BOOTSTRAPPING COMPLEXITY

LEARNING FROM SELF-ORGANIZING SYSTEMS IN NATURE AND TECHNOLOGY

- a remix of Kevin Kelly's book "Out of Control" (1994)

by
Andreas Lloyd
2009

Introduction

“Good morning, self-organizing systems!”

The cheerful speaker smiled with a polished ease and adjusted his tie. "I am indeed very happy to find the Office of Naval Research joining with the Armour Research Foundation in organizing this conference on what I personally consider an exceedingly important topic, and at such a well-chosen time."

It was a spring day in early May, 1959. Four hundred men from an astoundingly diverse group of scientific backgrounds had gathered in Chicago for what promised to be an electrifying meeting. Almost every major branch of science was represented: psychology, linguistics, engineering, embryology, physics, information theory, mathematics, astronomy, and social sciences. No one could remember a conference before this where so many top scientists in different fields were about to spend two days talking about one thing. Certainly there had never been a large meeting about this particular one thing.

It was a topic that only a young country flush with success and confident of its role in the world would even think about: self-organizing systems -- how organization bootstraps itself to life. Bootstrapping! It was the American dream put into an equation.

"The choice of time is particularly significant in my personal life, too," the speaker continued. "For the last nine months the Department of Defense of the United States of America has been in the throes of an organizational effort which shows reasonably clearly that we are still a long way from understanding what makes a self-organizing system."

Hearty chuckles from the early morning crowd just settling into their seats. At the podium Dr. Joachim Weyl, Research Director of the Office of Naval Research, beamed. The conference he hosted was a public rendezvous of scientists who had been convening in smaller meetings since 1942. These intimate, invitation-only gatherings were organized by the Josiah Macy, Jr. Foundation, and became known as the Macy Conferences. In the spirit of wartime urgency, the small gatherings were interdisciplinary, elite, and emphasized thinking big. Among the several dozen visionaries invited over the nine years of the conference were Gregory Bateson, Norbert Wiener, Margaret Mead, Lawrence Frank, John von Neumann, Warren McCulloch, and Arturo Rosenblueth. This stellar congregation later became known as the cybernetic group for the perspective they pioneered -- cybernetics, the art and science of control.

As has been noted by many writers, cybernetics derives from the Greek for "steersman" -- a pilot that steers a ship. In order to steer the ship, the pilot is constantly dependent on constant feedback. The ship and its sails, the wind and waves affecting it can be seen as a whole, closed self-sustaining system, of which the pilot is a vital part. Just as the pilot is dependent on the ship, the ship is dependent on the pilot actively steering to avoid sinking the ship.

In short, cybernetics is the study of the functions and processes of systems which participate in circular, causal chains that move from action to sensing to comparison with desired goal, and again to action. As cybernetician Louis Kauffman has defined it, "cybernetics is the study of systems and processes that interact with themselves and produce themselves from themselves."

The term Cybernetics became widespread because of Norbert Wiener’s book "Cybernetics", first published in 1948. The sub-title of the book was "control and communication in the animal and machine". This was important because it connects control (actions taken in *hope* of achieving goals) with communication (connection and information flow between the actor and the environment). The sub-title thus contains two central points. One: that effective action requires communication. Two: that both animals (biological systems) and machines (non-biological or "artificial" systems) can operate according to cybernetic principles - an explicit recognition that both living and non-living systems can hav*e purpose*.

Some beginnings are inconspicuous; this one wasn't. From the very first Macy Conference, the participants could imagine the alien vista they were opening. Despite their veteran science background and natural skepticism, they saw immediately that this new view would change their life's work. Anthropologist Margaret Mead recalled she was so excited by the ideas set loose in the first meeting that "I did not notice that I had broken one of my teeth until the Conference was over."

The core group consisted of key thinkers in biology, social science, and what we would now call computer science, although this group were only beginning to invent the concept of computers at the time. Their chief achievement was to articulate a language of control and design that worked for biology, social sciences, and computers. Much of the brilliance of these conferences came by the then unconventional approach of rigorously considering living things as machines and machines as living things. Von Neumann quantitatively compared the speed of brain neurons and the speed of vacuum tubes, boldly implying the two could be compared. Wiener reviewed the history of machine automata segueing into human anatomy. Rosenblueth, the doctor, saw homeostatic circuits in the body and in cells.

What brought all of these thinkers together was a shared quest for understanding what makes a self-organizing system. Fundamentally, these thinkers sought to find out how to make something out of nothing. Nature does this every day: First there is hard rock planet; then there is life, lots of it. First barren hills; then brooks with fish and cattails and red-winged blackbirds. First an acorn; then an oak tree forest. Bootstrapping systems that interact with themselves and produce themselves from themselves.

How do you make something from nothing? Although nature knows this trick, we haven't learned much just by watching her. We need to make our own mistakes through our own experiments.

Unfortunately, the cybernetic group lacked the funding and computing technology necessary to model and test their theories. Unable to find answers, they spent their efforts preparing an agenda for questions. So in spite of its bold and fresh ideas, which sparked a breakthroughs in a wide range of disciplines, the field of cybernetics itself soon withered away. By the late 1970s, cybernetics as an academic discipline had all but died out, partly due to lack of funding, partly due to the lack of computers powerful enough to model the complexity of self-organizing systems.

In the fabric of knowledge we call science, there was a rent here, a hole. It was only once computing technology had matured enough to make cybernetic experiments feasible that this hole could be bridged. And by then, the original cybernetic group had passed on, leaving their field to be filled by young enthusiasts not burdened by wise old men.

This book is an exploration of the heritage of the cybernetic group present at that conference in 1959 as it has been carried on by an unlikely group of young, ambitious scientists studying chaos, complexity, artificial life, subsumption architecture, artificial evolution, simulations, ecosystems, and bionic machines. All of these scientists across such diverse fields have found a common framework for their questions in cybernetics as they have continued the quest to understand what makes a self-organizing system. But even though cybernetics pervades every part of this book, references to the original cybernetic group and their work will be few and far between. Particularly since this new generation of scientists have come into cybernetics on their own, unencumbered by an academic tradition, they rarely describe their work in cybernetic terms.

These new cyberneticians are extracting the logical principle of both life and machines, and applying each to the task of building extremely complex systems, thus conjuring up contraptions that are at once both made and alive. In these efforts to create complex mechanical things, again and again they return to nature for directions. They have learned more by their failures in creating complexity and by combining these lessons with small successes in imitating and understanding natural systems than the original cybernetic group could have hoped for. And in doing so, they are fulfilling the notion first presented in the Whole Earth Catalog, itself inspired by the original cybernetics group: “We are as gods and might as well get good at it.”

As this book will show, these experiments are stretching the meanings of "mechanical" and "life" to the point where all complicated things can be perceived as machines, and all self-sustaining machines can be perceived as alive. Human-made things are behaving more lifelike, and biological life is becoming more engineered. I call those examples of self-organizing systems, both made and born, "vivisystems" for the lifelikeness each kind of system holds.

In the following chapters I survey these new cybernetic frontiers of computer science, the edges of biological research, and the odd corners of interdisciplinary experimentation, where researchers are seeking to understand existing self-organizing biological vivisystems as well as to experimentally implement and create self-sustaining, self-improving vivisystems of their own. Creating something from nothing, learning how to be good at being gods.

Because of the very diverse fields of research involved, describing different paths along the same theme of self-organization, the chapters of this book is not organised by some grand sweeping narrative, but rather the opposite: Each chapter tells its own story: One of rebuilding a natural eco-system, another of designing robots, yet another of the notion of co-evolution, and yet another describes research on how the mind works. All of these describe self-organizing, whole systems to some extent, but in completely different realms - some biological, some technical. At the end of this book, I will sum up the recurring patterns in all of these experiments, and extrapolate the insights of this new generation of cyberneticians in what I call “Nine Laws of God.”

But it is only by reading and experiencing the juxtaposition of these radically different expressions of cybernetic ideas that you will be able to fully appreciate the wonder of self-organizing systems. It is my hope that the reader will be able to apply these insights not only in the realm of biological and technical evolution, which they describe, but also in the many forms of human organization. Louis Couffignal, one of the early cyberneticians, characterized cybernetics as “the art of ensuring the efficacy of action” - I find that there is still much potential in ensuring efficacy in the way we organize ourselves.

Hive mind

The beehive beneath my office window quietly exhales legions of busybodies and then inhales them. On summer afternoons, when the sun seeps under the trees to backlight the hive, the approaching sunlit bees zoom into their tiny dark opening like curving tracer bullets. I watch them now as they haul in the last gleanings of nectar from the final manzanita blooms of the year. Soon the rains will come and the bees will hide. I will still gaze out the window as I write; they will still toil, but now in their dark home. Only on the balmiest day will I be blessed by the sight of their thousands in the sun.

Over years of beekeeping, I've tried my hand at relocating bee colonies out of buildings and trees as a quick and cheap way of starting new hives at home. One fall I gutted a bee tree that a neighbor felled. I took a chain saw and ripped into this toppled old tupelo. The poor tree was cancerous with bee comb. The further I cut into the belly of the tree, the more bees I found. The insects filled a cavity as large as I was. It was a gray, cool autumn day and all the bees were home, now agitated by the surgery. I finally plunged my hand into the mess of comb. Hot! Ninety-five degrees at least. Overcrowded with 100,000 cold-blooded bees, the hive had become a warm-blooded organism. The heated honey ran like thin, warm blood. My gut felt like I had reached my hand into a dying animal.

The idea of the collective hive as an animal was an idea late in coming. The Greeks and Romans were famous beekeepers who harvested respectable yields of honey from homemade hives, yet these ancients got almost every fact about bees wrong. Blame it on the lightless conspiracy of bee life, a secret guarded by ten thousand fanatically loyal, armed soldiers. Democritus thought bees spawned from the same source as maggots. Xenophon figured out the queen bee but erroneously assigned her supervisory responsibilities she doesn't have. Aristotle gets good marks for getting a lot right, including the semiaccurate observation that "ruler bees" put larva in the honeycomb cells. (They actually start out as eggs, but at least he corrects Democritus's misguided direction of maggot origins.) Not until the Renaissance was the female gender of the queen bee proved, or beeswax shown to be secreted from the undersides of bees. No one had a clue until modern genetics that a hive is a radical matriarchy and sisterhood: all bees, except the few good-for-nothing drones, are female and sisters. The hive was a mystery as unfathomable as an eclipse.

I've seen eclipses and I've seen bee swarms. Eclipses are spectacles I watch halfheartedly, mostly out of duty, I think, to their rarity and tradition, much as I might attend a Fourth of July parade. Bee swarms, on the other hand, evoke another sort of awe. I've seen more than a few hives throwing off a swarm, and never has one failed to transfix me utterly, or to dumbfound everyone else within sight of it.

A hive about to swarm is a hive possessed. It becomes visibly agitated around the mouth of its entrance. The colony whines in a centerless loud drone that vibrates the neighborhood. It begins to spit out masses of bees, as if it were emptying not only its guts but its soul. A poltergeist-like storm of tiny wills materializes over the hive box. It grows to be a small dark cloud of purpose, opaque with life. Boosted by a tremendous buzzing racket, the ghost slowly rises into the sky, leaving behind the empty box and quiet bafflement. The German theosophist Rudolf Steiner writes lucidly in his otherwise kooky Nine Lectures on Bees: "Just as the human soul takes leave of the body...one can truly see in the flying swarm an image of the departing human soul."

For many years Mark Thompson, a beekeeper local to my area, had the bizarre urge to build a Live-In Hive -- an active bee home you could visit by inserting your head into it. He was working in a yard once when a beehive spewed a swarm of bees "like a flow of black lava, dissolving, then taking wing." The black cloud coalesced into a 20-foot-round black halo of 30,000 bees that hovered, UFO-like, six feet off the ground, exactly at eye level. The flickering insect halo began to drift slowly away, keeping a constant six feet above the earth. It was a Live-In Hive dream come true.

Mark didn't waver. Dropping his tools he slipped into the swarm, his bare head now in the eye of a bee hurricane. He trotted in sync across the yard as the swarm eased away. Wearing a bee halo, Mark hopped over one fence, then another. He was now running to keep up with the thundering animal in whose belly his head floated. They all crossed the road and hurried down an open field, and then he jumped another fence. He was tiring. The bees weren't; they picked up speed. The swarm-bearing man glided down a hill into a marsh. The two of them now resembled a superstitious swamp devil, humming, hovering, and plowing through the miasma. Mark churned wildly through the muck trying to keep up. Then, on some signal, the bees accelerated. They unhaloed Mark and left him standing there wet, "in panting, joyful amazement." Maintaining an eye-level altitude, the swarm floated across the landscape until it vanished, like a spirit unleashed, into a somber pine woods across the highway.

"Where is 'this spirit of the hive'...where does it reside?" asks the author Maurice Maeterlinck as early as 1901. "What is it that governs here, that issues orders, foresees the future...?" We are certain now it is not the queen bee. When a swarm pours itself out through the front slot of the hive, the queen bee can only follow. The queen's daughters manage the election of where and when the swarm should settle. A half-dozen anonymous workers scout ahead to check possible hive locations in hollow trees or wall cavities. They report back to the resting swarm by dancing on its contracting surface. During the report, the more theatrically a scout dances, the better the site she is championing. Deputy bees then check out the competing sites according to the intensity of the dances, and will concur with the scout by joining in the scout's twirling. That induces more followers to check out the lead prospects and join the ruckus when they return by leaping into the performance of their choice.

It's a rare bee, except for the scouts, who has inspected more than one site. The bees see a message, "Go there, it's a nice place." They go and return to dance/say, "Yeah, it's really nice." By compounding emphasis, the favorite sites get more visitors, thus increasing further visitors. As per the law of increasing returns, them that has get more votes, the have-nots get less. Gradually, one large, snowballing finale will dominate the dance-off. The biggest crowd wins.

It's an election hall of idiots, for idiots, and by idiots, and it works marvelously. This is the true nature of democracy and of all distributed governance. At the close of the curtain, by the choice of the citizens, the swarm takes the queen and thunders off in the direction indicated by mob vote. The queen who follows, does so humbly. If she could think, she would remember that she is but a mere peasant girl, blood sister of the very nurse bee instructed (by whom?) to select her larva, an ordinary larva, and raise it on a diet of royal jelly, transforming Cinderella into the queen. By what karma is the larva for a princess chosen? And who chooses the chooser?

"The hive chooses," is the disarming answer of William Morton Wheeler, a natural philosopher and entomologist of the old school, who founded the field of social insects. Writing in a bombshell of an essay in 1911 ("The Ant Colony as an Organism" in the Journal of Morphology), Wheeler claimed that an insect colony was not merely the analog of an organism, it is indeed an organism, in every important and scientific sense of the word. He wrote: "Like a cell or the person, it behaves as a unitary whole, maintaining its identity in space, resisting dissolution...neither a thing nor a concept, but a continual flux or process."

It was a mob of 20,000 united into oneness.

In a darkened Las Vegas conference room, a cheering audience waves cardboard wands in the air. Each wand is red on one side, green on the other. Far in back of the huge auditorium, a camera scans the frantic attendees. The video camera links the color spots of the wands to a nest of computers set up by graphics wizard Loren Carpenter. Carpenter's custom software locates each red and each green wand in the auditorium. Tonight there are just shy of 5,000 wandwavers. The computer displays the precise location of each wand (and its color) onto an immense, detailed video map of the auditorium hung on the front stage, which all can see. More importantly, the computer counts the total red or green wands and uses that value to control software. As the audience wave the wands, the display screen shows a sea of lights dancing crazily in the dark, like a candlelight parade gone punk. The viewers see themselves on the map; they are either a red or green pixel. By flipping their own wands, they can change the color of their projected pixels instantly.

Loren Carpenter boots up the ancient video game of Pong onto the immense screen. Pong was the first commercial video game to reach pop consciousness. It's a minimalist arrangement: a white dot bounces inside a square; two movable rectangles on each side act as virtual paddles. In short, electronic ping-pong. In this version, displaying the red side of your wand moves the paddle up. Green moves it down. More precisely, the Pong paddle moves as the average number of red wands in the auditorium increases or decreases. Your wand is just one vote.

Carpenter doesn't need to explain very much. Every attendee at this 1991 conference of computer graphic experts was probably once hooked on Pong. His amplified voice booms in the hall, "Okay guys. Folks on the left side of the auditorium control the left paddle. Folks on the right side control the right paddle. If you think you are on the left, then you really are. Okay? Go!"

The audience roars in delight. Without a moment's hesitation, 5,000 people are playing a reasonably good game of Pong. Each move of the paddle is the average of several thousand players' intentions. The sensation is unnerving. The paddle usually does what you intend, but not always. When it doesn't, you find yourself spending as much attention trying to anticipate the paddle as the incoming ball. One is definitely aware of another intelligence online: it's this hollering mob.

The group mind plays Pong so well that Carpenter decides to up the ante. Without warning the ball bounces faster. The participants squeal in unison. In a second or two, the mob has adjusted to the quicker pace and is playing better than before. Carpenter speeds up the game further; the mob learns instantly.

"Let's try something else," Carpenter suggests. A map of seats in the auditorium appears on the screen. He draws a wide circle in white around the center. "Can you make a green '5' in the circle?" he asks the audience. The audience stares at the rows of red pixels. The game is similar to that of holding a placard up in a stadium to make a picture, but now there are no preset orders, just a virtual mirror. Almost immediately wiggles of green pixels appear and grow haphazardly, as those who think their seat is in the path of the "5" flip their wands to green. A vague figure is materializing. The audience collectively begins to discern a "5" in the noise. Once discerned, the "5" quickly precipitates out into stark clarity. The wand-wavers on the fuzzy edge of the figure decide what side they "should" be on, and the emerging "5" sharpens up. The number assembles itself.

"Now make a four!" the voice booms. Within moments a "4" emerges. "Three." And in a blink a "3" appears. Then in rapid succession, "Two... One...Zero." The emergent thing is on a roll.

Loren Carpenter launches an airplane flight simulator on the screen. His instructions are terse: "You guys on the left are controlling roll; you on the right, pitch. If you point the plane at anything interesting, I'll fire a rocket at it." The plane is airborne. The pilot is...5,000 novices. For once the auditorium is completely silent. Everyone studies the navigation instruments as the scene outside the windshield sinks in. The plane is headed for a landing in a pink valley among pink hills. The runway looks very tiny.

There is something both delicious and ludicrous about the notion of having the passengers of a plane collectively fly it. The brute democratic sense of it all is very appealing. As a passenger you get to vote for everything; not only where the group is headed, but when to trim the flaps.

But group mind seems to be a liability in the decisive moments of touchdown, where there is no room for averages. As the 5,000 conference participants begin to take down their plane for landing, the hush in the hall is ended by abrupt shouts and urgent commands. The auditorium becomes a gigantic cockpit in crisis. "Green, green, green!" one faction shouts. "More red!" a moment later from the crowd. "Red, red! REEEEED!" The plane is pitching to the left in a sickening way. It is obvious that it will miss the landing strip and arrive wing first. Unlike Pong, the flight simulator entails long delays in feedback from lever to effect, from the moment you tap the aileron to the moment it banks. The latent signals confuse the group mind. It is caught in oscillations of overcompensation. The plane is lurching wildly. Yet the mob somehow aborts the landing and pulls the plane up sensibly. They turn the plane around to try again.

How did they turn around? Nobody decided whether to turn left or right, or even to turn at all. Nobody was in charge. But as if of one mind, the plane banks and turns wide. It tries landing again. Again it approaches cockeyed. The mob decides in unison, without lateral communication, like a flock of birds taking off, to pull up once more. On the way up the plane rolls a bit. And then rolls a bit more. At some magical moment, the same strong thought simultaneously infects five thousand minds: "I wonder if we can do a 360?"

Without speaking a word, the collective keeps tilting the plane. There's no undoing it. As the horizon spins dizzily, 5,000 amateur pilots roll a jet on their first solo flight. It was actually quite graceful. They give themselves a standing ovation.

The conferees did what birds do: they flocked. But they flocked self- consciously. They responded to an overview of themselves as they co-formed a "5" or steered the jet. A bird on the fly, however, has no overarching concept of the shape of its flock. "Flockness" emerges from creatures completely oblivious of their collective shape, size, or alignment. A flocking bird is blind to the grace and cohesiveness of a flock in flight.

At dawn, on a weedy Michigan lake, ten thousand mallards fidget. In the soft pink glow of morning, the ducks jabber, shake out their wings, and dunk for breakfast. Ducks are spread everywhere. Suddenly, cued by some imperceptible signal, a thousand birds rise as one thing. They lift themselves into the air in a great thunder. As they take off they pull up a thousand more birds from the surface of the lake with them, as if they were all but part of a reclining giant now rising. The monstrous beast hovers in the air, swerves to the east sun, and then, in a blink, reverses direction, turning itself inside out. A second later, the entire swarm veers west and away, as if steered by a single mind. In the 17th century, an anonymous poet wrote: "...and the thousands of fishes moved as a huge beast, piercing the water. They appeared united, inexorably bound to a common fate. How comes this unity?"

A flock is not a big bird. Writes the science reporter James Gleick, "Nothing in the motion of an individual bird or fish, no matter how fluid, can prepare us for the sight of a skyful of starlings pivoting over a cornfield, or a million minnows snapping into a tight, polarized array....High-speed film [of flocks turning to avoid predators] reveals that the turning motion travels through the flock as a wave, passing from bird to bird in the space of about one-seventieth of a second. That is far less than the bird's reaction time." The flock is more than the sum of the birds.

In the film Batman Returns a horde of large black bats swarmed through flooded tunnels into downtown Gotham. The bats were computer generated. A single bat was created and given leeway to automatically flap its wings. The one bat was copied by the dozens until the animators had a mob. Then each bat was instructed to move about on its own on the screen following only a few simple rules encoded into an algorithm: don't bump into another bat, keep up with your neighbors, and don't stray too far away. When the algorithmic bats were run, they flocked like real bats.

The flocking rules were discovered by Craig Reynolds, a computer scientist working at Symbolics, a graphics hardware manufacturer. By tuning the various forces in his simple equation -- a little more cohesion, a little less lag time -- Reynolds could shape the flock to behave like living bats, sparrows, or fish. Even the marching mob of penguins in Batman Returns were flocked by Reynolds's algorithms. Like the bats, the computer-modeled 3-D penguins were cloned en masse and then set loose into the scene aimed in a certain direction. Their crowdlike jostling as they marched down the snowy street simply emerged, out of anyone's control.

So realistic is the flocking of Reynolds's simple algorithms that biologists have gone back to their hi-speed films and concluded that the flocking behavior of real birds and fish must emerge from a similar set of simple rules. A flock was once thought to be a decisive sign of life, some noble formation only life could achieve. Via Reynolds's algorithm it is now seen as an adaptive trick suitable for any distributed vivisystem, organic or made.

Wheeler, the ant pioneer, started calling the bustling cooperation of an insect colony a "superorganism" to clearly distinguish it from the metaphorical use of "organism." He was influenced by a philosophical strain at the turn of the century that saw holistic patterns overlaying the individual behavior of smaller parts. The enterprise of science was on its first steps of a headlong rush into the minute details of physics, biology, and all natural sciences. This pell-mell to reduce wholes to their constituents, seen as the most pragmatic path to understanding the wholes, would continue for the rest of the century and is still the dominant mode of scientific inquiry. Wheeler and colleagues were an essential part of this reductionist perspective, as the 50 Wheeler monographs on specific esoteric ant behaviors testify. But at the same time, Wheeler saw "emergent properties" within the superorganism superseding the resident properties of the collective ants. Wheeler said the superorganism of the hive "emerges" from the mass of ordinary insect organisms. And he meant emergence as science -- a technical, rational explanation -- not mysticism or alchemy.

Wheeler held that this view of emergence was a way to reconcile the reduce-it-to-its parts approach with the see-it-as-a-whole approach. The duality of body/mind or whole/part simply evaporated when holistic behavior lawfully emerged from the limited behaviors of the parts. The specifics of how superstuff emerged from baser parts was very vague in everyone's mind. And still is.

What was clear to Wheeler's group was that emergence was a common natural phenomena. It was related to the ordinary kind of causation in everyday life, the kind where A causes B which causes C, or 2 + 2 = 4. Ordinary causality was invoked by chemists to cover the observation that sulfur atoms plus iron atoms equal iron sulfide molecules. According to fellow philosopher C. Lloyd Morgan, the concept of emergence signaled a different variety of causation. Here 2 + 2 does not equal 4; it does not even surprise with 5. In the logic of emergence, 2 + 2 = apples. "The emergent step, though it may seem more or less saltatory [a leap], is best regarded as a qualitative change of direction, or critical turning-point, in the course of events," writes Morgan in Emergent Evolution, a bold book in 1923. Morgan goes on to quote a verse of Browning poetry which confirms how music emerges from chords:

And I know not if, save in this, such gift be allowed to man
That out of three sounds he frame, not a fourth sound, but a star.

We would argue now that it is the complexity of our brains that extracts music from notes, since we presume oak trees can't hear Bach. Yet "Bachness" -- all that invades us when we hear Bach -- is an appropriately poetic image of how a meaningful pattern emerges from musical notes and generic information.

The organization of a tiny honeybee yields a pattern for its tinier one-tenth of a gram of wing cells, tissue, and chitin. The organism of a hive yields integration for its community of worker bees, drones, pollen and brood. The whole 50-pound hive organ emerges with its own identity from the tiny bee parts. The hive possesses much that none of its parts possesses. One speck of a honeybee brain operates with a memory of six days; the hive as a whole operates with a memory of three months, twice as long as the average bee lives.

Ants, too, have hive mind. A colony of ants on the move from one nest site to another exhibits the Kafkaesque underside of emergent control. As hordes of ants break camp and head west, hauling eggs, larva, pupae -- the crown jewels -- in their beaks, other ants of the same colony, patriotic workers, are hauling the trove east again just as fast, while still other workers, perhaps acknowledging conflicting messages, are running one direction and back again completely empty-handed. A typical day at the office. Yet, the ant colony moves. Without any visible decision making at a higher level, it chooses a new nest site, signals workers to begin building, and governs itself.

The marvel of "hive mind" is that no one is in control, and yet an invisible hand governs, a hand that emerges from very dumb members. The marvel is that more is different. To generate a colony organism from a bug organism requires only that the bugs be multiplied so that there are many, many more of them, and that they communicate with each other. At some stage the level of complexity reaches a point where new categories like "colony" can emerge from simple categories of "bug." Colony is inherent in bugness, implies this marvel. Thus, there is nothing to be found in a beehive that is not submerged in a bee. And yet you can search a bee forever with cyclotron and fluoroscope, and you will never find the hive.

This is a universal law of vivisystems: higher-level complexities cannot be inferred by lower-level existences. Nothing -- no computer or mind, no means of mathematics, physics, or philosophy -- can unravel the emergent pattern dissolved in the parts without actually playing it out. Only playing out a hive will tell you if a colony is immixed in a bee. The theorists put it this way: running a system is the quickest, shortest, and only sure method to discern emergent structures latent in it. There are no shortcuts to actually "expressing" a convoluted, nonlinear equation to discover what it does. Too much of its behavior is packed away.

That leads us to wonder what else is packed into the bee that we haven't seen yet? Or what else is packed into the hive that has not yet appeared because there haven't been enough honeybee hives in a row all at once? And for that matter, what is contained in a human that will not emerge until we are all interconnected by wires and politics? The most unexpected things will brew in this bionic hivelike supermind.

The most inexplicable things will brew in any mind.

Because the body is plainly a collection of specialist organs-heart for pumping, kidneys for cleaning -- no one was too surprised to discover that the mind delegates cognitive matters to different regions of the brain.

In the late 1800s, physicians noted correlations in recently deceased patients between damaged areas of the brain and obvious impairments in their mental abilities just before death. The connection was more than academic: might insanity be biological in origin? At the West Riding Lunatic Asylum, London, in 1873, a young physician who suspected so surgically removed small portions of the brain from two living monkeys. In one, his incision caused paralysis of the right limbs; in the other he caused deafness. But in all other respects, both monkeys were normal. The message was clear: the brain must be compartmentalized. One part could fail without sinking the whole vessel.

If the brain was in departments, in what section were recollections stored? In what way did the complex mind divvy up its chores? In a most unexpected way.

In 1888, a man who spoke fluently and whose memory was sharp found himself in the offices of one Dr. Landolt, frightened because he could no longer name any letters of the alphabet. The perplexed man could write flawlessly when dictated a message. However, he could not reread what he had written nor find a mistake if he had made one. Dr. Landolt recorded, "Asked to read an eye chart, [he] is unable to name any letter. However he claims to see them perfectly....He compares the A to an easel, the Z to a serpent, and the P to a buckle."

The man's word-blindness degenerated to a complete aphasia of both speech and writing by the time of his death four years later. Of course, in the autopsy, there were two lesions: an old one near the occipital (visual) lobe and a newer one probably near the speech center.

Here was remarkable evidence of the bureaucratization of the brain. In a metaphorical sense, different functions of the brain take place in different rooms. This room handles letters, if spoken; that room, letters, if read. To speak a letter (outgoing), you need to apply to yet another room. Numbers are handled by a different department altogether, in the next building. And if you want curses, as the Monty Python Flying Circus skit reminds us, you'll need to go down the hall.

An early investigator of the brain, John Hughlings-Jackson, recounts a story about a woman patient of his who lived completely without speech. When some debris, which had been dumped across the street from the ward where she lived, ignited into flames, the patient uttered the first and only word Hughlings-Jackson had ever heard her say: "Fire!"

How can it be, he asked somewhat incredulous, that "fire" is the only word her word department remembers? Does the brain have its own "fire" department, so to speak?

As investigators probed the brain further, the riddle of the mind revealed itself to be deeply specific. The literature on memory features people ordinary in their ability to distinguish concrete nouns -- tell them "elbow" and they will point to their elbow -- but extraordinary in their inability to distinguish abstract nouns -- ask them about "liberty" or "aptitude" and they stare blankly and shrug. Contrarily, the minds of other apparently normal individuals have lost the ability to retain concrete nouns, while perfectly able to identify abstract things. In his wonderful and overlooked book The Invention of Memory, Israel Rosenfield writes:

One patient, when asked to define hay, responded, "I've forgotten"; and when asked to define poster, said, "no idea." Yet given the word supplication, he said, "making a serious request for help," and pact drew "friendly agreement."

Memory is a palace, say the ancient philosophers, where every room parks a thought. Yet with every clinical discovery of yet another form of specialized forgetfulness, the rooms of memory exploded in number. Down this road there is no end. Memory, already divided into a castle of chambers, balkanizes into a terrifying labyrinth of tiny closets.

One study pointed to four patients who could discern inanimate objects (umbrella, towel), but garbled living things, including foods! One of these patients could converse about nonliving objects without suspicion, but a spider to him was defined as "a person looking for things, he was a spider for a nation." There are records of aphasias that interfere with the use of the past tense. I've heard of another report (one that I cannot confirm, but one that I don't doubt) of an ailment that allows a person to discern all foods except vegetables.

The absurd capriciousness underlying such a memory system is best represented by the categorization scheme of an ancient Chinese encyclopedia entitled Celestial Emporium of Benevolent Knowledge, as interpreted by the South American fiction master J. L. Borges.

On those remote pages it is written that animals are divided into (a) those that belong to the Emperor, (b) embalmed ones, (c) those that are trained, (d) suckling pigs, (e) mermaids, (f) fabulous ones, (g) stray dogs, (h) those that are included in this classification, (i) those that tremble as if they were mad, (j) innumerable ones, (k) those drawn with a very fine camel's hair brush, (l) others, (m) those that have just broken a flower vase, (n) those that resemble flies from a distance.

As farfetched as the Celestial Emporium system is, any classification process has its logical problems. Unless there is a different location for every memory to be filed in, there will need to be confusing overlaps, say for instance, of a talking naughty pig, that may be filed under three different categories above. Filing the thought under all three slots would be highly inefficient, although possible.

The system by which knowledge is sequestered in our brain became more than just an academic question as computer scientists tried to build an artificial intelligence. What is the architecture of memory in a hive mind?

In the past most researchers leaned toward the method humans intuitively use for their own manufactured memory stashes: a single location for each archived item, with multiple cross-referencing, such as in libraries. The strong case for a single location in the brain for each memory was capped by a series of famously elegant experiments made by Wilder Penfield, a Canadian neurosurgeon working in the 1930s. In daring open-brain surgery, Penfield probed the living cerebellum of conscious patients with an electrical stimulant, and asked them to report what they experienced. Patients reported remarkably vivid memories. The smallest shift of the stimulant would generate distinctly separate thoughts. Penfield mapped the brain location of each memory while he scanned the surface with his probe.

His first surprise was that these recollections appeared repeatable, in what years later would be taken as a model of a tape recorder -- as in: "hit replay." Penfield uses the term "flash-back" in his account of a 26-year-old woman's postepileptic hallucination: "She had the same flash-back several times. These had to do with her cousin's house or the trip there -- a trip she has not made for ten to fifteen years but used to make often as a child."

The result of Penfield's explorations into the unexplored living brain produced the tenacious image of the hemispheres as fabulous recording devices, ones that seemed to rival the fantastic recall of the newly popular phonograph. Each of our memories was delicately etched into its own plate, catalogued and filed faithfully by the temperate brain, and barring violence, could be retrieved like a jukebox song by pushing the right buttons.

Yet, a close scrutiny of Penfield's raw transcripts of his probing experiments shows memory to be a less mechanical process. As one example, here are some of the responses of a 29-year-old woman to Penfield's pricks in her left temporal lobe: "Something coming to me from somewhere. A dream." Four minutes later, in exactly the same spot: "The scenery seemed to be different from the one just before..." In a nearby spot: "Wait a minute, something flashed over me, something I dreamt." In a third spot: further inside the brain, "I keep having dreams." The stimulation is repeated in the same spot: "I keep seeing things -- I keep dreaming of things."

These scripts tell of dreamlike glimpses, rather than disorienting reruns dredged up from the basement cubbyholes of the mind's archives. The owners of these experiences recognize them as fragmentary semimemories. They ramble with that awkward "assembled" flavor that dreams grow by -- unfocused tales of bits and pieces of the past reworked into a collage of a dream. The emotional charge of a déjá vu was absent. No overwhelming sense of "it was exactly like this was then" pushed against the present. The replays should have fooled nobody.

Human memories do crash. They crash in peculiar ways, by forgetting vegetables on a list of things to buy at the grocery or by forgetting vegetables in general. Memories often bruise in tandem with a physical bruise of the brain, so we must expect that some memory is bound in time and space to some degree, since being bound to time and space is one definition of being real.

But the current view of cognitive science leans more toward a new image: memories are like emergent events summed out of many discrete, unmemory-like fragments stored in the brain. These pieces of half-thoughts have no fixed home; they abide throughout the brain. Their manner of storage differs substantially from thought to thought-learning to shuffle cards is organized differently than learning the capital of Bolivia -- and the manner differs subtly from person to person, and equally subtly from time to time.

There are more possible ideas/experiences than there are ways to combine neurons in the brain. Memory, then, must organize itself in some way to accommodate more possible thoughts than it has room to store. It cannot have a shelf for every thought of the past, nor a place reserved for every potential thought of the future.

I remember a night in Taiwan twenty years ago. I was in the back of an open truck on a dirt road in the mountains. I had my jacket on; the hill air was cold. I was hitching a ride to arrive at a mountain peak by dawn. The truck was grinding up the steep, dark road while I looked up to the stars in the clear alpine air. It was so clear that I could see tiny stars near the horizon. Suddenly a meteor zipped across low, and because of my angle in the mountains, I could see it skip across the atmosphere. Skip, skip, skip, like a stone.

As I just now remembered this, the skipping meteor was not a memory tape I replayed, despite its ready vividness. The skipping meteor image doesn't exist anywhere in particular in my mind. When I resurrected my experience, I assembled it anew. And I assemble it anew each time I remember it. The parts are tiny bits of evidence scattered sparsely through the hive of my brain: a record of cold shivering, of a bumpy ride somewhere, of many sightings of stars, of hitchhiking. The records are even finer grained than that: cold, bump, points of light, waiting. They are the same raw impressions our minds receive from our senses and with which it assembles our perceptions of the present.

Our consciousness creates the present, just as it creates the past, from many distributed clues scattered in our mind. Standing before an object in a museum, my mind associates its parallel straight lines with the notion of a "chair," even though the thing has only three legs. My mind has never before seen such a chair, but it compiles all the associations -- upright, level seat, stable, legs- and creates the visual image. Very fast. In fact, I will be aware of the general "chairness" of the chair before I can perceive its unique details.

Our memories (and our hive minds) are created in the same indistinct, haphazard way. To find the skipping meteor, my consciousness grabbed a thread with streaks of light and gathered a bunch of feelings associated with stars, cold, bumps. What I created depended on what else I had thrown into my mind recently, including what other thing I was doing/feeling last time I tried to assemble the skipping meteor memory. That's why the story is slightly different each time I remember it, because each time it is, in a real sense, a completely different experience. The act of perceiving and the act of remembering are the same. Both assemble an emergent whole from many distributed pieces.

"Memory," says cognitive scientist Douglas Hofstadter, "is highly reconstructive. Retrieval from memory involves selecting out of a vast field of things what's important and what is not important, emphasizing the important stuff, downplaying the unimportant." That selection process is perception. "I am a very big believer," Hofstadter told me, "that the core processes of cognition are very, very tightly related to perception."

In the last two decades, a few cognitive scientists have contemplated ways to create a distributed memory. Psychologist David Marr proposed a novel model of the human cerebellum in the early 1970s by which memory was stored randomly throughout a web of neurons. In 1974, Pentti Kanerva, a computer scientist, worked out the mathematics of a similar web by which long strings of data could be stored randomly in a computer memory. Kanerva's algorithm was an elegant method to store a finite number of data points in a very immense potential memory space. In other words, Kanerva showed a way to fit any perception a mind could have into a finite memory mechanism. Since there are more ideas possible in the universe than there are atoms or minutes, the actual ideas or perceptions that a human mind can ever get to are relatively sparse within the total possibilities; therefore Kanerva called his technique a "sparse distributed memory" algorithm.

In a sparse distributed network, memory is a type of perception. The act of remembering and the act of perceiving both detect a pattern in a very large choice of possible patterns. When we remember, we re-create the act of the original perception; that is, we relocate the pattern by a process similar to the one we used to perceive the pattern originally.

Kanerva's algorithm was so mathematically clean and crisp that it could be roughly implemented by a hacker into a computer one afternoon. At the NASA Ames Research Center, Kanerva and colleagues fine-tuned his scheme for a sparse distributed memory in the mid-1980s by designing a very robust practical version in a computer. Kanerva's memory algorithm could do several marvelous things that parallel what our own minds can do. The researchers primed the sparse memory with several degraded images of numerals (1 to 9) drawn on a 20-by-20 grid. The memory stored these. Then they gave the memory another image of a numeral more degraded than the first samples to see if it could "recall" what the digit was. The memory could. It honed in on the prototypical shape that was behind all the degraded images. In essence it remembered a shape it had never seen before!

The breakthrough was not just being able to find or replay something from the past, but to find something in a vast hive of possibilities when only the vaguest clues are given. It is not enough to retrieve your grandmother's face; a memory must identify it when you see her profile in a wholly different light and from a different angle.

A hive mind is a distributed memory that both perceives and remembers. It is possible that a human mind may be chiefly distributed, yet, it is in artificial minds where distributed mind will certainly prevail. The more computer scientists thought about distributing problems into a hive mind, the more reasonable it seemed. They figured that most personal computers are not in actual use most of the time they are turned on! While composing a letter on a computer you may interrupt the computer's rest with a short burst of key pounding and then let it return to idleness as you compose the next sentence. Taken as a whole, the turned-on computers in an office are idle a large percentage of the day. The managers of information systems in large corporations look at the millions of dollars of personal computer equipment sitting idle on workers' desks at night and wonder if all that computing power might not be harnessed. All they would need is a way to coordinate work and memory in a very distributed system.

But merely combating idleness is not what makes distributing computing worth doing. Distributed being and hive minds have their own rewards, such as greater immunity to disruption. At Digital Equipment Corporation's research lab in Palo Alto, California, an engineer demonstrated this advantage of distributed computation by opening the door of the closet that held the company's own computer network and dramatically yanking a cable out of its guts. The network instantly routed around the breach and didn't falter a bit.

There will still be crashes in any hive mind, of course. But because of the nonlinear nature of a network, when it does fail we can expect glitches like an aphasia that remembers all foods except vegetables. A broken networked intelligence may be able to calculate pi to the billionth digit but not forward e-mail to a new address. It may be able to retrieve obscure texts on, say, the classification procedures for African zebra variants, but be incapable of producing anything sensible about animals in general. Forgetting vegetables in general, then, is less likely a failure of a local memory storage place than it is a systemwide failure that has, as one of its symptoms, the failure of a particular type of vegetable association -- just as two separate but conflicting programs on your computer hard disk may produce a "bug" that prevents you from printing words in italic. The place where the italic font is stored is not broken; but the system's process of rendering italic is broken.

Some of the hurdles that stand in the way of fabricating a distributed computer mind are being overcome by building the network of computers inside one box. This deliberately compressed distributed computing is also known as parallel computing, because the thousands of computers working inside the supercomputer are running in parallel. Parallel supercomputers don't solve the idle-computer-on-the-desk problem, nor do they aggregate widespread computing power; it's just that working in parallel is an advantage in and of itself, and worth building a million-dollar stand-alone contraption to do it.

Parallel distributed computing excels in perception, visualization, and simulation. Parallelism handles complexity better than traditional supercomputers made of one huge, incredibly fast serial computer. But in a parallel supercomputer with a sparse, distributed memory, the distinction between memory and processing fades. Memory becomes an reenactment of perception, indistinguishable from the original act of knowing. Both are a pattern that emerges from a jumble of interconnected parts.

A sink brims with water. You pull the plug. The water stirs. A vortex materializes. It blooms into a tiny whirlpool, growing as if it were alive. In a minute the whirl extends from surface to drain, animating the whole basin. An ever changing cascade of water molecules swirls through the tornado, transmuting the whirlpool's being from moment to moment. Yet the whirlpool persists, essentially unchanged, dancing on the edge of collapse. "We are not stuff that abides, but patterns that perpetuate themselves," wrote Norbert Wiener.

As the sink empties, all of its water passes through the spiral. When finally the basin of water has sunk from the bowl to the cistern pipes, where does the form of the whirlpool go? For that matter, where did it come from?

The whirlpool appears reliably whenever we pull the plug. It is an emergent thing, like a flock, whose power and structure are not contained in the power and structure of a single water molecule. No matter how intimately you know the chemical character of H2O, it does not prepare you for the character of a whirlpool. Like all emergent entities, the essence of a vortex emanates from a messy collection of other entities; in this case, a pool of water molecules. One drop of water is not enough for a whirlpool to appear in, just as one pinch of sand is not enough to hatch an avalanche. Emergence requires a population of entities, a multitude, a collective, a mob, more.

More is different. One grain of sand cannot avalanche, but pile up enough grains of sand and you get a dune that can trigger avalanches. Certain physical attributes such as temperature depend on collective behavior. A single molecule floating in space does not really have a temperature. Temperature is more correctly thought of as a group characteristic that a population of molecules has. Though temperature is an emergent property, it can be measured precisely, confidently, and predictably. It is real.

It has long been appreciated by science that large numbers behave differently than small numbers. Mobs breed a requisite measure of complexity for emergent entities. The total number of possible interactions between two or more members accumulates exponentially as the number of members increases. At a high level of connectivity, and a high number of members, the dynamics of mobs takes hold. More is different.

There are two extreme ways to structure "moreness." At one extreme, you can construct a system as a long string of sequential operations, such as we do in a meandering factory assembly line. The internal logic of a clock as it measures off time by a complicated parade of movements is the archetype of a sequential system. Most mechanical systems follow the clock.

At the other far extreme, we find many systems ordered as a patchwork of parallel operations, very much as in the neural network of a brain or in a colony of ants. Action in these systems proceeds in a messy cascade of interdependent events. Instead of the discrete ticks of cause and effect that run a clock, a thousand clock springs try to simultaneously run a parallel system. Since there is no chain of command, the particular action of any single spring diffuses into the whole, making it easier for the sum of the whole to overwhelm the parts of the whole. What emerges from the collective is not a series of critical individual actions but a multitude of simultaneous actions whose collective pattern is far more important. This is the swarm model.

