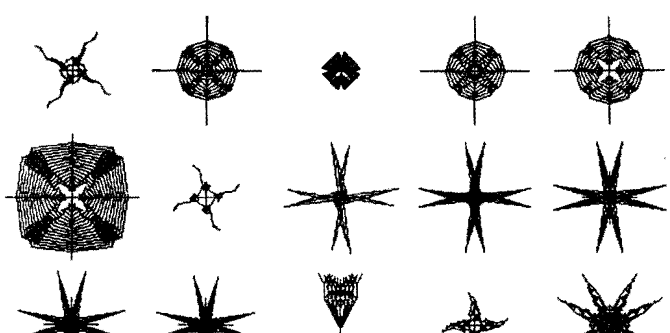

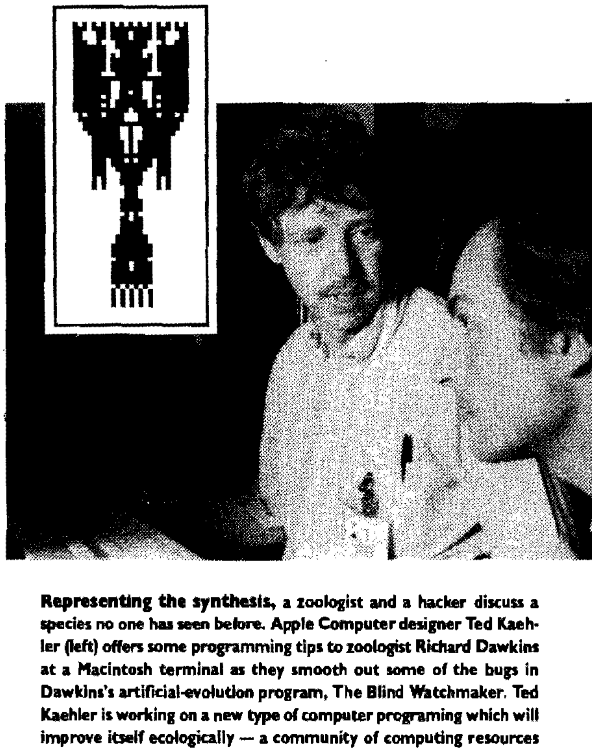

Representing the synthesis, a zoologist and a hacker discuss a species no one has seen before. Apple Computer designer Ted Kaehler (left) offers some programming tips to zoologist Richard Dawkins at a Macintosh terminal as they smooth out some of the bugs in Dawkins's artificial-evolution program, The Blind Watchmaker. Ted Kaehler is working on a new type of computer programing which will improve itself ecologically — a community of computing resources which compete to find an answer.

they must be addressed honestly, and openly.

Artificial Life is more than a scientific endeavor, it is a challenge to our most fundamental social, moral, philosophical, religious, and even cosmological beliefs. Like the Copernican revolution, Artificial Life will force us to re-examine our place in the universe and our role in nature. ☐

RESOURCES

1. Queries for future Artificial Life Conferences and published proceedings from the first one should contact Chris Langton at the Center for Nonlinear Studies, MS B258, Los Alamos National Laboratory, Los Alamos, New Mexico 87545; 505/667-1444.

2. News of current battles, upcoming contests, and technical tips for Core Wars is published in *The Core War Newsletter*, edited by William R. Buckley. Published quarterly by AMRAN, 5712 Kern Drive, Huntington Beach, CA 92649-4535. The International Core War Society is located at 8619 Wassell, Wichita, KS 67210-1934.

3. *Vehicles: Experiments in Synthetic Psychology:* Valentino Braitenberg, 1984. **$6.95** ($8.45 postpaid) from MIT Press, Cambridge, MA 02142; 617/253-2884.

4. *The Blind Watchmaker:* Richard Dawkins, 1986; 332 pp. **$7.95** postpaid (order # 3448) from W.W. Norton, 500 5th Avenue, New York, NY 10110; 212/345-5500. The Blind Watchmaker software program is available for $10.95 postpaid with a coupon from the paperback's appendix, which also constitutes the program's manual. It requires a Fat Macintosh or larger to run. Entries for the Holy Grail search should be mailed to W.W. Norton and Co.

5. There's a two-month waiting list for Cellular Automata Machines, version 6 (CAM 6) add-on boards which slip into IBM PC clones (PC, XT, and AT). They go for $1,500. Call or write Systems Concepts, 55 Francisco St., San Francisco, CA 94133; 415/985-1000. ■