These two poles of the organization of moreness exist only in theory because all systems in real life are mixtures of these two extremes. Some large systems lean to the sequential model (the factory); others lean to the web model (the telephone system).

It seems that the things we find most interesting in the universe are all dwelling near the web end. We have the web of life, the tangle of the economy, the mob of societies, and the jungle of our own minds. As dynamic wholes, these all share certain characteristics: a certain liveliness, for one.

We know these parallel-operating wholes by different names. We know a swarm of bees, or a cloud of modems, or a network of brain neurons, or a food web of animals, or a collective of agents. The class of systems to which all of the above belong is variously called: networks, complex adaptive systems, swarm systems, vivisystems, or collective systems. I use all these terms in this book.

Organizationally, each of these is a collection of many (thousands) of autonomous members. "Autonomous" means that each member reacts individually according to internal rules and the state of its local environment. This is opposed to obeying orders from a center, or reacting in lock step to the overall environment.

These autonomous members are highly connected to each other, but not to a central hub. They thus form a peer network. Since there is no center of control, the management and heart of the system are said to be decentrally distributed within the system, as a hive is administered.

There are four distinct facets of distributed being that supply vivisystems their character:

- The absence of imposed centralized control

- The autonomous nature of subunits
- The high connectivity between the subunits
- The webby nonlinear causality of peers influencing peers.

The relative strengths and dominance of each factor have not yet been examined systematically.

One theme of this book is that distributed artificial vivisystems, such as parallel computing, silicon neural net chips, or the grand network of online networks commonly known as the Internet, provide people with some of the attractions of organic systems, but also, some of their drawbacks. I summarize the pros and cons of distributed systems here:

Benefits of Swarm Systems

- Adaptable -- It is possible to build a clockwork system that can adjust to predetermined stimuli. But constructing a system that can adjust to new stimuli, or to change beyond a narrow range, requires a swarm -- a hive mind. Only a whole containing many parts can allow a whole to persist while the parts die off or change to fit the new stimuli.
- Evolvable -- Systems that can shift the locus of adaptation over time from one part of the system to another (from the body to the genes or from one individual to a population) must be swarm based. Noncollective systems cannot evolve (in the biological sense).
- Resilient -- Because collective systems are built upon multitudes in parallel, there is redundancy. Individuals don't count. Small failures are lost in the hubbub. Big failures are held in check by becoming merely small failures at the next highest level on a hierarchy.
- Boundless -- Plain old linear systems can sport positive feedback loops -- the screeching disordered noise of PA microphone, for example. But in swarm systems, positive feedback can lead to increasing order. By incrementally extending new structure beyond the bounds of its initial state, a swarm can build its own scaffolding to build further structure. Spontaneous order helps create more order. Life begets more life, wealth creates more wealth, information breeds more information, all bursting the original cradle. And with no bounds in sight.
- Novelty -- Swarm systems generate novelty for three reasons: (1) They are "sensitive to initial conditions" -- a scientific shorthand for saying that the size of the effect is not proportional to the size of the cause -- so they can make a surprising mountain out of a molehill. (2) They hide countless novel possibilities in the exponential combinations of many interlinked individuals. (3) They don't reckon individuals, so therefore individual variation and imperfection can be allowed. In swarm systems with heritability, individual variation and imperfection will lead to perpetual novelty, or what we call evolution.

Apparent Disadvantages of Swarm Systems

- Nonoptimal -- Because they are redundant and have no central control, swarm systems are inefficient. Resources are allotted higgledy-piggledy, and duplication of effort is always rampant. What a waste for a frog to lay so many thousands of eggs for just a couple of juvenile offspring! Emergent controls such as prices in free-market economy -- a swarm if there ever was one -- tend to dampen inefficiency, but never eliminate it as a linear system can.

- Noncontrollable -- There is no authority in charge. Guiding a swarm system can only be done as a shepherd would drive a herd: by applying force at crucial leverage points, and by subverting the natural tendencies of the system to new ends (use the sheep's fear of wolves to gather them with a dog that wants to chase sheep). An economy can't be controlled from the outside; it can only be slightly tweaked from within. A mind cannot be prevented from dreaming, it can only be plucked when it produces fruit. Wherever the word "emergent" appears, there disappears human control.

- Nonpredictable-The complexity of a swarm system bends it in unforeseeable ways. "The history of biology is about the unexpected," says Chris Langton, a researcher now developing mathematical swarm models. The word emergent has its dark side. Emergent novelty in a video game is tremendous fun; emergent novelty in our airplane traffic -- control system would be a national emergency.

- Nonunderstandable -- As far as we know, causality is like clockwork. Sequential clockwork systems we understand; nonlinear web systems are unadulterated mysteries. The latter drown in their self-made paradoxical logic. A causes B, B causes A. Swarm systems are oceans of intersecting logic: A indirectly causes everything else and everything else indirectly causes A. I call this lateral or horizontal causality. The credit for the true cause (or more precisely the true proportional mix of causes) will spread horizontally through the web until the trigger of a particular event is essentially unknowable. Stuff happens. We don't need to know exactly how a tomato cell works to be able to grow, eat, or even improve tomatoes. We don't need to know exactly how a massive computational collective system works to be able to build one, use it, and make it better. But whether we understand a system or not, we are responsible for it, so understanding would sure help.

- Nonimmediate -- Light a fire, build up the steam, turn on a switch, and a linear system awakens. It's ready to serve you. If it stalls, restart it. Simple collective systems can be awakened simply. But complex swarm systems with rich hierarchies take time to boot up. The more complex, the longer it takes to warm up. Each hierarchical layer has to settle down; lateral causes have to slosh around and come to rest; a million autonomous agents have to acquaint themselves. I think this will be the hardest lesson for humans to learn: that organic complexity will entail organic time.

The tradeoff between the pros and cons of swarm logic is very similar to the cost/benefit decisions we would have to make about biological vivisystems, if we were ever asked to. But because we have grown up with biological systems and have had no alternatives, we have always accepted their costs without evaluation.

We can swap a slight tendency for weird glitches in a tool in exchange for supreme sustenance. In exchange for a swarm system of 17 million computer nodes on the Internet that won't go down (as a whole), we get a field that can sprout nasty computer worms, or erupt inexplicable local outages. But we gladly trade the wasteful inefficiencies of multiple routing in order to keep the Internet's remarkable flexibility. On the other hand, when we construct autonomous robots, I bet we give up some of their potential adaptability in exchange for preventing them from going off on their own beyond our full control.

As our inventions shift from the linear, predictable, causal attributes of the mechanical motor, to the crisscrossing, unpredictable, and fuzzy attributes of living systems, we need to shift our sense of what we expect from our machines. A simple rule of thumb may help:

- For jobs where supreme control is demanded, good old clockware is the way to go.

- Where supreme adaptability is required, out-of-control swarmware is what you want.

For each step we push our machines toward the collective, we move them toward life. And with each step away from the clock, our contraptions lose the cold, fast optimal efficiency of machines. Most tasks will balance some control for some adaptability, and so the apparatus that best does the job will be some cyborgian hybrid of part clock, part swarm. The more we can discover about the mathematical properties of generic swarm processing, the better our understanding will be of both artificial complexity and biological complexity.

Swarms highlight the complicated side of real things. They depart from the regular. The arithmetic of swarm computation is a continuation of Darwin's revolutionary study of the irregular populations of animals and plants undergoing irregular modification. Swarm logic tries to comprehend the out-of-kilter, to measure the erratic, and to time the unpredictable. It is an attempt, in the words of James Gleick, to map "the morphology of the amorphous" -- to give a shape to that which seems to be inherently shapeless. Science has done all the easy tasks -- the clean simple signals. Now all it can face is the noise; it must stare the messiness of life in the eye.

Zen masters once instructed novice disciples to approach zen meditation with an unprejudiced "beginner's mind." The master coached students, "Undo all preconceptions." The proper awareness required to appreciate the swarm nature of complicated things might be called hive mind. The swarm master coaches, "Loosen all attachments to the sure and certain."

A contemplative swarm thought: The Atom is the icon of 20th century science.

The popular symbol of the Atom is stark: a black dot encircled by the hairline orbits of several other dots. The Atom whirls alone, the epitome of singleness. It is the metaphor for individuality: atomic. It is the irreducible seat of strength. The Atom stands for power and knowledge and certainty. It is as dependable as a circle, as regular as round.

The image of the planetary Atom is printed on toys and on baseball caps. The swirling Atom works its way into corporate logos and government seals. It appears on the back of cereal boxes, in school books, and stars in TV commercials.

The internal circles of the Atom mirror the cosmos, at once a law-abiding nucleus of energy, and at the same time the concentric heavenly spheres spinning in the galaxy. In the center is the animus, the It, the life force, holding all to their appropriate whirling stations. The symbolic Atoms' sure orbits and definite interstices represent the understanding of the universe made known. The Atom conveys the naked power of simplicity.

Another Zen thought: The Atom is the past. The symbol of science for the next century is the dynamical Net.

The Net icon has no center -- it is a bunch of dots connected to other dots -- a cobweb of arrows pouring into each other, squirming together like a nest of snakes, the restless image fading at indeterminate edges. The Net is the archetype -- always the same picture -- displayed to represent all circuits, all intelligence, all interdependence, all things economic and social and ecological, all communications, all democracy, all groups, all large systems. The icon is slippery, ensnaring the unwary in its paradox of no beginning, no end, no center. Or, all beginning, all end, pure center. It is related to the Knot. Buried in its apparent disorder is a winding truth. Unraveling it requires heroism.

When Darwin hunted for an image to end his book Origin of Species -- a book that is one long argument about how species emerge from the conflicting interconnected self-interests of many individuals -- he found the image of the tangled Net. He saw "birds singing on bushes, with various insects flitting about, with worms crawling through the damp earth"; the whole web forming "an entangled bank, dependent on each other in so complex a manner."

The Net is an emblem of multiples. Out of it comes swarm being -- distributed being -- spreading the self over the entire web so that no part can say, "I am the I." It is irredeemably social, unabashedly of many minds. It conveys the logic both of Computer and of Nature -- which in turn convey a power beyond understanding.

Hidden in the Net is the mystery of the Invisible Hand -- control without authority. Whereas the Atom represents clean simplicity, the Net channels the messy power of complexity.

The Net, as a banner, is harder to live with. It is the banner of noncontrol. Wherever the Net arises, there arises also a rebel to resist human control. The network symbol signifies the swamp of psyche, the tangle of life, the mob needed for individuality.

The inefficiencies of a network -- all that redundancy and ricocheting vectors, things going from here to there and back just to get across the street -- encompasses imperfection rather than ejecting it. A network nurtures small failures in order that large failures don't happen as often. It is its capacity to hold error rather than scuttle it that makes the distributed being fertile ground for learning, adaptation, and evolution.

The only organization capable of unprejudiced growth, or unguided learning, is a network. All other topologies limit what can happen.

A network swarm is all edges and therefore open ended any way you come at it. Indeed, the network is the least structured organization that can be said to have any structure at all. It is capable of infinite rearrangements, and of growing in any direction without altering the basic shape of the thing, which is really no outward shape at all. Craig Reynolds, the synthetic flocking inventor, points out the remarkable ability of networks to absorb the new without disruption: "There is no evidence that the complexity of natural flocks is bounded in any way. Flocks do not become 'full' or 'overloaded' as new birds join. When herring migrate toward their spawning grounds, they run in schools extending as long as 17 miles and containing millions of fish." How big a telephone network could we make? How many nodes can one even theoretically add to a network and still have it work? The question has hardly even been asked.

There are a variety of swarm topologies, but the only organization that holds a genuine plurality of shapes is the grand mesh. In fact, a plurality of truly divergent components can only remain coherent in a network. No other arrangement -- chain, pyramid, tree, circle, hub -- can contain true diversity working as a whole. This is why the network is nearly synonymous with democracy or the market.

A dynamic network is one of the few structures that incorporates the dimension of time. It honors internal change. We should expect to see networks wherever we see constant irregular change, and we do.

A distributed, decentralized network is more a process than a thing. In the logic of the Net there is a shift from nouns to verbs. Economists now reckon that commercial products are best treated as though they were services. It's not what you sell a customer, its what you do for them. It's not what something is, it's what it is connected to, what it does. Flows become more important than resources. Behavior counts.

Network logic is counterintuitive. Say you need to lay a telephone cable that will connect a bunch of cities; let's make that three for illustration: Kansas City, San Diego, and Seattle. The total length of the lines connecting those three cities is 3,000 miles. Common sense says that if you add a fourth city to your telephone network, the total length of your cable will have to increase. But that's not how network logic works. By adding a fourth city as a hub (let's make that Salt Lake City) and running the lines from each of the three cities through Salt Lake City, we can decrease the total mileage of cable to 2,850 or 5 percent less than the original 3,000 miles. Therefore the total unraveled length of a network can be shortened by adding nodes to it! Yet there is a limit to this effect. Frank Hwang and Ding Zhu Du, working at Bell Laboratories in 1990, proved that the best savings a system might enjoy from introducing new points into a network would peak at about 13 percent. More is different.

On the other hand, in 1968 Dietrich Braess, a German operations researcher, discovered that adding routes to an already congested network will only slow it down. Now called Braess's Paradox, scientists have found many examples of how adding capacity to a crowded network reduces its overall production. In the late 1960s the city planners of Stuttgart tried to ease downtown traffic by adding a street. When they did, traffic got worse; then they blocked it off and traffic improved. In 1992, New York City closed congested 42nd Street on Earth Day, fearing the worst, but traffic actually improved that day.

Then again, in 1990, three scientists working on networks of brain neurons reported that increasing the gain -- the responsivity -- of individual neurons did not increase their individual signal detection performance, but it did increase the performance of the whole network to detect signals.

Nets have their own logic, one that is out-of-kilter to our expectations. And this logic will quickly mold the culture of humans living in a networked world. What we get from heavy-duty communication networks, and the networks of parallel computing, and the networks of distributed appliances and distributed being is Network Culture.

Alan Kay, a visionary who had much to do with inventing personal computers, says that the personally owned book was one of the chief shapers of the Renaissance notion of the individual, and that pervasively networked computers will be the main shaper of humans in the future. It's not just individual books we are leaving behind, either. Global opinion polling in real-time 24 hours a day, seven days a week, ubiquitous telephones, asynchronous e-mail, 500 TV channels, video on demand: all these add up to the matrix for a glorious network culture, a remarkable hivelike being.

The tiny bees in my hive are more or less unaware of their colony. By definition their collective hive mind must transcend their small bee minds. As we wire ourselves up into a hivish network, many things will emerge that we, as mere neurons in the network, don't expect, don't understand, can't control, or don't even perceive. That's the price for any emergent hive mind.

Machines with an attitude

Despite millions of dollars in research funding, no hacker has been able to coax a machine to walk across a room under its own intellect. A few robots cross in the unreal time of days, or they bump into furniture, or conk out after three-quarters of the way. In December 1990, after a decade of effort, graduate students at Carnegie Mellon University's Field Robotics Center wired together a robot that slowly walked all the way across a courtyard. Maybe 100 feet in all. They named him Ambler.

The 19-foot-tall iron Ambler weighed 2 tons, not counting its brain which was so heavy it sat on the ground off to the side. It was funded to explore distant planets and cost several million dollars of tax money to construct. This huge machine toddled in a courtyard, deliberating at each step. It did nothing else. Walking without tripping was enough after such a long wait. Ambler's parents applauded happily at its first steps.

Moving its six crablike legs was the easiest part for Ambler. The giant had a harder time trying to figure out where it was. Simply representing the terrain so that it could calculate how to traverse it turned out to be Ambler's curse. Ambler spends its time, not walking, but worrying about getting the layout of the yard right. "This must be a yard," it says to itself. "Here are possible paths I could take. I'll compare them to my mental map of the yard and throw away all but the best one." Ambler works from a representation of its environment that it creates in its mind and then navigates from that symbolic chart, which is updated after each step. A thousand-line software program in the central computer manages Ambler's laser vision, sensors, pneumatic legs, gears, and motors. Despite its two-ton, two-story-high hulk, this poor robot is living in its head. And a head that is only connected to its body by a long cable.

Contrast that to a tiny, real ant just under one of Ambler's big padded feet. It crosses the courtyard twice during Ambler's single trip. An ant weighs, brain and body, 1/100th of a gram -- a pinpoint. It has no image of the courtyard and very little idea of where it is. Yet it zips across the yard without incident, without even thinking in one sense.

Ambler was built huge and rugged in order to withstand the extreme cold and grit conditions on Mars, where it would not be so heavy. But ironically Ambler will never make it to Mars because of its bulk, while robots built like ants may.

The ant approach to mobots is Rodney Brooks's idea. Rather than waste his time making one incapacitated genius, Brooks, an MIT professor, wants to make an army of useful idiots. He figures we would learn more from sending a flock of mechanical can-do cockroaches to a planet, instead of relying on the remote chance of sending a solitary overweight dinosaur with pretensions of intelligence.

In a widely cited 1989 paper entitled "Fast, Cheap and Out of Control: A Robot Invasion of the Solar System," Brooks claimed that "within a few years it will be possible at modest cost to invade a planet with millions of tiny robots." He proposed to invade the moon with a fleet of shoe-box-size, solar-powered bulldozers that can be launched from throwaway rockets. Send an army of dispensable, limited agents coordinated on a task, and set them loose. Some will die, most will work, something will get done. The mobots can be built out of off-the-shelf parts in two years and launched completely assembled in the cheapest one-shot, lunar-orbit rocket. In the time it takes to argue about one big sucker, Brooks can have his invasion built and delivered.

There was a good reason why some NASA folks listened to Brooks's bold ideas. Control from Earth didn't work very well. The minute-long delay in signals between an Earth station and a faraway robot teetering on the edge of a crevice demand that the robot be autonomous. A robot cannot have a remotely linked head, as Ambler did. It has to have an onboard brain operating entirely by internal logic and guidance without much communication from Earth. But the brains don't have to be very smart. For instance, to clear a landing pad on Mars an army of bots can dumbly spend twelve hours a day scraping away soil in the general area. Push, push, push, keep it level. One of them wouldn't do a very even job, but a hundred working as a colony could clear a building site. When an expedition of human visitors lands later, the astronauts can turn off any mobots still alive and give them a pat.

Most of the mobots will die, though. Within several months of landing, the daily shock of frigid cold and oven heat will crack the brain chips into uselessness. But like ants, individual mobots are dispensable. Compared to Ambler, they are cheaper to launch into space by a factor of 1000; thus, sending hundreds of mobots is a fraction of the cost of one large robot.

Brooks's original crackpot idea has now evolved into an official NASA program. Engineers at the Jet Propulsion Laboratory are creating a microrover. The project began as a scale model for a "real" planet rover, but as the virtues of small, distributed effort began to dawn on everyone, microrovers became real things in themselves. NASA's prototype tiny bot looks like a very flashy six-wheeled, radio-controlled dune buggy for kids. It is, but it is also solar-powered and self-guiding. A flock of these microrovers will probably end up as the centerpiece of the Mars Environmental Survey scheduled to land in 1997.

Microbots are fast to build from off-the-shelf parts. They are cheap to launch. And once released as a group, they are out of control, without the need for constant (and probably misleading) supervision. This rough-and-ready reasoning is upside-down to the slow, thorough, in-control approach most industrial designers bring to complex machinery. Such radical engineering philosophy was reduced to a slogan: Fast, cheap, and out of control. Engineers envisioned fast, cheap, and out-of-control robots ideal for: (1) Planet exploration; (2) Collection, mining, harvesting; and (3) Remote construction.

"Fast, cheap, and out of control" began appearing on buttons of engineers at conferences and eventually made it to the title of Rodney Brooks's provocative paper. The new logic offered a completely different view of machines. There is no center of control among the mobots. Their identity was spread over time and space, the way a nation is spread over history and land. Make lots of them; don't treat them so precious.

Rodney Brooks grew up in Australia, where like a lot of boys round the world, he read science fiction books and built toy robots. He developed a Downunder perspective on things, wanting to turn views on their heads. Brooks followed up on his robot fantasies by hopscotching around the prime robot labs in the U.S., before landing a permanent job as director of mobile robots at MIT.

There, Brooks began an ambitious graduate program to build a robot that would be more insect than dinosaur. "Allen" was the first robot Brooks built. It kept its brains on a nearby desktop, because that's what all robot makers did at the time in order to have a brain worth keeping. The multiple cables leading to the brain box from Allen's bodily senses of video, sonar, and tactile were a neverending source of frustration for Brooks and crew. There was so much electronic background interference generated on the cables that Brooks burnt out a long string of undergraduate engineering students attempting to clear the problem. They checked every known communication media, including ham radio, police walkie-talkies and cellular phones, as alternatives, but all failed to find a static-free connection for such diverse signals. Eventually the undergraduates and Brooks vowed that on their next project they would incorporate the brains inside a robot -- where no significant wiring would be needed -- no matter how tiny the brains might have to be.

They were thus forced to use very primitive logic steps, and very short and primitive connections in "Tom" and "Jerry," the next two robots they built. But to their amazement they found that the dumb way their onboard neural circuit was organized worked far better than a brain in getting simple things done. When Brooks reexamined the abandoned Allen in light of their modest success with dumb neurons, he recalled that "it turned out that in Allen's brain, there really was not much happening."

The success of this profitable downsizing sent Brooks on a quest to see how dumb he could make a robot and still have it do something useful. He ended up with a type of reflex-based intelligence, and robots as dumb as ants. But they were as interesting as ants, too.

Brooks's ideas gelled in a cockroachlike contraption the size of a football called "Genghis." Brooks had pushed his downsizing to an extreme. Genghis had six legs but no "brain" at all. All of its 12 motors and 21 sensors were distributed in a decomposable network without a centralized controller. Yet the interaction of these 12 muscles and 21 sensors yielded an amazingly complex and lifelike behavior.

Each of Genghis's six tiny legs worked on its own, independent of the others. Each leg had its own ganglion of neural cells -- a tiny microprocessor -- that controlled the leg's actions. Each leg thought for itself! Walking for Genghis then became a group project with at least six small minds at work. Other small semiminds within its body coordinated communication between the legs. Entomologists say this is how ants and real cockroaches cope -- they have neurons in their legs that do the leg's thinking.

In the mobot Genghis, walking emerges out of the collective behavior of the 12 motors. Two motors at each leg lift, or not, depending on what the other legs around them are doing. If they activate in the right sequence -- Okay, hup! One, three, six, two, five, four! -- walking "happens."

No one place in the contraption governs walking. Without a smart central controller, control can trickle up from the bottom. Brooks called it "bottom-up control." Bottom-up walking. Bottom-up smartness. If you snip off one leg of a cockroach, it will shift gaits with the other five without losing a stride. The shift is not learned; it is an immediate self-reorganization. If you disable one leg of Genghis, the other legs organize walking around the five that work. They find a new gait as easily as the cockroach.

In one of his papers, Rod Brooks first laid out his instructions on how to make a creature walk without knowing how:

There is no central controller which directs the body where to put each foot or how high to lift a leg should there be an obstacle ahead. Instead, each leg is granted a few simple behaviors and each independently knows what to do under various circumstances. For instance, two basic behaviors can be thought of as "If I'm a leg and I'm up, put myself down, " or "If I'm a leg and I'm forward, put the other five legs back a little." These processes exist independently, run at all times, and fire whenever the sensory preconditions are true. To create walking then, there just needs to be a sequencing of lifting legs (this is the only instance where any central control is evident). As soon as a leg is raised it automatically swings itself forward, and also down. But the act of swinging forward triggers all the other legs to move back a little. Since those legs happen to be touching the ground, the body moves forward.

Once the beast can walk on a flat smooth floor without tripping, other behaviors can be added to improve the walk. For Genghis to get up and over a mound of phone books on the floor, it needs a pair of sensing whiskers to send information from the floor to the first set of legs. A signal from a whisker can suppress a motor's action. The rule might be, "If you feel something, I'll stop; if you don't, I'll keep going."

While Genghis learns to climb over an obstacle, the foundational walking routine is never fiddled with. This is a universal biological principle that Brooks helped illuminate -- a law of god: When something works, don't mess with it; build on top of it. In natural systems, improvements are "pasted" over an existing debugged system. The original layer continues to operate without even being (or needing to be) aware that it has another layer above it.

When friends give you directions on how to get to their house, they don't tell you to "avoid hitting other cars" even though you must absolutely follow this instruction. They don't need to communicate the goals of lower operating levels because that work is done smoothly by a well-practiced steering skill. Instead, the directions to their house all pertain to high-level activities like navigating through a town.

Animals learn (in evolutionary time) in a similar manner. As do Brooks's mobots. His machines learn to move through a complicated world by building up a hierarchy of behaviors, somewhat in this order:

Avoid contact with objects

Wander aimlessly

Explore the world

Build an internal map

Notice changes in the environment

Formulate travel plans

Anticipate and modify plans accordingly

The Wander-Aimlessly Department doesn't give a hoot about obstacles, since the Avoidance Department takes such good care of that.

The grad students in Brooks's mobot lab built what they cheerfully called "The Collection Machine" -- a mobot scavenger that collected empty soda cans in their lab offices at night. The Wander-Aimlessly Department of the Collection Machine kept the mobot wandering drunkenly through all the rooms; the Avoidance Department kept it from colliding with the furniture while it wandered aimlessly.

The Collection Machine roamed all night long until its video camera spotted the shape of a soda can on a desk. This signal triggered the wheels of the mobot and propelled it to right in front of the can. Rather than wait for a message from a central brain (which the mobot did not have), the arm of the robot "learned" where it was from the environment. The arm was wired so that it would "look" at its wheels. If it said, "Gee, my wheels aren't turning," then it knew, "I must be in front of a soda can." Then the arm reached out to pick up the can. If the can was heavier than an empty can, it left it on the desk; if it was light, it took it. With a can in hand the scavenger wandered aimlessly (not bumping into furniture or walls because of the avoidance department) until it ran across the recycle station. Then it would stop its wheels in front of it. The dumb arm would "look" at its hand to see if it was holding a can; if it was it would drop it. If it wasn't, it would begin randomly wandering again through offices until it spotted another can.

That crazy hit-or-miss system based on random chance encounters was one heck of an inefficient way to run a recycling program. But night after night when little else was going on, this very stupid but very reliable system amassed a great collection of aluminum.

The lab could grow the Collection Machine into something more complex by adding new behaviors over the old ones that worked. In this way complexity can be accrued by incremental additions, rather than basic revisions. The lowest levels of activities are not messed with. Once the wander-aimlessly module was debugged and working flawlessly, it was never altered. Even if wander-aimlessly should get in the way of some new higher behavior, the proven rule was suppressed, rather than deleted. Code was never altered, just ignored. How bureaucratic! How biological!

Furthermore, all parts (departments, agencies, rules, behaviors) worked -- and worked flawlessly -- as stand-alones. Avoidance worked whether or not Reach-For-Can was on. Reach-For-Can worked whether or not Avoidance was on. The frog's legs jumped even when removed from the circuits of its head.

The distributed control layout for robots that Brooks devised came to be known as "subsumption architecture" because the higher level of behaviors subsumed the roles of lower levels of behaviors when they wished to take control.

If a nation were a machine, here's how you could build it using subsumption architecture:

You start with towns. You get a town's logistics ironed out: basic stuff like streets, plumbing, lights, and law. Once you have a bunch of towns working reliably, you make a county. You keep the towns going while adding a layer of complexity that will take care of courts, jails, and schools in a whole district of towns. If the county apparatus were to disappear, the towns would still continue. Take a bunch of counties and add the layer of states. States collect taxes and subsume many of the responsibilities of governing from the county. Without states, the towns would continue, although perhaps not as effectively or as complexly. Once you have a bunch of states, you can add a federal government. The federal layer subsumes some of the activities of the states, by setting their limits, and organizing work above the state level. If the feds went away the thousands of local towns would still continue to do their local jobs -- streets, plumbing and lights. But the work of towns subsumed by states and finally subsumed by a nation is made more powerful. That is, towns organized by this subsumption architecture can build, educate, rule, and prosper far more than they could individually. The federal structure of the U.S. government is therefore a subsumption architecture.

A brain and body are made the same way. From the bottom up. Instead of towns, you begin with simple behaviors -- instincts and reflexes. You make a little circuit that does a simple job, and you get a lot of them going. Then you overlay a secondary level of complex behavior that can emerge out of that bunch of working reflexes. The original layer keeps working whether the second layer works or not. But when the second layer manages to produce a more complex behavior, it subsumes the action of the layer below it.

Here is the generic recipe for distributed control that Brooks's mobot lab developed. It can be applied to most creations:

1) Do simple things first.

2) Learn to do them flawlessly.

3) Add new layers of activity over the results of the simple tasks.

4) Don't change the simple things.

5) Make the new layer work as flawlessly as the simple.

6) Repeat, ad infinitum.

This script could also be called a recipe for managing complexity of any type, for that is what it is.

What you don't want is to organize the work of a nation by a centralized brain. Can you imagine the string of nightmares you'd stir up if you wanted the sewer pipe in front of your house repaired and you had to call the Federal Sewer Pipe Repair Department in Washington, D.C., to make an appointment?

The most obvious way to do something complex, such as govern 100 million people or walk on two skinny legs, is to come up with a list of all the tasks that need to be done, in the order they are to be done, and then direct their completion from a central command, or brain. The former Soviet Union's economy was wired in this logical but immensely impractical way. Its inherent instability of organization was evident long before it collapsed.

Central-command bodies don't work any better than central-command economies. Yet a centralized command blueprint has been the main approach to making robots, artificial creatures, and artificial intelligences. It is no surprise to Brooks that braincentric folks haven't even been able to raise a creature complex enough to collapse.

Brooks has been trying to breed systems without central brains so that they would have enough complexity worth a collapse. In one paper he called this kind of intelligence without centrality "intelligence without reason," a delicious yet subtle pun. For not only would this type of intelligence -- one constructed layer by layer from the bottom up -- not have the architecture of "reasoning," it would also emerge from the structure for no apparent reason at all.

The USSR didn't collapse because its economy was strangled by a central command model. Rather it collapsed because any central-controlled complexity is unstable and inflexible. Institutions, corporations, factories, organisms, economies, and robots will all fail to thrive if designed around a central command.

Yes, I hear you say, but don't I as a human have a centralized brain?

Humans have a brain, but it is not centralized, nor does the brain have a center. "The idea that the brain has a center is just wrong. Not only that, it is radically wrong," claims Daniel Dennett. Dennett is a Tufts University professor of philosophy who has long advocated a "functional" view of the mind: that the functions of the mind, such as thinking, come from non-thinking parts. The semimind of a insectlike mobot is a good example of both animal and human minds. According to Dennett, there is no place that controls behavior, no place that creates "walking," no place where the soul of being resides. Dennett: "The thing about brains is that when you look in them, you discover that there's nobody home."

Dennett is slowly persuading many psychologists that consciousness is an emergent phenomenon arising from the distributed network of many feeble, unconscious circuits. Dennett told me, "The old model says there is this central place, an inner sanctum, a theater somewhere in the brain where consciousness comes together. That is, everything must feed into a privileged representation in order for the brain to be conscious. When you make a conscious decision, it is done in the summit of the brain. And reflexes are just tunnels through the mountain that avoid the summit of consciousness."

From this logic (very much the orthodox dogma in brain science) it follows, says Dennett, that "when you talk, what you've got in your brain is a language output box. Words are composed by some speech carpenters and put in the box. The speech carpenters get directions from a sub-system called the 'conceptualizer' which gives them a preverbal message. Of course the conceptualizer has to gets its message from some source, so it all goes on to an infinite regress of control."

Dennett calls this view the "Central Meanor." Meaning descends from some central authority in the brain. He describes this perspective applied to language -- making as the "idea that there is this sort of four-star general that tells the troops, 'Okay, here's your task. I want to insult this guy. Make up an English insult on the appropriate topic and deliver it.' That's a hopeless view of how speech happens."

Much more likely, says Dennett, is that "meaning emerges from distributed interaction of lots of little things, no one of which can mean a damn thing." A whole bunch of decentralized modules produce raw and often contradictory parts -- a possible word here, a speculative word there. "But out of the mess, not entirely coordinated, in fact largely competitive, what emerges is a speech act."

We think of speech in literary fashion as a stream of consciousness pouring forth like radio broadcasts from a News Desk in our mind. Dennett says, "There isn't a stream of consciousness. There are multiple drafts of consciousness; lots of different streams, no one of which will be singled out as the stream." In 1874, pioneer psychologist William James wrote, "...the mind is at every stage a theatre of simultaneous possibilities. Consciousness consists in the comparisons of these with each other, the selection of some, and the suppression of the rest...."

The idea of a cacophony of alternative wits combining to form what we think of as a unified intelligence is what Marvin Minsky calls "society of mind." Minsky says simply "You can build a mind from many little parts, each mindless by itself." Imagine, he suggests, a simple brain composed of separate specialists each concerned with some important goal (or instinct) such as securing food, drink, shelter, reproduction, or defense. Singly, each is a moron; but together, organized in many different arrangements in a tangled hierarchy of control, they can create thinking. Minsky emphatically states, "You can't have intelligence without a society of mind. We can only get smart things from stupid things."

The society of mind doesn't sound very much different from a bureaucracy of mind. In fact, without evolutionary and learning pressures, the society of mind in a brain would turn into a bureaucracy. However, as Dennett, Minsky, and Brooks envision it, the dumb agents in a complex organization are always both competing and cooperating for resources and recognition. There is a very lax coordination among the vying parts. Minsky sees intelligence as generated by "a loosely-knitted league of almost separate agencies with almost independent goals." Those agencies that succeed are preserved, and those that don't vanish over time. In that sense, the brain is no monopoly, but a ruthless cutthroat ecology, where competition breeds an emergent cooperation.

The slightly chaotic character of mind goes even deeper, to a degree our egos may find uncomfortable. It is very likely that intelligence, at bottom, is a probabilistic or statistical phenomenon -- on par with the law of averages. The distributed mass of ricocheting impulses which form the foundation of intelligence forbid deterministic results for a given starting point. Instead of repeatable results, outcomes are merely probabilistic. Arriving at a particular thought, then, entails a bit of luck.

Dennett admits to me, "The thing I like about this theory is that when people first hear about it they laugh. But then when they think about it, they conclude maybe it is right! Then the more they think about it, they realize, no, not maybe right, some version of it has to be right!"

As Dennett and others have noted, the odd occurrence of Multiple Personalities Syndrome (MPS) in humans depends at some level on the decentralized, distributed nature of human minds. Each personality -- Billy vs. Sally -- uses the same pool of personality agents, the same community of actors and behavior modules to generate visibly different personas. Humans with MPS present a fragmented facet (one grouping) of their personality as a whole being. Outsiders are never sure who they are talking to. The patient seems to lack an "I."

But isn't this what we all do? At different times of our life, and in different moods, we too shift our character. "You are not the person I used to know," screams the person we hurt by manifesting a different cut on our inner society. The "I" is a gross extrapolation that we use as an identity for ourselves and others. If there wasn't an "I" or "Me" in every person then each would quickly invent one. And that, Minsky says, is exactly what we do. There is no "I" so we each invent one.

There is no "I" for a person, for a beehive, for a corporation, for an animal, for a nation, for any living thing. The "I" of a vivisystem is a ghost, an ephemeral shroud. It is like the transient form of a whirlpool held upright by a million spinning atoms of water. It can be scattered with a fingertip.

But a moment later, the shroud reappears, driven together by the churning of a deep distributed mob. Is the new whirlpool a different form, or the same? Are you different after a near-death experience, or only more mature? If the chapters in this book were arranged in a different order, would it be a different book or the same? When you can't answer that question, then you know you are talking about a distributed system.

Inside every solitary living creature is a swarm of non-creature things. Inside every solitary machine one day will be a swarm of non-mechanical things. Both types of swarms have an emergent being and their own agenda.

Brooks writes: "In essence subsumption architecture is a parallel and distributed computation for connecting sensors to actuators in robots." An important aspect of this organization is that complexity is chunked into modular units arranged in a hierarchy. Many observers who are delighted with the social idea of decentralized control are upset to hear that hierarchies are paramount and essential in this new scheme. Doesn't distributed control mean the end of hierarchy?

As Dante climbed through a hierarchy of heavens, he ascended a hierarchy of rank. In a rank hierarchy, information and authority travels one way: from top down. In a subsumption or web hierarchy, information and authority travel from the bottom up, and from side to side. No matter what level an agent or module works at, as Brooks points out, "all modules are created equal....Each module merely does its thing as best it can."

In the human management of distributed control, hierarchies of a certain type will proliferate rather than diminish. That goes especially for distributed systems involving human nodes -- such as huge global computer networks. Many computer activists preach a new era in the network economy, an era built around computer peer-to-peer networks, a time when rigid patriarchal networks will wither away. They are right and wrong. While authoritarian "top-down" hierarchies will retreat, no distributed system can survive long without nested hierarchies of lateral "bottom-up" control. As influence flows peer to peer, it coheres into a chunk -- a whole organelle -- which then becomes the bottom unit in a larger web of slower actions. Over time a multi-level organization forms around the percolating-up control: fast at the bottom, slow at the top.

The second important aspect of generic distributed control is that the chunking of control must be done incrementally from the bottom. It is impossible to take a complex problem and rationally unravel the mess into logical interacting pieces. Such well-intentioned efforts inevitably fail. For example, large companies created ex nihilo, as in joint ventures, have a remarkable tendency to flop. Large agencies created to solve another department's problems become problem departments in themselves.

Chunking from the top down doesn't work for the same reason why multiplication is easier than division in mathematics. To multiply several prime numbers into a larger product is easy; any elementary school kid can do it. But the world's supercomputers choke while trying to unravel a product into its simple primes. Top-down control is very much like trying to decompose a product into its factors, while the large product is very easy to assemble from its factors up.

The law is concise: Distributed control has to be grown from simple local control. Complexity must be grown from simple systems that already work.

As a test bed for bottom-up, distributed control, Brian Yamauchi, a University of Rochester graduate student, constructed a juggling seeing-eye robot arm. The arm's task was to repeatedly bounce a balloon on a paddle. Rather than have one big brain try to figure out where the balloon was and then move the paddle to the right spot under the balloon and then hit it with the right force, Yamauchi decentralized these tasks both in location and in power. The final balancing act was performed by a committee of dumb "agents."

For instance, the extremely complex question of Where is the balloon? was dispersed among many tiny logic circuits by subdividing the problem into several standalone questions. One agent was concerned with the simple query: Is the balloon anywhere within reach? -- an easier question to act on. The agent in charge of that question didn't have any idea of when to hit the balloon, or even where the balloon was. Its single job was to tell the arm to back up if the balloon was not within the arm's camera vision, and to keep moving until it was. A network, or society, of very simpleminded decision-making centers like these formed an organism that exhibited remarkable agility and adaptability.

Yamauchi said, "There is no explicit communication between the behavior agents. All communication occurs through observing the effects of actions that other agents have on the external world." Keeping things local and direct like this allows the society to evolve new behavior while avoiding the debilitating explosion in complexity that occurs with hardwired communication processes. Contrary to popular business preaching, keeping everybody informed about everything is not how intelligence happens.

"We take this idea even further," Brooks said, "and often actually use the world as the communication medium between distributed parts." Rather than being notified by another module of what it expects to happen, a reflex module senses what happened directly in the world. It then sends its message to the others by acting upon the world. "It is possible for messages to get lost -- it actually happens quite often. But it doesn't matter because the agent keeps sending the message over and over again. It goes 'I see it. I see it. I see it' until the arm picks the message up, and does something in the world to alter the world, deactivating the agent."

Centralized communication is not the only problem with a central brain. Maintaining a central memory is equally debilitating. A shared memory has to be updated rigorously, timely, and accurately -- a problem that many corporations can commiserate with. For a robot, central command's challenge is to compile and update a "world model," a theory, or representation, of what it perceives -- where the walls are, how far away the door is, and, by the way, beware of the stairs over there.

What does a brain center do with conflicting information from many sensors? The eye says something is coming, the ear says it is leaving. Which does the brain believe? The logical way is to try to sort them out. A central command reconciles arguments and recalibrates signals to be in sync. In presubsumption robots, most of the great computational resources of a centralized brain were spent in trying to make a coherent map of the world based on multiple-vision signals. Different parts of the system believed wildly inconsistent things about their world derived from different readings of the huge amount of data pouring in from cameras and infrared sensors. The brain never got anything done because it never got everything coordinated.

So difficult was the task of coordinating a central world view that Brooks discovered it was far easier to use the real world as its own model: "This is a good idea as the world really is a rather good model of itself." With no centrally imposed model, no one has the job of reconciling disputed notions; they simply aren't reconciled. Instead, various signals generate various behaviors. The behaviors are sorted out (suppressed, delayed, activated) in the web hierarchy of subsumed control.

In effect, there is no map of the world as the robot sees it (or as an insect sees it, Brooks might argue). There is no central memory, no central command, no central being. All is distributed. "Communication through the world circumvents the problem of calibrating the vision system with data from the arm," Brooks wrote. The world itself becomes the "central" controller; the unmapped environment becomes the map. That saves an immense amount of computation. "Within this kind of organization," Brooks said, "very small amounts of computation are needed to generate intelligent behaviors."

With no central organization, the various agents must perform or die. One could think of Brooks's scheme as having, in his words, "multiple agents within one brain communicating through the world to compete for the resources of the robot's body." Only those that succeed in doing get the attention of other agents.

Astute observers have noticed that Brooks's prescription is an exact description of a market economy: there is no communication between agents, except that which occurs through observing the effects of actions (and not the actions themselves) that other agents have on the common world. The price of eggs is a message communicated to me by hundreds of millions of agents I have never met. The message says (among many other things): "A dozen eggs is worth less to us than a pair of shoes, but more than a two-minute telephone call across the country." That price, together with other price messages, directs thousands of poultry farmers, shoemakers, and investment bankers in where to put their money and energy.

Brooks's model, for all its radicalism in the field of artificial intelligence, is really a model of how complex organisms of any type work. We see a subsumption, web hierarchy in all kinds of vivisystems. He points out five lessons from building mobots. What you want is:

- Incremental construction -- grow complexity, don't install it
- Tight coupling of sensors to actuators -- reflexes, not thinking
- Modular independent layers -- the system decomposes into viable subunits
- Decentralized control -- no central planning
- Sparse communication -- watch results in the world, not wires

When Brooks crammed a bulky, headstrong monster into a tiny, featherweight bug, he discovered something else in this miniaturization. Before, the "smarter" a robot was to be, the more computer components it needed, and the heavier it got. The heavier it got, the larger the motors needed to move it. The heavier the motors, the bigger the batteries needed to power it. The heavier the batteries, the heavier the structure needed to move the bigger batteries, and so on in an escalating vicious spiral. The spiral drove the ratio of thinking parts to body weight in the direction of ever more body.

But the spiral worked in the other direction even nicer. The smaller the computer, the lighter the motors, the smaller the batteries, the smaller the structure, and the stronger the frame became relative to its size. This also drove the ratio of brains to body towards a mobot with a proportionally larger brain, small though its brain was. Most of Brooks's mobots weighed less than ten pounds. Genghis, assembled out of model car parts, weighed only 3.6 pounds. Within three years Brooks would like to have a 1-mm (pencil-tip-size) robot. "Fleabots" he calls them.

Brooks calls for an infiltration of robots not just on Mars but on Earth as well. Rather than try to bring as much organic life into artificial life, Brooks says he's trying to bring as much artificial life into real life. He wants to flood the world (and beyond) with inexpensive, small, ubiquitous semi-thinking things. He gives the example of smart doors. For only about $10 extra you could put a chip brain in a door so that it would know you were about to go out, or it could hear from another smart door that you are coming, or it could notify the lights that you left, and so on. If you had a building full of these smart doors talking to each other, they could help control the climate, as well as help traffic flow. If you extend that invasion to all kinds of other apparatus we now think of as inert, putting fast, cheap, out-of-control intelligence into them, then we would have a colony of sentient entities, serving us, and learning how to serve us better.

When prodded, Brooks predicts a future filled with artificial creatures living with us in mutual dependence -- a new symbiosis. Most of these creatures will be hidden from our senses, and taken for granted, and engineered with an insect approach to problems -- many hands make light work, small work done ceaselessly is big work, individual units are dispensable. Their numbers will outnumber us, as do insects. And in fact, his vision of robots is less that they will be R2D2s serving us beers, than that they will be an ecology of unnamed things just out of sight.

One student in the Mobot Lab built a cheap, bunny-size robot that watches where you are in a room and calibrates your stereo so it is perfectly adjusted as you move around. Brooks has another small robot in mind that lives in the corner of your living room or under the sofa. It wanders around like the Collection Machine, vacuuming at random whenever you aren't home. The only noticeable evidence of its presence is how clean the floors are. A similar, but very tiny, insectlike robot lives in one corner of your TV screen and eats off the dust when the TV isn't on.

Everybody wants programmable animals. "The biggest difference between horses and cars," says Keith Hensen, a popular techno-evangelist, "is that cars don't need attention every day, and horses do. I think there will be a demand for animals that can be switched on and off."

"We are interested in building artificial beings," Brooks wrote in a manifesto in 1985. He defined an artificial being as a creation that can do useful work while surviving for weeks or months without human assistance in real environment. "Our mobots are Creatures in the sense that on power-up they exist in the world and interact with it, pursuing multiple goals. This is in contrast to other mobile robots that are given programs or plans to follow for a specific mission." Brooks was adamant that he would not build toy (easy, simple) environments for his beings, as most other robotists had done, saying "We insist on building complete systems that exist in the real world so that we won't trick ourselves into skipping hard problems."

To date, one hard problem science has skipped is jump-starting a pure mind. If Brooks is right, it probably never will. Instead it will grow a mind from a dumb body. Almost every lesson from the

Mobot Lab seems to teach that there is no mind without body in a real unforgiving world. "To think is to act, and to act is to think," said Heinz von Foerster, gadfly of the 1950s cybernetic movement. "There is no life without movement."

Ambler's dinosaur troubles began because we humans, with our attendant minds, think we are more like Ambler than ants. Since the vital physiological role of the brain has become clear to medicine, the vernacular sense of our center has migrated from the ancient heart to newfangled mind.

We twentieth century humans live entirely in our heads. And so we build robots that live in their heads. Scientists -- humans too -- think of themselves as beings focused onto a spot just south of their forehead behind their eyeballs. There breathes us. In fact, in 1968, brain death became the deciding threshold for human life. No mind, no life.

Powerful computers birthed the fantasy of a pure disembodied intelligence. We all know the formula: a mind inhabiting a brain submerged in a vat. If science would assist me, the contemporary human says, I could live as a brain without a body. And since computers are big brains, I could live in a computer. In the same spirit a computer mind could just as easily use my body.

One of the tenets in the gospel of American pop culture is the widely held creed of transferability of mind. People declare that mind transfer is a swell idea, or an awful idea, but not that it is a wrong idea. In modern folk-belief, mind is liquid to be poured from one vessel to another. From that comes Terminator 2, Frankenstein, and a huge chunk of science fiction.

For better or worse, in reality we are not centered in our head. We are not centered in our mind. Even if we were, our mind has no center, no "I." Our bodies have no centrality either. Bodies and minds blur across each others' supposed boundaries. Bodies and minds are not that different from one another. They are both composed of swarms of sublevel things.

We know that eyes are more brain than camera. An eyeball has as much processing power as a supercomputer. Much of our visual perception happens in the thin retina where light first strikes us, long before the central brain gets to consider the scene. Our spinal cord is not merely a trunk line transmitting phone calls from the brain. It too thinks. We are a lot closer to the truth when we point to our heart and not our head as the center of behaviors. Our emotions swim in a soup of hormones and peptides that percolate through our whole body. Oxytocin discharges thoughts of love (and perhaps lovely thoughts) from our glands. These hormones too process information. Our immune system, by science's new reckoning, is an amazing parallel, decentralized perception machine, able to recognize and remember millions of different molecules.

For Brooks, bodies clarify, simplify. Intelligences without bodies and beings without form are spectral ghosts guaranteed to mislead. Building real things in the real world is how you'll make complex systems like minds and life. Making robots that have to survive in real bodies, day to day on their own, is the only way to find artificial intelligence, or real intelligence. If you don't want a mind to emerge, then unhinge it from the body.

Tedium can unhinge a mind.

Forty years ago, Canadian psychologist D. O. Hebbs was intrigued by the bizarre delusions reported by the ultrabored. Radar observers and long-distance truck drivers often reported blips that weren't there, and stopped for hitchhikers that didn't exist. During the Korean War, Hebbs was contacted by the Canadian Defense Research Board to investigate another troublesome product of monotony and boredom: confessions. Seems that captured UN soldiers were renouncing the West after being brainwashed (a new word) by the communists. Isolation tanks or something.

So in 1954 Hebbs built a dark, soundproof cell at McGill University in Montreal. Volunteers entered the tiny cramped room, donned translucent goggles, padded their arms in cardboard, gloved their hands with cotton mittens, covered their ears with earphones playing a low noise, and laid in bed, immobile, for two to three days. They heard a steady hum, which soon melted into a steady silence. They felt nothing but a dull ache in their backs. They saw nothing but a dim grayness, or was it blackness? The amazonian flow of colors, signals, urgent messages that had been besieging their brains since birth evaporated. Slowly, each of their minds unhitched from its moorings in the body and spun.

Half of the subjects reported visual sensations, some within the first hour: "a row of little men, a German helmet...animated integrated scenes of a cartoonlike character." In the innocent year of 1954 the Canadian scientists reported: "Among our early subjects there were several references, rather puzzling at first, to what one of them called 'having a dream while awake.' Then one of us, while serving as a subject, observed the phenomenon and realized its peculiarity and extent." By the second day of stillness the subjects might report "loss of contact with reality, changes in body image, speech difficulties, reminiscence and vivid memories, sexual preoccupation, inefficiencies of thought, complex dreams, and a higher incident of worry and fright." They didn't say "hallucinations" because that wasn't a word in their vocabulary. Yet.

Hebb's experiments were taken up a few years later by Jack Vernon, who built a "black room" in the basement of the psychology hall at Princeton. He recruited graduate students who hoped to spend four days or so in the dark "getting some thinking done." One of the initial students to stay in the numbing room told the debriefing researchers later, "I guess I was in there about a day or so before you opened the observation window. I wondered why you waited so long to observe me." There was, of course, no observation window.

In the silent coffin of disembodiment, few subjects could think of anything in particular after the second day. Concentration crumbled. The pseudobusyness of daydreaming took over. Worse were thoughts of an active mind that got stuck in an inactive loop. "One subject made up a game of listing, according to the alphabet, each chemical reaction that bore the name of the discoverer. At the letter n he was unable to think of an example. He tried to skip n and go on, but n kept doggedly coming up in his mind, demanding an answer. When this became tiresome, he tried to dismiss the game altogether, only to find that he could not. He endured the insistent demand of his game for a short time, and, finding that he was unable to control it, he pushed the panic button."

The body is the anchor of the mind, and of life. Bodies are machines to prevent the mind from blowing away under a wind of its own making. The natural tendency of neural circuitry is to play games with itself. Left on its own, without a direct link to "outside," a brainy network takes its own machinations as reality. A mind cannot possibly consider anything beyond what it can measure or calculate; without a body it can only consider itself. Given its inherent curiosity, even the simplest mind will exhaust itself devising solutions to challenges it confronts. Yet if most of what it confronts is its own internal circuitry and logic, then it spends its days tinkering with its latest fantasy.

The body -- that is, any bundle of senses and activators -- interrupts this natural mental preoccupation with an overload of urgent material that must be considered right now! A matter of survival! Should we duck?! The mind no longer needs to invent its reality -- the reality is in its face, rapidly approaching dead-on. Duck! it decides by a new and wholly original insight it had never tried before, and would have never thought to try.

Without senses, the mind mentally masturbates, engendering a mental blindness. Without the interruptions of hellos from the eye, ear, tongue, nose, and finger, the evolving mind huddles in the corner picking its navel. The eye is most important because being half brain itself (chock-full of neurons and biochips) it floods the mind with an impossibly rich feed of half-digested data, critical decisions, hints for future steps, clues of hidden things, evocative movements, and beauty. The mind grinds under the load, and behaves. Cut loose from its eyes suddenly, the mind will rear up, spin, retreat.

The cataracts that afflict elderly men and women after a life of sight can be removed, but not without a brief journey into a blindness even darker than what cataracts bring. Doctors surgically remove the lens growths and then cover patients' eyes with a black patch to shield them from light and to prevent the eyeballs from moving, as they unconsciously do whenever they look. Since the eyes move in tandem, both are patched. To further reduce eye movement, patients lie in bed, quiet, for up to a week. At night, when the hospital bustle dies down, the stillness can match the blackness under the blindfold. In the early 1900s when this operation was first commonly performed, there was no machinery in hospitals, no TV or radio, few night shifts, no lights burning. Eyes wrapped in bandages in the cataract ward, the world as hushed and black as the deepest forever.

The first day was dim but full of rest and still. The second day was darker. Numbing. Restless. The third day was black, black, black, silent, and filled with red bugs crawling on the walls.

"During the third night following surgery [the 60-year-old woman] tore her hair and the bedclothes, tried to get out of bed, claimed that someone was trying to get her, and said that the house was on fire. She subsided when the bandage was removed from the unoperated eye," stated a hospital report in 1923.

In the early 1950s, doctors at Mount Sinai Hospital in New York studied a sample of 21 consecutive admissions to the cataract ward. "Nine patients became increasingly restless, tore off the masks, or tried to climb over the siderails. Six patients had paranoid delusions, four had somatic complaints, four were elated [!!], three had visual hallucinations, and two had auditory hallucinations."

"Black patch psychosis" is now something ophthalmologists watch for on the wards. I think universities should keep an eye out for it too. Every philosophy department should hang a pair of black eye patches in a red firealarm-like box that says, "In case of argument about mind/body, break glass, put on."

In an age of virtual everything, the importance of bodies cannot be overemphasized. Rod Brooks have advanced further than most in creating personas for machines, because the creatures are fully embodied. They insist that their robots be situated in real environments.

It is a rare mobile robot that has an "on" lifetime of more than dozens of hours. For the most part, automatons are improved while they are off. In essence, robotists are trying to evolve things while dead, a curious situation that hasn't escaped some researchers' notice. "You know, I'd like to build a robot that could run 24 hours a day for weeks. That's the way for a robot to learn," says Maja Mataric, one of Brooks's robot builders at MIT.

When I visited the Mobot Lab at MIT, Genghis lay sprawled in disassembled pieces on a lab bench. New parts lay nearby. "He's learning," quipped Brooks.

Genghis was learning, but not in any ultimately useful manner. He had to rely on the busy schedules of Brooks and his busy grad students. How much better to learn while alive. That is the next big step for machines. To learn over time, on their own. To not only adapt, but evolve.

Evolution proceeds in steps. Genghis is an insect-equivalent. Its descendants someday will be rodents, and someday further, as smart and nimble as apes.

But we need to be a little patient in our quest for machine evolution, Brooks cautions. From day one of Genesis, it took billions of years for life to reach plant stage, and another billion and a half before fish appeared. A hundred million years later insects made the scene. "Then things really started moving fast," says Brooks. Reptiles, dinosaurs, and mammals appeared within the next 100 million years. The great, brainy apes, including man, arrived in the last 20 million years.

The relatively rapid complexification in most recent geological history suggests to Brooks "that problem solving behavior, language, expert knowledge and reason, are all pretty simple once the essence of being and reacting are available." Since it took evolution 3 billion years to get from single cells to insects, but only another half billion years from there to humans, "this indicates the nontrivial nature of insect level intelligence."

So insect life -- the problem Brooks is sweating over -- is really the hard part. Get artificial insects down, and artificial apes will soon follow. This points to a second advantage to working with fast, cheap, and out-of-control mobots: the necessity of mass numbers for evolution. One Genghis can learn. But evolution requires a seething population of Genghises to get anything done.

To evolve machines, we'll need huge flocks of them. Gnatbots might be perfect. Brooks ultimately dreams of engineering vivisystems full of machines that both learn (adjust to variations in environment) and evolve (populations of critters undergoing "gazillions of trials").

When democracy was first proposed for (and by) humans, many reasonable people rightly feared it as worse than anarchy. They had a point. A democracy of autonomous, evolving machines will be similarly feared as Anarchy Plus. This fear too has some truth.

The greatest social consequence of the Darwinian revolution was the grudging acceptance by humans that humans were random descendants of monkeys, neither perfect nor engineered. The greatest social consequence of neo-biological civilization will be the grudging acceptance by humans that humans are the random ancestors of machines, and that as machines we can be engineered ourselves.

I'd like to condense that further: Natural evolution insists that we are apes; artificial evolution insists that we are machines with an attitude.

I believe that humans are more than the combination of ape and machine (we have a lot going for us!), but I also believe that we are far more ape and machine than we think. That leaves room for an unmeasured but discernible human difference, a difference that inspires great literature, art, and our lives as a whole. I appreciate and indulge in those sentiments. But what I have encountered in the rather mechanical process of evolution, and in the complex but knowable interconnections underpinning living systems, and in the reproducible progress in manufacturing reliable behaviors in robots, is a singular unity between simple life, machines, complex systems, and us. This unity can stir lofty inspirations the equal of any passion in the past.

Machines are a dirty word now. This is because we have withheld from them the full elixir of life. But we are poised to remake them into something that one day may be taken as a compliment.

As humans, we find spiritual refuge in knowing that we are a branch in the swaying tree of life spread upon this blue ball. Perhaps someday we will find spiritual wholesomeness in knowing we are a link in a complex machine layered on top of the green life. Perhaps we'll sing hymns rhapsodizing our role as an ornate node in a vast network of new life that is spawning on top of the old.

Assembling complexity

As an autumn gray settles, I stand in the middle of one of the last wildflower prairies in America. A slight breeze rustles the tan grass. I close my eyes and say a prayer to Jesus, the God of rebirth and resurrection. Then I bend at the waist, and with a strike of a match, I set the last prairie on fire. It burns like hell.

"The grass of the field alive today is thrown into the oven tomorrow," says the rebirth man. The Gospel passage comes to mind as an eight-foot-high wall of orange fire surges downwind crackling loudly and out of control. The heat from the wisps of dead grass is terrific. I am standing with a flapping rubber mat on a broom handle trying to contain the edges of the wall of fire as it marches across the buff-colored field. I remember another passage: "The new has come, the old is gone."

While the prairie burns, I think of machines. Gone is the old way of machines; come is the reborn nature of machines, a nature more alive than dead.

I've come to this patch of fire-seared grass because in its own way this wildflower field is another item of human construction, as I can explain in a moment. The burnt field makes a case that life is becoming manufactured, just as the manufactured is becoming life, just as both are becoming something wonderful and strange.

The future of machines lies in the tangled weeds underfoot. Machines have steadily plowed under wildflower prairies until none are left except the tiny patch I'm standing in. But in a grand irony, this patch holds the destiny of machines, for the future of machines is biology.

My guide to the grassy inferno is Steve Packard, an earnest, mid-thirties guy, who fondles bits of dry weeds -- their Latin names are intimately familiar to him -- as we ramble through the small prairie. Almost two decades ago, Packard was captured by a dream he couldn't shake. He imagined a suburban dumping ground blooming again in its original riotous prairie-earth colors, an oasis of life giving soulful rest to harried cosmopolitans. He dreamt of a prairie gift that would "pay for itself in quality-of-life dollars," as he was fond of telling supporters. In 1974 Packard began working on his vision. With the mild help of skeptical conservation groups, he began to recreate a real prairie not too far from the center of the greater city of Chicago.

Packard knew that the godfather of ecology, Aldo Leopold, had successfully recreated a prairie of sorts in 1934. The University of Wisconsin, where Leopold worked, had purchased an old farm, called the Curtis place, to make an arboretum out of it. Leopold convinced the University to let the Curtis farm revert to prairie again. The derelict farm would be plowed one last time, then sown with disappearing and all but unknown prairie seeds, and left to be.

This simple experiment was not undoing the clock; it was undoing civilization.

Until Leopold's innocent act, every step in civilization had been another notch in controlling and retarding nature. Houses were designed to keep nature's extreme temperatures out. Gardens contrived to divert the power of botanical growth into the tame artifacts of domesticated crops. Iron mined in order to topple trees for lumber.

Respites from this march of progress were rare. Occasionally a feudal lord reserved a wild patch of forest from destruction for his game hunting. Within this sanctuary a gamekeeper might plant wild grain to attract favored animals for his lord's hunt. But until Leopold's folly no one had ever deliberately planted wilderness. Indeed, even as Leopold oversaw the Curtis project, he wondered if anyone could plant wilderness. As a naturalist, he figured it must be largely a matter of letting nature reclaim the spot. His job would be protecting whatever gestures nature made. With the help of colleagues and small bands of farm boys hired by the Civilian Conservation Corps during the Depression, Leopold nursed 300 acres of young emerging prairie plants with buckets of water and occasional thinning of competitors for the first five years.

The prairie plants flourished; but so did the nonprairie weeds. Whatever was carpeting this meadow, it was not the prairie that once did. Tree seedlings, Eurasian migrants, and farm weeds all thrived along with the replanted prairie species. Ten years after the last plowing, it was evident to Leopold that the reborn Curtis prairie was only a half-breed wilderness. Worse, it was slowly becoming an overgrown weedy lot. Something was missing.

A key species, perhaps. A missing species which once reintroduced, would reorder the whole community of ecology of plants. In the mid-1940s that species was identified. It was a wary animal, once ubiquitous on the tall grass prairies, that roamed widely and interacted with every plant, insect, and bird making a home over the sod. The missing member was fire.

Fire made the prairie work. It hatched certain fire-triggered seeds, it eliminated intruding tree saplings, it kept the fire-intolerant urban competitors down. The rediscovery of fire's vital function in tall grass prairie ecology coincided with the rediscovery of fire in the role of almost all the other ecologies in North America. It was a rediscovery because fire's effects on nature had been recognized and used by the aboriginal researchers of the land. The ubiquitous prevalence of fire on the pre-whiteman prairie was well documented by European settlers.

While evident to us now, the role of fire as a key ingredient of the prairie was not clear to ecologists and less clear to conservationists, or what we would now call environmentalists. Ironically, Aldo Leopold, the greatest American ecologist, argued fiercely against letting wildfire burn in wilderness. He wrote in 1920, "The practice of [light-burning] would not only fail to prevent serious fires but would ultimately destroy the productivity of the forests on which western industries depend for their supply of timber." He gave five reasons why fire was bad, none of them valid. Railing against the "light-burning propagandists," Leopold wrote, "It is probably a safe prediction to state that should light-burning continue for another fifty years, our existing forest areas would be further curtailed to a very considerable extent."

A decade later, when more was known about the interdependencies of nature, Leopold finally conceded the vital nature of organic fire. When he reintroduced fire into the synthetic plots of the Wisconsin field grass arboretum, the prairie flourished like it had not for centuries. Species that were once sparse started to carpet the plots.

Still, even after 50 years of fire and sun and winter snows, the Curtis prairie today is not completely authentic in the diversity of its members. Around the edges especially, where ecological diversity is usually the greatest, the prairie suffers from invasions of monopolistic weeds -- the same few ones that thrive on forgotten lots.

The Wisconsin experiment proved one could cobble together a fair approximation of a prairie. What in the world would it take to make a pure prairie, authentic in every respect, an honest-to-goodness recreated prairie? Could one grow a real prairie from the ground up? Is there a way to manufacture a self-sustaining wilderness?

In the fall of 1991, I stood with Steve Packard in one of his treasures -- what he called a "Rembrandt found in the attic" -- at the edge of a suburban Chicago woods. This was the prairie we would burn. Several hundred acres of rustling, wind-blown grass swept over our feet and under scattered oak trees. We swam in a field far richer, far more complete, and far more authentic than

Leopold had seen. Dissolved into this pool of brown tufts were hundreds of uncommon species. "The bulk of the prairie is grass," Packard shouted to me in the wind, "but what most people notice is the advertising of the flowers." At the time of my visit, the flowers were gone, and the ordinary-looking grass and trees seemed rather boring. That "barrenness" turned out to be a key clue in the rediscovery of an entire lost ecosystem.

To arrive at this moment, Packard spent the early 1980s locating small, flowery clearings in the thickets of Illinois woods. He planted prairie wildflower seeds in them and expanded their size by clearing the brush at their perimeters. He burnt the grass to discourage nonnative weeds. At first he hoped the fire would do the work of clearing naturally. He would let it leap from the grass into the thicket to burn the understory shrubs. Then, because of the absence of fuel in the woods, the fire would die naturally. Packard told me, "We let the fires blast into the bush as far as they would go. Our motto became 'Let the fires decide.' "

But the thickets would not burn as he hoped, so Packard and his crews interceded with axes in hand and physically cleared the underbrush. Within two years, they were happy with their results. Thick stands of wild rye grass mingled with yellow coneflower in the new territory. The restorers manually hacked back the brush each season and planted the choicest prairie flower seed they could find.

But by the third year, it was clear something was wrong. The plantings were doing poorly in the shade, producing poor fuel for the season's fires. The grasses that did thrive were not prairie species; Packard had never seen them before. Gradually, the replanted areas reverted to brush.

Packard began to wonder if anyone, including himself, would go through the difficulties of burning an empty plot for decades if they had nothing to show for it. He felt yet another ingredient must be missing which prevented a living system from snapping together. He started reading the botanical history of the area and studying the oddball species.

When he identified the unknown species flourishing so well in the new oak-edge patches, he discovered they didn't belong to a prairie, but to a savanna ecosystem -- a prairie with trees. Researching the plants that were associated with savanna, Packard soon came up with a list of other associated species -- such as thistles, cream gentians, and yellow pimpernels -- that he quickly realized peppered the fringes of his restoration sites. Packard had even found a blazing star flower a few years before. He had brought the flowering plant to a university expert because varieties of blazing star defy nonexpert identification. "What the heck is this?" he'd asked the botanist. "It's not in the books, it's not listed in the state catalogue of species. What is it?" The botanist had said, "I don't know. It could be a savanna blazing star, but there aren't any savannas here, so it couldn't be that. Don't know what is." What one is not looking for, one does not see. Even Packard admitted to himself that the unusual wildflower must have been a fluke, or misidentified. As he recalls, "The savanna species weren't what I was looking for at first so I had sort of written them off."

But he kept seeing them. He found more blazing star in his patches. The oddball species, Packard was coming to realize, were the main show of the clearings. There were many other species associated with savannas he did not recognize, and he began searching for samples of them in the corners of old cemeteries, along railway right-of-ways, and old horse paths -- anywhere a remnant of an earlier ecosystem might survive. Whenever he could, he collected their seed.

An epiphany of sorts overtook Packard when he watched the piles of his seed accumulate in his garage. The prairie seed mix was dry and fluffy-like grass seed. The emerging savanna seed collection, on the other hand, was "multicolored handfuls of lumpy, oozy, glop," ripe with pulpy seeds and dried fruits. Not by wind, but by animals and birds did these seeds disperse. The thing --

the system of coevolved, interlocking organisms -- he was seeking to restore was not a mere prairie, but a prairie with trees: a savanna.

The pioneers in the Midwest called a prairie with trees a "barren." Weedy thickets and tall grass grew under occasional trees. It was neither grassland nor forest -- therefore barren to the early settlers. An almost entirely different set of species kept it a distinct biome from the prairie. The savanna barrens were particularly dependent on fire, more so than the prairies, and when farmers arrived and stopped the fires, the barrens very quickly collapsed into a woods. By the turn of this century the barrens were almost extinct, and the list of their constituent species hardly recorded. But once Packard got a "search image" of the savanna in his mind, he began to see evidence of it everywhere.

Packard sowed the mounds of mushy oddball savanna species, and within two years the fields were ablaze with rare and forgotten wildflowers: bottlebrush grass, blue-stem goldenrod, starry champion, and big-leafed aster. In 1988, a drought shriveled the non-native weeds as the reseeded natives flourished and advanced. In 1989, a pair of eastern bluebirds (which had not been seen in the county for decades) settled into their familiar habitat -- an event that Packard took as "an endorsement." The university botanists called back. Seems like there were early records of savanna blazing star in the state. The biologists were putting it on the endangered list. Oval milkweed somehow returned to the restored barren although it grows nowhere else in the state. Rare and endangered plants like the white-fringed orchid and a pale vetchling suddenly sprouted on their own. The seed might have lain dormant -- and between fire and other factors found the right conditions to hatch -- or been brought in by birds such as the visiting blue birds. Just as miraculously, the silvery-blue butterfly, which had not been seen anywhere in Illinois for a full decade, somehow found its way to suburban Chicago where its favorite food, vetchling, was now growing in the emerging savanna.

"Ah," said the expert entomologists. "The classic savanna butterfly is Edwards hairstreak. But we don't see any. Are you sure this is a savanna?" But by the fifth year of restoration, the Edwards hairstreak butterfly was everywhere on the site.

If you build it, they will come. That's what the voice said in the Field of Dreams. And it's true. And the more you build it, the more that come. Economists call it the "law of increasing returns" -- the snowballing effect. As the web of interrelations is woven tighter, it becomes easier to add the next piece.

Yet there was still an art to it. As Packard knotted the web, he noticed that it mattered what order he added the pieces in. And he learned that other ecologists had discovered the same thing. A colleague of Leopold had found that he got closer to a more authentic prairie by planting prairie seed in a weedy field, rather than in a newly plowed field, as Leopold had first done. Leopold had been concerned that the aggressive weeds would strangle the wildflowers, but a weedy field is far more like a prairie than a plowed field. Some weeds in an old weedy lot are latecomers, and a few of these latecomers are prairie members; their early presence in the conversion quickens the assembly of the prairie system. But the weeds that immediately sprout in a plowed, naked field are very aggressive, and the beneficial late-arriving weeds come into the mix too late. It's like having the concrete reinforcement bars arrive after you've poured the cement foundation for your house. Succession is important.

Stuart Pimm, an ecologist at the University of Tennessee, compares succession paths -- such as the classic series of fire, weed, pine, broadleaf trees -- to well-rehearsed assembly sequences that "the players have played many times. They know, in an evolutionary sense, what the sequence is."

Evolution not only evolves the functioning community, but it also finely tunes the assembly process of the gathering until the community practically falls together. Restoring an ecosystem community is coming at it from the wrong side. "When we try to restore a prairie or wetland, we are trying to assemble an ecosystem along a path that the community has no practice in," says Pimm. We are starting with an old farm, while nature may have started with a glacial moraine ten thousand years ago. Pimm began asking himself: Can we assemble a stable ecosystem by taking in the parts at random? Because at random was exactly how humans were trying to restore ecosystems.

In a laboratory at the University of Tennessee, ecologists Pimm and Jim Drake had been assembling ingredients of microecosystems in different random orders to chart the importance of sequence. Their tiny worlds were microcosms. They started with 15 to 40 different pure strains of algae and microscopic animals, and added these one at a time in various combinations and sequences to a large flask. After 10 to 15 days, if all went well, the aquatic mixture formed a stable, self-reproducing slime ecology -- a distinctive mix of species surviving off of each other. In addition Drake set up artificial ecologies in aquaria and in running water for artificial stream ecologies. After mixing them, they let them run until they were stable. "You look at these communities and you don't need to be a genius to see that they are different," Pimm remarks. "Some are green, some brown, some white. But the interesting thing is that there is no way to tell in advance which way a particular combination of species will go. Like most complex systems, you have to set them up and run them to find out."

It was also not clear at the start whether finding a stable system would be easy. A randomly made ecosystem was likely, Pimm thought, "to just wander around forever, going from one state to the next and back again without ever coming to a persistent state." But the artificial ecosystems didn't wander. Instead, much to their surprise, Pimm found "all sorts of wonderful wrinkles. For one, these random ecosystems have absolutely no problem in stabilizing. Their most common feature is that they always come to a persistent state, and typically it's one state per system."

It was very easy to arrive at a stable ecosystem, if you didn't care what system you arrived at. This was surprising. Pimm said, "We know from chaos theory that many deterministic systems are exquisitely sensitive to initial conditions -- one small difference will send it off into chaos. This stability is the opposite of that. You start out in complete randomness, and you see these things assemble towards something that is a lot more structured than you had any reason to believe could be there. This is anti-chaos."

To complement their studies in vitro, Pimm also set up experiments "in silico"-simplified ecological models in a computer. He created artificial "species" of code that required the presence of certain other species to survive, and also gave them a pecking order so that species B might drive out species A if and when the population of B reached a certain density. (Pimm's models of random ecologies bear some resemblance to Stuart Kauffman's models of random genetic networks; see chapter 20). Each species was loosely interconnected to the others in a kind of vast distributed network. Running thousands of random combinations of the same list of species, Pimm mapped how often the resulting system would stabilize so that minor perturbations, such as introductions or removals of a few species, would not destabilize the collective mix. His results mirrored the results from his bottled living microworlds.

In Pimm's words, the computer models showed that "with just 10 to 20 components in the mix, the number of peaks [or stabilities] may be in the tens, twenties or hundreds. And if you play the tape of life again, you get to a different peak." In other words, after dropping in the same inventory of species, the mess headed toward a dozen final arrangements, but changing the entry sequence of even one of the species was enough to divert the system from one of the end-points to another. The system was sensitive to initial conditions, but it was usually attracted to order.

Pimm saw Packard's work in restoring the Illinois prairie/savanna as validating his findings: "When Packard first tried to assemble the community, it didn't work in the sense that he couldn't get the species he wanted to stick and he had a lot of trouble taking out things he didn't want. But once he introduced the oddball, though proper, species it was close enough to the persistent state that it easily moved there and will probably stay there."

Pimm and Drake discovered a principle that is a great lesson to anyone concerned about the environment, and anyone interested in building complex systems. "To make a wetland you can't just flood an area and hope for the best," Pimm told me. "You are dealing with systems that have assembled over hundreds of thousand, or millions of years. Nor is compiling a list of what's there in terms of diversity enough. You also have to have the assembly instructions."

Steve Packard set out to extend the habitat of authentic prairie. On the way he resurrected a lost ecosystem, and perhaps acquired the assembly instructions for a savanna. Working in an ocean of water instead an ocean of grass, David Wingate in Bermuda set out thirty years ago to nurse a rare species of shorebird back from extinction. On the way, he recreated the entire ecology of a subtropical island, and illuminated a further principle of assembling large functioning systems.

The Bermuda tale involves an island suffering from an unhealthy, ad hoc, artificial ecosystem. By the end of World War II, Bermuda was ransacked by housing developers, exotic pests, and a native flora wrecked by imported garden species. The residents of Bermuda and the world's scientific community were stunned, then, in 1951 by the announcement that the cahow -- a gull-size seabird -- had been rediscovered on the outer cliffs of the island archipelago. The cahow was thought to be extinct for centuries. It was last seen in the 1600s, around the time the dodo had gone extinct. But by a small miracle, a few pairs of breeding cahows hung for generations on some remote sea cliffs in the Bermuda archipelago. They spent most of their life on water, only coming ashore to nest underground, so they went unnoticed for four centuries.

As a schoolboy with a avid interest in birds, David Wingate was present in 1951 when a Bermudan naturalist succeeded in weaseling the first cahow out of its deep nesting crevice. Later, Wingate became involved in efforts to reestablish the bird on a small uninhabited island near Bermuda called Nonsuch. He was so dedicated to the task that he moved -- newly married -- to an abandoned building on the uninhabited, unwired outer island.

It quickly became apparent to Wingate that the cahow could not be restored unless the whole ecosystem of which it was part was also restored. Nonsuch and Bermuda itself were once covered by thick groves of cedar, but the cedars had been wiped out by an imported insect pest in a mere three years between 1948 and 1952. Only their huge white skeletons remained. In their stead were a host of alien plants, and on the main island, many tall ornamental trees that Wingate was sure would never survive a once-in-fifty-year hurricane.

The problem Wingate faced was the perennial paradox that all whole systems makers confront: where do you start? Everything requires everything else to stay up, yet you can't levitate the whole thing at once. Some things have to happen first. And in the correct order.

Studying the cahows, Wingate determined that their underground nesting sites had been diminished by urban sprawl and subsequently by competition with the white-tailed tropicbird for the few remaining suitable sites. The aggressive tropicbird would peck a cahow chick to death and take over the nest. Drastic situations require drastic measures, so Wingate instituted a "government housing program" for the cahow. He built artificial nest sites -- sort of underground birdhouses. He couldn't wait until Nonsuch reestablished a forest of trees, which tip slightly in hurricanes to uproot just the right-sized crevice, too small for the tropic bird to enter, but just perfect for the cahow. So he created a temporary scaffolding to get one piece of the puzzle going.

Since he needed a forest, he planted 8,000 cedar trees in the hope that a few would be resistant to the blight, and a few were. But the wind smothered them. So Wingate planted a scaffold species -- a fast-growing non-native evergreen, the casuarinas -- as a windbreak around the island. The casuarinas grew rapidly, and let the cedars grow slowly, and over the years, the better-adapted cedars displaced the casuarinas. The resown forest made the perfect home for a night heron which had not been seen on Bermuda for a hundred years. The heron gobbled up land crabs which, without the herons, had become a pest on the islands. The exploded population of land crabs had been feasting on the succulent sprouts of wetland vegetation. The crab's reduced numbers now allowed rare Bermudan sedges to grow, and in recent years, to reseed. It was like the parable of "For Want of a Nail, The Kingdom Was Lost," but in reverse: By finding the nail, the kingdom was won. Notch by notch, Wingate was reassembling a lost ecosystem.

Ecosystems and other functioning systems, like empires, can be destroyed much faster than they can be created. It takes nature time to grow a forest or marsh because even nature can't do everything at once. The kind of assistance Wingate gave is not unnatural. Nature commonly uses interim scaffolding to accomplish many of her achievements. Danny Hillis, an artificial intelligence expert, sees a similar story in the human thumb as a platform for human intelligence. A dexterous hand with a thumb-grasp made intelligence advantageous (for now it could make tools), but once intelligence was established, the hand was not as important. Indeed, Hillis claims, there are many stages needed to build a large system that are not required once the system is running. "Much more apparatus is probably necessary to exercise and evolve intelligence than to sustain it," Hillis wrote. "One can believe in the necessity of the opposable thumb for the development of intelligence without doubting a human capacity for thumbless thought."

When we lie on our backs in an alpine meadow tucked on the perch of high mountains, or wade into the mucky waters of a tidal marsh, we are encountering the "thumbless thoughts" of nature. The intermediate species required to transform the proto-meadow into a regenerating display of flowers are now gone. We are only left now with the thought of flowers and not the helpful thumbs that chaperoned them into being.

You may have heard the heartwarming account of "The Man Who Planted Trees and Grew Happiness." It's about how a forest and happiness were created out of almost nothing. The story is told by a young European man who hikes into a remote area of the Alps in 1910.

The young man wanders into a windy, treeless region, a harsh place whose remaining inhabitants are a few mean, poor, discontented charcoal burners huddled in a couple of dilapidated villages. The hiker meets the only truly happy inhabitant in the area, a lone shepherd hermit. The young man watches in wonder as the hermit wordlessly and idiotically spends his days poking acorns one by one into the moonscape. Every day the silent hermit plants 100 acorns. The hiker departs, eager to leave such desolation, only to return many years later by accident, after the interruption of World War I. He now finds the same village almost unrecognizable in its lushness. The hills are flush with trees and vegetation, brimming with streams, and full of wildlife and a new population of content villagers. Over three decades the hermit had planted 90 square miles thick with oak, beech, and birch trees. His single -- handed work-a mere nudge in the world of nature -- had remodeled the local climate and restored the hopes of many thousands of people.

The only unhappy part of the story is that it is not true. Although it has been reprinted as a true story all over the world, it is, in fact, a fantasy written by a Frenchman for Vogue magazine. There are, however, genuine stories of idealists recreating a forest environment by planting trees in the thousands. And their results confirm the Frenchman's intuition: tiny plants grown on a large scale can divert a local ecosystem in a positive loop of increasing good.

As one true example, in the early 1960s, an eccentric Englishwoman, Wendy Campbell-Purdy, journeyed to North Africa to combat the encroaching sand dunes by planting trees in the desert. She planted a "green wall" of 2,000 trees on 45 acres in Tiznit, Morocco. In six years time, the trees had done so well, she founded a trust to finance the planting of 130,000 more trees on a 260-acre dump in the desert wastes at Bou Saada, Algeria. This too took off, creating a new minihabitat that was suitable for growing citrus, vegetables, and grain.

Given a slim foothold, the remarkable latent power in interconnected green things can launch the law of increasing returns: "Them that has, gets more." Life encourages an environment that encourages yet more life. On Wingate's island the presence of herons enables the presence of sedges. In Packard's prairie the toehold of fire enables the existence of wildflowers which enable the existence of butterflies. In Bou Saada, Algeria, some trees alter the climate and soil to make them fit for more trees. More trees make a space for animals and insects and birds, which prepare a place for yet more trees. Out of acorns, nature makes a machine that provides a luxurious home for people, animals, and plants.

The story of Nonsuch and the other forests of increasing returns, as well as the data from Stuart Pimm's microcosms overlap into a powerful lesson that Pimm calls the Humpty Dumpty Effect. Can we put the Humpty Dumpty of a lost ecosystem together again? Yes, we can if we have all the pieces. But we don't know if we do. There may be chaperone species that catalyze the assembly of an ecosystem in some early stage -- the thumb for intelligence -- that just aren't around the neighborhood anymore. Or, in a real tragedy, a key scaffold species may be globally extinct. One could imagine a hypothetical small, prolific grass essential to creating the matrix out of which the prairie arose, which was wiped out by the last ice age. With it gone, Humpty Dumpty can't be put back together again. "Keep in mind you can't always get there from here," Pimm says.

Packard has contemplated this sad idea. "One of the reasons the prairie may never be fully restored is that some parts are forever gone. Perhaps without the megaherbivores like the mastodon of old or even the bison of yesteryear, the prairie won't come back." Even more scary is yet another conclusion of Pimm's and Drake's work: that it is not just the presence of the right species, in the right order, but the absence of the right species at the right time as well. A mature ecology may be able to tolerate species X easily; but during its assembly, the presence of species X will divert the system onto some other path leading toward a different ecosystem. "That's why," Packard sighs, "it may take a million years to make an ecosystem." Which species now rooted on Nonsuch island or dwelling in the Chicago suburbs might push the reemerging savanna ecosystem away from its original destination?

The rule for machines is counterintuitive but clear: Complex machines must be made incrementally and often indirectly. Don't try to make a functioning mechanical system all at once, in one glorious act of assembly. You have to first make a working system that serves as a platform for the system you really want. To make a mechanical mind, you need to make the equivalent of a mechanical thumb -- a lateral approach that few appreciate. In assembling complexity, the bounty of increasing returns is won by multiple tries over time -- a process anyone would call growth.

Ecologies and organisms have always been grown. Today computer networks and intricate silicon chips are grown too. Even if we owned the blueprints of the existing telephone system, we could not assemble a replacement as huge and reliable as the one we had without in some sense recapitulating its growth from many small working networks into a planetary web.

Creating extremely complex machines, such as robots and software programs of the future, will be like restoring prairies or tropical islands. These intricate constructions will have to be assembled over time because that is the only way to make sure they work from top to bottom. Unripe machinery let out before it is fully grown and fully integrated with diversity will be a common complaint. "We ship no hardware before its time," will not sound funny before too long.

Co-evolution

What color is a chameleon placed on the mirror?

Stewart Brand posed that riddle to Gregory Bateson in the early 1970s. Bateson, together with Norbert Wiener, was a founding father of the modern cybernetic movement. Bateson had a most orthodox Oxford education and a most unorthodox career. He filmed Balinese dance in Indonesia; he studied dolphins; he developed a useful theory of schizophrenia. While in his sixties, he taught at the University of California at Santa Barbara, where his eccentric brilliant views on mental health and evolutionary systems caught the attention of holistically minded hippies.

Stewart Brand, a student of Bateson's, was himself a legendary promoter of cybernetic holism. Brand published his chameleon koan in his Whole Earth Catalog, in 1974. Writes Brand of his riddle: "I asked the question of Gregory Bateson at a point in our interview when we were lost in contemplation of the function, if any, of consciousness -- self-consciousness. Both of us being biologists, we swerved to follow the elusive chameleon. Gregory asserted that the creature would settle at a middle value in its color range. I insisted that the poor beast trying to disappear in a universe of itself would endlessly cycle through a number of its disguises."

The mirror is a clever metaphor for informational circuits. Two ordinary mirrors facing each other will create a fun-house hall that ricochets an image back and forth until it vanishes into an infinite regress. Any message loosed between the two opposing mirrors bounces to exhaustion without changing its form. But what if one side is a responsive mirror, just as the chameleon is, in part reflecting, in part generating? The very act of accommodating itself to its own reflection would disturb it anew. Could it ever settle into a pattern persistent enough to call it something?

Bateson felt the system -- perhaps like self-consciousness -- would quickly settle out at an equilibrium determined by the pull of the creature's many extremes in color. The conflicting colors (and conflicting viewpoints in a society of mind) would compromise upon a "middle value," as if it were a democracy voting. On the other hand, Brand opined that equilibrium of any sort was next to impossible, and that the adaptive system would oscillate without direction or end. He imagined the colors fluctuating chaotically in a random, psychedelic paisley.

The chameleon responding to its own shifting image is an apt analog of the human world of fashion. Taken as a whole, what are fads but the response of a hive mind to its own reflection?

In a 21st-century society wired into instantaneous networks, marketing is the mirror; the collective consumer is the chameleon. What color is the consumer when you put him on the marketplace? Does he dip to the state of the lowest common denominator -- a middle average consumer? Or does he oscillate in mad swings of forever trying to catch up with his own moving reflection?

Bateson was tickled by the depth of the chameleon riddle and passed it on to his other students. One of them, Gerald Hall, proposed a third hypothesis for the final color of the mirror visitor: "The chameleon will stay whatever color he was at the moment he entered the mirror domain."

This is the most logical answer in my view. The coupling between mirror and chameleon is probably so tight and immediate that almost no adaptation is possible. In fact, it may be that once the chameleon bellies up to the mirror, it can't budge from its color unless a change is induced from outside or from an erroneous drift in the chameleon's coloration process. Otherwise, the mirror/chameleon system freezes solidly onto whatever initial value it begins with.

For the mirrored world of marketing, this third answer means the consumer freezes. He either locks onto whatever brand he began with, or he stops purchasing altogether.

There are other possible answers, too. While conducting interviews for this book, I sometimes posed the chameleon riddle to my interviewees. The scientists understood it for the archetypal case of adaptive feedback it was. Their answers ranged over the map. Some examples:

MATHEMATICIAN JOHN HOLLAND: It goes kaleidoscopic! There's a lag time, so it'll flicker all over the place. The chameleon won't ever be a uniform color.

COMPUTER SCIENTIST MARVIN MINSKY: It might have a number of eigenvalues or colors, so it will zero in on a number of colors. If you put it in when it's green it might stay green, and if it was red it might stay red, but if you put it in when it was brown it might tend to go to green.

NATURALIST PETER WARSHALL: A chameleon changes color out of a fright response so it all depends on its emotional state. It might be frightened by its image at first, but then later "warm up" to it, and so change colors.

Putting a chameleon on a mirror seemed a simple enough experiment that I thought that even a writer could perform it. So I did. I built a small, mirrored box, and I bought a color-changing lizard and placed it inside. Although Brand's riddle had been around for 20 years, this was the first time, as far as I know, anyone had actually tried it.

On the mirror the lizard stabilized at one color of green -- the green of young leaves on trees in the spring -- and returned to that one color each time I tried the experiment. But it would spend periods being brown before returning to green. Its resting color in the box was not the same dark brown it seemed to like when out of the mirrored box.

Although I performed this experiment, I place very little confidence in my own results for the following important reasons: the lizard I used was not a true chameleon, but an anole, a species with a far more limited range of color adaptation than a true chameleon. (A true chameleon may cost several hundred dollars and requires a terrarium of a quality I did not want to possess.) More importantly, according to the little literature I read, anoles change colors for other reasons in addition to trying to match their background. As Warshall said, they also alter in response to fright. And frightened it was. The anole did not want to go into the mirrored box. The color green it presented in the box is the same color it uses when it is frightened. It may be that the chameleon in the mirror is merely in a constant state of fright at its own amplified strangeness now filling its universe. I certainly would go nutty in a mirrored box. Finally, there is this observer problem: I can only see the lizard when my face is peeking into the mirrored box, an act which inserts a blue eye and red nose into the anole's universe, a disturbance I could not circumvent.