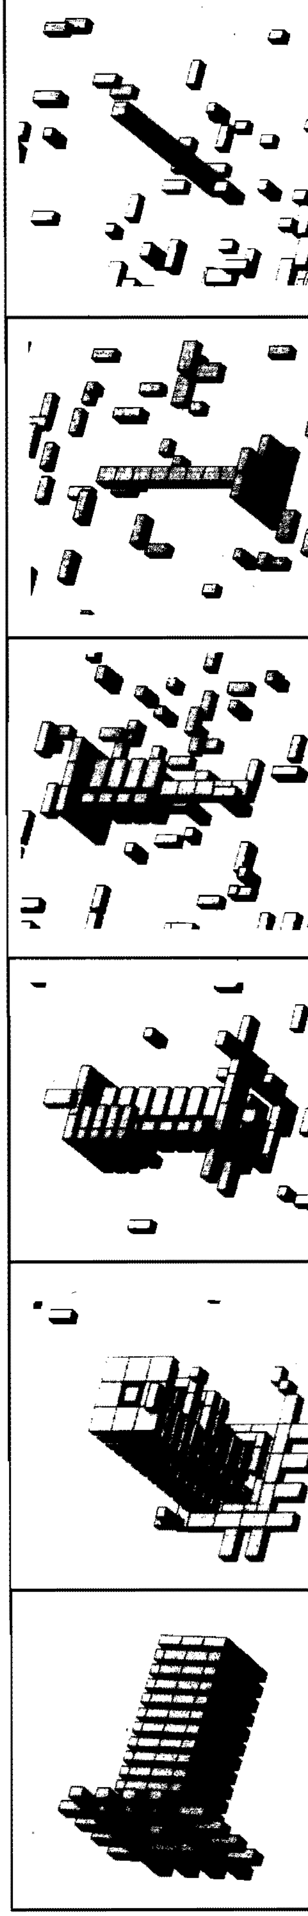

Self-organizing virus. Start with various building blocks (proteins in real life) that can bind to other blocks only if several sides of the block bind at once. In other words, all the attractive sites on a block must be joined at once in order for any of them to be joined. This is called "configuration bonding" in technical literature. Proteins follow this pattern as they combine into the complexity of a virus. Can you make a virus by designing elementary blocks so that they self-assemble by their own attractions when you put all the pieces into a bag and shake them?

Narendra Goel, at SUNY in Binghamton, NY, designed a computer bacteriophage that would assemble itself by the configurational bonding energy of its pieces. Once built, the virus (a simulation of a T4 bacteriophage), would attach its bottom flange against a simulated protein membrane and puncture it, as in real life.

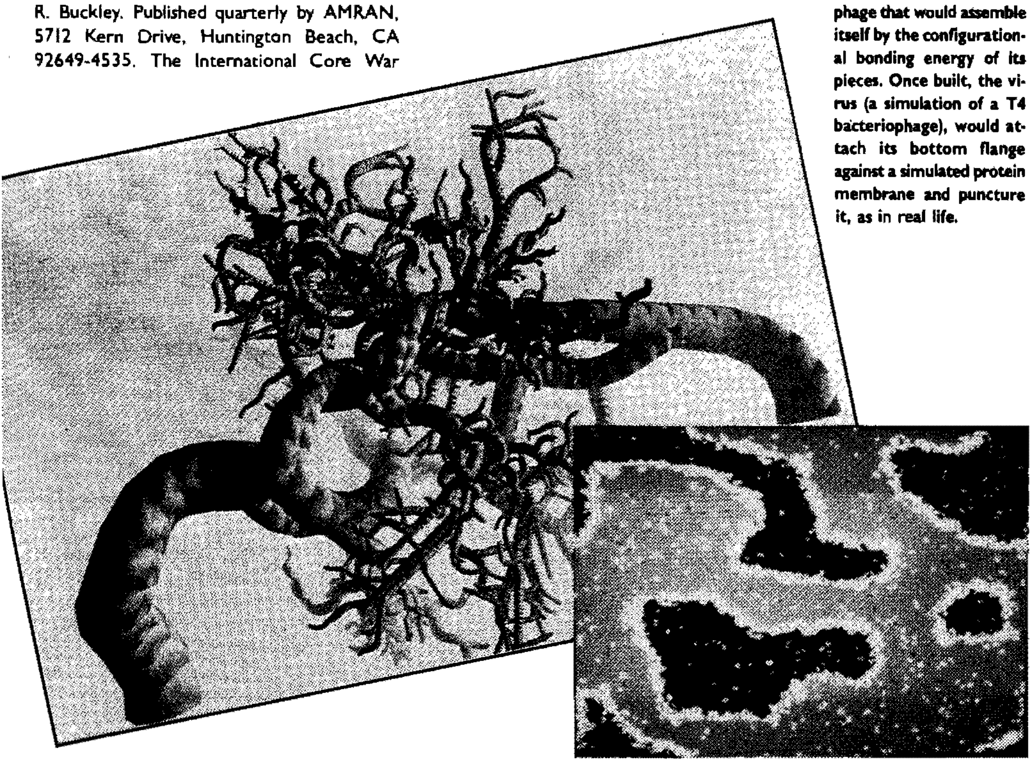

Twisting phenotype in real time. These trees, entwining deeper by the minute, are controlled by knobs on the graphics machine of Peter Oppenheimer at New York Institute of Technology. The knobs determine the phenotype (the physical manifestation fixed by genes) of a tree-like structure displayed on the screen. By twirling their settings he can send the bark of the tree into shagginess, or deepen its color, or compress the stature of its trunk or the spacing of its branches, or alter its curliness. The knobs rotate through species in a marvelous continuum, juniper to cedar to Ponderosa pine. In between, the trees often turn into trees that aren't, but could be. It's the modeling of counterfeit life. "Controlling nature is addictive, even obsessive," says Oppenheimer as he smiles and spins the knobs again.

A membrane of coalitions. Science-fiction author and mathematician Rudy Rucker started cellular automata brewing on a CAM 6 board in his PC clone.[5] Based on the game of Life, invented by John Conway, his rules generate elaborate patterns of populations that are governed by "voting" coalitions. Rucker's world is red and black. At the boundary between colors, cellular-automata life thrives on the "shoreline." They vote on which neighbors should live, and the survivors then vote again, and so on. Rucker set up a world where a near-tie vote does the unexpected. "You win if you get 40 percent of the vote, but not more than 49 percent. It's sort of a radical political tidepool where has-beens can get elected. I wanted to try something different than majority-wins, which just freezes up into a crystal structure. It's a gerrymander life." The pattern of the edge of living and dying forms a throbbing, fluid membrane that pulsates across the screen like a fat amoeba.

Rudy Rucker came up with the most expansive definition of artificial life I have heard. "Right now you can spend a year writing up a program that will only take a few minutes to run. Artificial life is about writing down a few lines of programming that will take decades to run."