It may be that an exact answer to the riddle requires future experiments with an authentic chameleon and many more controls than I had. But I doubt it. True chameleons are full-bodied animals just as anoles are, with more than one reason for changing colors. The chameleon on a mirror riddle is best kept in idealized form as a thought experiment.

Even in the abstract, the "real" answer depends on such specific factors as the reaction time of the chameleon's color cells, their sensitivity to a change in hue, and whether other factors influence the signals -- all the usual critical values in feedback circuits. If one could alter these functions in a real chameleon, one could then generate each of the chameleon-on-the-mirror scenarios mentioned above. This, in fact, is what engineers do when they devise electronic control circuits to guide spaceships or steer robot arms. By tweaking delay times, sensitivity to signals, dampening values, etc., they can tailor a system to seek either a wide-ranging equilibrium (say, keeping the temperature between 68 and 70 degrees), or constant change, or some homeostatic point in between.

We see this happening in networked markets. A sweater manufacturer will try to rig a cultural mirror that encourages wild fluctuations in the hopes of selling many styles of sweaters, while a dishwasher manufacture will try to focus the reflections onto the common denominators of only a few dishwasher images, since making varieties of sweaters is much cheaper than making varieties of dishwashers. The type of market is determined by quantity and speed of feedback signals.

The important point about the chameleon on the mirror riddle is that the lizard and glass become one system. "Lizardness" and "mirrorness" are encompassed into a larger essence -- a "lizard-glass" -- which acts differently than either a chameleon or a mirror.

Medieval life was remarkably unnarcissistic. Common folk had only vague notions of their own image in the broad sense. Their individual and social identities were informed by participating in rituals and traditions rather than by reflection. On the other hand, the modern world is being paved with mirrors. We have ubiquitous TV cameras, and ceaseless daily polling ("63 Percent of Us Are Divorced") to mirror back to us every nuance of our collective action. A steady paper trail of bills, grades, pay stubs, and catalogs helps us create our individual identity. Pervasive digitalization of the approaching future promises clearer, faster, and more omnipresent mirrors. Every consumer becomes both a reflection and reflector, a cause and an effect.

The Greek philosophers were obsessed with the chain of causality, how the cause of an effect should be traced back in a relay of hops until one reached the Prime Cause. That backward path is the foundation of Western, linear logic. The lizard-glass demonstrates an entirely different logic -- the circular causality of the Net. In the realm of recursive reflections, an event is not triggered by a chain of being, but by a field of causes reflecting, bending, mirroring each other in a fun-house nonsense. Rather than cause and control being dispensed in a straight line from its origin, it spreads horizontally, like creeping tide, influencing in roundabout, diffuse ways. Small blips can make big splashes, and big blips no splashes. It is as if the filters of distance and time were subverted by the complex connecting of everything to everything.

Computer scientist Danny Hillis has noted that computation, particularly networked computation, exhibits a nonlinear causality field. He wrote:

In the physical universe the effect that one event has on another tends to decrease with the distance in time or in space between them. This allows us to study the motions of the Jovian moons without taking into account the motion of Mercury. It is fundamental to the twin concepts of object and action. Locality of action shows itself in the finite speed of light, in the inverse square law of fields, and in macroscopic statistical effects, such as rates of reaction and the speed of sound.

In computation, or at least in our old models of computation, an arbitrarily small event can and often does cause an arbitrarily large effect. A tiny program can clear all of memory. A single instruction can stop the machine. In computation there is no analog of distance. One memory location is as easily influenced as another.

The lines of control in natural ecologies also dissolve into a causality horizon. Control is not only distributed in space, but it is also blurred in time as well. When the chameleon steps onto the mirror, the cause of his color dissolves into a field of effects spinning back on themselves. The reasons for things do not proceed like an arrow, but rather spread to the side like a wind.

Stewart Brand majored in biology at Stanford, where his teacher was Paul Ehrlich, a population biologist. Ehrlich too was fascinated by the rubbery chameleon-on-the-mirror paradox. He saw it most vividly in the relationship between a butterfly and its host plant. Fanatical butterfly collectors had long ago figured out that the best way to get perfect specimens was to encase a caterpillar, along with a plant it feeds on, in a box while waiting for the larvae to metamorphose. After transformation, the butterfly would emerge in the box sporting flawless unworn wings. It would be immediately killed and mounted.

This method required that collectors figure out which plants butterflies ate. With the prospect of perfect specimens, they did this thoroughly. The result was a rich literature of plant/butterfly communities, whose summary indicated that many butterflies in the larvae stage chomp on only one specific plant. Monarch caterpillars, for instance, devour only milkweeds. And, it seemed, the milkweed invited only the monarch to dine on it.

Ehrlich noticed that in this sense the butterfly was reflected in the plant, and the plant was reflected in the butterfly. Every step the milkweed took to keep the monarch larvae at bay so the worm wouldn't devour it completely, forced the monarch to "change colors" and devise a way to circumvent the plant's defenses. The mutual reflections became a dance of two chameleons belly to belly. In defending itself so thoroughly against the monarch, the milkweed became inseparable from the butterfly. And vice versa. Any long-term antagonistic relationship seemed to harbor this kind of codependency. In 1952, W. Ross Ashby, a cybernetician interested in how machines could learn, wrote, "[An organism's gene-pattern] does not specify in detail how a kitten shall catch a mouse, but provides a learning mechanism and a tendency to play, so that it is the mouse which teaches the kitten the finer points of how to catch mice."

Ehrlich came across a word to describe this tightly coupled dance in the title of a 1958 paper by C. J. Mode in the journal Evolution. It was called "coevolution," as in "A mathematical model for the co-evolution of obligate parasites and their hosts." Like most biological observations, the notion of coevolution was not new. The amazing Darwin himself wrote of "coadaptions of organic beings to each other..." in his 1859 masterpiece Origin of Species.

The formal definition of coevolution runs something like this: "Coevolution is reciprocal evolutionary change in interacting species," says John Thompson in Interaction and Coevolution. But what actually happens is more like a tango. The milkweed and monarch, shoulder to shoulder, lock into a single system, an evolution toward and with each other. Every step of coevolutionary advance winds the two antagonists more inseparably, until each is wholly dependent on the other's antagonism. The two become one. Biochemist James Lovelock writes of this embrace, "The evolution of a species is inseparable from the evolution of its environment. The two processes are tightly coupled as a single indivisible process."

Brand picked up the term and launched a magazine called CoEvolution Quarterly. It was devoted to the larger notion of all things -- biological, societal, and technological -- adapting to and creating each other, and at the same time weaving into one whole system. As an introduction Brand penned a definition: "Evolution is adapting to meet one's needs. Coevolution, the larger view, is adapting to meet each other's needs."

The "co" in coevolution is the mark of the future. In spite of complaints about the steady demise of interpersonal relationships, the lives of modern people are increasingly more codependent than ever. All politics these days means global politics and global politics means copolitics. The new online communities built between the spaces of communication networks are coworlds. Marshall McLuhan was not quite right. We are not hammering together a cozy global village. We are weaving together a crowded global hive -- a coworld of utmost sociality and mirrorlike reciprocation. In this environment, all evolution, including the evolution of manufactured entities, is coevolution. Nothing changes without also moving closer to its changing neighbors.

Nature is chock-a-block with coevolution. Every green corner sports parasites, symbionts, and tightly coupled dances. Biologist P. W. Price estimated that over 50 percent of today's species are parasitic. (The figure has risen from the deep paleologic past and is expected to keep rising.) Here's news: half of the living world is codependent! Business consultants commonly warn their clients against becoming a symbiont company dependent upon a single customer-company, or a single supplier. But many do, and as far as I can tell, live profitable lives, no shorter on average than other companies. The surge of alliance-making in the 1990s among large corporations -- particularly among those in the information and network industries -- is another facet of an increasing coevolutionary economic world. Rather than eat or compete with a competitor, the two form an alliance -- a symbiosis.

The parties in a symbiosis don't have to be symmetrical or even at parity. In fact, biologists have found that almost all symbiotic alliances in nature entail a greater advantage for one party -- in effect some hint of parasitism -- in every codependency. But even though one side gains at the expense of the other, both sides gain over all, and so the pact continues.

In his magazine CoEvolution Brand began collecting stories of coevolutionary games. One of the most illustrative examples of alliance making in nature is the following:

In eastern Mexico live a variety of acacia shrubs and marauding ants. Most acacias have thorns, bitter leaves, and other protection against a hungry world. One, the "swollen thorn acacia," learned to encourage a species of ant to monopolize it as a food source and kill or run off all other predators. Enticements gradually included nifty water-proof swollen thorns to live in, handy nectar fountains, and special ant-food buds at the leaf tips. The ants, whose interests increasingly coincided with the acacia's, learned to inhabit the thorns, patrol the acacia day and night, attack every acacia-hungry organism, and even prune away invading plants such as vines and tree seedlings that might shade Mother Acacia. The acacia gave up its bitter leaves, sharp thorns, and other devices and now requires the acacia-ant for survival. And the ant colonies can no longer live without the acacia. Together they're unbeatable.

In evolutionary time, the instances of coevolution have increased as sociability in life has increased. The more copious life's social behaviors are, the more likely they are to be subverted into mutually beneficial interactions. The more mutually responsive we construct our economic and material world, the more coevolutionary games we'll see.

Parasitic behavior itself is a new territory for organisms to make a living in. Thus we find parasites upon parasites. Ecologist John Thompson notes that "just as the richness of social behaviors may increase mutualism with other species, so may some mutualisms allow for the evolution of new social behaviors." In true coevolutionary fashion, coevolution breeds coevolution.

A billion years from now life on Earth may be primarily social, and stuffed with parasites and symbionts; and the world economy may be primarily a crowded network of alliances. What happens, then, when coevolution saturates a complete planet? What does a sphere of reflecting, responsive, coadapting, and recursive bits of life looping back upon itself do?

The butterfly and the milkweed constantly dance around each other, and by this ceaseless crazed ballet they move far beyond the forms they would have if they were at peace with each other. The chameleon on the mirror flipping without rest slips into some deranged state far from sanity. There is a sort of madness in pursuing self-reflections, that same madness we sensed in the nuclear arms race of post-World War II. Coevolution moves things to the absurd. The butterfly and the milkweed, although competitors in a way, cannot live apart. Paul Ehrlich sees coevolution pushing two competitors into "obligate cooperation." He wrote, "It's against the interests of either predator or prey to eliminate the enemy." That is clearly irrational, yet that is clearly a force that drives nature.

When a human mind goes off the deep end and gets stuck in the spiral of watching itself watching a mirror, or becomes so dependent upon its enemies that it apes them, then we declare it insane. Yet there is a touch of insanity -- a touch of the off-balance -- in intelligence and consciousness itself. To some extent a mind, even a primitive mind, must watch itself. Must any consciousness stare at its own navel?

This was the point in the conversation when Stewart Brand pointed out to Gregory Bateson his fine riddle of the chameleon on the mirror, and the two biologists swerved to follow it. The chase arrives at the odd conclusion that consciousness, life, intelligence, coevolution are off-balanced, unexpected, even unreasonable, given the resting point of everything else. We find intelligence and life spooky because they maintain a precarious state far from equilibrium. Compared to the rest of the universe, intelligence and consciousness and life are stable instabilities.

They are held together, poised upright like a pencil standing on its point, by the recursive dynamics of coevolution. The butterfly pushes the milkweed, and the milkweed pushes the butterfly, and the harder they push the more impossible it becomes for them to let go, until the whole butterfly/milkweed thing emerges as its own being -- a living insect/plant system-pulling itself up by its bootstraps.

Rabid mutualism doesn't just happen in pairs. Threesomes can meld into an emergent, coevolutionarily wired symbiosis. Whole communities can be coevolutionary. In fact, any organism that adapts to organisms around it will act as an indirect coevolutionary agent to some degree. Since all organisms adapt that means all organisms in an ecosystem partake in a continuum of coevolution, from direct symbiosis to indirect mutual influence. The force of coevolutionism flows from one creature to its most intimate neighbors, and then ripples out in fainter waves until it immeasurably touches all living organisms. In this way the loose network of a billion species on this home planet are knit together so that unraveling the coevolutionary fabric becomes impossible, and the parts elevate themselves into some aggregate state of spooky, stable instability.

The network of life on Earth, like all distributed being, transcends the life of its ingredients. But bully life reaches deeper and ties up the entire planet in the web of its network, also roping in the nonliving matrix of rock and gas into its coevolutionary antics.

Thirty years ago, biologists asked NASA to shoot a couple of unmanned probes towards the two likeliest candidates for extraterrestrial life, Mars and Venus, and poke a dipstick into their soil to check for vital signs.

The life-meter that NASA came up with was a complicated, delicate (and expensive) contraption that would, upon landing, be sprinkled with a planet's soil and check for evidence of bacterial life. One of the consultants hired by NASA was a soft-spoken British biochemist, James Lovelock, who found that he had a better way of checking for life on planets, a method that did not require a multimillion-dollar gadget, or even a rocket at all.

Lovelock was very rare breed in modern science. He practiced science as a maverick, working out of a stone barn among the rural hedgerows in Cornwall, England. He maintained a spotless scientific reputation, yet he had no formal institutional affiliation, a rarity in the heavily funded world of science. His stark independence both nurtured and demanded free thinking. In the early 1960s Lovelock came up with a radical proposal that irked the rest of the folks on the NASA probe team. They really wanted to land a meter on a another planet. He said they didn't have to bother.

Lovelock told them he could determine whether there was life on a planet by looking through a telescope. He could measure the spectrum of a planet's atmosphere, and thereby determine its composition. The makeup of the bubble of gases surrounding a planet would yield the secret of whether life inhabited the sphere. You therefore didn't need to hurl an expensive canister across the solar system to find out. He already knew the answer.

In 1967, Lovelock wrote two papers predicting that Mars would be lifeless based on his interpretation of its atmosphere. The NASA orbiters that circled Mars later in the decade, and the spectacular Mars soft landings the decade following made it clear to everyone that Mars was indeed as dead as Lovelock had forecasted. Equivalent probes to Venus brought back the same bad news: the solar system was barren outside of Earth.

How did Lovelock know?

Chemistry and coevolution. When the compounds in the Martian atmosphere and soil were energized by the sun's rays, and heated by the planetary core, and then contained by the Martian gravity, they settled into a dynamic equilibrium after millions of years. The ordinary laws of chemistry permit a scientist to make calculations of their reactions as if the planet were a large flask of matter. When a chemist derives the approximate formulas for Mars, Venus, and the other planets, the equations roughly balance: energy, compounds in; energy, compounds out. The measurements from the telescopes, and later the probes, matched the results predicted by the equations.

Not so the Earth. The mixture of gases in the atmosphere of the Earth are way out of whack. And they are out of whack, Lovelock was to find out, because of the curious accumulative effects of coevolution.

Oxygen in particular, at 21 percent, makes the Earth's atmosphere unstable. Oxygen is a highly reactive gas, combining with many elements in a fierce explosive union we call fire or burning. Thermodynamically, the high oxygen content of Earth's atmosphere should fall quickly as the gas oxidizes surface solids. Other reactive trace gases such as nitrous oxide and methyl iodide also remain at elevated and aberrant levels. Both oxygen and methane coexist, yet they are profoundly incompatible, or rather too compatible since they should burn each other up. Carbon dioxide is inexplicably a mere trace gas when it should be the bulk of the air, as it is on other planets. In addition to its atmosphere, the temperature and alkalinity of the Earth's surface also exhibits a queer level. The entire surface of the Earth seems to be a vast unstable chemical anomaly.

It seemed to Lovelock as if an invisible power, an invisible hand, pushed the interacting chemical reactions into a raised state that should at any minute swing back to a balanced rest. The chemistry of Mars and Venus was as balanced as the periodic table, and as dead. The chemistry of the Earth was out of kilter, wholly unbalanced by the periodic table, and alive. From this, Lovelock concluded that any planet that has life would reveal a chemistry that held odd imbalances. A life-friendly atmosphere might not be oxygen-rich, but it should buck textbook equilibria.

That invisible hand was coevolutionary life.

Life in coevolution, which has the remarkable knack of generating stable instability, moved the chemical circuitry of the Earth's atmosphere into what Lovelock calls a "persistent state of disequilibrium." At any moment, the atmosphere should fall, but for millions of years it doesn't. Since high oxygen levels are needed for most microbial life, and since microbial fossils are billions of years old, this odd state of discordant harmony has been quite persistent and stable.

The Earth's atmosphere seeks a steady oxygen level much as a thermostat hones in on a steady temperature. The uniform 20 percent oxygen level it has found turns out to be "fortuitous" as one scientist put it. Lower oxygen would be anemic, while greater oxygen would be too flammable. George R. Williams at the University of Toronto writes: "An O2 content of about 20 percent seems to ensure a balance between almost complete ventilation of the oceans without incurring greater risks of toxicity or increased combustibility of organic material." But where are the sensors and the thermostatic control mechanisms? For that matter, where is the furnace?

Dead planets find equilibrium by geological circuits. Gases, such as carbon dioxide, dissolve in liquids and can precipitate out as solids. Only so much gas will dissolve before it reaches a natural saturation. Solids can release gases back into the atmosphere when heated and pressed by volcanic activity. Sedimentation, weathering, uplift -- all the grand geological forces -- also act as strong chemical agents, breaking and making the bonds of materials. Thermodynamic entropy draws all chemical reactions down to their minimal energy level. The furnace metaphor breaks down. Equilibrium on a dead planet is less like a thermostat and more like the uniform level of water in a bowl; it simply levels out when it can't get any lower.

But the Earth has the self of a thermostat. A spontaneous circuit, provided by the coevolutionary tangle of life, which guides the chemicals of the planet toward some elevated potential. Presumably if all life on Earth were extinguished, the Earth's atmosphere would fall back to a persistent equilibrium, and become as boringly predictable as Mars and Venus. But as long as the distributed hand of life dominates, it will keep the chemicals of Earth off key.

Yet the off-balance is itself balanced. The persistent disequilibrium that coevolutionary life generates, and that Lovelock seeks as an acid test for its presence, is stable in its own way. As far as we can tell Earth's atmospheric oxygen has remained at about 20 percent for hundreds of millions of years. The atmosphere acts not merely as an acrobat on a tightrope pitched far from the vertical, but as an acrobat teetering between tilting and falling, and poised there for millions of years. She never falls, but never gets out of falling. It's a state of permanent almost-fell.

Lovelock recognized that persistent almost-fell is a hallmark of life. Recently complexity investigators have recognized that persistent almost-fell is a hallmark of any vivisystem: an economy, a natural ecosystem, a deep computer simulation, an immune system, or an evolutionary system. All share that paradoxical quality of working best when they remain poised in an Escher-like state of forever descending without ever being lowered. They remain poised in the act of collapsing.

David Layzer, writing in his semiscientific book Cosmogenesis, argues that "the central property of life is not reproductive invariance, but reproductive instability." The key of life is its ability to reproduce slightly out of kilter rather than with exactitude. This almost-falling into chaos keeps life proliferating.

A little noticed but central character of such vivisystems is that this paradoxical essence is contagious. Vivisystems spread their poised instability into whatever they touch, and they reach for everything. On Earth, life elbows its way into solid, liquid, gas. No rocks, to our knowledge, are untouched by life in former times. Tiny oceanic microorganisms solidify carbon and oxygen gases dissolved in sea water to produce a salt which settles on the sea floor. The deposits eventually become pressed under sedimentary weight into stone. Tiny plant organisms transport carbon from the air into soil and lower into the sea bottom, to be submerged and fossilized into oil. Life generates methane, ammonia, oxygen, hydrogen, carbon dioxide, and many other gases. Iron -- and metal-concentrating bacteria create metallic ores. (Iron, the very emblem of nonlife, born of life!) Upon close inspection, geologists have concluded that all rocks residing on the Earth's surface (except perhaps volcanic lava) are recycled sediments, and therefore all rocks are biogenic in nature, that is, in some way affected by life. The relentless push and pull of coevolutionary life eventually brings into its game the abiotic stuff of the universe. It makes even the rocks part of its dancing mirror.

One of the first to articulate the transcendent view that life directly shaped the physicality of this planet was the Russian geologist Vladimir Vernadsky, writing in 1926. Vernadsky tallied up the billions of organisms on Earth and considered their collective impact upon the material resources of the planet. He called this grand system of resources the "biosphere," (although Eduard Suess had coined the term a few years earlier) and set out to measure it quantitatively in his book The Biosphere, a volume only recently translated into English.

In articulating life as a chameleon on a rocky mirror, Vernadsky committed heresy on two counts. He enraged biologists by considering the biosphere of living creatures as a large chemical factory. Plants and animals were mere temporary chemical storage units for the massive flow of minerals around the world. "Living matter is a specific kind of rock...an ancient and, at the same time, an eternally young rock," Vernadsky wrote. Living creatures were delicate shells to hold these minerals. "The purpose of animals," he once said of their locomotion and movement, "is to assist the wind and waves to stir the brewing biosphere."

At the same time, Vernadsky enraged geologists by considering rocks as if they were half-alive. Since the genesis of every rock was in life, their gradual interaction with living organisms meant that rocks were the part of life that moved the slowest. The mountains, the waters of the ocean, and the gases of the sky were very slow life. Naturally, geologists balked at this apparent mysticism.

The two heresies melded into a beautiful symmetry. Life as ever-renewing mineral, and minerals as slow life. They could only be opposite sides of a single coin. The two sides of this equation cannot be mathematically unraveled; they are one system: lizard-mirror, plant/insect, rock-life, and now in modern times, human/machine. The organism behaves as environment, the environment behaves as organism.

This has been a venerable idea at the edge of science for at least several hundred years. Many evolutionary biologists in the last century such as T. H. Huxley, Herbert Spencer, and Darwin, too, understood it intuitively -- that the physical environment shapes its creatures and the creatures shape their environment, and if considered in the long view, the environment is the organism and the organism is the environment. Alfred Lotka, an early theoretical biologist, wrote in 1925, "It is not so much the organism or the species that evolves, but the entire system, species plus environment. The two are inseparable." The entire system of evolving life and planet was coevolution, the dance of the chameleon on the mirror.

If life were to vanish from Earth, Vernadsky realized, not only would the planet sink back into the "chemical calm" of an equilibrium state, but the clay deposits, limestone caves, ores in mine, chalk cliffs, and the very structure of all that we consider the Earth's landscape would retreat. "Life is not an external and accidental development on the terrestrial surface. Rather, it is intimately related with the constitution of the Earth's crust," Vernadsky wrote in 1929. "Without life, the face of the Earth would become as motionless and inert as the face of the moon."

Three decades later, free-thinker James Lovelock arrived at the same conclusions based on his telescopic analysis of other planets. Lovelock observed, "In no way do organisms simply 'adapt' to a dead world determined by physics and chemistry alone. They live in a world that is the breath and bones of their ancestors and that they are now sustaining." Lovelock had more complete knowledge of early Earth than was available to Vernadsky, and a slightly better understanding of the global patterns of gases and material flows on Earth. All this led him to suggest in complete seriousness that "the air we breathe, the oceans, and the rocks are all either the direct products of living organisms or else have been greatly modified by their presence."

Such a remarkable conclusion was foreshadowed by the French natural philosopher, Jean Baptiste Lamarck, who in 1800 had even less information about planetary dynamics than Vernadsky did. As a biologist, Lamarck was equal to Darwin. He, not Darwin, was the true discoverer of evolution, but Lamarck is stuck with an undeserved reputation as a loser, in part because he relied a little too much on intuition rather than the modern notion of detailed facts. Lamarck made an intuitive guess about the biosphere and again was prescient. Since there wasn't a shred of scientific evidence to support Lamarck's claims at the time, his observations were not influential. He wrote in 1802, "Complex mineral substances of all kinds that constitute the external crust of the Earth occurring in the form of individual accumulations, ore bodies, parallel strata, etc., and forming lowlands, hills, valleys, and mountains, are exclusively products of the animals and plants that existed within these areas of the Earth's surface."

The bold claims of Lamarck, Vernadsky, and Lovelock seem ludicrous at first, but in the calculus of lateral causality make fine sense: that all we can see around us -- the snow-covered Himalayas, the deep oceans east and west, vistas of rolling hills, awesome painted desert canyons, game-filled valleys -- are all as much the product of life as the honeycomb.

Lovelock kept gazing into the mirror and finding that it was nearly bottomless. As he examined the biosphere in succeeding years, he added more complex phenomena to the list of life-made. Some examples: plankton in the oceans release a gas (DMS) which oxidizes to produce submicroscopic aerosols of sulfate salts which form nuclei for the condensation of cloud droplets. Thus perhaps even clouds and rain may be biogenic. Summer thunderstorms may be life raining on itself. Some studies hinted that a majority of nuclei in snow crystals may be decayed vegetation, bacteria, or fungi spores; and so snow may be largely life-triggered. Only very little could escape life's imprint. "It may be that the core of our planet is unchanged as a result of life; but it would be unwise to assume it," Lovelock said.

"Living matter is the most powerful geological force," Vernadsky claimed, "and it is growing with time." The more life, the greater its material force. Humans intensify life further. We harness fossil energy and breathe life into machines. Our entire manufactured infrastructure -- as an extension of our own bodies -- becomes part of a wider, global-scale life. As the carbon dioxide from our industry pours into the air and alters the global air mix, the realm of our artificial machines also becomes part of the planetary life. Jonathan Weiner writing in The Next One Hundred Years then can rightly say, "The Industrial Revolution was an astonishing geological event." If rocks are slow life, then our machines are quicker slow life.

The Earth as mother was an old and comforting notion. But the Earth as mechanical device has been a harder idea to swallow. Vernadsky came very close to Lovelock's epiphany that the Earth's biosphere exhibits a regulation beyond chemical equilibrium. Vernadsky noted that "organisms exhibit a type of self-government" and that the biosphere seemed to be self-governed, but Vernadsky didn't press further because the crucial concept of self-government as a purely mechanical process had not yet been uncovered. How could a mere machine control itself?

We now know that self-control and self-governance are not mystical vital spirits found only in life because we have built machines that contain them. Rather, control and purpose are purely logical processes that can emerge in any sufficiently complex medium, including that of iron gears and levers, or even complex chemical pathways. If a thermostat or a steam engine can own self-governance, the idea of a planet evolving such graceful feedback circuits is not so alien.

Lovelock brought an engineer's sensibilities to the analysis of Mother Earth. He was a tinkerer, inventor, patent holder, and had worked for the biggest engineering firm of all time, NASA. In 1972, Lovelock offered a hypothesis of where the planet's self-government lay. He wrote, "The entire range of living matter on Earth, from whales to viruses, from oaks to algae, could be regarded as constituting a single living entity, capable of manipulating the Earth's atmosphere to suit its overall needs and endowed with faculties and powers far beyond those of its constituent part." Lovelock called this view Gaia. Together with microbiologist Lynn Margulis, the two published the view in 1972 so that it could be critiqued on scientific terms. Lovelock says, "The Gaia theory is a bit stronger than coevolution," at least as biologists use the word.

A pair of coevolutionary creatures chasing each other in an escalating arms race can only seem to veer out of control. Likewise, a pair of cozy coevolutionary symbionts embracing each other can only seem to lead to stagnant solipsism. But Lovelock saw that if you had a vast network of coevolutionary impulses, such that no creatures could escape creating its own substrate and the substrate its own creatures, then the web of coevolution spread around until it closed a circuit of self-making and self-control. The "obligate cooperation" of Ehrlich's coevolution -- whether of mutual enemies or mutual partners -- cannot only raise an emergent cohesion out of the parts, but this cohesion can actively temper its own extremes and thereby seek its own survival. The solidarity produced by a planetary field of creatures mirrored in a coevolving environment and each other is what Lovelock means by Gaia.

Many biologists (including Paul Ehrlich) are unhappy with the idea of Gaia because Lovelock expanded the definition of life without asking their permission. He unilaterally enlarged life's scope to include a predominantly mechanical apparatus. In one easy word, a solid planet became "the largest manifestation of life" that we know. It is an odd beast: 99.9 percent rock, a lot of water, and a little air, wrapped up in the thinnest green film that would stretch around it.

But if Earth is reduced to the size of a bacteria, and inspected under high-powered optics, would it seem stranger than a virus? Gaia hovers there, a blue sphere under the stark light, inhaling energy, regulating its internal states, fending off disturbances, complexifying, and ready to transform another planet if given a chance.

While Lovelock backs off earlier assertions that Gaia is an organism, or acts as if it is one, he maintains that it really is a system that has living characteristics. It is a vivisystem. It is a system that is alive, whether or not it possesses all the attributes needed for an organism.

That Gaia is made up of many purely mechanical circuits shouldn't deter us from applying the label of life. After all, cells are mostly chemical cycles. Some ocean diatoms are mostly inert, crystallized calcium. Trees are mostly dead pulp. But they are still living organisms.

Gaia is a bounded whole. As a living system, its inert, mechanistic parts are part of its life. Lovelock: "There is no clear distinction anywhere on the Earth's surface between living and nonliving matter. There is merely a hierarchy of intensity going from the material environment of the rocks and atmosphere to the living cells." Somewhere at the boundary of Gaia, either in the rarefied airs of the stratosphere or deep in the Earth's molten core, the effects of life fade. No one can say where that line is, if there is a line.

The trouble with Gaia, as far as most skeptics are concerned, is that it makes a dead planet into a "smart" machine. We already are stymied in trying to design an artificial learning machine from inert computers, so the prospect of artificial learning evolving unbidden at a planetary scale seems ludicrous.

But learning is overrated as something difficult to evolve. This may have to do with our chauvinistic attachment to learning as an exclusive mark of our species. There is a strong sense, which I hope to demonstrate in this book, in which evolution itself is a type of learning. Therefore learning occurs wherever evolution is, even if artificially.

The dethronement of learning is one of the most exciting intellectual frontiers we are now crossing. In a virtual cyclotron, learning is being smashed into its primitives. Scientists are cataloguing the elemental components for adaptation, induction, intelligence, evolution, and coevolution into a periodic table of life. The particles for learning lie everywhere in all inert media, waiting to be assembled (and often self-assembled) into something that surges and quivers.

Coevolution is a variety of learning. Stewart Brand wrote in CoEvolution Quarterly: "Ecology is a whole system, alright, but coevolution is a whole system in time. The health of it is forward-systemic self-education which feeds on constant imperfection. Ecology maintains. Coevolution learns."

Colearning might be a better term for what coevolving creatures do. Coteaching also works, for the participants in coevolution are both learning and teaching each other at the same time. (We don't have a word for learning and teaching at the same time, but our schooling would improve if we did.)

The give and take of a coevolutionary relationship-teaching and learning at once-reminded many scientists of game playing. A simple child's game such as "Which hand is the penny in?" takes on the recursive logic of a chameleon on a mirror as the hider goes through this open-ended routine: "I just hid the penny in my right hand, and now the guesser will think it's in my left, so I'll move it into my right. But she also knows that I know she knows that, so I'll keep it in my left."

Since the guesser goes through a similar process, the players form a system of mutual second-guessing. The riddle "What hand is the penny in?" is related to the riddle, "What color is the chameleon on a mirror?" The bottomless complexity which grows out of such simple rules intrigued John von Neumann, the mathematician who developed programmable logic for a computer in the early 1940s, and along with Wiener and Bateson launched the field of cybernetics.

Von Neumann invented a mathematical theory of games. He defined a game as a conflict of interests resolved by the accumulative choices players make while trying to anticipate each other. He called his 1944 book (coauthored by economist Oskar Morgenstern) Theory of Games and Economic Behavior because he perceived that economies possessed a highly coevolutionary and gamelike character, which he hoped to illuminate with simple game dynamics. The price of eggs, say, is determined by mutual second-guessing between seller and buyer-how much will he accept, how much does he think I will offer, how much less than what I am willing to pay should I offer? The aspect von Neumann found amazing was that this infinite regress of mutual bluffing, codeception, imitation, reflection, and "game playing" would commonly settle down to a definite price, rather than spiral on forever. Even in a stock market made of thousands of mutual second-guessing agents, the group of conflicting interests would quickly settle on a price that was fairly stable.

Von Neumann was particularly interested in seeing if he could develop optimal strategies for these kinds of mutual games, because at first glance they seemed almost insolvable in theory. As an answer he came up with a theory of games. Researchers at the U.S. government-funded RAND corporation, a think tank based in Santa Monica, California, extended von Neumann's initial work and eventually catalogued four basic varieties of mutual second-guessing games. Each variety had a different structure of rewards for winning, losing, or drawing. The four simple games were called "social dilemmas" in the technical literature, but could be thought of as the four building blocks of complicated coevolutionary games. They were: Chicken, Stag Hunt, Deadlock, and the Prisoner's Dilemma

Chicken is the game played by teenage daredevils. Two cars race toward a cliff's edge; the driver who jumps out last, wins. Stag Hunt is the dilemma faced by a bunch of hunters who must cooperate to kill a stag, but may do better sneaking off by themselves to hunt a rabbit if no one cooperates. Do they gamble on cooperation (high payoff) or defection (low, but sure payoff)? Deadlock is a boring game where mutual defection pays best. The last one, the Prisoner's Dilemma, is the most illuminating, and became the guinea pig model for over 200 published social psychology experiments in the late 1960s.

The Prisoner's Dilemma, invented in 1950 by Merrill Flood at RAND, is a game for two separately held prisoners who must independently decide whether to deny or confess to a crime. If both confess, each will be fined. If neither confesses, both go free. But if only one should confess, he is rewarded while the other is fined. Cooperation pays, but so does betrayal, if played right. What would you do?

Played only once, betrayal of the other is the soundest choice. But when two "prisoners" played the game over and over, learning from each other-a game known as the Iterated Prisoner Dilemma-the dynamics of the game shifted. The other player could not be dismissed; he demanded to be attended to, either as obligate enemy or obligate colleague. This tight mutual destiny closely paralleled the coevolutionary relationship of political enemies, business competitors, or biological symbionts. As study of this simple game progressed, the larger question became, What were the strategies of play for the Iterated Prisoner's Dilemma that resulted in the highest scores over the long term? And what strategies succeeded when played against many varieties of players, from the ruthless to the kind?

In 1980, Robert Axelrod, a political science professor at University of Michigan, ran a tournament pitting 14 submitted strategies of Prisoner's Dilemma against each other in a round robin to see which one would triumph. The winner was a very simple strategy crafted by psychologist Anatol Rapoport called Tit-For-Tat. The Tit-For-Tat strategy prescribed reciprocating cooperation for cooperation, and defection for defection, and tended to engender periods of cooperation. Axelrod had discovered that "the shadow of the future," cast by playing a game repeatedly rather than once, encouraged cooperation, because it made sense for a player to cooperate now in order to ensure cooperation from others later. This glimpse of cooperation set Axelrod on this quest: "Under what conditions will cooperation emerge in a world of egoists without central authority?"

For centuries, the orthodox political reasoning originally articulated by Thomas Hobbes in 1651 was dogma: that cooperation could only develop with the help of a benign central authority. Without top-down government, Hobbes claimed, there would be only collective selfishness. A strong hand had to bring forth political altruism, whatever the tone of economics. But the democracies of the West, beginning with the American and French Revolutions, suggested that societies with good communications could develop cooperative structures without heavy central control. Cooperation can emerge out of self-interest. In our postindustrial economy, spontaneous cooperation is a regular occurrence. Widespread industry-initiated standards (both of quality and protocols such as 110 volts or ASCII) and the rise of the Internet, the largest working anarchy in the world, have only intensified interest in the conditions necessary for hatching coevolutionary cooperation.

This cooperation is not a new age spiritualism. Rather it is what Axelrod calls "cooperation without friendship or foresight"-cold principles of nature that work at many levels to birth a self-organizing structure. Sort of cooperation whether you want it or not.

Games such as Prisoner's Dilemma can be played by any kind of adaptive agent-not just humans. Bacteria, armadillos, or computer transistors can make choices according to various reward schemes, weighing immediate sure gain over future greater but riskier gain. Played over time with the same partners, the results are both a game and a type of coevolution.

Every complex adaptive organization faces a fundamental tradeoff. A creature must balance perfecting a skill or trait (building up legs to run faster) against experimenting with new traits (wings). It can never do all things at once. This daily dilemma is labeled the tradeoff between exploration and exploitation. Axelrod makes an analogy with a hospital: "On average you can expect a new medical drug to have a lower payoff than exploiting an established medication to its limits. But if you gave every patient the current best drug, you'd never get proven new drugs. From an individual's point of view you should never do the exploration. But from the society of individuals' point of view, you ought to try some experiments." How much to explore (gain for the future) versus how much to exploit (sure bet now) is the game a hospital has to play. Living organisms have a similar tradeoff in deciding how much mutation and innovation is needed to keep up with a changing environment. When they play the tradeoff against a sea of other creatures making similar tradeoffs, it becomes a coevolutionary game.

Axelrod's 14-player Prisoner's Dilemma round robin tournament was played on a computer. In 1987, Axelrod extended the computerization of the game by setting up a system in which small populations of programs played randomly generated Prisoner's Dilemma strategies. Each random strategy would be scored after a round of playing against all the other strategies running; the ones with the highest scores got copied the most to the next generation, so that the most successful strategies propagated. Because many strategies could succeed only by "preying" on other strategies, they would thrive only as long as their prey survived. This leads to the oscillating dynamics found everywhere in the wilds of nature; how fox and hare populations rise and fall over the years in coevolutionary circularity. When the hares increase the foxes boom; when the foxes boom, the hares die off. But when there are no hares, the foxes starve. When there are less foxes, the hares increase. And when the hares increase the foxes do too, and so on.

In 1990, Kristian Lindgren, working at the Niels Bohr Institute in Copenhagen, expanded these coevolutionary experiments by increasing the population of players to 1,000, introducing random noise into the games, and letting this artificial coevolution run for up to 30,000 generations. Lindgren found that masses of dumb agents playing Prisoner's Dilemma not only reenacted the ecological oscillations of fox and hare, but the populations also created many other natural phenomenon such as parasitism, spontaneously emerging symbiosis, and long-term stable coexistence between species, as if they were an ecology. Lindgren's work excited some biologists because his very long runs displayed long periods when the mix of different "species" of strategy was very stable. These historical epochs were interrupted by very sudden, short-lived episodes of instability, when old species went extinct and new ones took root. Quickly a new stable arrangement of new species of strategies arose and persisted for many thousands of generations. This motif matches the general pattern of evolution found in earthly fossils, a pattern known in the evolutionary trade as punctuated equilibrium, or "punk eek" for short.

One marvelous result from these experiments bears consideration by anyone hoping to manage coevolutionary forces. It's another law of the gods. It turns out that no matter what clever strategy you engineer or evolve in a world laced by chameleon-on-a-mirror loops, if it is applied as a perfectly pure rule that you obey absolutely, it will not be evolutionary resilient to competing strategies. That is, a competing strategy will figure out how to exploit your rule in the long run. A little touch of randomness (mistakes, imperfections), on the other hand, actually creates long-term stability in coevolutionary worlds by allowing some strategies to prevail for relative eons by not being so easily aped. Without noise-wholly unexpected and out-of-character choices-the opportunity for escalating evolution is lost because there are not enough periods of stability to keep the system going. Error keeps the glue of coevolutionary relationships from binding too tightly into runaway death spirals, and therefore error keeps a coevolutionary system afloat and moving forward. Honor thy error.

Playing coevolutionary games in computers has provided other lessons. One of the few notions from game theory to penetrate the popular culture was the distinction of zero-sum and nonzero-sum games. Chess, elections, races, and poker are zero-sum games: the winner's earnings are deducted from the loser's assets. Natural wilderness, the economy, a mind, and networks on the other hand, are nonzero-sum games. Wolverines don't have to lose just because bears live. The highly connected loops of coevolutionary conflict mean the whole can reward (or at times cripple) all members. Axelrod told me, "One of the earliest and most important insights from game theory was that nonzero-sum games had very different strategic implications than zero-sum games. In zero-sum games whatever hurts the other guy is good for you. In nonzero-sum games you can both do well, or both do poorly. I think people often take a zero-sum view of the world when they shouldn't. They often say, 'Well I'm doing better than the other guy, therefore I must be doing well.' In a nonzero-sum you could be doing better than the other guy and both be doing terribly."

Axelrod noticed that the champion Tit-For-Tat strategy always won without exploiting an opponent's strategy-it merely mirrored the other's actions. Tit-For-Tat could not beat anyone's strategy one on one, but in a nonzero-sum game it would still win a tournament because it had the highest cumulative score when played against many kinds of rules. As Axelrod pointed out to William Poundstone, author of Prisoner's Dilemma, "That's a very bizarre idea. You can't win a chess tournament by never beating anybody." But with coevolution-change changing in response to itself-you can win without beating others. Hard-nosed CEOs in the business world now recognize that in the era of networks and alliances, companies can make billions without beating others. Win-win, the cliché is called.

Win-win is the story of life in coevolution.

Sitting in his book-lined office, Robert Axelrod mused on the consequences of understanding coevolution and then added, "I hope my work on the evolution of cooperation helps the world avoid conflict. If you read the citation which the National Academy of Science gave me," he said pointing to a plaque on the wall, "they think it helped avoid nuclear war." Although von Neumann was a key figure in the development of the atom bomb, he did not formally apply his own theories to the gamelike politics of the nuclear arms race. But after von Neumann's death in 1957, strategists in military think tanks began using his game theory to analyze the cold war, which had taken on the flavor of a coevolutionary "obligate cooperation" between two superpower enemies. Gorbachev had a fundamental coevolutionary insight, says Axelrod. "He saw that the Soviets could get more security with fewer tanks rather than with more tanks. Gorbi unilaterally threw away 10,000 tanks, and that made it harder for US and Europe to have a big military budget, which helped get this whole process going that ended the cold war."

Perhaps the most useful lesson of coevolution for "wannabe" gods is that in coevolutionary worlds control and secrecy are counterproductive. You can't control, and revelation works better than concealment. "In zero-sum games you always try to hide your strategy," says Axelrod. "But in nonzero-sum games you might want to announce your strategy in public so the other players need to adapt to it." Gorbachev's strategy was effective because he did it publicly; unilaterally withdrawing in secret would have done nothing.

The chameleon on the mirror is a completely open system. Neither the lizard nor the glass has any secrets. The grand closure of Gaia keeps cycling because all its lesser cycles inform each other in constant coevolutionary communication. From the collapse of Soviet command-style economies, we know that open information keeps an economy stable and growing.

Coevolution can be seen as two parties snared in the web of mutual propaganda. Coevolutionary relationships, from parasites to allies, are in their essence informational. A steady exchange of information welds them into a single system. At the same time, the exchange-whether of insults or assistance or plain news-creates a commons from which cooperation, self-organization, and win-win endgames can spawn.

In the Network Era-that age we have just entered-dense communication is creating artificial worlds ripe for emergent coevolution, spontaneous self-organization, and win-win cooperation. In this Era, openness wins, central control is lost, and stability is a state of perpetual almost-falling ensured by constant error.

The natural flux

Tonight is the Chinese Lunar Festival. Downtown in San Francisco's Chinatown, immigrants are exchanging moon cakes and telling tales of the Ghost Maiden who escaped as an orb in the sky. Twelve miles away where I live, I can walk in a cloud. The fog of the Golden Gate has piled up along the steep bank behind our house, engulfing our neighborhood in vapor. Under the light of Lady Moon, I take a midnight hike.

I wade chest-high in bleached ryegrass murmuring in the wind, and spy down the rugged coast of California. It is a disruptive land. For most purposes it is a mountainous desert that meets a generous ocean which cannot provide rain. Instead the sea sneaks in the water of life by rolling out blankets of fog at night. Come morning, the mist condenses into drops on the edges of twig and leaf, which tinkle to the earth. Much water is transported this way over a summer, bypassing the monopoly thunderclouds have on water delivery elsewhere. On this stingy substitute rain, the behemoth of all living things, the redwood, thrives.

The advantage of rain is that it is massive and indiscriminate. When it rains, it will wet a wide, diverse constituency. Fog on the other hand, is local. It relies on low-powered convection currents to ramble wherever it is easiest to drift to, and is then trapped by gentle, patient cul-de-sacs in the hills. In this way, the shape of the land steers the water, and indirectly, life. The correctly shaped hill can catch fog, or funnel drip into a canyon. A sunny south-facing mound will lose more precious moisture to evaporation than a shadier northern slope. Certain outcroppings of soil retain water better than others. Play these variables on top of each other and you have a patchwork of habitats. In a desert land, water decides life. And in a desert land where water is not delivered democratically, but parochially, on a whim, the land itself decides life.

The result is a patchwork landscape. The hills behind my house are cloaked with three separate quilts. A community of low-lying grass -- and of mice, owl, thistle, and poppy -- runs to the sea on one slope. On the crest of the hill, gnarly juniper and cypress trees preside over a separate association of deer, fox, and lichen. And on the other side of the rise, an endless impenetrable thicket of poison oak and coyote brush hides quail and other members of its guild.

The balance of these federations is kinetic. Their mutual self-supporting pose is continuously almost-falling, like a standing wave in a spring creek. When the mass of nature's creatures push against each other in coevolutionary embrace, their interactions among the uneven terrain of land and weather breaks their aggregate into local enclaves of codependency. And these patches roam over the land in time.

Wind and spring floods erode soils, exposing underlying layers and premiering new compositions of humus and minerals on the surface. As the mix of soil churns on the land, the mix of plants and animals coupled to it likewise churn. A thick stand of cactus, such as a Saguaro forest, can migrate onto or off of a patch of southwestern desert in little as 100 years. In a time-lapse film, a Saguaro grove would seem to creep across the desertscape like a pool of mercury. And it's not just cactus that would roam. Under the same time-lapse view, the wildflower prairie savanna of the midwest would flow around stands of oaks like an incoming tide, sometimes dissolving the woods into prairie, and sometimes, if the wildfires died out, retreating from the spreading swell of oak groves. Ecologist Dan Botkin speaks of forests "marching slowly across the landscape to the beat of the changing climate."

"Without change, deserts deteriorate," claims Tony Burgess, a burly ecologist with a huge red beard. Burgess is in love with deserts. He inhales desert lore and data all his waking hours. Out in the stark sun near Tucson, Arizona, he has been monitoring a desert plot that several generations of scientists have continuously measured and photographed for 80 years; the plot is the longest uninterrupted ecological observation anywhere. From studying the data of 80 years of desert change, Burgess has concluded that "variable rainfall is the key to the desert. Every year it should be a slightly different ball game to keep every species slightly out of equilibrium. If rainfall is variant then the mixture of species increases by two or three orders of magnitude. Whereas if you have a constant schedule of rainfall with respect to the annual temperature cycle, the beautiful desert ecology will almost always collapse into something simpler."

"Equilibrium is dead," Burgess states matter-of-factly. This opinion has not been held very long by the ecological science community. "Until the mid-1970s we were all working under a legacy which said that communities are on a trajectory towards an unchanging equilibrium, the climax. But now we see that it is turbulence and variance that really gives the richness to nature."

A major reason why ecologists favored equilibrium end points in nature was exactly the same reason why economists favored equilibrium end points in the economy: the mathematics of equilibria were possible. You could write an equation for a process that you could actually solve. But if you said that the system was perpetually in disequilibrium, you were saying it followed a model you couldn't solve and therefore couldn't explore. You were saying almost nothing. It is no coincidence, therefore, that a major shift in ecological (and economic) understanding occurred in the era when cheap computers made nonequilibrial and nonlinear equations easy to program. It was suddenly no problem to model a chaotic, coevolutionary ecosystem on a personal computer, and see that, hey, it acts very much like the odd behavior of a Saguaro forest or a prairie savanna on the march.

A thousand varieties of nonequilibrial models have blossomed in recent years; in fact there is now a small cottage industry of makers of chaotic and nonlinear mathematics, differential equations, and complexity theory, all this activity lending a hand in overturning the notion that nature or an economy seeks a stable balance. This new perspective -- that a certain unremitting flux is the norm -- has illuminated past data for reinterpretation. Burgess can display old photographs of the desert that show in a relatively short time -- over a few decades -- patches of Saguaro drifting over the Tucson basin. "What we found from our desert plot," Burgess said, "is that these patches are not in sync in terms of development and that by not being in sync, they make the whole desert richer because if something catastrophic wipes out one patch, another patch at a different stage of its natural history can export organisms and seeds to the decimated patch. Even ecosystems, such as tropical rain forests, which don't have variable rainfall, also have patch dynamics due to periodic storms and tree falls."

"Equilibrium is not only dead, it is death," Burgess emphasizes. "To enrich a system you need variance in time and space. But too much change will kill you too. You go from an ecocline to ecotone."

Burgess finds nature's reliance on disturbances and variance to be a practical issue. "In nature, it is no problem if you have very erratic production [of vegetation, seeds, or meat] from year to year. Nature actually increases her richness from this variance. But when people try to sustain themselves on the production from an ecosystem like a desert that is so variance driven, they can only do it by simplifying the system into what we call agriculture -- which gives a constant production for a variable environment." Burgess hopes the flux of the desert can teach us how to live with a variable environment without simplifying it. It is not a completely foolish dream. Part of what an information-driven economy provides us with is an adaptable infrastructure that can bend and work around irregular production; this is the basis for flexible and "just-in-time" manufacturing. It is theoretically possible that we could use information networks to coordinate the investment and highly irregular output of a rich, fluxing ecosystem that provides food and organic resources. But, as Burgess admits, "At the moment we have no industrial economic models that are variance driven, except gambling."

If it is true that nature is fundamentally in constant flux, then instability may cause the richness of biological forms in nature. But the idea that the elements of instability are the root of diversity runs counter to one of the hoariest dictums of environmentalism: that stability begets diversity, and diversity begets stability. If natural systems do not settle into a neat balance, then we should make instability our friend.

Biologists finally got their hands on computers in the late 1960s and began to model kinetic ecologies and food webs on silicon networks. One of the first questions they attempted to answer was, Where does stability come from? If you create predator/prey relationships in silico, what conditions cause the virtual organisms to settle into a long-term coevolutionary duet, and what conditions cause them to crash?

Among the earliest studies of simulated stability was a paper published in 1970 by Gardner and Ashby. Ashby was an engineer interested in nonlinear control circuits and the virtues of positive feedback loops. Ashby and Gardner programmed simple network circuits in hundreds of variations into a computer, systematically changing the number of nodes and the degrees of connectivity between nodes. They discovered something startling: that beyond a certain threshold, increasing the connectivity would suddenly decrease the ability of the system to rebound after disturbances. In other words, complex systems were less likely to be stable than simple ones.

A similar conclusion was published the following year by theoretical biologist Robert May, who ran model ecologies on computers populated with large multitudes of interacting species, and some virtual ecologies populated with few. His conclusions contradicted the common wisdom of stability/diversity, and he cautioned against the "simple belief" that stability is a consequence of increasing complexity of the species mix. Rather, May's simulated ecologies suggested that neither simplicity nor complexity had as much impact on stability as the pattern of the species interaction.

"In the beginning, ecologists built simple mathematical models and simple laboratory microcosms. They were a mess. They lost species like crazy," Stuart Pimm told me. "Later ecologists built more complex systems in the computer and in the aquarium. They thought these complex ones would be good. They were wrong. They were an even worse mess. Complexity just makes things very difficult -- the parameters have to be just right. So build a model at random and, unless it's really simple (a one-prey-one-resource population model) it won't work. Add diversity, interactions, or increase the food chain lengths and soon these get to the point where they will also fall apart. That's the theme of Gardner, Ashby, May and my early work on food webs. But keep on adding species, keep on letting them fall apart and, surprisingly, they eventually reach a mix that will not fall apart. Suddenly one gets order for free. It takes a lot of repeated messes to get it right. The only way we know how to get stable, persistent, complex systems is to repeatedly assemble them. And as far as I know, no one really understands why that works."

In 1991 Stuart Pimm, together with colleagues John Lawton and Joel Cohen, reviewed all the field measurements of food webs in the wild and by analyzing them mathematically concluded that "the rate at which populations recovered from disasters...depends on food chain length," as well as the number of prey and predators a species had. An insect eating a leaf is a chain of one. A turtle eating the insect that eats the leaf makes a chain of two. A wolf may sit many links away from a leaf. In general, the longer the chain, the less stable the interacting web to environmental disruption.

The other important point one can extract from May's simulations was best articulated in an observation made a few years earlier by the Spanish ecologist Ramon Margalef. Margalef noticed, as May did, that systems with many components would have weak relations between them, while systems that had few components would have tightly coupled relationships. Margalef put it this way: "From empirical evidence it seems that species that interact freely with others do so with a great number of other species. Conversely, species with strong interactions are often part of a system with a small number of species." This apparent tradeoff in an ecosystem between many loosely coupled members or few tightly coupled members is nicely paralleled by the now well-known tradeoff which biological organisms must choose in reproduction strategies. They can either produce a few well-protected offspring or a zillion unprotected ones.

Biology suggests that in addition to regulating the numbers of connections per "node" in a network, a system tends to also regulate the "connectance" (the strength of coupledness) between each pair of nodes in a network. Nature seems to conserve connectance. We should thus expect to find a similar law of the conservation of connectance in cultural, economic, and mechanical systems, although I am not aware of any studies that have attempted to show this. If there is such a law in all vivisystems, we should also expect to find this connectance being constantly adjusted, perpetually in flux.

"An ecosystem is a network of living creatures," says Burgess. The creatures are wired together in various degrees of connectance by food webs and by smells and vision. Every ecosystem is a dynamic web always in flux, always in the processes of reshaping itself. "Wherever we seek to find constancy we discover change," writes Botkin.

When we make a pilgrimage to Yellowstone National Park, or to the California Redwood groves, or to the Florida Everglades, we are struck by the reverent appropriateness of nature's mix in that spot. The bears seem to belong in those Rocky Mountain river valleys; the redwoods seem to belong on those coastal hills, and the alligators seem to belong in those plains. Thus our spiritual urge to protect them from disturbance. But in the long view, they are natural squatters who haven't been there long and won't always be there. Botkin writes, "Nature undisturbed is not constant in form, structure, or proportion, but changes at every scale of time and space."

A study of pollen lifted from holes drilled at the bottom of African lakes shows that the African landscape has been in a state of flux for the past several million years. Depending on when you looked in, the African landscape would look vastly different from now. In the recent geological past, the Sahara desert vastness of northern Africa was tropical forest. It's been many ecological types between then and now. We hold wilderness to be eternal; in reality, nature is constrained flux.

Complexity poured into the artificial medium of machines and silicon chips will only be in further flux. We see, too, that human institutions -- those ecologies of human toil and dreams -- must also be in a state of constant flux and reinvention, yet we are always surprised or resistant when change begins. (Ask a hip postmodern American if he would like to change the 200-year-old rule book known as the Constitution. He'll suddenly become medieval.)

Change, not redwood groves or parliaments, is eternal. The questions become: What controls change? How can we direct it? Can the distributed life in such loose associations as governments, economies, and ecologies be controlled in any meaningful way? Can future states of change even be predicted?

Let's say you purchase a worn-out 100-acre farm in Michigan. You fence the perimeter to keep out cows and people. Then you walk away. You monitor the fields for decades. That first summer, garden weeds take over the plot. Each year thereafter new species blow in from outside the fence and take root. Some newcomers are eventually overrun by newer newcomers. An ecological combo self-organizes itself on the land. The mix fluxes over the years. Would a knowledgeable ecologist watching the fencing-off be able to predict which wildlife species would dominate the land a century later?

"Yes, without a doubt he could," says Stuart Pimm. "But his prediction is not as interesting as one might think."

The final shape of the Michigan plot is found in every standard ecology college textbook in the chapter on the concept of succession. The first year's weeds on the Michigan plot are annual flowering plants, followed by tougher perennials like crabgrass and ragweed. Woodier shrubs will shade and suppress the flowers, followed by pines, which suppress the shrubs. But the shade of the pine trees protect hardwood seedlings of beech and maple, which in turn steadily elbow out the pines. One hundred years later the land is almost completely owned by a typical northern hardwood forest.

It is as if the brown field itself is a seed. The first year it sprouts a hair of weeds, a few years later it grows a shrubby beard, and then later it develops into a shaggy woods. The plot unfolds in predictable stages just as a tadpole unfolds out of a frog's egg.

Yet, the curious thing about this development is that if you start with a soggy 100-acre swamp, rather than a field, or with the same size lot of Michigan dry sandy dunes, the initial succession species are different (sedges in the swamp, raspberries on the sand), but the mix of species gradually converges to the same end point of a hardwood forest. All three seeds hatch the same adult. This convergence led ecologists to the notion of an omega point, or a climax community. For a given area, all ecological mixtures will tend to shift until they reach a mature, ultimate, stable harmony.

What the land "wants" to be in the temperate north is a hardwood forest. Give it enough time and that's what a drying lake or a windblown sand bog will become. If it ever warmed up a little, that's what an alpine mountaintop wants to be also. It is as if the ceaseless strife in the complicated web of eat-or-be-eaten stirs the jumble of species in the region until the mixture arrives at the hardwood climax (or the specific climax in other climates), at which moment it quietly settles into a tolerable peace. The land coming to a rest in the climax blend.

Mutual needs of diverse species click together so smartly in the climax arrangement that the whole is difficult to disrupt. In the space of 30 years the old-growth chestnut forest in North America lost every specimen of a species -- the mighty chestnut -- that formerly constituted a significant hunk of the forest's mass. Yet, there weren't any huge catastrophes in the rest of the forest; it still stands. This persistent stability of a particular composite of species -- an ecosystem -- speaks of some basin of efficiency that resembles the coherence belonging to an organism. Something whole, something alive dwells in that mutual support. Perhaps a maple forest is but a grand organism composed of lesser organisms.

On the other hand, Aldo Leopold writes, "In terms of conventional physics, the grouse represents only a millionth of either the mass or the energy of an acre. Yet subtract the grouse and the whole thing is dead."

In 1916, Frederic Clements, one of the founding fathers of ecology, called a community of creatures such as the beech hardwood forest an emergent superorganism. In his words, a climax formation is a superorganism because it "arises, grows, matures, and dies....comparable in its chief features with the life history of an individual plant." Since a forest could reseed itself on an abandoned Michigan field, Clements portrayed that act as reproduction, a further characteristic of an organism. To any astute observer, a beech-maple forest displays an integrity and identity as much as a crow does. What else but a (super)organism could reproduce itself so reliably, propagating on empty fields and sandy barrens?

Superorganism was a buzz word among biologists in the 1920s. They used it to describe the then novel idea that a collection of agents could act in concert to produce phenomena governed by the collective. Like a slime mold that assembled itself from moldy spots into a thrusting blob, an ecosystem coalesced into a stable superorganization -- a hive or forest. A Georgia pine forest did not act like a pine tree, nor a Texas sagebrush desert like a sagebrush, just as a flock is not a big bird. They were something else, a loose federation of animals and plants united into an emergent superorganism exhibiting distinctive behavior.

A rival of Clements, biologist H. A. Gleason, the other father of modern ecology, thought the superorganism federation was too flabby and too much the product of a human mind looking for patterns. In opposition to Clements, Gleason proposed that the climax community was merely a fortuitous association of organisms that came and went depending on climate and geological conditions. An ecosystem was more like a conference than a community -- indefinite, pluralistic, tolerant, and in constant flux.

The wilds of nature hold evidence for both views. In places the boundary between communities is decisive, much as one expects if ecosystems are superorganisms. Along the rocky coast of the Pacific Northwest, for instance, the demarcation between the high tide seaweed community and the watery edge of the spruce forest is an extreme no-man's-land of barren beach. One can stand on this yard-wide strip of salty desert and sense the two superorganisms on either side, fidgeting in their separate lives. As another example, the border between deciduous forest and wildflower prairie in the midwest is remarkably impermeable.

In search of an answer to the riddle of ecological superorganisms, biologist William Hamilton began modeling ecosystems on computers in the 1970s. He found that in his models (as well as in real life) very few systems were able to self-organize into any kind of lasting coherence. My examples above are a few exceptions in the wild. He found a few others: a sphagnum moss peat bog can repel the invasion of pine trees for thousands of years. Ditto for the tundra steppes. But most ecological communities stumble along into a mongrel mixture of species that offers no outstanding self-protection to the group as a team. Most ecological communities, both simulated and real, can be easily invaded in the longer run.

Gleason was right. The couplings between members of an ecosystem are far more flexible and transient than the couplings between members of an organism. The cybernetic difference between an organism such as a pollywog and an ecosystem such as a fresh-water bog is that an organism is tightly bound, and strict; an ecosystem is loosely bound, and lax.

In the long view, ecologies are temporary networks. Although some links become hardwired and nearly symbiotic, most species are promiscuous in evolutionary time, shacking up with a different partners as the partners themselves evolve.

In this light of evolutionary time, ecology can be seen as one long dress rehearsal. It's an identity workshop for biological forms. Species try out different roles with one another and explore partnerships. Over time, roles and performance are assimilated by an organism's genes. In poetic language, the gene is reluctant to assimilate into its code any interactions and functions directly based upon its neighbors' ways because the neighborhood can shift at any evolutionary moment. It pays to stay flexible, unattached, and uncommitted.

At the same time Clements was right. There is a basin of efficiency that, all things being equal, will draw down a certain mix of parts into a stable harmony. As a metaphor, consider the way rocks make their way to the valley floor. Not all rocks will land at the bottom; a particular rock may get stuck on a small hill somewhere. In the same way, stable intermediate less-than-climax mixtures of species can be found in places on the landscape. For extremely short periods of geological time -- hundreds of thousands of years -- ecosystems form an intimate troupe of players, who brook no interference and need no extras. These associations are far briefer than even the brief life of individual species, which typically flame-out after a million years or two.

Evolution requires a certain connectance among its participants to express its power; and so evolutionary dynamics exert themselves most forcefully in tightly coupled systems. In systems connected loosely, such as ecosystems, economic systems, and cultural systems, a less structured adaptation takes place. We know very little about the general dynamics of loosely coupled systems because this kind of distributed change is messy and infinitely indirect. Howard Pattee, an early cybernetician, defined hierarchical structure as a spectrum of connectance. He said, "To a Platonic mind, everything in the world is connected to everything else -- and perhaps it is. Everything is connected, but some things are more connected than others." Hierarchy for Pattee was the product of differential connectedness within one system. Members that were so loosely connected as to be "flat" would tend to form a separate organizational level distinct from areas where members were tightly connected. The range of connectance created a hierarchy.

In the most general terms, evolution is a tight web and ecology a loose one. Evolutionary change seems a strongly bound process very similar to mathematical computation, or even to thinking. In this way it is "cerebral." Ecological change, on the other hand, seems a weak-minded, circuitous process, centered in bodies shoved against wind, water, gravity, sunlight, and rock. "Community [ecological] attributes are more the product of environment than the product of evolutionary history," writes ecologist Robert Ricklefs. While evolution is governed by the straightforward flow of symbolic information issuing from the gene or computer chips, ecology is governed by the far less abstract, far more untidy complexity embodied by flesh.

Because evolution is such a symbolic process, we now can artificially create it and attempt to govern it. But because ecological change is so body bound, we cannot synthesize it well until we can more easily simulate bodies and richer artificial environments.

Where does diversity come from? In 1983, microbiologist Julian Adams discovered a clue when he brewed up a soup of cloned E. coli bacteria. He purified the broth until he had a perfectly homogenized pool of identical creatures. He put this soup of clones into a specially constructed chemostat that provided a uniform environment for them -- every E. coli bug had the same temperature and nutrient bath. Then he let the soup of identical bugs replicate and ferment. At the end of 400 generations, the E. coli bacteria had bred new strains of itself with slightly different genes. Out of a starting point in a constant featureless environment, life spontaneously diversified.

A surprised Adams dissected the genes of the variants (they weren't new species) to find out what happened. One of the original bugs had undergone a mutation that caused it to excrete acetate, an organic chemical. A second bug experienced a mutation that allowed it to make use of the acetate excreted from the first. Suddenly a symbiotic codependence of acetate maker and acetate eater had emerged from the uniformity, and the pool diverged into an ecology.

Although uniformity can yield diversity, variance does better. If the Earth were as smooth as a shiny ball bearing -- a perfect spherical chemostat spread evenly with uniform climate and homogeneous soils -- then the diversity of ecological communities on it would be far reduced from what it is now. In a constant environment, all variation and all diversity must be driven by internal forces. The only constraints on life would be other coevolutionary life.

If evolution had its way, with no interference from geographical and geological dynamics -- that is, without the clumsiness of a body -- then mindlike evolution would feed upon itself and breed heavily recursive relationships. On a globe without mountains or storms or unexpected droughts, evolution would wind life into a ever-tightening web of coevolution, a smooth world stuffed with parasites, parasites upon parasites (hyperparasites), mimics, and symbionts, all caught up in accelerating codependence. But each species would be so tightly coupled with the others that it would be difficult to distinguish where the identity of one began and the other left off. Eventually evolution on a ball-bearing planet would mold everything into a single, massive, ultradistributed planetwide superorganism.

Creatures born in the rugged environments of arctic climes must deal with the unpredictable variations that nature is always throwing at them. Freezing at night, baking during the day, ice storms after spring thaw, all create a rugged habitat. Habitats in the tropics and in the very deep sea are relatively "smooth" because of their constant temperature, rainfall, lightfall, and nutrients. Thus the smoothness of tropical or benthic environments allows species there to relinquish the need to adapt in physiological ways and allows them room to adapt in purely biological ways. In these steady habitats we should expect to see many instances of weird symbiotic and parasitic relationships -- parasites preying upon parasites, males living inside of females, and creatures mimicking and mirroring other creatures -- and that's what we do find.

Without a rugged environment life can only play off itself. It will still produce variation and novelty. But far more diversity can be manufactured in natural and artificial worlds by setting creatures in a rugged and vastly differentiated environment.

This lesson has not been lost on the wannabe gods trying to create lifelike behavior in computer worlds. When self-replicating and self-mutating computer viruses are loosed into a computer memory uniformly distributed with processing resources, the computer viruses quickly evolve a host of wildly recursive varieties including parasites, hyperparasites, and hyper-hyperparasites. David Ackley, one computer life researcher, told me, "I finally figured out that the way to get wonderfully lifelike behavior is not to try to make a really complex creature, but to make a wonderfully rich environment for a simple creature."

It's two o'clock on a blustery afternoon, six months after my midnight hike, when I climb the hill behind my house again. The windblown grass is green from the winter's rain. Up near the ridge I stop at a circle where the deer have matted the soft grass into a cushion. The trampled stems are weathered, buff with a tinge of violet, as if the color has rubbed off the deer's bellies. I rest in this recess. The wind swipes overhead.

I can see wildflowers crouched among the blown grass blades. For some reason every species is blue-violet: lupine, blue-eyed grass, thistle, gentian. Between me, the bent grass, and the ocean there are shrubs, squat creatures outfitted with silvery olive leaves -- standard desert issue.

Here's a stem of Queen Anne's lace. Its furrowed leaves are mind- bogglingly intricate. Each leaf has two dozen minileaves arrayed on it, and each of those minileaves has a dozen microleaves arrayed on it. The recursive shape is the result of some obsessive process, no doubt. Its bunched flower head, 30 miniature cream white florets surrounding a single tiny purple floret in the center, is equally unexpected. On this one slope where I rest, the diversity of living forms is overwhelming in its detail and unlikeliness.

I should be impressed. But what strikes me as I sit among two million grass plants and several thousand juniper shrubs, is how similar life on Earth is. For all the possible shapes and behaviors animated matter could take, only a few -- in wide variation -- are tried out. Life can't fool me. It's all the same, like those canned goods in grocery stores with different labels but all manufactured by the same food conglomerate. Life on Earth obviously all comes from one transnational conglomerate.

The grass pushing up on my seat, the scraggly thistle stem rubbing my shirt, the brown-breasted swallow swooping downhill: they are a single thing stretching out in many directions. I recognized it because I am stretched into it too.

Life is a networked thing -- a distributed being. It is one organism extended in space and time. There is no individual life. Nowhere do we find a solo organism living. Life is always plural. (And not until it became plural -- cloning itself -- could life be called life.) Life entails interconnections, links, and shared multiples. "We are of the same blood, you and I," coos the poet Mowgli. Ant, we are of the same blood, you and I. Tyrannosaurus, we are of the same blood, you and I. AIDS virus, we are of the same blood, you and I.

The apparent individuals that life has dispersed itself into are illusions. "Life is [primarily] an ecological property, and an individual property for only a fleeting moment," writes microbiologist Clair Folsome, a man who dabbled in making superorganisms inside bottles. We live one life, distributed. Life is a transforming flood that fills up empty containers and then spills out of them on its way to fill up more. The shape and number of vessels submerged by the flood doesn't make a bit of difference.

Life works as an extremist, a fanatic without moderation. It infiltrates everywhere. It saturates the atmosphere, covers the Earth's surface and wheedles its way into bedrock cracks. It will not be refused. As Lovelock noted, we have dug up no ancient rocks without also digging up ancient life preserved in them. John von Neumann, who thought of life in mathematical terms, said, "living organisms are...by any reasonable theory of probability or thermodynamics, highly improbable... [However] if by any peculiar accident there should ever be one of them, from there on the rules of probability do not apply, and there will be many of them." Life once made, filled the Earth immediately, commandeering matter from all the realms -- gas, liquid, solid -- into its schemes. "Life is a planetary-scale phenomenon," said James Lovelock. "There cannot be sparse life on a planet. It would be as unstable as half of an animal."

A thin membrane of whole life now covers the entire Earth. It is a coat that cannot be taken off. Rip one seam and the coat will patch itself on the spot. Abuse it, and the coat will metamorphose itself to thrive on the abuse. Not a threadbare green, it is a lush technicolor coat, a flamboyant robe surrounding the colossal corporeality of the planet.

In practice, it is an everlasting coat. The great secret which life has kept from us is that once born, life is immortal. Once launched, it cannot be eradicated.

Despite the rhetoric of radical environmentalists, it is beyond the power of human beings to wipe the whole flood of life off the planet. Mere nuclear bombs would do little to halt life in general, and might, in fact, increase the nonhuman versions.

There must have been a time billions of years ago when life crossed the threshold of irreversibility. Let's call that the I-point (for irreversible, or immortal). Before the I-point life was tenuous; indeed it faced a steep uphill slope. Frequent meteor impacts, fierce radiation, and harsh temperature fluctuations on Earth four billion years ago created an incredibly hostile environment for any half-formed, about-to-replicate complexity. But then, as Lovelock tells the story, "very early in the history of the planet, the climate conditions formed a window of opportunity just about right for life. Life had a short period in which to establish itself. If it failed, the whole system for future life failed."

But once established, life stuck fast. And once past the I-point life turned out to be neither delicate nor fragile, but hardy and irrepressible. Single cell bacteria are astonishingly indomitable, living in every possible antagonistic environment one could imagine, including habitats doused with heavy radiation. As hospitals know, it is frustratingly difficult to rid a few rooms of bacterial life. The Earth? Ha!

We should heed the unstoppable nature of life, because it has much to do with the complexity of vivisystems. We are about to make machines as complex as grasshoppers and let them loose in the world. Once born, they won't go away. Of the thousands of computer viruses cataloged by virus hunters so far, not one species of them has gone extinct. According to the companies that write antiviral software there are several dozens of new computer viruses created per week. They'll be with us for as long as we have computers.

The reason life cannot be halted is that the complexity of life's dynamics has exceeded the complexity of all known destructive forces. Life is far more complex than nonlife. While life can serve as an agent of death -- predator chomping on prey -- the consumption of one life form by another generally does not diminish complexity in the whole system and may even add to it.

It takes, on average, all the diseases and accidents of the world working 24 hours a day, 7 days a week, with no vacations, about 621,960 hours to kill a human organism. That's 70 years of full-time attack to break the bounds of human life -- barring the intervention of modern medicine (which may either accelerate or hinder death, depending on your views). This stubborn persistence in life is directly due to the complexity of the human body.

In contrast, a well-built car that managed to puff its way to an upper limit of 200,000 miles before blowing a valve would have run for about 5,000 hours. A jet turbine engine may run for 40,000 hours before being rebuilt. A simple light bulb with no moving parts is good for 2,000 hours. The longevity of nonliving complexity isn't even in the same league as the persistence of life.

The museum at the Harvard Medical School dedicates a display case to the "crowbar skull." This skull reveals a hole roughly gouged by a speeding iron bar. The skull belonged to Phineas Gage, a 19th-century quarry foreman who was packing a black powder charge into a hole with the iron bar when the powder exploded. The iron bar pierced his head. His crew sawed off the protruding bar before taking him to an ill-equipped doctor. According to anecdotes from those who knew him, Gage lived for another 13 years, more or less functional, except that after the accident he became short-tempered and peevish. Which is understandable. But the machine kept going.

People who lack a pancreas, a second kidney, a small intestine, may not run marathons, but they live. While debasement of many small components of the body -- glands in particular -- can cause death to the whole, these parts are heavily buffered from easy disruption. Indeed, warding off disruption is the principal property of complex systems.

Animals and plants in the wild regularly survive drastic violence and injury. The only study I know that has tried to measure the rate of injury in the wild focused on Brazilian lizards and concluded that 12 percent of them were missing at least one toe. Elk survive gunshot wounds, seals heal after shark bites, oak trees resprout after decapitation. In one experiment gastropods whose shells were deliberately crushed by researchers and returned to the wild lived as long as uninjured controls. The heroic achievement in nature is not the little fish that gets away, but that old man death is ever able to crash a system.

Networked complexity inverts the usual relation of reliability in things. As an example, individual switch parts in a modern camera may have 90 percent dependability. Linked dumbly in a series, not in a distributed way, the hundreds of switches would have great unreliability as a group -- let's say they have 75 percent dependability. Connected right -- each part informing the others -- as they are in advanced point-'n'-shoots, the reliability of the camera counter intuitively rises as a whole to 99 percent, exceeding the reliability of the individual parts (90 percent).

But the camera now has new subgroups of parts which act like parts themselves. More virtual parts means the total possibility for unpredictable behavior at the component level increases. There are now novel ways to go wrong. So while the camera as a whole is utterly more dependable, when it does surprise, it can often be a very surprising surprise. The old cameras were easy to fail, easy to repair. The new cameras fail creatively.

Failing creatively is the hallmark of vivisystems. Dying is difficult, but there are a thousand ways to do it. It took two hundred overpaid engineers two weeks of emergency alert work to figure out why the semi-alive American telephone switching system repeatedly failed in 1990. And these are the guys who built it. It had never failed this way, and probably won't fail this way again.

While every human is born pretty much the same, every death is different. If coroner's cause-of-death certificates were exact, each one would be unique. Medicine finds it more instructive to round off the causes and classify them generally, so the actual idiosyncratic nature of each death is not recorded.

A complex system cannot die simply. The members of a system have a bargain with the whole. The parts say, "We are willing to sacrifice to the whole, because together we are greater than our sum." Complexity locks in life. The parts may die, but the whole lives. As a system self-organizes into greater complexity, it increases its life. Not the length of its life, but its lifeness. It has more lives.

We tend to think of life and death as binary; a creature is either off or on. The self-organizing subsystems in organisms suggest, though, that some things are more alive than others. Biologist Lynn Margulis and others have pointed out that even a cell has lives in plural, as each cell is a historical marriage of at least three vestigial forms of bacteria.

"I am the most alive among the living," crows the Russian poet A. Tarkovsky (father of the filmmaker). That's politically incorrect, but probably true. There may be no real difference between the aliveness of a sparrow and a horse, but there is a difference of aliveness between a horse and a willow tree, or between a virus and a cricket. The greater the complexity of a vivisystem, the more life it may harbor. As long as the universe continues to cool down, life will build up in more curious varieties and in further mutual networks.

I head up the hill behind my house one more time. I ramble over to a grove of eucalyptus trees, where the local 4-H club used to keep its beehives. The grove snoozes in moist shade this time of day; the west-facing hill it stands on blocks the warm morning sun.

I imagine the valley all rock and barren at history's start-a hill of naked flint and feldspar, desolate and shiny. A billion years flicker by. Now the rock is clothed with a woven mat of grass. Life has filled a space in the grove with wood reaching higher than I can. Life is trying to fill the whole valley in. For the next billion years, it will keep trying new forms, erupting in whatever crevice or emptiness it can find.

Before life, there was no complex matter in the universe. The entire universe was utterly simple. Salts. Water. Elements. Very boring. After life, there was much complex matter. According to astrochemists, we can't find complex molecules in the universe outside of life. Life tends to hijack any and all matter it comes in contact with and complexify it. By some weird arithmetic, the more life stuffs itself into the valley, the more spaces it creates for further life. In the end, this small valley along the northern coast of California will become a solid block of life. In the end, left to its own drift, life may infiltrate all matter.

Why isn't the Earth a solid green from space? Why doesn't life cover the oceans and fill the air? I believe the answer is that if left alone, the Earth will be solid green someday. The conquest of air by living organisms is a relatively recent event, and one not yet completed. The complete saturation of the oceans may have to wait for rugged mats of kelp to evolve, ones able to withstand storm waves. But in the end, life will dominate; the oceans will be green.

The galaxy may be green someday too. Distant planets now toxic to life won't always remain so. Life can evolve representations of itself capable of thriving in environments that seem hostile now. But more importantly, once one variety of life has a toehold in a place, the inherently transforming nature of life modifies the environment until it is fit for other species of life.

In the 1950s, the physicist Erwin Schrödinger called the life force "negentropy" to indicate its opposite direction from the push of thermal decay. In the 1990s, an embryonic subculture of technocrats thriving in the U.S. calls the life force "extropy."

"Extropians," as promoters of extropy call themselves, issued a seven-point lifestyle manifesto based on the vitalism of life's extropy. Point number three is a creed that states their personal belief in "boundless expansion"-the faith that life will expand until it fills the universe. Those who don't believe this are tagged "deathists." In the context of their propaganda, this creed could be read as mere pollyanna self-inspiration, as in: We can do anything!

But somewhat perversely I take their boast as a scientific proposition: life will fill the universe. Nobody knows what the theoretical limits to the infection of matter by life would be. Nor does anybody know what the maximum amount of life-enhanced matter that our sun could support is.

In the 1930s, the Russian geochemist/biologist Vernadsky wrote, "The property of maximum expansion is inherent to living matter in the same manner as it is characteristic of heat to transfer from more heated to less heated bodies, of a soluble substance to dissolve in a solvent, and of a gas to dissipate in space." Vernadsky called it "pressure of life" and measured this expansion as velocity. His record for the velocity of life expansion was a giant puffball, which, he said, produced spores at such a rate that if materials were provided fast enough for the developing fungus, in only three generations puffballs would exceed the volume of Earth. He calculated by some obscure method that the life force's "speed of transmission" in bacteria is about 1,000 kilometers per hour. Life won't get far in filling up the universe at that rate.

When reduced to its essentials, life is very close to a computational function. For a number of years Ed Fredkin, a maverick thinker once associated with MIT, has been spinning out a heretical theory that the universe is a computer. Not metaphorically like a computer, but that matter and energy are forms of information processing of the same general class as the type of information processing that goes on inside a Macintosh. Fredkin disbelieves in the solidity of atoms and says flatly that "the most concrete thing in the world is information." Stephen Wolfram, a mathematical genius who did pioneering work on the varieties of computer algorithms agrees. He was one of the first to view physical systems as computational processes, a view that has since become popular in some small circles of physicists and philosophers. In this outlook the minimal work accomplished by life resembles the physics and thermodynamics of the minimal work done in a computer. Fredkin and company would say that knowing the maximum amount of computation that could be done in the universe (if we considered all its matter as a computer) would tells us whether life will fill the universe, given the distribution of matter and energy we see in the cosmos. I do not know if anyone has made that calculation.

One of the very few scientists to have thought in earnest about the final destiny of life is the theoretical physicist Freeman Dyson. Dyson did some rough calculations to estimate whether life and intelligence could survive until the ultimate end of the universe. He concluded it could, writing: "The numerical results of my calculations show that the quantities of energy required for permanent survival and communication are surprisingly modest....[T]hey give strong support to an optimistic view of the potentialities of life. No matter how far we go into the future, there will always be new things happening, new information coming in, new worlds to explore, a constantly expanding domain of life, consciousness and memory."

Dyson has taken this further than I would have dared. I was merely concerned about the dynamics of life, and how it infiltrates all matter, and how nothing known can halt it. But just as life irretrievably conquers matter, the lifelike higher processing power we call mind irrevocably conquers life and thus also all matter. Dyson writes in his lyrical and metaphysical book, Infinite in All Directions:

It appears to me that the tendency of mind to infiltrate and control matter is a law of nature....The infiltration of mind into the universe will not be permanently halted by any catastrophe or by any barrier that I can imagine. If our species does not choose to lead the way, others will do so, or may have already done so. If our species is extinguished, others will be wiser or luckier. Mind is patient. Mind has waited for 3 billion years on this planet before composing its first string quartet. It may have to wait for another 3 billion years before it spreads all over the galaxy. I do not expect that it will have to wait so long. But if necessary, it will wait. The universe is like a fertile soil spread out all around us, ready for the seeds of mind to sprout and grow. Ultimately, late or soon, mind will come into its heritage. What will mind choose to do when it informs and controls the universe? That is a question which we cannot hope to answer.

About a century ago, the common belief that life was a mysterious liquid that infused living things was refined into a modern philosophy called vitalism. The position which vitalism held was not very far from the meaning in the everyday phrase, "She lost her life." We all imagine some invisible substance seeping away at death. The vitalists took this vernacular meaning seriously. They held that while the essential spirit stirring in creatures was not itself alive, neither was it wholly an inanimate material or mechanism either. It was something else: a vital impulse that existed outside of the creature it animated.

My description of the aggressive character of life is not meant to be a postmodern vitalism. It is true that defining life as "an emergent property contingent upon the organization of inanimate parts but not reducible to them" (the best that science can do right now), comes very close to sounding like a metaphysical doctrine. But it is intended to be testable.

I take the view that life is a nonspiritual, almost mathematical property that can emerge from networklike arrangements of matter. It is sort of like the laws of probability; if you get enough components together, the system will behave like this, because the law of averages dictates so. Life results when anything is organized according to laws only now being uncovered; it follows rules as strict as those that light obeys.

This lawful process coincidentally clothes life in a spiritual looking garb. One reason is that this organization must, by law, produce the unpredictable and novel. Secondly, the result of organization must replicate at every opportunity, giving it a sense of urgency and desire. And thirdly, the result can easily loop around to protect its own existence, and thus it acquires an emergent agenda. Altogether, these principals might be called the "emergent" doctrine of life. This doctrine is radical because it entails a revised notion of what laws of nature mean: irregularity, circular logic, tautology, surprise.

Vitalism, like every wrong idea, contains a useful sliver of truth. Hans Driesch, the arch twentieth-century vitalist, defined vitalism in 1914 as "the theory of the autonomy of the process of life," and in certain respects he was right. Life in our dawning new view can be divorced from both living bodies and mechanical matrix, and set apart as a real, autonomous process. Life can be copied from living bodies as a delicate structure of information (spirit or gene?) and implanted in new lifeless bodies, whether they are of organic parts or machine parts.

In the history of ideas, we have progressively eliminated discontinuities from our perception of our role as humans. Historian of science David Channell summarizes this progression in his book The Vital Machine: A Study of Technology and Organic Life.

First, Copernicus eliminated the discontinuity between the terrestrial world and the rest of the physical universe. Next, Darwin eliminated the discontinuity between human beings and the rest of the organic world. And most recently, Freud eliminated the discontinuity between the rational world of the ego and the irrational world of the unconscious. But as [historian and psychologist Bruce] Mazlish has argued, there is one discontinuity that faces us yet. This "fourth discontinuity" is between human beings and the machine.

We are now crossing the fourth discontinuity. No longer do we have to choose between the living or the mechanical because that distinction is no longer meaningful. Indeed, the most meaningful discoveries in this coming century are bound to those that celebrate, explore, and exploit the unified quality of technology and life.

The bridge between the worlds of the born and the manufactured is the perpetual force of radical disequilibrium -- a law called life. In the future, the essence that both living creatures and machines will have in common -- that which will distinguish them from all other matter in the universe -- is that they both will have the dynamics of self-organized change.

We can now take the premise that life is a something in flux that is obeying laws which humans can uncover and recognize, even if we can't understand them fully. As a way to discover the commonalty between machines and creatures in this book, I've found it useful to ask, What does life want? I also consider evolution in the same way. What does evolution want? Or to be more precise, What does the world look like from life and evolution's point of view? If we consider life and evolution as "autonomous processes," then what are their selfish goals? Where are they headed? What are they becoming?

Gretel Ehrlich writes in her lyrical book Montana Spaces : "Wildness has no conditions, no sure routes, no peaks or goals, no source that is not instantly becoming something more than itself, then letting go of that, always becoming. It cannot be stripped to its complexity by cat scan or telescope. Rather, it is a many-pointed truth, almost a bluntness, a sudden essence like the wild strawberries strung along the ground on scarlet runners under my feet. Wildness is source and fruition at once, as if every river circled round, the mouth eating the tail -- and the tail, the source..."

There is no purpose, other than itself, to wildness. It is both "source and fruition," the mingling of cause and effect in circular logic. What Ehrlich calls wildness, I call a network of vital life, an outpouring of a nearly mechanic force that seeks only to enlarge itself, and that pushes its disequilibrium into all matter, erupting in creatures and machines alike.

Wildness/life is always becoming, Ehrlich says. Becoming what? Becoming becoming. Life is on its way to further complications, further deepness and mystery, further processes of becoming and change. Life is circle of becoming, an autocatalytic set, inflaming itself with its own sparks, breeding upon itself more life and more wildness and more "becomingness." Life has no conditions, no moments that are not instantly becoming something more than life itself.

As Ehrlich hints, wild life resembles that strange loop of the Uroborus biting its tail, consuming itself. But in truth, wild life is the far stranger loop of a snake releasing itself from its own grip, unmouthing an ever fattening tail tapering up to an ever increasingly larger mouth, birthing an ever larger tail, filling the universe with this strangeness.

The emergence of control

The invention of autonomous control, like most inventions, has roots in ancient China. There, on a dusty windswept plain, a small wooden statue of a man in robes teeters upon a short pole. The pole is carried between a pair of turning wagon wheels, pulled by two red horses outfitted in bronze finery.

The statue man, carved in the flowing dresses of 9th-century China, points with outstretched hand towards a distant place. By the magic of noisy gears connecting the two wooden wheels, as the cart races along the steppes, the wooden man perched on the stick invariably, steadily, without fail, points south. When the cart turns left or right, the geared wheels calculate the change and swing the wooden man's (or is it a god's?) arm a corresponding amount in the opposite direction, negating the cart's shift and keeping the guide forever pointing to the south. With an infallible will, and on his own accord, the wooden figure automatically seeks south. The south-pointing chariot precedes a lordly procession, preventing the party from losing its way in the desolate countryside of old China.

How busy was the ingenious medieval mind of China! Peasant folk in the backwaters of southwestern China, wishing to temper the amount of wine downed in the course of a fireside toast, came upon a small device which, by its own accord, would control the rowdy spirits of the wine. Chou Ch'u-Fei, a traveler among the Ch'i Tung natives then, reported that drinking bouts in this kingdom had been perfected by means of a two-foot-long bamboo straw which automatically regulated wine consumption, giving large-throated and small-mouthed drinkers equal advantage. A "small fish made of silver" floated inside the straw. The downward weight of the internal metal float restricted the flow of warm plum wine if the drinker sucked too feebly (perhaps through intoxication), thereby calling an end for his evening of merriment. If he inhaled too boisterously, he also got nothing, as the same float became wedged upwards by force of the suction. Only a temperate, steady draw was profitable.

Upon inspection, neither the south-pointing carriage nor the wine straw are truly automatic in a modern (self-steering) sense. Both devices merely tell their human masters, in the most subtle and unconscious way, of the adjustment needed to keep the action constant, and leave the human to make the change in direction of travel or power of lung. In the lingo of modern thinking, the human is part of the loop. To be truly automatic, the south-pointing statue would have to turn the cart itself, to make it a south-heading carriage. Or a carrot would have to be dangled from the point of his finger so that the horses (now in the loop) followed it. Likewise the drinking straw would have to regulate its volume no matter how hard one sucked. Although not automatic, the south-pointing cart is based on the differential gear, a thousand-year-old predecessor to the automobile transmission, and an early prototype of modern self-pointing guns on an armored tank which aid the drivers inside where a magnetic compass is useless. Thus, these clever devices are curious stillbirths in our genealogy of automation. The very first truly automatic devices had actually been built long before, a millennia earler.

Ktesibios was a barber who lived in Alexandria in the first half of the third century B.C. He was obsessed with mechanical devices, for which he had a natural genius. He eventually became a proper mechanician -- a builder of artifactual creations -- under King Ptolemy II. He is credited with having invented the pump, the water organ, several kinds of catapults, and a legendary water clock. At the time, Ktesibios's fame as an inventor rivaled that of the legendary engineer Archimedes. Today, Ktesibios is credited with inventing the first honest-to-goodness automatic device.

Ktesibios's clock kept extraordinarily good time (for then) by self-regulating its water supply. The weakness of most water clocks until that moment was that as the reservoir of water propelling the drive mechanism emptied, the speed of emptying would gradually decrease (because a shallow level of water provides less pressure than a high level), slowing down the clock's movements. Ktesibios got around this perennial problem by inventing a regulating valve (regula) comprised of a float in the shape of a cone which fit its nose into a mating inverted funnel. Within the regula, water flowed from the funnel stem, over the cone, and into the bowl the cone swam in. The cone would then float up into the concave funnel and constrict the water passage, thus throttling its flow. As the water diminished, the float would sink, opening the passage again and allowing more water in. The regula would immediately seek a compromise position where it would let "just enough" water for a constant flow through the metering valve vessel.

Ktesibios's regula was the first nonliving object to self-regulate, self-govern, and self-control. Thus, it became the first self to be born outside of biology. It was a true auto thing -- directed from within. We now consider it to be the primordial automatic device because it held the first breath of lifelikeness in a machine.

It truly was a self because of what it displaced. A constant autoregulated flow of water translated into a constant autoregulated clock and relieved a king of the need for servants to tend the water clock's water vessels. In this way, "auto-self" shouldered out the human self. From the very first instance, automation replaced human work.

Ktesibios's invention is first cousin to that all-American 20th-century fixture, the flush toilet. Readers will recognize the Ktesibios floating valve as the predecessor to the floating ball in the upper chamber of the porcelain throne. After a flush, the floating ball sinks with the declining water level, pulling open the water valve with its metal arm. The incoming water fills the vessel again, raising the ball triumphantly so that its arm closes the flow of water at the precise level of "full." In a medieval sense, the toilet yearns to keep itself full by means of this automatic plumbing. Thus, in the bowels of the flush toilet we see the archetype for all autonomous mechanical creatures.

About a century later, Heron, working in the same city of Alexandria, came up with a variety of different automatic float mechanisms, which look to the modern eye like a series of wildly convoluted toilet mechanisms. In actuality, these were elaborate party wine dispensers, such as the "Inexhaustible Goblet" which refilled itself to a constant level from a pipe fitted into its bottom. Heron wrote a huge encyclopedia (the Pneumatica) crammed with his incredible (even by today's standards) inventions. The book was widely translated and copied in the ancient world and was influential beyond measure. In fact, for 2,000 years (that is, until the age of machines in the 18th century), no feedback systems were invented that Heron had not already fathered.

The one exception was dreamed up in the 17th century by a Dutch alchemist, lens grinder, pyromaniac, and hobby submariner by the name of Cornelis Drebbel. (Drebbel made more than one successful submarine dive around 1600!) While tinkering in his search for gold, Drebbel invented the thermostat, the other universal example of a feedback system. As an alchemist, Drebbel suspected that the transmutation of lead into gold in a laboratory was inhibited by great temperature fluctuations of the heat sources cooking the elements. In the 1620s he jerry-rigged a minifurnace which could bake the initial alchemic mixture over moderate heat for a very long time, much as might happen to gold-bearing rock bordering the depths of Hades. On one side of his ministove, Drebbel attached a glass tube the size of a pen filled with alcohol. The liquid would expand when heated, pushing mercury in a connecting second tube, which in turn would push a rod that would close an air draft on the stove. The hotter the furnace, the futher the draft would close, decreasing the fire. The cooling tube retracted the rod, thus opening the draft and increasing the fire. An ordinary suburban tract home thermostat is conceptually identical -- both seek a constant temperature. Unfortunately, Drebbel's automatic stove didn't make gold, nor did Drebbel ever publish its design, so his automatic invention perished without influence, and its design had to be rediscovered a hundred years later by a French gentleman farmer, who built one to incubate his chicken eggs.

James Watt, who is credited with inventing the steam engine, did not. Working steam engines had been on the job for decades before Watt ever saw one. As a young engineer, Watt was once asked to repair a small-scale model of an early working, though inefficient, Newcomen steam engine. Frustrated by its awkwardness, Watt set out to improve it. At about the time of the American Revolution, he added two things to the existing engines; one of them evolutionary, the other revolutionary. His key evolutionary innovation was separating the heating chamber from the cooling chamber; this made his engine extremely powerful. So powerful that he needed to add a speed regulator to moderate this newly unleashed machine power. As usual Watt turned to what already existed. Thomas Mead, a mechanic and miller, had invented a clumsy centrifugal regulator for a windmill that would lower the millstone onto the grain only when stone's speed was sufficient. It regulated the output but not the power of a millstone.

Watt contrived a radical improvement. He borrowed Mead's regulator from the mill and revisioned it into a pure control circuit. By means of his new regulator the steam machine gripped the throat of its own power. His completely modern regula automatically stabilized his now ferocious motor at a constant speed of the operator's choice. By adjusting the governor, Watt could vary the steam engine to run at any rate. This was revolutionary.

Like Heron's float and Drebbel's thermostat, Watt's centrifugal governor is transparent in its feedback. Two leaden balls, each at the end of a stiff pendulum, swing from a pole. As the pole rotates the balls spin out levitating higher the faster the system spins. Linkages scissored from the twirling pendulums slide up a sleeve on the pole, levering a valve which controls the speed of rotation by adjusting the steam. The higher the balls spin, the more the linkages close the valve, reducing the speed, until an equilibrium point of constant rpms (and height of spinning balls) is reached. The control is thus as dependable as physics.

Rotation is an alien power in nature. But among machines, it is blood. The only known bearing in biology is at the joint of a sperm's spinning hair propeller. Outside of this micromotor, the axle and wheel are unknown to those with genes. To the ungened machine, whirling wheels and spinning shafts are reasons to live. Watt gave machines the secret to controlling their own revolutions, which was his revolution. His innovation spread widely and quickly. The mills of the industrial age were fueled by steam, and the engines earnestly regulated themselves with the universal badge of self-control: Watt's flyball governor. Self-powered steam begat machine mills which begat new kinds of engines which begat new machine tools. In all of them, self-regulators dwelt, fueling the principle of snowballing advantages. For every one person visibly working in a factory, thousands of governors and self-regulators toiled invisibly. Today, hundreds of thousands of regulators, unseen, may work in a modern plant at once. A single human may be their coworker.

Watt took the volcanic fury of expanding steam and tamed it with information. His flyball governor is undiluted informational control, one of the first non-biological circuits. The difference between a car and an exploding can of gasoline is that the car's information -- its design -- tames the brute energy of the gas. The same amount of energy and matter are brought together in a car burning in a riot and one speeding laps in the Indy 500. In the latter case, a critical amount of information rules over the system, civilizing the dragon of fire. The full heat of fire is housetrained by small amounts of self-perception. Furious energy is educated, brought in from the wilds to work in the yard, in the basement, in the kitchen, and eventually in living rooms.

The steam engine is an unthinkable contraption without the domesticating loop of the revolving governor. It would explode in the face of its inventors without that tiny heart of a self. The immense surrogate slave power released by the steam engine ushered in the Industrial Revolution. But a second, more important revolution piggybacked on it unnoticed. There could not have been an industrial revolution without a parallel (though hidden) information revolution at the same time, launched by the rapid spread of the automatic feedback system. If a fire-eating machine, such as Watt's engine, lacked self-control, it would have taken every working hand the machine displaced to babysit its energy. So information, and not coal itself, turned the power of machines useful and therefore desirable.

The industrial revolution, then, was not a preliminary primitive stage required for the hatching of the more sophisticated information revolution. Rather, automatic horsepower was, itself, the first phase of the knowledge revolution. Gritty steam engines, not teeny chips, hauled the world into the information age.

Heron's regulator, Drebbel's thermostat, and Watt's governor bestowed on their vessels a wisp of self-control, sensory awareness, and the awakening of anticipation. The governing system sensed its own attributes, noted if it had changed in a certain respect since it last looked, and if it had, it adjusted itself to conform to a goal. In the specific case of a thermostat, the tube of alcohol detected the system's temperature, and then took action or not to tweak the fire in order to align itself with the fixed goal of a certain temperature. It had, in a philosophical sense, a purpose.

Although it may strike us as obvious now, it took a long while for the world's best inventors to transpose even the simplest automatic circuit such as a feedback loop into the realm of electronics. The reason for the long delay was that from the moment of its discovery electricity was seen primarily as power and not as communication. The dawning distinction of the two-faced nature of the spark was acknowledged among leading German electrical engineers of the last century as the split between the techniques of strong current and the techniques of weak current. The amount of energy needed to send a signal is so astoundingly small that electricity had to be reimagined as something altogether different from power. In the camp of the wild-eyed German signalists, electricity was a sibling to the speaking mouth and the writing hand. The inventors (we would call them hackers now) of weak current technology brought forth perhaps the least precedented invention of all time -- the telegraph. With this device human communication rode on invisible particles of lightning. Our entire society was reimagined because of this wondrous miracle's descendants.

Telegraphers had the weak model of electricity firmly in mind, yet despite their clever innovations, it wasn't until August 1929, that telephone engineer H. S. Black, working at Bell Laboratories, tamed an electrical feedback loop. Black was hunting for a way to make durable amplifier relays for long-distance phone lines. Early amplifiers were made of crude materials that tended to disintegrate over use, causing the amp to "run away." Not only would an aging relay amplify the phone signal, it would mistakenly compound any tiny deviation from the range it expected until the mushrooming error filled and killed the system. What was needed was Heron's regula, a counter signal to rein in the chief signal, to dampen the effect of the perpetual recycling. Black came up with a negative feedback loop, which was designated negative in contrast to the snowballing positive loop of the amplifier. Conceptually, the electrical negative feedback loop is a toilet flusher or thermostat. This braking circuit keeps the amplifier honed in on a steady amplification in the same way a thermostat hones in on a steady temperature. But instead of metallic levers, a weak train of electrons talks to itself. Thus, in the byways of the telephone switching network, the first electrical self was born.

From World War I and after, the catapults that launched missiles had become so complicated, and their moving targets so sophisticated, that calculating ballistic trajectories taxed human talent. Between battles, human calculators, called computers, computed the settings for firing large guns under various wind, weather and altitude conditions. The results were sometimes printed in pocketsize tables for the gunmen on the front line, or if there was enough time and the missile-gun was common, the tables were mechanically encoded into an apparatus on the gun, known as the automaton. In the U.S., the firing calculations were compiled in a laboratory set up at the Navy's Aberdeen Proving Ground in Maryland, where rooms full of human computers (almost exclusively women) employed hand-cranked adding machines to figure the tables.

By World War II, the German airplanes which the big guns boomed at were flying as fast as the missiles themselves. Speedier on-the-spot calculations were needed, ideally ones that could be triggered from measurements of planes in flight made by the newly invented radar scanner. Besides, Navy gunmen had a weighty problem: how to move and aim these monsters with the accuracy the new tables gave them. The solution was as close at hand as the stern of the ship: a large ship controlled its rudder by a special type of automatic feedback loop known as a servomechanism.

Servomechanisms were independently and simultaneously invented a continent apart by an American and a Frenchman around 1860. It was the Frenchman, engineer Leon Farcot, who tagged the device with a name that stuck: moteur asservi, or servo-motor. As boats had increased in size and speed over time, human power at the tiller was no longer sufficient to move the rudder against the force of water surging beneath. Marine technicians came up with various oil-hydraulic systems that amplified the power of the tiller so that gently swinging the miniature tiller at the captain's helm would move the mighty rudder, kind of. A repeated swing of the minitiller would translate into different amounts of steerage of the rudder depending on the speed of the boat, waterline, and other similar factors. Farcot invented a linkage system that connected the position of the heavy rudder underwater back to the position of the easy-to-swing tiller -- the automatic feedback loop! The tiller then indicated the actual location of the rudder, and by means of the loop, moving the indicator moved the reality. In the jingo of current computerese, What you see is what you get!

The heavy gun barrels of World War II were animated the same way. A hydraulic hose of compressed oil connected a small pivoting lever (the tiller) to the pistons steering the barrel. As the shipmate's hand moved the lever to the desired location, that tiny turn compressed a small piston which would open a valve releasing pressurized oil, which would nudge a large piston moving the heavy gun barrel. But as the barrel swung it would push a small piston that, in return, moved the hand lever. As he tried to turn the tiller, the sailor would feel a mild resistance, a force created by the feedback from the rudder he wanted to move.

Bill Powers was a teenage Electronic Technician's Mate who worked with the Navy's automated guns, and who later pursued control systems as explanation for living things. He describes the false impression one gets by reading about servomechanism loops:

The sheer mechanics of speaking or writing stretches out the action so it seems that there is a sequence of well-separated events, one following the other. If you were trying to describe how a gun-pointing servomechanism works, you might start out by saying, "Suppose I push down on the gun-barrel to create a position error. The error will cause the servo motors to exert a force against the push, the force getting larger as the push gets larger." That seems clear enough, but it is a lie. If you really did this demonstration, you would say "Suppose I push down on the gun-barrel to create an error...wait a minute. It's stuck."

No, it isn't stuck. It's simply a good control system. As you begin to push down, the little deviation in sensed position of the gun-barrel causes the motor to twist the barrel up against your push. The amount of deviation needed to make the counteractive force equal to the push is so small that you can neither see nor feel it. As a result, the gun-barrel feels as rigid as if it were cast in concrete. It creates the appearance of one of those old-fashioned machines that is immovable simply because it weighs 200 tons, but if someone turned off the power the gun-barrel would fall immediately to the deck.

Servomechanisms have such an uncanny ability to aid steering that they are still used (in updated technology) to pilot boats, to control the flaps in airplanes, and to wiggle the fingers in remotely operated arms handling toxic and nuclear waste.

More than the purely mechanical self-hood of the other regulators like Heron's valve, Watt's governor, and Drebbel's thermostat, the servomechanism of Farcot suggested the possibility of a man-machine symbiosis -- a joining of two worlds. The pilot merges into the servomechanism. He gets power, it gets existence. Together they steer. These two aspects of the servomechanisms -- steering and symbiosis -- inspired one of the more colorful figures of modern science to recognize the pattern that connected these control loops.

Of all the mathematicians assigned during World War I to the human calculating lab in charge of churning out more accurate firing tables at the Aberdeen Proving Grounds, few were as overqualified as Private Norbert Wiener, a former math prodigy whose genius had an unorthodox pedigree.

The ancients recognized genius as something given rather than created. But America at the turn of the century was a place where the wisdom of the past was often successfully challenged. Norbert's father, Leo Wiener, had come to America to launch a vegetarian commune. Instead, he was distracted with other untraditional challenges, such as bettering the gods. In 1895, as a Harvard professor of Slavic languages, Leo Wiener decided that his firstborn son was going to be a genius. A genius deliberately made, not born.

Norbert Wiener was thus born into high expectations. By the age of three he was reading. At 18 he earned his Ph.D. from Harvard. By 19 he was studying metamathematics with Bertrand Russell. Come 30 he was a professor of mathematics at MIT and a thoroughly odd goose. Short, stout, splay-footed, sporting a goatee and a cigar, Wiener waddled around like a smart duck. He had a legendary ability to learn while slumbering. Numerous eyewitnesses tell of Wiener sleeping during a meeting, suddenly awakening at the mention of his name, and then commenting on the conversation that passed while he dozed, usually adding some penetrating insight that dumbfounded everyone else.

In 1948 he published a book for nonspecialists on the feasibility and philosophy of machines that learn. The book was initially published by a French publisher (for roundabout reasons) and went through four printings in the United States in its first six months, selling 21,000 copies in the first decade of its influence -- a best seller then. It rivaled the success of the Kinsey Report on sexual behavior, issued the same year. As a Business Week reporter observed in 1949, "In one respect Wiener's book resembles the Kinsey Report: the public response to it is as significant as the content of the book itself."

Wiener's startling ideas sailed into the public mind, even though few could comprehend his book, by means of the wonderfully colorful name he coined for both his perspective and the book: Cybernetics. As has been noted by many writers, cybernetics derives from the Greek for "steersman" -- a pilot that steers a ship. Wiener, who worked with servomechanisms during World War II, was struck by their uncanny ability to aid steering of all types. What is usually not mentioned is that cybernetics was also used in ancient Greece to denote a governor of a country. Plato attributes Socrates as saying, "Cybernetics saves the souls, bodies, and material possessions from the gravest dangers," a statement that encompasses both shades of the word. Government (and that meant self-government to these Greeks) brought order by fending off chaos. Also, one had to actively steer to avoid sinking the ship. The Latin corruption of kubernetes is the derivation of governor, which Watt picked up for his cybernetic flyball.

The managerial nature of the word has further antecedent to French speakers. Unbeknownst to Wiener, he was not the first modern scientist to reactivate this word. Around 1830 the French physicist Ampere (whence we get the electrical term amperes, and its shorthand "amp") followed the traditional manner of French grand scientists and devised an elaborate classification system of human knowledge. Ampere designated one branch the realm of "Noological Sciences," with the subrealm of Politics. Within political science, immediately following the sub-subcategory of Diplomacy, Ampere listed the science of Cybernetics, that is, the science of governance.

Wiener had in mind a more explicit definition, which he stated boldly in the full title of his book, Cybernetics: or control and communication in the animal and the machine. As Wiener's sketchy ideas were embodied by later computers and fleshed out by other theorists, cybernetics gradually acquired more of the flavor of Ampere's governance, but without the politics.

The result of Wiener's book was that the notion of feedback penetrated almost every aspect of technical culture. Though the central concept was both old and commonplace in specialized circumstances, Wiener gave the idea legs by generalizing the effect into a universal principle: lifelike self-control was a simple engineering job. When the notion of feedback control was packaged with the flexibility of electronic circuits, they married into a tool anyone could use. Within a year or two of Cybernetics's publication, electronic control circuits revolutionized industry.

The avalanche effects of employing automatic control in the production of goods were not all obvious. Down on the factory floor, automatic control had the expected virtue of moderating high-powered energy sources as mentioned earlier. There was also an overall speeding up of things because of the continuous nature of automatic control. But those were relatively minor compared to a completely unexpected miracle of self-control circuits: their ability to extract precision from grossness.

As an illustration of how the elemental loop generates precision of out imprecise parts, I follow the example suggested by the French writer Pierre de Latil in his 1956 book Thinking by Machine. Generations of technicians working in the steel industry pre-1948 had tried unsuccessfully to produce a roll of sheet metal in a uniform thickness. They discovered about a half-dozen factors that affected the thickness of the steel grinding out the rolling-mill -- such as speed of the rollers, temperature of the steel, and traction on the sheet -- and spent years strenuously perfecting the regulation of each of them, and more years attempting their synchronization. To no avail. The control of one factor would unintentionally disrupt the other factors. Slowing the speed would raise the temperature; lowering the temperature would raise the traction; increasing traction lowers the speed, and so on. Everything was influencing everything else. The control was wrapped up in some interdependent web. When the steel rolled out too thick or too thin, chasing down the culprit out of six interrelated suspects was inevitably a washout. There things stalled until Wiener's brilliant generalization published in Cybernetics. Engineers around the world immediately grasped the crucial idea and installed electronic feedback devices in their mills within the following year or two.

In implementation, a feeler gauge measures the thickness of the just-made sheet metal (the output) and sends this signal back to a servo-motor controlling the single variable of traction, the variable to affect the steel last, just before the rollers. By this meager, solo loop, the whole caboodle is regulated. Since all the factors are interrelated, if you can keep just one of them directly linked to the finished thickness, then you can indirectly control them all. Whether the deviation tendency comes from uneven raw metal, worn rollers, or mistakenly high temperatures doesn't matter much. What matters is that the automatic loop regulates that last variable to compensate for the other variables. If there is enough leeway (and there was) to vary the traction to make up for an overly thick source metal, or insufficiently tempered stock, or rollers contaminated with slag, then out would come consistently even sheets. Even though each factor is upsetting the others, the contiguous and near instantaneous nature of the loop steers the unfathomable network of relationships between them toward the steady goal of a steady thickness.

The cybernetic principle the engineers discovered is a general one: if all the variables are tightly coupled, and if you can truly manipulate one of them in all its freedoms, then you can indirectly control all of them. This principle plays on the holistic nature of systems. As Latil writes, "The regulator is unconcerned with causes; it will detect the deviation and correct it. The error may even arise from a factor whose influence has never been properly determined hitherto, or even from a factor whose very existence is unsuspected." How the system finds agreement at any one moment is beyond human knowing, and more importantly, not worth knowing.

The irony of this breakthrough, Latil claims, is that technologically this feedback loop was quite simple and "it could have been introduced some fifteen or twenty years earlier, if the problem had been approached with a more open mind..." Greater is the irony that twenty years earlier the open mind for this view was well established in economic circles. Frederick Hayek and the influential Austrian school of economics had dissected the attempts to trace out the routes of feedback in complex networks and called the effort futile. Their argument became known as the "calculation argument." In a command economy, such as the then embryonic top-down economy installed by Lenin in Russia, resources were allotted by calculation, tradeoffs, and controlled lines of communication. Calculating, even less controlling, the multiple feedback factors among distributed nodes in an economy was as unsuccessful as the engineer's failure in chasing down the fleeing interlinked factors in a steel mill. In a vacillating economy it is impossible to calculate resource allotment. Instead, Hayek and other Austrian economists of the 1920s argued that a single variable -- the price -- is used to regulate all the other variables of resource allotment. That way, one doesn't care how many bars of soap are needed per person, or whether trees should be cut for houses or for books. These calculations are done in parallel, on the fly, from the bottom up, out of human control, by the interconnected network itself. Spontaneous order.

The consequence of this automatic control (or human uncontrol) is that the engineers could relax their ceaseless straining for perfectly uniform raw materials, perfectly regulated processes. Now they could begin with imperfect materials, imprecise processes. Let the self-correcting nature of automation strain to find the optima which let only the premium through. Or, starting with the same quality of materials, the feedback loop could be set for a much higher quality setting, delivering increased precision for the next in line. The identical idea could be exported upstream to the suppliers of raw materials, who could likewise employ the automatic loop to extract higher quality products. Cascading further out in both directions in the manufacturing stream, the automatic self became an overnight quality machine, ever refining the precision humans can routinely squeeze from matter.

Radical transformations to the means of production had been introduced by Eli Whitney's interchangeable parts and Ford's idea of an assembly line. But these improvements demanded massive retooling and capital expenditures, and were not universally applicable. The homely auto-circuit, on the other hand -- a suspiciously cheap accessory -- could be implanted into almost any machine that already had a job. An ugly duckling, like a printing press, was transformed into a well-behaved goose laying golden eggs.

But not every automatic circuit yields the ironclad instantaneity that Bill Power's gun barrel enjoyed. Every unit added onto a string of connected loops increases the likelihood that the message traveling around the greater loop will arrive back at its origin to find that everything has substantially changed during its journey. In particularly vast networks in fast moving environments, the split second it takes to traverse the circuit is greater than the time it takes for the situation to change. In reaction, the last node tends to compensate by ordering a large correction. But this also is delayed by the long journey across many nodes, so that it arrives missing its moving mark, birthing yet another gratuitous correction. The same effect causes student drivers to zigzag down the road, as each late large correction of the steering wheel overreacts to the last late overcorrection. Until the student driver learns to tighten the feedback loop to smaller, quicker corrections, he cannot help but swerve down the highway hunting (in vain) for the center. This then is the bane of the simple auto-circuit. It is liable to "flutter" or "chatter," that is, to nervously oscillate from one overreaction to another, hunting for its rest. There are a thousand tricks to defeat this tendency of overcompensation, one trick each for the thousand advance circuits that have been invented. For the last 40 years, engineers with degrees in control theory have written shelffuls of treatises communicating their latest solution to the latest problem of oscillating feedback. Fortunately, feedback loops can be combined into useful configurations.

Let's take our toilet, that prototypical cybernetic example. We install a knob which allows us to adjust the water level of the tank. The self-regulating mechanism inside would then seek whatever level we set. Turn it down and it satisfies itself with a low level; turn it up and it hones in on a high level of water. (Modern toilets do have such a knob.) Now let's go further and add a self-regulating loop to turn the knob, so that we can let go of that, too. This second loop's job is to seek the goal for the first loop. Let's say the second mechanism senses the water pressure in the feed pipe and then moves the knob so that it assigns a high level to the toilet when there is high water pressure and a lower level when the pressure is low.

The second circuit is controlling the range of the first circuit which is controlling the water. In an abstract sense the second loop brings forth a second order of control -- the control of control -- or a metacontrol. Our newfangled second-order toilet now behaves "purposefully." It adapts to a shifting goal. Even though the second circuit setting the goal for the first is likewise mechanical, the fact that the whole is choosing its own goal gives the metacircuit a mildly biological flavor.

As simple as a feedback loop is, it can be stitched together in endless combinations and forever stacked up until it forms a tower of the most unimaginable complexity and intricacy of subgoals. These towers of loops never cease to amuse us because inevitably the messages circulating along them cross their own paths. A triggers B, and B triggers C, and C triggers A. In outright paradox, A is both cause and effect. Cybernetician Heinz von Foerster called this elusive cycle "circular causality." Warren McCulloch, an early artificial intelligence guru called it "intransitive preference," meaning that the rank of preferences would cross itself in the same self-referential way the children's game of Paper-Scissors-Stone endlessly intersects itself: Paper covers stone; stone breaks scissors; scissors cuts paper; and round again. Hackers know it as a recursive circuit. Whatever the riddle is called, it flies in the face of 3,000 years of logical philosophy. It undermines classical everything. If something can be both its own cause and effect, then rationality is up for grabs.

The compounded logic of stacked loops which doubles back on itself is the source of the strange counterintuitive behaviors of complex circuits. Made with care, circuits perform dependably and reasonably, and then suddenly, by their own drumbeat, they veer off without notice. Electrical engineers get paid well to outfox the lateral causality inherent in all circuits. But pumped up to the density required for a robot, circuit strangeness becomes indelible. Reduced back to its simplest -- a feedback cycle -- circular causality is a fertile paradox.

Where does self come from? The perplexing answer suggested by cybernetics is: it emerges from itself. It cannot appear any other way. Brian Goodwin, an evolutionary biologist, told reporter Roger Lewin, "The organism is the cause and effect of itself, its own intrinsic order and organization. Natural selection isn't the cause of organisms. Genes don't cause organisms. There are no causes of organisms. Organisms are self-causing agencies." Self, therefore, is an auto-conspired form. It emerges to transcend itself, just as a long snake swallowing its own tail becomes Uroborus, the mythical loop.

The Uroborus, according to C. G. Jung, is one of those resonant projections of the human soul that cluster around timeless forms. The ring of snake consuming its own tail first appeared as art adorning Egyptian statuary. Jung developed the idea that the nearly chaotic variety of dream images visited on humans tend to gravitate around certain stable nodes which form key and universal images, much as interlinked complex systems tend settle down upon "attractors," to use modern terminology. A constellation of these attracting, strange nodes form the visual vocabulary of art, literature, and some types of therapy. One of the most enduring attractors, and an early pattern to be named, was the Thing Eating Its Own Tail, often graphically simplified to a snakelike dragon swallowing its own tail in a perfect circle.

The loop of Uroborus is so obviously an emblem for feedback that I have trouble ascertaining who first used it in a cybernetic context. In the true manner of archetypes it was probably realized as a feedback symbol independently more than once. I wouldn't doubt that the faint image of snake eating its tail spontaneously hatches whenever, and wherever, the GOTO START loop dawns on a programmer.

Snake is linear, but when it feeds back into itself it becomes the archetype of nonlinear being. In the classical Jungian framework, the tail-biting Uroborus is the symbolic depiction of the self. The completeness of the circle is the self-containment of self, a containment that is at the same time made of one thing and made of competing parts. The flush toilet then, as the plainest manifestation of a feedback loop, is a mythical beast -- the beast of self.

The Jungians say that the self is taken to be "the original psychic state prior to the birth of ego consciousness," that is, "the original mandala-state of totality out of which the individual ego is born." To say that a furnace with a thermostat has a self is not to say it has an ego. The self is a mere ground state, an auto-conspired form, out of which the more complicated ego can later distinguish itself, should its complexity allow that.

Every self is a tautology: self-evident, self-referential, self-centered, and self-created. Gregory Bateson said a vivisystem was "a slowly self-healing tautology." He meant that if disturbed or disrupted, a self will "tend to settle toward tautology" -- it will gravitate to its elemental self-referential state, its "necessary paradox."

Every self is an argument trying to prove its identity. The self of a thermostat system has endless internal bickering about whether to turn the furnace up or down. Heron's valve system argues continuously around the sole, solitary action it can take: should it move the float or not?

A system is anything that talks to itself. All living systems and organisms ultimately reduce to a bunch of regulators -- chemical pathways and neuron circuits -- having conversations as dumb as "I want, I want, I want; no, you can't, you can't, you can't."

The sowing of selves into our built world has provided a home for control mechanisms to trickle, pool, spill, and gush. The advent of automatic control has come in three stages and has spawned three nearly metaphysical changes in human culture. Each regime of control is boosted by deepening loops of feedback and information flow.

The control of energy launched by the steam engine was the first stage. Once energy was controlled it became "free." No matter how much more energy we might release, it won't fundamentally change our lives. The amount of calories (energy) require to accomplish something continues to dwindle so that our biggest technological gains no longer hinge on further mastery of powerful energy sources.

Instead, our gains now derive from amplifying the accurate control of materials -- the second regime of control. Informing matter by investing it with high degrees of feedback mechanisms, as is done with computer chips, empowers the matter so that increasingly smaller amounts do the same work of larger uninformed amounts. With the advent of motors the size of dust motes (successfully prototyped in 1991), it seems as if you can have anything you want made in any size you want. Cameras the size of molecules? Sure, why not? Crystals the size of buildings? As you wish. Material is under the thumb of information, in the same handy way that energy now is -- just spin a dial. "The central event of the twentieth century is the overthrow of matter," says technology analyst George Gilder. This is the stage in the history of control in which we now dwell. Essentially, matter -- in whatever shape we want -- is no longer a barrier. Matter is almost "free."

The third regime of the control revolution, seeded two centuries ago by the application of information to coal steam, is the control of information itself. The miles of circuits and information looping from place to place that administers the control of energy and matter has incidentally flooded our environment with messages, bits, and bytes. This unmanaged data tide is at toxic levels. We generate more information than we can control. The promise of more information has come true. But more information is like the raw explosion of steam -- utterly useless unless harnessed by a self. To paraphrase Gilder's aphorism: "The central event of the twenty-first century will be the overthrow of information."

Genetic engineering (information which controls DNA information) and tools for electronic libraries (information which manages book information) foreshadow the subjugation of information. The impact of information domestication will be felt initially in industry and business, just as energy and material control did, and then later seep to the realm of individual.

The control of energy conquered the forces of nature (and made us fat); the control of matter brought material wealth within easy reach (and made us greedy). What mixed cornucopia will the blossoming of full information control bring about? Confusion, brilliance, impatience?

Without selves, very little happens. Motors, by the millions, bestowed with selves, now run factories. Silicon chips, by the billions, bestowed with selves, will redesign themselves smaller and faster and rule the motors. And soon, the fibrous networks, by the zillions, bestowed with selves, will rethink the chips and rule all that we let them. If we had tried to exploit the treasures of energy, material, and information by holding all the control, it would have been a loss.

As fast as our lives allow us, we are equipping our constructed world to bootstrap itself into self-governance, self-reproduction, self-consciousness, and irrevocable selfhood. The story of automation is the story of a one-way shift from human control to automatic control. The gift is an irreversible transfer from ourselves to the second selves.

The second selves are out of our control. This is the key reason, I believe, why the brightest minds of the Renaissance never invented another self-regulator beyond the obvious ones known to ancient Heron. The great Leonardo da Vinci built control machines, not out-of-control machines. German historian of technology Otto Mayr claims that great engineers in the Enlightenment could have built regulated steam power of some sort with the technology available to them at the time. But they didn't because they didn't have the ability to let go of their creation.

The ancient Chinese on the other hand, although they never got beyond the south-pointing cart, had the right no-mind about control. Listen to these most modern words from the hand of the mystical pundit Lao Tzu, writing in the Tao Teh King 2,600 years ago:

Intelligent control appears as uncontrol or freedom.

And for that reason it is genuinely intelligent control.

Unintelligent control appears as external domination.

And for that reason it is really unintelligent control.

Intelligent control exerts influence without appearing to do so.

Unintelligent control tries to influence by making a show of force.

Lao Tzu's wisdom could be a motto for a gung-ho 21st-century Silicon Valley startup. In an age of smartness and superintelligence, the most intelligent control methods will appear as uncontrol methods. Investing machines with the ability to adapt on their own, to evolve in their own direction, and grow without human oversight is the next great advance in technology.

The chief psychological chore of the 21st century will be letting go, with dignity. Until recently, all our artifacts, all our own handmade creations have been under our authority. But as we cultivate synthetic life in our artifacts, we cultivate the loss of our command. "Out of control," to be honest, is a great exaggeration of the state that our enlivened machines will take. They will remain indirectly under our influence and guidance but free of our domination.

Though I have searched everywhere, I could not find the word that describes this type of clout. We simply have no name for the loose relationship between an influential creator and a creation with a mind of its own -- a thing we shall see more of. The realm of parent and child should have such a word, but sadly doesn't. We do better with sheep where we have the notion of "shepherding." When we herd a flock of sheep, we know we are not in complete authority, yet neither are we without control. Perhaps we will shepherd artificial lives.

We also "husband" plants, as we assist them in their natural goals, or deflect them slightly for our own. "Manage" is probably the closest in meaning to the general type of control we will need for artificial lives, such as a virtual Mickey Mouse. A women can "manage" her difficult child, or a barking dog, or the 300-strong sales department under her authority.

"Manage" is close, but not perfect. Although we manage wilderness areas like the Everglades, we actually have little say in what goes on among the seaweed, snakes and marsh grass. Although we manage the national economy, it does what it wants. And although we manage a telephone network, we have no supervision on how a particular call is completed. The word "management" may imply more oversight then we really have in the examples above, and more than we will have in future very complex systems.

The word I'm looking for is more like "co-control." It's seen in some mechanical settings already. Keeping a 747 Jumbo Jet aloft and landing it in bad weather is a very complex task. Because of the hundreds of systems running simultaneously, the immediate reaction time required by the speed of the plane, and disorienting effects of sleepless long trips and hazardous weather, a computer can fly a jet better a human pilot. The sheer number of human lives at stake permits no room for errors or second best. Why not have a very smart machine control the jet?

So engineers wired together an autopilot, and it turns out be very capable. It flies and lands a Jumbo Jet oh so nicely. Flying-by-wire also fits very handily into the craving for order by the air traffic controllers -- everything is under digital control. The original idea was that human pilots would monitor the computer in case anything went wrong. The only problem is that humans are terrible at passive monitoring. They get bored. They daydream. Then they start missing critical details. Then an emergency pops up which they have to tackle cold.

So instead of having the pilot watch the computer, the new idea was to invert the relationship and have the computer watch the pilot. This approach was taken in the European Airbus A320, one of the most highly automated planes built to date. Introduced in 1988, the onboard computer supervises the pilot. When he pushes the control stick to turn the plane, the computer figures out how far to bank left or right, but it won't let the plane bank more than 67 degrees or nose up or down more than 30 degrees. This means, in the words of Scientific American, "the software spins an electronic cocoon that stops the aircraft from exceeding its structural limitations." It also means, pilots complain, that the pilot surrenders control. In 1989 British Airways pilots flying 747s experienced six different incidents where they had to override a computer-initiated power reduction. Had they not been able to override the erroneous automatic pilot -- which Boeing blamed on a software bug -- the error could have been fatal. The Airbus A320, however, provides no override of its autosystem.

Human pilots felt they were fighting for control of the plane. Should the computer be a pilot or navigator? The pilots joked that the computer was like putting a dog into the cockpit. The dog's job was to bite the pilot if he tries to touch the controls; and the pilot's only job was to feed the dog. In fact, in the emerging lingo of automated flying, pilots are called "system managers."

On one hand, the computer can be seen as an autonomous entity unto itself, an "artificial" colleague. On the other, the user-machine conversation can be seen as a sort of internal dialogue, as if the computer were a prosthesis of the brain. An extension of the thinking processes of the user.

There is much that computer already does outside of the reach of the pilot. Planes will fly by co-control. But the pilot will manage, or shepherd, the computer's behavior. Computers will be able to perform as powerful prostheses, coevolving with their users to enable new modes of creative thought, communication, and collaboration.

The future of control: Partnership, Co-control, Cyborgian control. What it all means is that the creator must share control, and his destiny, with his creations.

Closed systems

At one end of a long row of displays in the Steinhart Aquarium in San Francisco, a concentrated coral reef sits happily tucked under lights. The Aquarium's self-contained South Pacific ocean compresses the distributed life in a mile-long underwater reef into a few glorious yards behind glass.

The condensed reef's extraordinary hues and alien life forms cast a New Age vibe. To stand in front of this rectangular bottle is to stand on a harmonic node. Here are more varieties of living creatures crammed into a square meter than anywhere else on the planet. Life does not get any denser. The remarkable natural richness of the coral reef has been squeezed further into the hyper-natural richness of a synthetic reef.

A pair of wide plate glass windows peer into an Alician wonderland of exotic beings. Fish in hippie day-glo colors stare back-accents of orange- and white-banded clown fish or a minischool of iridescent turquoise damsels. The flamboyant creatures scoot between the feathery wands of chestnut-tinted soft corals or weave between the slowly pulsating fat lips of giant sea clams.

No mere holding pen, this is home for these creatures. They will eat, sleep, fight, and breed among each other, forever if they can. Given enough time, they will coevolve toward a shared destiny. Theirs is a true living community.

Behind the coral display tank, a clanking army of pumps, pipes, and gizmos vibrate on electric energy to support the toy reef's ultradiversity. A visitor treks to the pumps from the darkened viewing room of the aquarium by opening an unmarked door. Blinding E.T.-like light gushes out of the first crack. Inside, the white-washed room suffocates in warm moisture and stark brightness. An overhead rack of hot metal halide lamps pumps out 15 hours of tropical sun per day. Saltwater surges through a bulky 4-ton concrete tub of wet sand brimming with cleansing bacteria. Under the artificial sunlights, long, shallow plastic trays full of green algae thrive filtering out the natural toxins from the reef water.

Industrial plumbing fixtures are the surrogate Pacific for the reef. Sixteen thousand gallons of reconstituted ocean water swirl through the bionic system to provide the same filtration, turbulence, oxygen, and buffering that the miles of South Pacific algae gardens and sand beaches perform for a wild reef. The whole wired show is a delicate, hard-won balance requiring daily energy and attention. One wrong move and the reef could unravel in a day.

As the ancients knew, what can unravel in a day may take years or centuries to build. Before the Steinhart coral reef was constructed, no one was sure if a coral reef community could be assembled artificially, or how long it would take if it could. Marine scientists were pretty sure a coral reef, like any complex ecosystem, must be assembled in the correct order. But no one knew what that order was. Marine biologist Lloyd Gomez certainly didn't know when he first started puttering around in the dank basement of the Academy's aquarium building. Gomez mixed buckets of microorganisms together in large plastic trays, gradually adding species in different sequences in hopes of attaining a stable community. He built mostly failures.

He began each trial by culturing a thick pea-green soup of algae -- the scum of a pond out of whack -- directly under the bank of noon-lights. If the system started to drift away from the requirements of a coral reef, Gomez would flush the trays. Within a year, he eventually got the proto-reef soup headed in the right direction.

It takes time to make nature. Five years after Gomez launched the coral reef, it is only now configuring itself into self-sustenance. Until recently Gomez had to feed the fish and invertebrates dwelling on the synthetic reef with supplemental food. But now he thinks the reef has matured. "After five years of constant babying, I have a full food web in my tank so I no longer have to feed them anything." Except sunlight, which pours on the artificial reef in a steady burst of halide energy. Sunlight feeds the algae which feed the animals which feed the corals, sponges, clams, and fish. Ultimately this reef runs on electricity.

Gomez predicts further shifts as the reef community settles into its own. "I expect to see major changes until it is ten years old. That's when the reef fusing takes place. The footing corals start to anchor down on the loose rocks, and the subterranean sponges burrow underneath. It all combines into one large mass of animal life." A living rock grown from a few seed organisms.

Much to everyone's surprise, about 90 percent of the organisms that fuse the toy reef were stowaways that did not appear to be present in the original soup. A sparse but completely invisible population of the microbes were present, but not until five years down the road, when the reef had prepared itself to be fused, were the conditions right for the blossoming of the fuser microorganisms which had been floating unseen and patient.

During the same time, certain species dominating the initial reef disappeared. Gomez says, "I was not expecting that. It startled me. Organisms were dying off. I asked myself what did I do wrong? It turns out that I didn't do anything wrong. That's just the community cycle. Heavy populations of microalgae need to be present at first. Then within ten months, they've gone. Later, some initially abundant sponges disappeared, and another type popped up. Just recently a black sponge has taken up in the reef. I have no idea where it came from." As in the restorations of Packard's prairie and Wingate's Nonsuch Island, chaperone species were needed to assemble a coral but not to maintain it. Parts of the reef were "thumbs."

Lloyd Gomez's reef-building skills are in big demand at night school. Coral reefs are the latest challenge for obsessive hobbyists, who sign up to learn how to reduce oceanic monuments to 100 gallons. Miniature saltwater systems shrink miles of life into a large aquarium, plus paraphernalia. That's dosing pumps, halide lights, ozone reactors, molecular absorption filters, and so on, at a cool $15,000 per living room tank. The expensive equipment acts like the greater ocean, cleaning, filtering the reef's water. Corals demand a delicate balance of dissolved gases, trace chemicals, pH, microorganisms, light, wave action, temperature -- all of which are provided in an aquarium by an interconnected network of mechanical devices and biological agents. The common failure, Gomez says, is trying to stuff more species of life into the habitat than the system can carry, or not introducing them in the correct sequence, as Pimm and Drake discovered. How critical is the ordering? Gomez: "As critical as death."

The key to stabilizing a coral reef seemed to be getting the initial microbial matrix right. Clair Folsome, a microbiologist working at the University of Hawaii, had concluded from his own work with microbial soups in jars that "the foundation for stable closed ecologies of all types is basically a microbial one." He felt that microbes were responsible for "closing the bio-elemental loops" -- the flows of atmosphere and nutrients -- in any ecology. He found his evidence in random mixtures of microbes, similar to the experiments of Pimm and Drake, except that Folsome sealed the lid of the jars. Rather than model a tiny slice of life on Earth, Folsome modeled a self-contained self-recycling whole Earth. All matter on Earth is recycled (except for the insignificant escape of a trace of light gases and the fractional influx of meteorites). In system-science terms, we say Earth is materially closed. The Earth is also energetically/informationally open: sunlight pours in, and information comes and goes. Like Earth, Folsome's jars were materially closed, energetically open. He scooped up samples of brackish microbes from the bays of the Hawaiian Islands and funneled them into one- or two-liter laboratory glass flasks. Then he sealed them airtight and, by extracting microscopic amounts from a sampling port, measured their species ratios and energy flow until they stabilized.

Just as Pimm was stunned to find how readily random mixtures settled into self-organizing ecosystems, Folsome was surprised to see that even the extra challenge of generating closed nutrient recycling loops in a sealed flask didn't deter simple microbial societies from finding an equilibrium. Folsome said that he and another researcher, Joe Hanson, realized in the fall of 1983 that closed ecosystems "having even modest species-diversity, rarely if ever fail." By that time some of Folsome's original flasks had been living for 15 years. The oldest one, thrown together and sealed in 1968, is now 25 years old. No air, food, or nutrients have ever been added. Yet this and all of his other jar communities are still flourishing years later under florescent room lights.

No matter how long they lived, though, the bottled systems required an initial staging period, a time of fluctuation and precarious instability lasting between 60 and 100 days, when anything might happen. Gomez saw this in his coral microbes: the beginnings of complexity are rooted in chaos. But if a complex system is able to find a common balance after a period of give and take, thereafter not much will derail it.

How long can such closed complexity run? Folsome said his initial interest in making materially closed worlds was sparked by a legend that the Paris National Museum displayed a cactus sealed in a glass jar in 1895. He couldn't verify its existence, but it was claimed to be covered with recurrent blooms of algae and lichens that have cycled through a progression of colors from shades of green to hues of yellow for the past century. If the sealed jar had light and a steady temperature, there was theoretically no reason why the lichens couldn't live until the sun dies.

Folsome's sealed microbial miniworlds had their own living rhythms that mirrored our planet's. They recycled their carbon, from CO2 to organic matter and back again, in about two years. They maintained biological productivity rates similar to outside ecosystems. They produced stable oxygen levels slightly higher than on Earth. They registered energy efficiencies similar to larger ecosystems. And they maintained populations of organisms apparently indefinitely.

From his flask worlds, Folsome concluded that it was microbes -- tiny celled microbits of life, and not redwoods, crickets, orangutans -- which do the lion's share of breathing, generating air, and ultimately supporting the indefinite populations of other noticeable organisms on Earth. An invisible substrate of microbial life steers the course of life's whole and welds together the different nutrient loops. The organisms that catch our eye and demand our attention, Folsome suspected, were mere ornate, decorative placeholdings as far as the atmosphere was concerned. It was the microbes in the guts in mammals and the microbes that clung to tree roots that made trees and mammals valuable in closed systems, including our planet.

I once had a tiny living planet stationed on my desk. It even had a number: world #58262. I didn't have much to do to keep my planet happy. Just watch it every now and then.

World #58262 was smashed to smithereens at 5:04 P.M., October 17, during an abrupt heave of the 1989 San Francisco earthquake. A bookcase shook loose from my office wall during the tremor and spilled over my desk. In a blink, a heavy tome on ecosystems crushed the glass membrane of my living planet, irrevocably scrambling its liquid guts in a fatal Humpty Dumpty maneuver.

World #58262 was a human-made biosphere of living creatures, delicately balanced to live forever, and a descendent of Folsome's and Hanson's microbial jars. Joe Hanson, who worked at NASA's Advance Life-support Program in the Jet Propulsion Laboratory at Caltech, had come up with a more diverse world than Folsome's microbes. Hanson was the first to find a simple combination of self-sustaining creatures that included an animal. He put tiny brine shrimp and brine algae in an everlasting cosmos.

The basic commercial version of his closed world -- sold under the label of "Ecosphere" -- is a glass globe about the size of a large grapefruit. My world #58262 was one of these. Completely sealed inside the transparent ball were four tiny brine shrimp, a feathery mass of meadowgreen algae draped on a twig of coral, and microbes in the invisible millions. A bit of sand sat on the bottom. No air, water, or any other material entered or exited the globe. The thing ate only sunlight.

The oldest living Hanson-world so far is ten years old; that's as long as they have been manufactured. That's surprising since the average life-span of the shrimp swimming inside was thought to be about five years. Getting them to reproduce in their closed world has been problematic, although researchers know of no reason why they could not go on replicating forever. Individual shrimp and algae cells die, of course. What "lives forever" is the collective life, the aggregate life of a community.

You can buy an Ecosphere by mail order. It's like buying a Gaia or an experiment in emergent life. You unpack the orb from the heavy-duty insulation stuffed around it. The shrimp seem fine after their stormy ride. Then you hold the cannonball-size sphere in one hand up to the light; it sparkles with gemlike clarity. Here is a world blown into a bottle, the glass tidily pinched off at the top.

In its fragile immortality, the Ecosphere just sits there. Naturalist Peter Warshall, who owns one of the first Ecospheres, keeps it perched on his bookshelf. Warshall reads obscure dead poets and French philosophers in French and monographs on squirrel taxonomy. Nature is a kind of poetry for him; an Ecosphere is a book jacket blurb about the real thing. Warshall's Ecosphere lives under a regime of benign neglect, almost as a maintenance-free pet. He writes of his nonhobby: "You can't feed the shrimp. You can't snip off the decaying, dreary brown parts. You can't fiddle with the nonexistent filter, aerator, or pumps. You can't open it up and test the water's warmth with your finger. All you can do, if 'do' is an appropriate word, is to look and think."

The Ecosphere is a totem, a totem of all closed living systems. Tribesmen select totem creatures as a bridge between the separate worlds of spirit and dreams. Simply by being, the distinct world sealed behind an Ecosphere's clear glass invites us to meditate on such hard-to-grasp totemic ideas like "systems," "closed," and even "living."

"Closed" means separated from the flow. A manicured flower garden on the edge of the woods exists apart from the naturally structured wilderness surrounding, but the separateness of a garden mesocosm is partial -- more a division of mind than fact. Every garden is really a small slice of the larger biosphere we all are immersed in. Moisture and nutrients flow underground into it, and a harvest and oxygen come out. If the rest of the sustaining biosphere were absent, gardens would wither. A truly closed system does not partake in outside flows of elements; all its cycles are autonomous.

"System" means interconnected. Things in a system are intertwined, linked directly or indirectly into a common fate. In an ecospheric world, shrimp eat algae, algae live on the light, microbes survive on the "wastes" of both. If the temperature soars too high (above 90 degrees), the shrimp molt faster than they can eat; thus they consume themselves. Not enough light and the algae won't grow fast enough to satiate the shrimp. The flicking tails of the shrimp stir up the water, which stirs the microbes so that each bug has a chance to catch the sunlight. The whole has a life in addition to the individual lives.

"Living" means surprises. One ordinary Ecosphere managed to stay alive in a total darkness for six months, contrary to logical expectations. Another ecosystem erupted one day after two years of unwavering steady temperature and light in an office into a breeding panic, crowding the globe with 30 tiny descendants of shrimp.

But it is stasis that does an Ecosphere in. In an unguarded moment Warshall writes of his orb, "There is the feeling of too much peacefulness that comes from the Ecosphere. It contrasts sharply with our frantic, daily lives. I have felt like playing the abiotic God. Pick it up and shake it. How's that for an earthquake, you little shrimp!"

That would actually be a good thing for an Ecosphere world, as momentarily discombobulating as it might be for its citizens. In turbulence is the preservation of the world.

A forest needs the severe destruction of hurricanes to blow down the old and make space for the new. The turbulence of fire on the prairie unloosens bound materials that cannot be loosened unless ignited. A world without lightning and fire becomes rigid. An ocean has the fire of undersea thermal vents in the short run, and the fire of compressed seafloor and continental plates in the long geological run. Flash heat, volcanism, lightning, wind, and waves all renew the material world.

The Ecosphere has no fire, no flash, no high levels of oxygen, no serious friction -- even in its longest cycle. Over a period of years in its small space, phosphate, an essential element in all living cells, becomes tightly bound with other elements. In a sense, phosphate is taken out of circulation in the Ecosphere, diminishing the prospects of more life. Only the thick blob of blue-green algae will thrive in low phosphate environment, and so over time this species tends to dominate these stable systems.

A phosphate sink, and the inevitable takeover of blue-green algae, might be reversed by adding, say, a lightning-generating appendage to the glass globe. Several times a year, the calm world of the shrimp and algae would crackle and hiss and boil as calamity reigned for a few hours. Their vacations would be ruined, but their world would be rejuvenated.

In Peter Warshall's Ecosphere (which despite his idle thoughts has lain undisturbed for years), minerals have precipitated into a layer of solid crystals on the globe's inside. In a Gaian sense, the Ecosphere manufactured land. The "land" -- composed of silicates, carbonates, and metal salts -- built up on the glass because of an electric charge, a kind of natural electroplating. Don Harmony, the chief honcho at the small company making Ecospheres, was familiar with this tendency of tiny glass Gaia, and half in jest suggested that perhaps fusing an electrical ground wire onto the globe might keep the precipitates from forming.

Eventually the weight of the salt crystals peels them off the upper surface and they settle into the bottom of the liquid. On Earth, the deposit of sedimentary rock at the bottom of the ocean is part of larger geological cycles. Carbon and minerals circulate through air, water, land, rocks, and back again into life. Likewise in the Ecosphere. The elements it cradles are in a dynamic equilibrium with the cycling composition of the atmosphere and water and biosphere.

Most field ecologists were surprised by how simple such a self-sustaining closed world could be. With the advent of this toy biosphere, sustainable self-sufficiency appeared to be quite easy to create, especially if you didn't care what kind of life was being sustained. The Ecosphere was a mail-order proof of a remarkable assertion: self-sustained systems want to happen.

If simple and tiny was easy, how far could you expand the harmony and still have a sustainable world closed to all but energy input?

It turns out that ecospheres scale up well. A huge commercial Ecosphere can weigh in at 200 liters. That's about the volume of a large garbage can -- so big you can't reach your arms around it. Inside a stunning 30-inch-diameter glass globe, shrimp paddle between fronds of algae. But instead of the usual three or four spore-eating shrimp, the giant Ecosphere holds 3,000. It's a tiny moon with its own inhabitants. Here, the law of large numbers takes hold; more is different. More individual lives make the ecosystem more resilient. The larger an Ecosphere is, the longer it takes to stabilize, and the harder it is to kill it. But once in gear, the collective give and take of a vivisystem takes root and persists.

The next question is evident: How big a bottle closed to outside flows, filled with what kind of living organisms, would you need to support a human inside?

When human daredevils ventured beyond the soft bottle of the Earth's atmosphere, this once academic question took on practical meaning. Could you keep a person alive in space -- like shrimp in an Ecosphere -- by keeping plants alive? Could you seal a man up in a sunlit bottle with enough living things so that their mutual exhalations would balance? It was a question worth doing something about.

Every school child knows animals consume the oxygen and food that plants generate, while plants consume the carbon dioxide and nutrients that animals generate. It's a lovely mirror, one side producing what the other needs, just as the shrimp and algae serve each other. Perhaps the right mix of plants and mammals in their symmetrical demands could support each other. Perhaps a human could find its proper doppelganger of organisms in a closed bottle.

The first person crazy enough to experimentally try this was a Russian researcher at the Moscow Institute for Biomedical Problems. In 1961, during the heady early years of space research, Evgenii Shepelev welded together a steel casket big enough to hold himself and eight gallons of green algae. Shepelev's careful calculations showed that eight gallons of chlorella algae under sodium lights should supply enough oxygen for one man, and one man should generate enough carbon dioxide for eight gallons of chlorella algae. The two sides of the equation should cancel each other out into unity. In theory it should work. On paper it balanced. On the blackboard it made perfect sense.

Inside the airtight iron capsule, it was a different story. You can't breathe theories. If the algae faltered, the brilliant Shepelev would follow; or, if he succumbed, the algae would do likewise. In the box the two species would become nearly symbiotic allies entirely dependent on each other, and no longer dependent upon the vast planetary web of support outside -- the oceans, air, and creatures large and small. Man and algae sealed in the capsule divorced themselves from the wide net woven by the rest of life. They would be a separate, closed system. It was by an act of faith in his science that a trim Shepelev crawled into the chamber and sealed the door.

Algae and man lasted a whole day. For about 24 hours, man breathed into algae and algae breathed into man. Then the staleness of the air drove Shepelev out. The oxygen content initially produced by the algae plummeted rapidly by the close of the first day. In the final hour when Shepelev cracked open the sealed door to clamber out, his colleagues were bowled over by the revolting stench in his cabin. Carbon dioxide and oxygen had traded harmoniously, but other gases, such as methane, hydrogen sulfide, and ammonia, given off by algae and Shepelev himself, had gradually fouled the air. Like the mythological happy frog in slowly boiling water, Shepelev had not noticed the stink.

Shepelev's adventuresome work was taken up in seriousness by other Soviet researchers at a remote and secret lab in northern Siberia. Shepelev's own group was able to keep dogs and rats alive within the algae system for up to seven days. Unbeknownst to them, about the same time the United States Air Force School of Aviation Medicine linked a monkey to an algae-produced atmosphere for 50 hours. Later, by parking the tiny eight-gallon tub of chlorella in a larger sealed room, and tweaking the algae nutrients as well as the intensity of lights, Shepelev's lab found that a human could live in this airtight room for 30 days! At this extreme duration the researchers noticed that the respirations of man and algae were not exactly matched. To keep a balance of atmosphere, excess carbon dioxide needed to be removed by chemical filters. But the scientists were encouraged that stinky methane stabilized after 12 days.

By 1972, more than a decade later, the Soviet team, directed by Josepf Gitelson, constructed the third version of a small biologically based habitat that could support humans. The Russians called it Bios-3. It housed up to three men. The habitat was crowded inside. Four small airtight rooms enclosed tubs of hydroponically (soil-less) grown plants anchored under xenon lights. The men-in-a-box planted and harvested the kind of crops you might expect in Russia -- potatoes, wheat, beets, carrots, kale, radishes, onions and dill. From the harvest they prepared about half of their own food, including bread from the grain. In this cramped, stuffy, sealed greenhouse, the men and plants lived on each other for as long as six months.

The box was not perfectly closed. While its atmosphere was sealed to air exchanges, the setup recycled only 95 percent of its water. The Soviet scientists stored half of their food (meat and proteins) beforehand. In addition, the Bios-3 system did not recycle human fecal wastes or kitchen scraps; the Bios-dwellers ejected these from the container, thereby ejecting some trace elements and carbon.

In order not to lose all carbon from the cycle, the inhabitants burned a portion of the inedible dead plant matter rendering it into carbon dioxide and ash. Over weeks the rooms accumulated trace gases generated by a number of sources: the plants, the materials of the room, and the men themselves. Some of these vapors were toxic, and methods to recycle them unknown then, so the men burned off the gases by simply "burning" the air inside with a catalytic furnace.

NASA, of course, was interested in feeding and housing humans in space. In 1977 they launched the still-going CELSS program (Controlled Ecological Life Support Systems). NASA took the reductionist approach: find the simplest units of life that can produce the required oxygen, protein, and vitamins for human consumption. It was in messing around with elemental systems that NASA's Joe Hanson stumbled on the interesting, but to NASA's eyes, not very useful shrimp/algae combo.

In 1986 NASA initiated the Breadboard Project. The program's agenda was to take what was known from tabletop experiments and implement them at a larger scale. Breadboard managers found an abandoned cylinder left over from the Mercury space shots. This giant tubular container had been built to serve as pressure-testing chamber for the tiny astronaut capsule that would spearhead the Mercury rocket. NASA retrofitted the two-story cylinder with outside ductwork and plumbing, transforming the interior into a bottled home with racks of lights, plants, and circulating nutrients.

Just as the Soviet Bios-3 experiments did, Breadboard used higher plants to balance the atmosphere and provide food. But a human can only choke down so much algae each day. Even if algae was all one ate, chlorella only provides 10 percent of the daily nutrients a person needs. For this reason, NASA researchers drifted away from algae-based systems, and migrated toward plants that provided not only clean air but also food.

Ultra-intensive gardening seemed be what everyone was coming up with. Gardening could produce really edible stuff, like wheat. Among the most workable setups were various hydroponic contraptions that delivered aqueous nutrients to plants as a mist, a foam, or a thin film dripping through plastic holding racks matted with lettuce or other greens. This highly engineered plumbing produced concentrated plant growth in cramped spaces. Frank Salisbury of Utah State University discovered ways to plant spring wheat at 100 times its normal density by precisely controlling the wheat's optimal environment of light, humidity, temperature, carbon dioxide, and nutrients.

Extrapolating from his field results, Salisbury calculated the amount of calories one could extract from a square meter of ultradensely planted wheat sown, say, on enclosed lunar base. He concluded that "a moon farm about the size of an American football field would support 100 inhabitants of Lunar City."

One hundred people living off a football field-size truck farm! The vision was Jeffersonian! One could envision a nearby planet colonized by a network of Superdome villages, each producing its own food, water, air, people, and culture.

But NASA's approach to inventing a living in a closed system struck many as being overly cautious, strangulatingly slow, and intolerably reductionistic. The operative word for NASA's Controlled Ecological Life Support Systems was "Controlled."

What was needed was a little "out-of-control."

The appropriate out-of-controlness started on a ramshackle ranch near Santa Fe, New Mexico. During the commune heydays of the early 1970s, the ranch collected a typically renegade group of cultural misfits. Most communes then were freewheeling. This one, named Synergia Ranch, wasn't; it demanded discipline and hard work. Rather than lie back and whine while the apocalypse approached, the New Mexican commune worked on how it might build something to transcend the ills of society. They came up with several designs for giant arks of sanity. The more grandiose their mad ark visions got, the more interested in the whole idea they all became.

It was the commune's architect, Phil Hawes, who came up with the galvanizing idea. At a 1982 conference in France, Hawes presented a mock-up of a spherical, transparent spaceship. Inside the glass sphere were gardens, apartments, and a pool beneath a waterfall. "Why not look at life in space as a life instead of merely travel?" Hawes asked. "Why not build a spaceship like the one we've been traveling on?" That is, why not create a living satellite instead of hammering together a dead space station? Reproduce the holistic nature of Earth itself as a tiny transparent globe sailing through space. "We knew it would work," said John Allen, the ranch's charismatic leader, "because that's what the biosphere does every day. We just had to get the size right."

The Synergians stuck with the private vision of a living ark long after they left the ranch. In 1983, Ed Bass of Texas, one of the ranch's former members, used part of his extraordinary family oil fortune to finance a proof-of-concept prototype.

Unlike NASA, the Synergians wouldn't rely on technology as the solution. Their idea was to stuff as many biological systems -- plants, animals, insects, fish, and microorganisms -- as they possibly could into a sealed glass dome, and then rely on the emergent system's own self-stabilizing tendencies to self-organize a biospheric atmosphere. Life is in the business of making its environment agreeable for life. If you could get a bunch of life together and then give it enough freedom to cultivate the conditions it needed to thrive, it would go forever, and no one needed to understand how it worked.

Indeed, neither they nor biologists had any real idea of how one plant worked -- what's its exact needs and products were -- and no idea at all of how a distributed miniecosystem sealed in a hut would work. Instead, they would rely on decentralized, uncontrolled life to sort itself out and come to some self-enhancing harmony.

No one had ever built any living thing that large. Even Gomez hadn't built his coral reef yet. The Synergians had only a vague notion of Clair Folsome's ecospheres and even vaguer knowledge of the Russian Bios-3 experiments.

The group, now calling itself Space Biosphere Ventures (SBV), and financed to the tune of tens of millions of dollars by Ed Bass, designed and built a tiny cottage-size test unit during the mid-1980s. The hut was crammed with a greenhouse-worth of plants, some fancy plumbing for recycling water, black boxes of sensitive environmental monitoring equipment, a tiny kitchenette and bathroom, and lots of glass.

In September 1988, for three days, John Allen sealed himself in for the unit's first trial run. Much like Evgenii Shepelev's bold step, this was a act of faith. The plants had been selected by rational guess, but there was nothing controlled about how well they would work as a system. Contrary to Gomez's hard-won knowledge about sequencing, the SBV folks just threw everything in together, at once. The sealed home depended on at least some of the individual plants being able to keep up with the lungs of one man.

The test results were very encouraging. Allen wrote in his journal for September 12: "It appears we are getting close to equilibrium, the plants, soil, water, sun, night and me." In the confined loop of a 100 percent recycled atmosphere, 47 trace gases, "all of which were probably anthropogenic in origin," fell to minute levels when the air of the hut was sent through the plant soil -- an old technique modernized by SBV. Unlike Shepelev's case, when Allen stepped out, the air inside was fresh, ready for more human life. To someone outside, a whiff of the air inside was shockingly moist, thick, and "green."

The data from Allen's trial suggested a human could live in the hut for a while. Biologist Linda Leigh would later spend three weeks in the small glass shed. After her 21-day solo drive Leigh told me, "At first I was concerned whether I'd be able to stand breathing in there, but after two weeks I hardly noticed the moisture. In fact I felt invigorated, more relaxed, and healthier, probably because of the air-cleansing and oxygen-producing nature of close plants. The atmosphere even in that small space was stable. I felt that the test module could have gone on for the full two years and kept its atmosphere right."

During the three-week run, the sophisticated internal monitoring equipment indicated no buildup of gases either from building materials or biological sources. Although the atmosphere was stable overall, it was sensitive to perturbations which caused it to vacillate easily. While harvesting sweet potatoes out of their dirt beds in the hut, Leigh's digging disturbed CO_2-producing soil organisms. The rattled bugs temporarily altered the CO_2- concentration in the module's air. This was an illustration of the butterfly effect. In complex systems a small alteration in the initial conditions can amplify into wide-ranging effects throughout the rest of the system. The principle is usually illustrated by the fantasy of the flap of a butterfly's wings in Beijing triggering a hurricane in Florida. Here in SBV's sealed glass cottage the butterfly effect appeared in miniature: by wiggling her fingers Leigh upset the balance of the atmosphere.

John Allen and another Synergian, Mark Nelson, envisioned a near-future Mars station built as a mammoth closed-system bottle. Allen and Nelson gradually formulated a hybrid technology -- called ecotechnics -- based on a convergence of both machines and living organisms to support future human habitats.

They were dead serious about going to Mars and began working out the details. In order to journey to Mars or beyond, you needed a crew. How many people would you need? Military captains, expedition leaders, start-up managers, and crisis centers had long recognized that a team of eight was the ideal number for any complex hazardous project. More than eight people, and decisions got slow and squirrely; less than eight, accidents and ignorance became serious handicaps. Allen and Nelson settled on a crew of eight.

Next step: how big would you have to make a bottle-world to shelter, feed, water, and oxygenate eight people indefinitely?

Human requirements were well established. Each day a human adult needed about half a kilogram of food, a kilo of oxygen, 1.8 kilos of drinking water, FDA amounts of vitamins, and a couple of gallons of water for washing. Clair Folsome had extrapolated the results of his tiny ecospheres and calculated that you would need a sphere with a radius of 58 meters -- half air and half microbial soup -- to support the oxygen needs for one person indefinitely. Allen and Nelson then took the data from the Russian Bios-3 experiments and combined it with Folsome's, Salisbury's, and others' intensive farming harvest results. They estimated that right now -- with the knowledge and technology of 1980s -- they could support eight adults on...three acres of land.

Three acres! The transparent container would have to be the size of the Astrodome. Such a span would demand at least a 50-foot ceiling. Clothed in glass, it would be quite a sight. And quite expensive.

But it would be magnificent! They would build it! And they did, with the further help of Ed Bass - to the tune of $100 million. Hard-hat construction of the 8-person ark began in 1988. The Synergians called the grand project Biosphere 2 (Bio2), a bonsai version of Biosphere 1, our Earth. It took three years to build.

Small compared to Earth, the completed self-contained terrarium was awesome at the human scale. Bio2 was a gigantic glass ark the size of an airport hangar. Think of an inverted ocean liner whose hull is transparent. The gigantic greenhouse was superairtight, sealed at the bottom, too, with a stainless steel tray 25 feet under the soil to prevent seepage of air from its basement. No gas, water, or matter could enter or leave the ark. It was a stadium-size Ecosphere -- a big materially closed and energetically open system -- but far more complex. Bio2 was the second only to Biosphere 1 (the Earth) as largest closed vivisystem.

The challenge of creating a living system of any size is daunting. Creating a living wonder at the scale of Bio2 could only be described as an experiment in sustained chaos. The challenge included: Select a couple of thousand parts of out of several billion possibilities, and arrange them so that all the parts complemented and provided for each other, so that the whole mixture was self-sustaining over time, and that no single organism became dominant at the expense of others, so that the whole aggregate kept all the constituents in constant motion, without letting any ingredient become sequestered off to the side, while keeping the entire level of activity and atmospheric gases elevated at the point of perpetually almost-falling. Oh, and humans should be able to live, eat, and drink within and from it.

SBV decided to stake the survival of Bio2 on the design tenet that an extraordinarily diverse hodgepodge of living creatures would settle into a unified stability. If it proved nothing else, the experiment would at least shed some light on the almost universally held assumption in the last two decades: that diversity ensures stability. It would also test whether a certain level of complexity birthed self-sustainability.

As an architecture of maximum diversity, the final Bio2 floor plan had seven biomes (biogeographical habitats). Under the tallest part of the glass canopy, a rock-faced concrete mountain bulged. Planted with transplanted tropical trees and a misting system, the synthetic hill was transformed into a cloud forest -- a high altitude rain forest. The cloud forest drained into an elevated hot grassland (the size of a big patio, but stocked with waist-high wild grasses). One edge of the rain forest stopped before a rocky cliff which fell to a saltwater lagoon, complete with coral, colorful fishes, and lobsters. The high savanna lowered into a lower, drier savanna, dark with thorny, tangled thickets. This biome is called thornscrub and is one of the most common of all habitats on Earth. In real life it is nearly impenetrable to humans (and thus ignored), but in Bio2, it served as a little hideaway for both wildlife and humans. The thicket leads into a compact marshy wetland, the fifth biome, which finally emptied into the lagoon. The low end of Bio2 was a desert, as big as a gymnasium. Since it was pretty humid inside, the desert was planted with fog desert plants from Baja California and South America. Off to one side was the seventh biome -- an intensive agriculture and urban area where eight Homo sapiens grew all their own food. Like Noah's place, animals were aboard; some for meat, some for pets, and some on the loose: lizards, fish, birds roaming about the wild parts. There were honey bees, papaya trees, a beach, cable TV, a library, a gym, and a laundromat. Utopia!

The scale was stupendous. Once while I was visiting the construction site, an 18-wheeler semi-truck pulled up to the Bio2 office. The truck driver leaned out the window and asked where they wanted their ocean. He'd been hauling a full truckload of ocean salt and needed to unload it before dark. The office clerks pointed down to a very large hole in the center of the project. That's where Walter Adey from the Smithsonian Institution was building a one-million-gallon ocean, coral reef, and lagoon. There was enough elbow room in this gargantuan aquarium for all kinds of surprises to emerge.

Making an ocean is no cinch. Ask Gomez and the hobby saltwater aquarists. Adey had grown an artificial self-regenerating coral reef once before as a museum exhibit at the Smithsonian. But this one in Bio2 was huge; it had its own sandy beach. An expensive wave-making pump at one end would supply the turbulence coral love. The same machine created a half-meter tide on a lunar cycle.

The trucker unloaded the ocean: stacks of 50-pound bags of Instant-Ocean, the same stuff you buy at tropical aquarium stores. A starter solution harboring all the right microbeasties (sort of the yeast for the dough) was later hauled in on a different truck from the Pacific Ocean. Stir together well, and pour.

The ecologists building the wilderness areas of Bio2 were of the school that says: soil + bugs = ecology. To have the kind of tropical rainforest you want, you needed to have the right kind of jungle dirt. And to get that in Arizona you had to make it from scratch. Take a couple of bulldozer buckets of basalt, a few of sand, and a few of clay. Sprinkle in the right microorganisms. Mix in place. The underlying soils in each of the six wild biomes of Bio2 were manufactured in this painstaking way. "The thing we didn't realize at first," said Tony Burgess, "was that soils are alive. They breathe as fast as you do. You have to treat soil as a living organism. Ultimately it controls the biota."

Once you have soil, you can play Noah. Noah rounded up everything that moved for his ark, but that certainly wasn't going to work here. The designers of the Bio2 closed-system kept coming back to that most exasperating but thrilling question: what species should Bio2 include? No longer was it merely "Which organisms do we need to mirror the breath of eight humans?" The dilemma was "Which organisms do we need to mirror Gaia? Which combination of species would produce oxygen to breathe, plants to eat, plants to feed the animals to eat (if any), and species to support the food plants? How do we weave a self-supporting network out of random organisms? How do we launch a coevolutionary circuit?"

Take almost any creature as an example. Most fruit requires insects to pollinate it. So if you wanted blueberries in Bio2, you needed honeybees. But in order to have honeybees around when the blueberries are ready for pollination, you needed to provide the honeybees with flowers for the rest of the season. But in order to supply sufficient seasonal flowers to keep honeybees alive, there would be no room for other kinds of plants. So, perhaps another type of pollinating bee would work? You could use straw bees which can be supported with meager amounts of flowers, but they don't pollinate blueberry blossoms or several other fruits you wanted. How about moths? And so on down the catalog of living creatures. Termites are necessary to decompose old woody vegetation, but they were fond of eating the sealant around the windows. What's a benign termite substitute that would get along with the rest of the crowd?

"It's a sticky problem," said Peter Warshall, a consulting ecologist for the project. "It's a pretty impossible job to pick 100 living things, even from the same place, and put them together to make a 'wilderness'. And here we're taking them from all over the world to mix together since we have so many biomes."

To cobble together a synthetic biome, the half-dozen Bio2 ecologists sat down at a table together and played this ultimate jigsaw puzzle. Each scientist had expertise in either mammals, insects, birds, or plants. But while they knew something about sedges and pond frogs, very little of their knowledge was systematically accessible. Warshall sighed, "It would have been nice if somewhere there was a database of all known species listing their food and energy requirements, their habitat, their waste products, their companion species, their breeding needs, etc., but there isn't anything remotely like that. We know very little about even common species. In fact, what this project shows is how little we know about any species."

The burning question for the summer the biomes were designed was "Well, how many moths does a bat really eat?" In the end, selecting the thousand or so higher species came down to informed guesses and biodiplomacy. Each ecologist wrote up a long lists of possible candidates, including favorite species they thought would be the most versatile and flexible. Their heads were full of conflicting factors -- pluses and minuses, likes to be near this guy but can't stand this one. The ecologists projected the competitiveness of rival organisms. They bickered for water or sunlight rights. It was if they were ambassadors protecting the territory of their species from encroachments.

"I needed as much fruit as possible dropped from trees for my turtles to eat," said Bio2 desert ecologist Tony Burgess, "but the turtles would leave none for the fruit flies to breed on, which Warshall's hummingbirds needed to eat. Should we have more trees for leftover fruit, or use the space for bat habitat?"

So negotiations take place: If I can have this flower for the birds, you can keep the bats. Occasionally the polite diplomacy reverted to open subversion. The marsh-man wanted his pick of sawgrass, but Warshall didn't like his choice because he felt the species was too aggressive and would invade the dry land biome he was overseeing. In the end Warshall capitulated to the marsh-man's choice, but added, half in jest, "Oh, it doesn't make any difference because I'm just gonna plant taller elephant grass to shade out your stuff, anyway." The marsh-man retaliated by saying he was planting pine trees, taller than either. Warshall promised with a hearty laugh to plant a defense border of guava trees, which don't grow any taller, but grow much faster, staking out the niche early.

Everything was connected to everything. It made planning a nightmare. One approach the ecologists favored was building redundancy of pathways into the food webs. With multiple foodchains in every web, if the sand flies died off, then something else became second choice food for the lizards. Rather than fight the dense tangle of interrelationships, they exploited them. The key was to find organisms with as many alternative roles as possible, so that if one didn't work out, it had another way or two to complete somebody's loop.

"Designing a biome was an opportunity to think like God," recalled Warshall. You, as a god, could create something by nothing. You could create something -- some wonderful synthetic vibrant ecosystem -- but you had no control over precisely what something emerged. All you could do was gather all the parts and let them self-assemble into something that worked. Walter Adey said, "Ecosystems in the wild are made up of patches. You inject as many species as you can into the system and let it decide what patch of species it wants to be in." Surrendering control became one of the "Principles of Synthetic Ecology." Adey continued, "We have to accept the fact that the amount of information contained in an ecosystem far exceeds the amount contained in our heads. We are going to fail if we only try things we can control and understand." The exact details of an emerging Bio2 ecology, he warned, were beyond predicting.

But details counted. Eight human lives rested on the details fusing into a whole. Tony Burgess, one of the Bio2 gods, ordered dune sand to be trucked in for the desert biome because construction sand, the only kind on hand at the Bio2 site, was too sharp for the land turtles; it cut their feet. "You've got to take care of your turtles, so they can take care of you," he said in a priestly way.

The number of free-roaming animals taking care of the system was pretty thin for the first two years in Bio2 because there wasn't enough wild food to support very many of them. Warshall almost didn't put any monkeylike galagos from Africa in because he wasn't sure the young acacia trees could produce enough gum to satisfy them. In the end he released four galagos and stored a couple hundred pounds of emergency monkeychow in the basement of the ark. Other wild animal occupants of Bio2 included leopard tortoises, blue-tongued skinks ("because they are generalists" -- not picky what they eat), various lizards, small finches, and pygmy green hummingbirds, partially for pollination. "Most of the species will be pygmy," Warshall told a Discover reporter before closure, "because we really don't have that much space. In fact, ideally we'd have pygmy people, too."

The animals didn't go in two by two. "You want to have a higher ratio of females to males for reproduction insurance," Warshall told me. "Ideally we like to have at minimum five females per three males. I know director John Allen says that eight humans -- four female, four male -- is the minimum-size group needed for human colony start-up and reproduction, but from an ecologically correct rather than politically correct point of view, the Bio2 crew should be five females and three males."

For the first time biologists were being forced by the riddle of creating a biosphere to think like engineers: "Here is what we need, what materials will do that job?" At the same time, the engineers on the project were being forced to think like biologists: "That's not dirt, that's a living organism!"

A stubborn problem for the designers of Bio2 was making rain for the cloud forest. Rain is hard. The original plans optimistically called for cooling coils at the peak of the 85-foot glass roof over the jungle section. The coils would condense the jungle's moisture into gentle drops descending from the celestial heights -- real artificial rain. Early tests proved the drops to be scarce, too large and destructive when they landed, and not at all the constant gentle mist the plants wanted. Second plan was for the rain to be pumped up into sprinklers bolted to the frame structure high overhead, but that proved to be a maintenance nightmare since over a two-year period the fine-holed mist heads were sure to need unclogging or replacements. The design they ended up with was "rain" squirted from misting nozzles fitted on the ends of pipes stationed here and there on the slopes.

One unexpected consequence of living in a small materially closed system is that rather than water becoming precious, it's in virtual abundance. In about one week 100 percent of the water is recycled, cleansed by microbiological activity in wetland treatment areas. When you use more water, it just goes around the loop a little faster.

Any field of life is a cloth woven with countless separate loops. The loops of life -- the routes which materials, functions, and energy follow -- double up, cross over and interweave as knots until it is impossible to tell one thread from another. Only the larger pattern knitted by the loops emerges. Each circle strengthens the others, until the whole is hard to unravel.

That is not to say there will be no extinctions in a tightly wrapped ecosystem. A certain extinction rate is essential for evolution. Walter Adey had about 1 percent attrition rate in his previous partially closed coral reef. He expected about a 30 to 40 percent drop-off in species within the whole of Bio2 by the end of its first two-year run. (The biologists from Yale University who are currently counting the species after reopening have not finished their studies of species attrition as of my writing).

But Adey believes that he already has learned how to grow diversity: "What we are doing is cramming more species in than we expect to survive. So the numbers drop. Particularly the insects and lower organisms. Then, at the beginning of the next run we overstock it again, injecting slightly different species -- our second guesses. What will probably happen is that there will still be a large loss again, maybe one quarter, but we reinject again next closure. Each time the numbers of species will stabilize at a higher level than the first. The more complex the system, the more species it can hold. We keep doing that, building up the diversity. If you loaded up Biosphere 2 with all the species it ends up with, it would collapse at the start." The huge glass bottle is a diversity pump that grows complexity.

The Bio2 ecologists were left with the large question of how best to jump start the initial variety, upon which further diverse growth would be leveraged. This was very much related to the practical problem of how to load all the animals onto the ark. How do you get 3,000 interdependent creatures into a cage, alive? Adey proposed moving an entire natural biome into Bio2's relatively miniature space by compressing it in the manner of a condensed book: selecting choice highlights here and there, and fusing these bits into a sampler.

He selected a fine 30-mile stretch of a Florida Everglade mangrove swamp and had it surveyed into a grid. Every half mile or so along the salt gradient, a small cube (4-feet deep by 4-feet square) of mangrove roots was dug out. The block of leafy branches, roots, mud, and piggybacking barnacles was boxed and hauled ashore. The segments of the marsh, each one tuned to a slightly different salt content with slightly different microorganisms, were trucked to Arizona (after long negotiations with very confused agricultural custom agents who thought "mangroves" were "mangoes").

While the chunks of everglades were waiting to be placed in the Bio2 marsh, the Bio2 workers hooked the watertight boxes up into a network of pipes so that they became one distributed saltwater tide. Later the 30 or so cubes were reassembled into Bio2. Unboxed, the reconstituted marsh takes up only a micro 90-by-30 feet. But within this volleyball court-size everglade, each section harbors a gradually increasing salt-loving mixture of microorganisms. Thus, the flow of life from freshwater to brine is compressed into talking distance. The problem with the analog method is that scale is an important dimension of an ecosystem. As Warshall juggled the parts to manufacture a miniature savanna, he shook his head: "At best we are putting about one-tenth the variety of a system into Bio2. For the insect population it's more like one-hundredth. In a West African savanna there are 35 species of worms. At most we'll have three kinds. So the dilemma is: are we making a savanna or a lawn? It's surely better than a lawn...but how much better I don't know."

Constructing a wetlands or savanna by reassembling portions of a natural one is only one method of biome building -- which the ecologists call the "analog" way. It seemed to work fine. But as Tony Burgess pointed out, "You can go two ways with this. You can mimic an analog of a particular environment you find in nature, or you can invent a synthetic one based on many of them." Bio2 wound up being a synthetic ecosystem, with many analog parts, such as Adey's marshland.

"Bio2 is a synthetic ecosystem, but so is California by now," said Burgess. Warshall agrees: "What you see in California is a symbol of the future. A heavily synthetic ecology. It has hundreds of exotic species. A lot of Australia is going this way too. And the redwood/eucalyptus forest is also a new synthetic ecology." As are many other ecosystems in this world of jet travel, when species are jet-setted far from their home territories and introduced accidentally or deliberately in lands they would otherwise never reach. Warshall said, "Walter Adey first used the term synthetic ecology. Then I realized that there were already huge amounts of synthetic ecology in Biosphere One. And that I wasn't inventing a synthetic ecology in Bio2, I was merely duplicating what already existed." Edward Mills of Cornell University has identified 136 species of fish from Europe, the Pacific and elsewhere now thriving in the Great Lakes. "Probably most of the biomass in the Great Lakes is exotic," Mills claims. "It's a very artificial system now."

We might as well develop a science of synthetic ecosystem creation since we've been creating them anyway in a haphazard fashion. Many archeo- ecologists believe that the entire spectrum of early humanoid activities -- hunting, grazing, setting prairie fires, and selective herb gathering -- forged an "artificial" ecology upon the wilderness, that is, an ecology greatly shaped by human arts. In fact, all that we think of as natural virgin wilderness is abundant with artificiality and the mark of human activity. "Many rain forests are actually pretty heavily managed by indigenous Indians," Burgess says. "But the first thing we do when we come in is wipe out the indigenous people, so the management expertise disappears. We assumed that this growth of old trees is pristine rain forest because the only way we know how to manage a forest is to clear the trees, and these weren't clear-cut." Burgess believes that the mark of human activity runs so deep that it cannot be undone easily. "Once you alter the ecosystem, and you get the right seeds dispersed in the ground and the essential climate window, then the transformation starts and it's irreversible. This does not require the presence of man to keep the synthetic ecosystem going. It runs undisturbed. All the people in California could die and its current synthetic flora and fauna will remain. It's a new meta-stable state that remains as long as the self-reinforcing conditions stay the same."

"California, Chile and Australia are converging very rapidly to become the same synthetic ecology," Burgess claims. "They were established by the same people, and shaped by the same goal: removal of the ancient herbivores to be replaced by the production of bovines: cow meat." As a synthetic ecology, Bio2 is a foreshadowing of ecologies to come. It is clear that we are not retreating from our influence on nature. Perhaps the bottle of Bio2 can teach us how to artificially evolve useful, less disruptive synthetic ecosystems.

As the ecologists began to assemble the first deliberately synthesized ecology they made an attempt to devise guidelines they felt would be important in creating any living closed biosystem. The makers of Bio2 called these the Principles of Biospherics. When creating a biosphere remember that:

- Microorganisms do most of the work.
- Soil is an organism. It is alive. It breathes.
- Make redundant food webs.
- Increase diversity gradually.
- If you can't provide a physical function, you need to simulate it.
- The atmosphere communicates the state of the whole system.
- Listen to the system; see where it wants to go.

Rain forests, tundras, and everglades are not themselves natural closed systems; they are open to each other. There is only one natural closed system we know of: the Earth as a whole, or Gaia. In the end our interest in fashioning new closed systems rests on concocting second examples of living closed systems so that we may generalize their behavior and understand the system of Earth, our home.

Closed systems are a particularly intense variety of coevolution. Pouring shrimp into a flask and pinching off the throat of the flask is like putting a chameleon in a mirrored bottle and pinching closed the entrance. The chameleon responds to the image it has generated, just as the shrimp responds to the atmosphere it has generated. The closed bottle -- once the internal loops weave together and tighten -- accelerates change and evolution within. This isolation, like the isolation in terrestrial evolution, breeds variety and marked differences.

But eventually all closed systems are opened or at least leak. We can be certain that whatever artificial closed systems we fabricate will sooner or later be opened. Bio2 will be closed and unsealed every year or so. And in the heavens, on the scale of galactic time, the closed systems of planets will be penetrated and shared in a type of cross-panspermia -- a few exchanges of species here and there. The ecology of the cosmos is this type: a universe of isolated systems (planets), furiously inventing things in that mad way of a chameleon locked in a mirrored bottle. Every now and then marvels from one closed system will arrive with a shock into another.

On Gaia, the briefly closed miniature Gaias we construct are mostly instructional aides. They are models made to answer primarily one question: what influence do we, and can we, have over the unified system of life on Earth? Are there levels we can reach, or is Gaia entirely out of our control?

Artificial evolution

The first time Tom Ray released his tiny hand-made creature into his computer, it reproduced rapidly until hundreds of copies occupied the available memory space. Ray's creature was an experimental computer virus of sorts; it wasn't dangerous because the bugs couldn't replicate outside his computer. The idea was to see what would happen if they had to compete against each other in a confined world.

Ray cleverly devised his universe so that out of the thousands of clones from the first ancestral virus, about ten percent replicated with small variations. The initial creature was an "80" -- so named because it had 80 bytes of code. A number of 80s "flipped a bit" at random and became creatures 79 or 81 bytes long. Some of these new mutant viruses soon took over Ray's virtual world. In turn, they mutated into further varieties. Creature 80 was nearly overwhelmed to the point of extinction by the mushrooming ranks of new "organisms." But the 80s never completely died, and long after the new arrivals 79, 51, and 45 emerged and peaked in population, the 80s rebounded.

After a few hours of operation, Tom Ray's electric-powered evolution machine had evolved a soup of nearly a hundred types of computer viruses, all battling it out for survival in his isolated world. On his very first try, after months of writing code, Ray had brewed artificial evolution.

When he was a shy, soft-spoken Harvard undergraduate, Ray had collected ant colonies in Costa Rica for the legendary ant-man, E. O. Wilson. Wilson needed live leafcutting ant colonies for his Cambridge labs. Ray hired on in the lush tropics of Central America to locate and capture healthy colonies in the field, and then ship them to Harvard. He found that he was particularly good at the task. The trick was to dig into the jungle soil with the deftness of a surgeon in order to remove the guts of a colony. What was needed was the intact inner chamber of the queen's nest, along with the queen herself, her nurse ants, and a mini-ant-garden stocked with enough food to support the chamber for shipping. A young newborn colony was perfect. The heart of such a colony might fit into a tea cup. That was the other essential trick: to locate a really small nest hidden under the natural camouflaged debris of the forest floor. From a minuscule core that could be warmed in one's hands, the colony could grow in a few years to fill a large room.

While collecting ants in the rain forest, Ray discovered a obscure species of butterfly that would tag along the advancing lines of army ants. The army ants' ruthless eating habits -- devouring any animal life in their path -- would flush a cloud of flying insects eager to get out of the way. A kind of bird evolved to follow the pillaging army, happily picking off the agitated fleeing insects in the air. The butterfly, in turn, followed the birds who followed the army ants. The butterflies tagged along to feast on the droppings of the ant-birds -- a much needed source of nitrogen for egg laying. The whole motley crew of ants, ant-birds and ant-bird-butterflies, and who knew what else, would roam across the jungle like a band of gypsies in cahoots.

Ray was overwhelmed by such wondrous complexity. Here was an entirely nomadic community! Most attempts to understand ecological relations seemed laughable in light of these weird creations. How in the universe did these three groups of species (one ant, three butterflies, and about a dozen birds) ever wind up in this peculiar codependency? And why?

By the time he had finished his Ph.D., Ray felt that the science of ecology was moribund because it could not offer a satisfying answer to such big questions. Ecology lacked good theories to generalize the wealth of observations piling up from every patch of wilderness. It was stymied by extensive local knowledge: without an overarching theory, ecology was merely a library of fascinating just-so stories. The life cycles of barnacle communities, or the seasonal pattern of buttercup fields, or behavior of bobcat clans were all known, but what principles, if any, guided all three? Ecology needed a science of complexity that addressed the riddles of form, history, development -- all the really interesting questions -- yet was supported by field data.

Along with many other biologists, Ray felt that the best hope for ecology was to shift its focus from ecological time (the thousand-year lifetime of a forest) to evolutionary time (the million-year lifetime of a tree species). Evolution at least had a theory. Yet, the study of evolution too was caught up with the same fixation on specifics. "I was frustrated," Ray told me, "because I didn't want to study the products of evolution -- vines and ants and butterflies. I wanted to study evolution itself."

Tom Ray dreamed of making an electric-powered evolution machine. With a black box that contained evolution he could demonstrate the historical principles of ecology, how a rain forest descends from earlier woods, and how in fact ecologies emerge from the same primordial forces that spawn species. If he could develop an evolution engine, he'd have a test-bed with which to do real ecological experiments. He could take a community and run it over and over again in different combinations, making ponds without algae, woods without termites, grasslands without gophers, or just to cover the bases, jungles with gophers and grasslands with algae. He could start with viruses and see where it all would lead him.

Ray was a bird watcher, insect collector, plantsman -- the farthest thing from a computer nerd -- yet he was sure such a machine could be built. He remembered a moment ten years earlier when he was learning the Japanese game of Go from an MIT hacker who used biological metaphors to explain the rules. As Ray tells it, "He said to me, 'Do you know that it is possible to write a computer program that can self-replicate?' And right at that moment I imagined all the things I'm doing now. I asked him how to do it, and he said, 'Oh, it's trivial,' but I didn't remember what he said, or whether in fact he actually knew. When I remembered that conversation I stopped reading novels and started reading computer manuals."

Ray's solution to the problem of making an electronic evolution machine was to start with simple replicators and give them a cozy habitat and plenty of energy and places to fill. The closest real things to these creatures were bits of self-replicating RNA. But the challenge seemed doable. He would cook up a soup of computer viruses.

About this time in 1989, the news magazines were chock-full of cover stories pronouncing computer viruses worse than the plague and as evil as technology could get. Yet Ray saw in the simple codes of computer viruses the beginnings of a new science: experimental evolution and ecology.

To protect the outside world (and to keep his own computer from crashing), Ray devised a virtual computer to contain his experiments. A virtual computer is a bit of clever software that emulates a pretend computer deep within the operating subconscious of the real computer. By containing his tiny bits of replicating code inside this shadow computer, Ray sealed them from the outside world and gave himself room to mess with vital functions, such as computer memory, without jeopardizing the integrity of his host computer. "After a year of reading computer manuals, I sat down and wrote code. In two months the thing was running. And in the first two minutes of running without a crash, I had evolving creatures."

Ray seeded his world (which he called "Tierra") with a single creature he programmed by hand -- the 80-byte creature -- inserted into a block of RAM in his virtual computer. The 80 creature reproduced by finding an empty RAM block 80 bytes big and then filling it with a copy of itself. Within minutes the RAM was saturated with copies of 80.

But Ray had added two key features that modified this otherwise Xerox-like copying machine into an evolution machine: his program occasionally scrambled the digital bits during copying, and he assigned his creatures a priority tag for an executioner. In short he introduced variation and death.

Computer scientists had told him that if he randomly varied bits of a computer code (which is all his creatures really are), the resulting programs would break and then crash the computer. They felt that the probability of getting a working program by randomly introducing bugs into code was so low as to make his scheme a waste of time. This sentiment seemed in line with what Ray knew about the fragile perfection needed to keep computers going; bugs killed progress. But because his creature programs would run in his shadow computer, whenever a mutation would birth a creature that was seriously broken, his executioner program -- he named it "the Reaper" -- would kill it while the rest of his Tierra world kept running. In essence, Tierra spotted the buggy programs that couldn't reproduce and yanked them out of the virtual computer.

Yet, the Reaper would pass over the very rare mutants that worked, that is, those that happened to form a bona fide alternative program. These legitimate variations could multiply and breed other variants. If you ran Tierra for a billion computer cycles or so, as Ray did, a startling number of randomly generated creatures formed during those billion chances. And just to keep the pot boiling, Ray also assigned creatures an age stamp so that older creatures would die. "The Reaper kills either the oldest creature or the most screwed-up creature," Ray says with a smile.

On Ray's first run of Tierra, random variation, death, and natural selection worked. Within minutes Ray witnessed an ecology of newly created creatures emerge to compete for computer cycles. The competition rewarded creatures of smaller size since they needed less cycles, and in Darwinian ruthlessness, terminated the greedy consumers, the infirm, and the old. Creature 79 (one byte smaller than 80) was lucky. It worked productively and soon outpaced the 80s.

Ray also found something very strange: a viable creature with only 45 very efficient bytes which overran all other creatures. "I was amazed how fast this system would optimize," Ray recalls. "I could graph its pace as the system would generate organisms surviving on shorter and shorter genomes."

On close examination of 45's code, Ray was amazed to discover that it was a parasite. It contained only a part of the code it needed to survive. In order to reproduce, it "borrowed" the reproductive section from the code of an 80 and copied itself. As long as there were enough 80 hosts around, the 45s thrived. But if there were too many 45s in the limited world, there wouldn't be enough 80s to supply copy resources. As the 80s waned, so did the 45s. The pair danced the classic coevolutionary tango, back and forth endlessly, just like populations of foxes and rabbits in the north woods.

"It seems to be a universal property of life that all successful systems attract parasites," Ray reminds me. In nature parasites are so common that hosts soon coevolve immunity to them. Then eventually the parasites coevolve strategies to circumvent that immunity. And eventually the hosts coevolve defenses to repel them again. In reality, these actions are not alternating steps but two constant forces pressing against one another.

Ray learned to run ecological experiments in Tierra using parasites. He loaded his "soup" with 79s which he suspected were immune to the 45 parasite. They were. But as the 79s prospered, a second parasite evolved that could prey on them. This one was 51 bytes long. When Ray sequenced its genes he found that a single genetic event had transformed a 45 into a 51. "Seven instructions of unknown origin," Ray says, "had replaced one instruction somewhere near the middle of the 45," transforming a disabled parasite into a newly potent one. And so it went. A new creature evolved that was immune to 51s, and so on.

Poking around in the soups of long runs, Ray discovered parasites that preyed on other parasites -- hyperparasites: "Hyperparasites are like neighbors who steal power from your lines to the power plant. You sit in the dark while they use your power and you pay the bill." In Tierra, organisms such as the 45s discovered that they didn't need to carry a lot of code around to replicate themselves because their environment was full of code -- of other organisms. Quips Ray, "It's just like us using other animals' amino acids [when we eat them]." On further inspection Ray found hyper-hyperparasites thriving, parasites raised to the third. He found "social cheaters" -- creatures that exploit the code of two cooperating hyperparasites (the "cooperating" hyperparasites were stealing from each other!). Social cheaters require a fairly well developed ecology. They can't be seen yet, but there are probably hyper-hyper-hyperparasites and no end to elaborate freeloading games possible in his world.

And Ray found creatures that surpassed the programming skills of human software engineers.

"I started with a creature 80 bytes large," Ray remembers, "because that's the best I could come up with. I figured that maybe evolution could get it down to 75 bytes or so. I let the program run overnight and the next morning there was a creature -- not a parasite, but a fully self-replicating creature -- that was only 22 bytes! I was completely baffled how a creature could manage to self-replicate in only 22 instructions without stealing instructions from others, as parasites do. To share this novelty, I distributed its basic algorithm onto the Net. A computer science student at MIT saw my explanation, but somehow didn't get the code of the 22 creature. He tried to recreate it by hand, but the best he could do was get it to 31 instructions. He was quite distressed when he found out I came up with 22 instructions in my sleep!"

What humans can't engineer, evolution can. Ray puts it nicely as he shows off a monitor with traces of the 22s propagating in his soup: "It seems utterly preposterous to think that you could randomly alter a computer program and get something better than what you carefully crafted by hand, but here's living proof." It suddenly dawns on the observer that there is no end to the creativity that these mindless hackers can come up with.

Because creatures consume computer cycles, there is an advantage to smaller (shorter sets of instructions) creatures. Ray reprogrammed Tierra's code so the system assigned computer resources to creatures in proportion to their size; large ones getting more cycles. In this mode, Ray's creatures inhabited a size-neutral world, which seemed more suited for long runs since it wasn't biased to either the small or large. Once Ray ran a size-neutral world for 15 billion cycles of his computer. Somewhere around 11 billion cycles, a diabolically clever 36 creature evolved. It calculated its true size, then behind its back so to speak, shifted all the bits in the measurement to the left one bit, which in binary code is equal to doubling the number. So by lying about its size, creature 36 sneakily garnered the resources of a 72 creature, which meant that it got twice the usual CPU time. Naturally this mutation swept through the system.

Perhaps the most astounding thing about Tom Ray's electrically powered evolution machine is that it created sex. Nobody told it about sex, but it found it nonetheless. In an experiment to see what would happen if he turned the mutation function off, Ray let the soup run without deliberate error. He was flabbergasted to discover that even without programmed mutation, evolution pushed forward.

In real natural life, sex is a much more important source of variation than mutations. Sex, at the conceptual level, is genetic recombination -- a few genes from Dad and a few genes from Mom combined into a new genome for Junior. Sometimes in Tierra a parasite would be in the middle of asexual reproduction, "borrowing" the copy function of some other creature's code, when the Reaper would happen to kill the host midway in the process. When this happens the parasite uses some copy code of the new creature born in the old creature's space, and part of the "dead" creature's interrupted reproduction function. The resultant junior was a wild, new recombination created without deliberate mutation. (Ray also says this weird reproduction "amounts to sex with the dead!") Interrupted sex had happened all the time in his soup, but only when Ray turned off his "flip-a-bit" mutator did he notice its results. It turned out that inadvertent recombination alone was enough to fuel evolution. There was sufficient irregularity in the moment of death, and where creatures lived in RAM, that this complexity furnished the variety that evolution required. In one sense, the system evolved variation.

To scientists, the most exhilarating news to come out of Ray's artificial evolution machine is that his small worlds display what seems to be punctuated equilibrium. For relatively long periods of time, the ratio of populations remain in a steady tango of give and take with only the occasional extinction or birth of a new species. Then, in a relative blink, this equilibrium is punctuated by a rapid burst of roiling change with many newcomers and eclipsing of the old. For a short period change is rampant. Then things sort out and stasis and equilibrium reigns again. The current interpretation of fossil evidence on Earth is that this pattern predominates in nature. Stasis is the norm; change occurs in bouts. The same punctuated equilibrium pattern has been seen in other evolutionary computer models as well, such as Kristian Lindgren's coevolutionary Prisoner's Dilemma world. If artificial evolution mirrors organic evolution, one has to wonder what would happen if Ray let his world run forever? Would his viral creatures invent multicellularity?

Unfortunately, Ray has never turned his world on marathon mode just to see what would happen over months or years. He's still fiddling with the program, gearing it up to collect the immense store of data (50 megabytes per day) such a marathon run would generate. He admits that "sometimes we're like a bunch of boys with a car. We've always got the hood up and pieces of the engine out on the garage floor, but we hardly ever drive the car because we're too interested in souping it up."

In fact, Ray has his sights fixed on a new piece of hardware, a technology that ought to be. Ray figures that he could take his virtual computer and the fundamental language he wrote for it and "burn" it into a computer chip -- a slice of silicon that did evolution. This off-the-shelf Darwin Chip would then be a module you could plug into any computer, and it would breed stuff for you, fast. You could evolve lines of computer code, or subroutines, or maybe even entire software programs. "I find it rather peculiar," Ray confides, "that as a tropical plant ecologist I'm now designing computers."

The prospects that a Darwin Chip might serve up are delicious. Imagine you have one in your PC where you use Microsoft Word as a word processor. With resident Darwinism loaded into your operating system, Word would evolve as you worked. It would use your computer's idle CPU cycles to improve, and learn, in a slow evolutionary way, to fit itself to your working habits. Only those alterations that improved the speed or the accuracy would survive. However Ray feels strongly that messy evolution should happen away from the job. "You want to divorce evolution from the end user," he says. He imagines "digital husbandry" happening offline in back rooms, so to speak, so that the common failures necessary for evolution are never seen by its customer. Before an evolving application is turned over to an end user, it is "neutered" so that it can't evolve while in use.

Retail evolution is not so farfetched. Today you can buy a spreadsheet module that does something similar in software. It's called, naturally enough, "Evolver." Evolver is a template for spreadsheets on the Macintosh -- very complicated spreadsheets spilling over with hundreds of variables and "what-if" functions. Engineers and database specialists use it.

Let's say you have the medical records of thirty thousand patients. You'd probably like to know what a typical patient looks like. The larger the database, the harder it is to see what you have in there. Most software can do averaging, but that does not extract a "typical" patient. What you would like to know is what set of measurements -- out of the thousands of categories collected by the records -- have similar values for the maximum number of people? It's a problem of optimizing huge numbers of interacting variables. The task is familiar to any living species: how does it maximize the results of thousands of variables? Raccoons have to ensure their own survival, but there are a thousand variables (foot size, night vision, heart rate, skin color, etc.) that can be changed over time, and altering one parameter will alter another. The only way to tread through this vast space of possible answers, and retain some hope of reaching a peak, is by evolution.

The Evolver software optimizes the broadest possible profile for the largest number of patients by trying a description of a typical patient, then testing how many fit that description, then tweaking the profile in a multitude of directions to see if more patients fit it, and then varying, selecting, and varying again, until a maximum number of patients fit the profile. It's a job particularly suited for evolution.

"Hill climbing," computer scientists call the process. Evolutionary programs attempt to scale the peak in the libraries of form where the optimal solution resides. By relentlessly pushing the program toward better solutions, the programs climb up until they can't climb any higher. At that point, they are on a peak -- a maximum -- of some sort. The question always is: is their summit the tallest peak around, or is the program stuck on a local peak adjacent to a much taller peak across the valley, with no way to retreat?

Finding a solution -- a peak -- is not difficult. What evolution in nature and evolutionary programs in computers excel at is hill climbing to global summits -- the highest peaks around -- when the terrain is rugged with many false summits.

John Holland is a gnomic figure of indeterminate age who once worked on the world's earliest computers, and who now teaches at the University of Michigan. He was the first to invent a mathematical method of describing evolution's optimizing ability in a form that could be easily programmed on a computer. Because of the way his math mimicked the effects of genetic information, Holland called them genetic algorithms, or GAs for short.

Holland, unlike Tom Ray, started with sex. Holland's genetic algorithms took two strings of DNA-like computer code that did a job fairly well and recombined the two at random in a sexual swap to see if the new offspring code might do a little better. In designing his system, Holland had to overcome the same looming obstacle that Ray faced: any random generation of a computer program would most likely produce not a program that was either slightly better or slightly worse, but one that was not sensible at all. Statistically, successive random mutations to a working code were bound to produce successive crashes.

Mating rather than mutating was discovered by theoretical biologists in the early 1960s to make a more robust computer evolution -- one that birthed a higher ratio of sensible entities. But sexual mating alone was too restrictive in what it could come up with. In the mid-1960s Holland devised his GAs; these relied chiefly on mating and secondarily on mutation as a background instigator. With sex and mutation combined, the system was both flexible and wide.

Like many other systems thinkers, Holland sees the tasks of nature and the job of computers as similar. "Living organisms are consummate problem solvers," Holland wrote in a summary of his work. "They exhibit a versatility that puts the best computer programs to shame. This observation is especially galling for computer scientists, who may spend months or years of intellectual effort on an algorithm, whereas organisms come by their abilities through the apparently undirected mechanism of evolution and natural selection."

The evolutionary approach, Holland wrote, "eliminates one of the greatest hurdles in software design: specifying in advance all the features of a problem." Anywhere you have many conflicting, interlinked variables and a broadly defined goal where the solutions may be myriad, evolution is the answer.

Just as evolution deals in populations of individuals, genetic algorithms mimic nature by evolving huge churning populations of code, all processing and mutating at once. GAs are swarms of slightly different strategies trying to simultaneously hill-climb over a rugged landscape. Because a multitude of code strings "climb" in parallel, the population visits many regions of the landscape concurrently. This ensures it won't miss the Big Peak.

Implicit parallelism is the magic by which evolutionary processes guarantee you climb not just any peak but the tallest peak. How do you locate the global optima? By testing bits of the entire landscape at once. How do you optimally balance a thousand counteracting variables in a complex problem? By sampling a thousand combinations at once. How do you develop an organism that can survive harsh conditions? By running a thousand slightly varied individuals at once.

In Holland's scheme, the highest performing bits of code anywhere on the landscape mate with each other. Since high performance increases the assigned rate of mating in that area, this focuses the attention of the genetic algorithm system on the most promising areas in the overall landscape. It also diverts computational cycles away from unpromising areas. Thus parallelism sweeps a large net over the problem landscape while reducing the number of code strings that need manipulating to locate the peaks.

Parallelism is one of the ways around the inherent stupidity and blindness of random mutations. It is the great irony of life that a mindless act repeated in sequence can only lead to greater depths of absurdity, while a mindless act performed in parallel by a swarm of individuals can, under the proper conditions, lead to all that we find interesting.

John Holland invented genetic algorithms while studying the mechanics of adaptation in the 1960s. His work was ignored until the late 1980s by all but a dozen wild-eyed computer grad students. A couple of other researchers, such as the engineers Lawrence Fogel and Hans Bremermann, independently played around with mechanical evolution of populations in the 1960s; they enjoyed equal indifference from the science community. Michael Conrad, a computer scientist now at Wayne State University, Michigan, also drifted from the study of adaptation to modeling evolving populations in computers in the 1970s, and met the same silence that Holland did a decade earlier. The totality of this work was obscure to computer science and completely unknown in biology.

No more than a couple of students wrote theses on GA until Holland's book Adaptation in Natural and Artificial Systems about GAs and evolution appeared in 1975. The book sold only 2500 copies until it was reissued in 1992. Between 1972 and 1982, no more than two dozen articles on GAs were published in all of science. You could not even say computational evolution had a cult following.

The lack of interest from biology was understandable (but not commendable); biologists reasoned that nature was far too complex to be meaningfully represented by computers of that time. The lack of interest from computer science is more baffling. I was often perplexed in my research for this book why such a fundamental process as computational evolution could be so wholly ignored? I now believe the disregard stems from the messy parallelism inherent in evolution and the fundamental conflict it presented to the reigning dogma of computers: the von Neumann serial program.

The first functioning electronic computer was the ENIAC, which was booted up in 1945 to solve ballistic calculations for the U.S. Army. The ENIAC was an immense jumble of 18,000 hot vacuum tubes, 70,000 resistors, and 10,000 capacitors. The instructions for the machine were communicated to it by setting 6,000 switches by hand and then turning the program on. In essence the machine calculated all its values simultaneously in a parallel fashion. It was a bear to program.

The genius von Neumann radically altered this awkward programming system for the EDVAC, the ENIAC's successor and the first general-purpose computer with a stored program. Von Neumann had been thinking about systemic logic since the age of 24 when he published his first papers (in 1927) on mathematical logic systems and game theory. Working with the EDVAC computer group, he invented a way to control the slippery calculations needed to program a machine that could solve more than one problem. Von Neumann proposed that a problem be broken into discrete logical steps, much like the steps in a long division problem, and that intermediate values in the task be stored temporarily in the computer in such a way that those values could be considered input for the next portion of the problem. By feeding back the calculation through a coevolutionary loop (or what is now called a subroutine), and storing the logic of the program in the machine so that it could interact with the answer, von Neumann was able to take any problem and turn it into a series of steps that could be comprehended by a human mind. He also invented a notation for describing this step-wise circuit: the now familiar flow chart. Von Neumann's serial architecture for computation -- where one instruction at a time was executed -- was amazingly versatile and extremely suited to human programming. He published the general outlines for the architecture in 1946, and it immediately became the standard for every commercial computer thereafter, without exception.

In 1949, John Holland worked on Project Whirlwind, a follow-up to the EDVAC. In 1950 he joined the logical design team on what was then called IBM's Defense Calculator, later to become the IBM 701, the world's first commercial computer. Computers at that point were room-size calculators consuming a lot of electricity. But in the mid-fifties Holland participated in the legendary circle of thinkers who began to map out the possibility of artificial intelligence.

While luminaries such as Herbert Simon and Alan Newall thought of learning as a noble, high-order achievement, Holland thought of it as a polished type of lowly adaptation. If we could understand adaptation, especially evolutionary adaptation, Holland believed, we might be able to understand and maybe imitate conscious learning. But although the others could appreciate the parallels between evolution and learning, evolution was the low road in a fast-moving field.

Browsing for nothing in particular in the University of Michigan math library in 1953, Holland had an epiphany. He stumbled upon a volume, The Genetical Theory of Natural Selection, written by R. A. Fisher in 1929. It was Darwin who led the consequential shift from thinking about creatures as individuals to thinking about populations of individuals, but it was Fisher who transformed this population-thinking into a quantitative science. Fisher took what appeared to be a community of flittering butterflies evolving over time and saw them as a whole system transmitting differentiated information in parallel through a population. And he worked out the equations that governed that diffusion of information. Fisher single-handedly opened a new world of human knowledge by subjugating nature's most potent force -- evolution -- with humankind's most potent tool -- mathematics. "That was the first time I realized that you could do significant mathematics on evolution," Holland recalled of the encounter. "The idea appealed to me tremendously." Holland was so enamored of treating evolution as a type of math that in a desperate attempt to get a copy of the out-of-print text (in the days before copiers) he begged the library (unsuccessfully) to sell it to him. Holland absorbed Fisher's vision and then leaped to a vision of his own: butterflies as coprocessors in a field of computer RAM.

Holland felt artificial learning at its core was a special case of adaptation. He was pretty sure he could implement adaptation on computers. Taking the insights of Fisher -- that evolution was a class of probability -- Holland began the job of trying to code evolution into a machine.

Very early in his efforts, he confronted the dilemma that evolution is a parallel processor while all available electronic computers were von Neumann serial processors.

In his eagerness to wire up a computer as a platform for evolution, Holland did the only reasonable thing: he designed a massively parallel computer to run his experiments. During parallel computing, many instructions are executed concurrently, rather than one at a time. In 1959 he presented a paper which, as its title says, describes "A Universal Computer Capable of Executing an Arbitrary Number of Sub-programs Simultaneously," a contraption that became known as a "Holland Machine." It was almost thirty years before one was built.

In the interim, Holland and the other computational evolutionists had to rely on serial computers to grow evolution. By various tricks they programmed their fast serial CPUs to simulate a slow parallelism. The simulations worked well enough to hint at the power of true parallelism.

It wasn't until the mid-1980s that Danny Hillis began building the first massively parallel computer. Just a few years earlier Hillis had been a wunderkind computer science student. His pranks and hacks at MIT were legendary, even on the campus that invented hacking. With his usual clarity, Hillis summed up for writer Steven Levy the obstacle the von Neumann bottleneck had become in computers: "The more knowledge you gave them, the slower computers got. Yet with a person, the more knowledge you give him, the faster he gets. So we were in this paradox that if you tried to make computers smart, they got stupider."

Hillis really wanted to be a biologist, but his knack for understanding complex programs drew him to the artificial intelligence labs of MIT, where he wound up trying to build a thinking computer "that would be proud of me." He attributes to John Holland the seminal design notions for a swarmy, thousand-headed computing beast. Eventually Hillis led a group that invented the first parallel processing computer, the Connection Machine. In 1988 it sold for a cool $1 million apiece, fully loaded. Now that the machines are here, Hillis has taken up computational biology in earnest.

"There are only two ways we know of to make extremely complicated things," says Hillis. "One is by engineering, and the other is evolution. And of the two, evolution will make the more complex." If we can't engineer a computer that will be proud of us, we may have to evolve it.

Hillis's first massively parallel Connection Machine had 64,000 processors working in unison. He couldn't wait to get evolution going. He inoculated his computer with a population of 64,000 very simple software programs. As in Holland's GA or in Ray's Tierra, each individual was a string of symbols that could be altered by mutation. But in Hillis's Connection Machine, each program had an entire computer processor dedicated to running it. The population, therefore, would react extremely quickly and in numbers that were simply not possible for serial computers to handle.

Each bug in his soup was initially a random sequence of instructions, but over tens of thousands of generations they became a program that sorted a long string of numbers into numerical order. Such a sort routine is an integral part of most larger computer programs; over the years many hundreds of man hours have been spent in computer science departments engineering the most efficient sort algorithms. Hillis let thousands of his sorters proliferate in his computer, mutate at random, and occasionally sexually swap genes. Then in the usual evolutionary maneuver, his system tested them and terminated the less fit so that only the shortest (the best) sorting programs would be given a chance to reproduce. Over ten thousand generations of this cycle, his system bred a software program that was nearly as short as the best sorting programs written by human programmers.

Hillis then reran the experiment but with this important difference: He allowed the sorting test itself to mutate while the evolving sorter tried to solve it. The string of symbols in the test varied to become more complicated in order to resist easy sorting. Sorters had to unscramble a moving target, while tests had to resist a moving arrow. In effect Hillis transformed the test list of numbers from a harsh passive environment into an active organism. Like foxes and hares or monarchs and milkweed, sorters and tests got swept up by a textbook case of coevolution.

A biologist at heart, Hillis viewed the mutating sorting test as a parasitic organism trying to disrupt the sorter. He saw his world as an arms race -- parasite attack, host defense, parasite counterattack, host counter -- defense, and so on. Conventional wisdom claimed such locked arms races are a silly waste of time or an unfortunate blind trap to get stuck in. But Hillis discovered that rather than retard the advance of the sorting organisms, the introduction of a parasite sped up the rate of evolution. Parasitic arms races may be ugly, but they turbocharged evolution.

Just as Tom Ray would discover, Danny Hillis also found that evolution can surpass ordinary human skills. Parasites thriving in the Connection Machine prodded sorters to devise a solution more efficient than the ones they found without parasites. After 10,000 cycles of coevolution, Hillis's creatures evolved a sorting program previously unknown to computer scientists. Most humbling, it was only a step short of the all-time shortest algorithm engineered by humans. Blind dumb evolution had designed an ingenious, and quite useful, software program.

A single processor in the Connection Machine is very stupid. It might be as smart as an ant. On its own, a single processor could not come up with an original solution to anything, no matter how many years it spent. Nor would it come up with much if 64,000 processors were strung in a row.

But 64,000 dumb, mindless, ant-brains wired up into a vast interconnected network become a field of evolving populations and, at the same time, look like a mass of neurons in a brain. Out of this network of dumbness emerge brilliant solutions to problems that tax humans. This "order-emerging-out-of-massive-connections" approach to artificial intelligence became known as "connectionism."

Connectionism rekindled earlier intuitions that evolution and learning were deeply related. The connectionists who were reaching for artificial learning latched onto the model of vast webs interconnecting dumb neurons, and then took off with it. They developed a brand of connected concurrent processing -- running in either virtual or hardwired parallel computers -- that performed simultaneous calculations en masse, similar to genetic algorithms but with more sophisticated (smarter) accounting systems. These smartened up networks were called neural networks. So far neural nets have achieved only limited success in generating partial "intelligence," although their pattern-recognition abilities are useful.

But that anything at all emerges from a field of lowly connections is startling. What kind of magic happens inside a web to give it an almost divine power to birth organization from dumb nodes interconnected, or breed software from mindless processors wired to each other? What alchemic transformation occurs when you connect everything to everything? One minute you have a mob of simple individuals, the next, after connection, you have useful, emergent order.

There was a fleeting moment when the connectionists imagined that perhaps all you needed to produce reason and consciousness was a sufficiently large field of interlinked neurons out of which rational intelligence would assemble itself. That dream vanished as soon as they tried it.

But in an odd way, the artificial evolutionists still pursue the dream of connectionism. Only they, in sync with the slow pace of evolution, would be more patient. But it is the slow, very slow, pace of evolution that bothers me. I put my concern to Tom Ray this way: "What worries me about off-the-shelf evolution chips and parallel evolutionary processing machines is that evolution takes an incredible amount of time. Where is this time going to come from? Look at the speed at which nature is working. Consider all the little molecules that have just been snapped together as we talk here. Nature is incredibly speedy and vast and humongously parallel, and here we are going to try to beat it. It seems to me there's simply not enough time to do it.

Ray replied: "Well, I worry about that too. On the other hand, I'm amazed at how fast evolution has occurred in my system with only one virtual processor churning it. Besides, time is relative. In evolution, a generation sets the time scale. For us a generation is thirty years, but for my creatures it is a fraction of a second. And, when I play god I can crank up the global mutation rate. I'm not sure, but I may be able to get more evolution on a computer."

There are other reasons for doing evolution in a computer. For instance, Ray can record the sequence of every creature's genome and keep a complete demographic and genealogic record of every creature's birth and death. It produces an avalanche of data that is impossible to compile in the real world. And though the complexity and cost of extracting the information will surge as the complexity of the artificial worlds surge, it will probably remain easier to do than in the unwired organic world. As Ray told me, "Even if my world gets as complex as the real world, I'm god. I'm omniscient. I can get information on whatever attracts my attention without disturbing it, without walking around crushing plants. That's a crucial difference."

Back in the 18th century, Benjamin Franklin had a hard time convincing his friends that the mild electrical currents produced in his lab were identical in their essence to the thundering lightning that struck in the wild. The difference in scale between his artificially produced microsparks and the sky-splitting, tree-shattering, monstrous bolts generated in the heavens was only part of the problem. Primarily, observers found it unnatural that Franklin could re-create nature, as he claimed.

Today, Tom Ray has trouble convincing his colleagues that the evolution he has synthesized in his lab is identical in essence to the evolution shaping the animals and plants in nature. The difference in time scale between the few hours his world has evolved and the billions of years wild nature has evolved is only part of the problem. Primarily, skeptics find it unnatural that Ray can re-create such an intangible and natural process as he claims.

Two hundred years after Franklin, artificially generated lightning -- tamed, measured, and piped through wires into buildings and tools -- is the primary organizing force in our society, particularly our digital society. Two hundred years from now, artificial adaptation -- tamed, measured and piped into every type of mechanical apparatus we have -- will become the central organizing force in our society.

No computer scientist has yet synthesized an artificial intelligence -- as desirable and immensely powerful and life-changing as that would be. Nor has any biochemist created an artificial life. But evolution captured, as Ray and others have done, and re-created on demand, is now seen by many technicians as the subtle spark that can create both our dreams of artificial life and artificial intelligence, unleashing their awesome potential. We can grow rather than make them.

We have built machines as complicated as is possible with unassisted engineering. The kind of projects we now have on the drawing boards -- software programs reckoned in tens of millions of lines of code, communication systems spanning the planet, factories that must adapt to rapidly shifting global buying habits and retool in days, cheap Robbie the Robots -- all demand a degree of complexity that only evolution can coordinate.

Because it is slow, invisible, and diffuse, evolution has the air of a hardly believable ghost in this fast-paced, in-your-face world of humanmade machines. But I prefer to think of evolution as a natural technology that is easily moved into computer code. It is this supercompatibility between evolution and computers that will propel artificial evolution into our digital lives.

Artificial evolution is not merely confined to silicon, however. Evolution will be imported wherever engineering balks. Synthetic evolution technology is already employed in the frontier formerly called bioengineering.

Here's a real-world problem. You need a drug to combat a disease whose mechanism has just been isolated. Think of the mechanism as a lock. All you need is the right key molecule -- a drug -- that triggers the active binding sites of the lock.

Organic molecules are immensely complex. They consist of thousands of atoms that can be arranged in billions of ways. Simply knowing the chemical ingredients of a protein does not tell us much about its structure. Extremely long chains of amino acids are folded up into a compact bundle so that the hot spots -- the active sites of the protein -- are held on the outside at just the right position. Folding a protein is similar to the task of pushing a mile-long stretch of string marked in blue at six points, and trying to fold the string up into a bundle so that the six points of blue all land on different outside faces of the bundle. There are uncountable ways you could proceed, of which only a very few would work. And usually you wouldn't know which sequence was even close until you had completed most of it. There is not enough time in the universe to try all of the variations.

Drug makers have had two traditional manners for dealing with this complexity. In the past, pharmacists relied on hit or miss. They tried all existing chemicals found in nature to see if any might work on a given lock. Often, one or two natural compounds activated a couple of sites -- a sort of partial key. But now in the era of engineering, biochemists try to decipher the pathways between gene code and protein folding to see if they can engineer the sequence of steps needed to create a molecular shape. Although there has been some limited success, protein folding and genetic pathways are still far too complex to control. Thus this logical approach, called "rational drug design," has bumped the ceiling of how much complexity we can engineer.

Beginning in the late 1980s, though, bioengineering labs around the world began perfecting a new procedure that employs the only other tool we have for creating complex entities: evolution.

In brief, the evolutionary system generates billions of random molecules which are tested against the lock. Out of the billion humdrum candidates, one molecule contains a single site that matches one of, say, six sites on the lock. That partial "warm" key sticks to the lock and is retained. The rest are washed down the drain. Then, a billion new variations of that surviving warm key are made (retaining the trait that works) and tested against the lock. Perhaps another warm key is found that now has two sites correct. That key is kept as a survivor while the rest die. A billion variations are made of it, and the most fit of that generation will survive to the next. In less than ten generations of repeating the wash/mutate/bind sequence, this molecular breeding program will find a drug -- perhaps a lifesaving drug -- that keys all the sites of the lock.

Almost any kind of molecule might be evolved. An evolutionary biotechnician could evolve an improved version of insulin, say, by injecting insulin into a rabbit and harvesting the antibodies that the rabbit's immune system produced in reaction to this "toxin." (Antibodies are the complementary shape to a toxin.) The biotechnician then puts the extracted insulin antibodies into an evolutionary system where the antibodies serve as a lock against which new keys are tested. After several generations of evolution, he would have a complementary shape to the antibody, or in effect, an alternative working shape to the insulin shape. In short, he'd have another version of insulin. Such an alternative insulin would be extremely valuable. Alternative versions of natural drugs can offer many advantages: they might be smaller; more easily delivered in the body; produce fewer side effects; be easier to manufacture; or be more specific in their targets.

Of course, the bioevolutionists could also harvest an antibody against, say, a hepatitis virus and then evolve an imitation hepatitis virus to match the antibody. Instead of a perfect match, the biochemist would select for a surrogate molecule that lacked certain activation sites that cause the disease's fatal symptoms. We call this imperfect, impotent surrogate a vaccine. So vaccines could also be evolved rather than engineered.

All the usual reasons for creating drugs lend themselves to the evolutionary method. The resulting molecule is indistinguishable from rationally designed drugs. The only difference is that while an evolved drug works, we have no idea of how or why it does so. All we know is that we gave it a thorough test and it passed. Cloaked from our understanding, these invented drugs are "irrationally designed."

Evolving drugs allows a researcher to be stupid, while evolution slowly accumulates the smartness. Andrew Ellington, an evolutionary biochemist at Indiana University, told Science that in evolving systems "you let the molecule tell you about itself, because it knows more about itself than you do."

Breeding drugs would be a medical boon. But if we can breed software and then later turn the system upon itself so that software breeds itself, leading to who knows what, can we set molecules too upon the path of open-ended evolution?

Yes, but it's a difficult job. Tom Ray's electric-powered evolution machine is heavy on the heritable information but light on bodies. Molecular evolution programs are heavy on bodies but skimpy on heritable information. Naked information is hard to kill, and without death there is no evolution. Flesh and blood greatly assist the cause of evolution because a body provides a handy way for information to die. Any system that can incorporate the two threads of heritable information and mortal bodies has the ingredients for an evolutionary system.

Gerald Joyce, a biochemist at San Diego whose background is the chemistry of very early life, devised a simple way to incorporate the dual nature of information and bodies into one robust artificial evolutionary system. He accomplished this by recreating a probable earlier stage of life on Earth -- "RNA world" -- in a test tube.

RNA is a very sophisticated molecular system. It was not the very first living system, but life on Earth at some stage almost certainly became RNA life. Says Joyce, "Everything in biology points to the fact that 3.9 billion years ago, RNA was running the show."

RNA has a unique advantage that no other system we know about can boast. It acts at once as both body and info, phenotype and genotype, messenger and message. An RNA molecule is at once the flesh that must interact in the world and the information that must inherit the world, or at least be transmitted to the next generation. Though limited by this uniqueness, RNA is a wonderfully compact system in which to begin open-ended artificial evolution.

Gerald Joyce runs a modest group of graduates and postdocs at Scripps Institute, a sleek modern lab along the California coast near San Diego. His experimental RNA worlds are tiny drops that pool in the bottom of plastic micro-test tubes hardly the volume of thimbles. At any one time dozens of these pastel-colored tubes, packed in ice in styrofoam buckets, await being warmed up to body temperature to start evolving. Once warmed, RNA will produce a billion copies in one hour.

"What we have here," Joyce says pointing to one of the tiny tubes, "is a huge parallel processor. One of the reasons I went into biology instead of doing computer simulations of evolution is that no computer on the face of the Earth, at least for the near future, can give me 1015 microprocessors in parallel." The drops in the bottom of the tubes are about the size of the smart part of computer chips. Joyce polishes the image: "Actually, our artificial system is even better than playing with natural evolution because there aren't too many natural systems that come close to letting us turn over 1015 individuals in a hour, either."

In addition to the intellectual revolution a self-sustaining life system would launch, Joyce sees evolution as a commercially profitable way to create useful chemicals and drugs. He imagines molecular evolution systems that run 24 hours, 365 days a year: "You give it a task, and say don't come out of your closet until you've figured out how to convert molecule A to molecule B."

Joyce rattles off a list of biotech companies that are today dedicated solely to research in directed molecular evolution (Gilead, Ixsys, Nexagen, Osiris, Selectide, and Darwin Molecule). His list does not include established biotech companies, such as Genentech, which are doing advanced research into directed evolutionary techniques, but which also practice rational drug design. Darwin Molecule, whose principal patent holder is complexity researcher Stuart Kauffman, raised several million dollars to exploit evolution's power to design drugs. Manfred Eigen, Nobel Prize-winning biochemist, calls directed evolution "the future of biotechnology."

But is this really evolution? Is this the same vital spirit that brought us insulin, eyelashes, and raccoons in the first place? It is. "We approach evolution with a capital D for Darwin," Joyce told me. "But since the selection pressure is determined by us, rather than nature, we call this directed evolution."

Directed evolution is another name for supervised learning, another name for the Method of traversing the Library, another name for breeding. Instead of letting the selection emerge, the breeder directs the choice of varieties of dogs, pigeons, pharmaceuticals, or graphic images.

David Ackley is a researcher of neural nets and genetic algorithms at Bellcore, the R&D labs for the Baby Bells. Ackley has some of the most original ways of looking at evolutionary systems that I've come across.

Ackley is a bear of a guy with a side-of-the-mouth wisecracking delivery. He broke up 250 serious scientists at the 1990 Second Artificial Life Conference with a wickedly funny video of a rather important artificial life world he and colleague Michael Littman had made. His "creatures" were actually bits of code not too different from a classical GA, but he dressed them up with moronic smiley faces as they went about chomping each other or bumping into walls in his graphical world. The smart survived, the dumb died. As others had, Ackley found that his world was able to evolve amazingly fit organisms. Successful individuals would live Methuselahian lifetimes -- 25,000 day-steps in his world. These guys had the system all figured out. They knew how to get what they needed with minimum effort. And how to stay out of trouble. Not only would individuals live long, but the populations that shared their genes would survive eons as well.

Noodling around with the genes of these streetwise creatures, Ackley uncovered a couple of resources they hadn't taken up. He saw that he could improve their chromosomes in a godlike way to exploit these resources, making them even better adapted to the environment he had set up for them. So in an early act of virtual genetic engineering, he modified their evolved code and set them back again into his world. As individuals, they were superbly fitted and flourished easily, scoring higher on the fitness scale than any creatures before them.

But Ackley noticed that their population numbers were always lower than the naturally evolved guys. As a group they were anemic. Although they never died out, they were always endangered. Ackley felt their low numbers wouldn't permit the species to last more than 300 generations. So while handcrafted genes suited individuals to the max, they lacked the robustness of organically grown genes, which suited the species to the max. Here, in the home-brewed world of a midnight hacker, was the first bit of testable proof for hoary ecological wisdom: that what is best for an individual ain't necessarily best for the species.

"It's tough accepting that we can't figure out what's best in the long run," Ackley told the Artificial Life conference to great applause, "but, hey, I guess that's life!"

Bellcore allowed Ackley to pursue his microgod world because they recognized that evolution is a type of computation. Bellcore was, and still is, interested in better computational methods, particularly those based on distributed models, because ultimately a telephone network is a distributed computer. If evolution is a useful type of distributed computation, what might some other methods be? And what improvements or variations, if any, can we make to evolutionary techniques? Taking up the usual library/space metaphor, Ackley gushes, "The space of computational machinery is unbelievably vast and we have only explored very tiny corners of it. What I'm doing, and what I want to do more of, is to expand the space of what people recognize as computation."

Of all the possible types of computation, Ackley is primarily interested in those procedures that underpin learning. Strong learning methods require smart teachers; that's one type of learning. A smart teacher tells a learner what it should know, and the learner analyzes the information and stores it in memory. A less smart teacher can also teach by using a different method. It doesn't know the material itself, but it can tell when the learner guesses the right answer -- as a substitute teacher might grade tests. If the learner guesses a partial answer the weak teacher can give a hint of "getting warm," or "getting cold" to help the learner along. In this way, a weaker teacher can potentially generate information that it itself doesn't own. Ackley has been pushing the edge of weak learning as a way of maximizing computation: leveraging the smallest amount of information in, to get the maximum information out. "I'm trying to come up with the dumbest, least informative teacher as possible," Ackley told me. "And I think I found it. My answer is: death."

Death is the only teacher in evolution. Ackley's mission was to find out: what can you learn using only death as a teacher? We don't know for sure, but some candidates are: soaring eagles, or pigeon navigation systems, or termite skyscrapers. It takes a while, but evolution is clever. Yet it is obviously blind and dumb. "I can't imagine any dumber type of learning than natural selection," says Ackley.

In the space of all possible computation and learning, then, natural selection holds a special position. It occupies the extreme point where information transfer is minimized. It forms the lowest baseline of learning and smartness, below which learning doesn't happen and above which smarter, more complicated learning takes place. Even though we still do not fully understand the nature of natural selection in coevolutionary worlds, natural selection remains the elemental melting point of learning. If we could measure degrees of evolution (we can't yet) we would have a starting benchmark against which to rate other types of learning.

Natural selection plays itself out in many guises. Ackley was right; computer scientists now realize that many modes of computation exist -- many of them evolutionary. For all anyone knows, there may be hundreds of styles of evolution and learning. All such strategies, however, perform a search routine through a library or space. "Discovering the notion of the 'search' was the one and only brilliant idea that traditional AI research ever had," claims Ackley. A search can be accomplished in many ways. Natural selection -- as it is run in organic life -- is but one flavor.

Biological life is wedded to a particular hardware: carbon-based DNA molecules. This hardware limits the versions of search-by-natural-selection that can successfully operate upon it. With the new hardware of computers, particularly parallel computers, a host of other adaptive systems can be conjured up, and entirely different search strategies set out to shape them. For instance, a chromosome of biological DNA cannot broadcast its code to DNA molecules in other organisms in order for them to receive the message and alter their code. But in a computer environment you can do that.

David Ackley and Michael Littman, both of Bellcore's Cognitive Science Research Group, set out to fabricate a non-Darwinian evolutionary system in a computer. They chose a most logical alternative: Lamarckian evolution -- the inheritance of acquired traits. Lamarckism is very appealing. Intuitively such a system would seem deeply advantageous over the Darwinian version, because presumably useful mutations would be adopted into the gene line more quickly. But a look at its severe computational requirements quickly convinces the hopeful engineer how unlikely such a system would be in real life.

If a blacksmith acquires bulging biceps, how does his body reverse- engineer the exact changes in his genes needed to produce this improvement? The drawback for a Lamarckian system is its need to trace a particular advantageous change in the body back through embryonic development into the genetic blueprints. Since any change in an organism's form may be caused by more than one gene, or by many instructions interacting during the body's convoluted development, unraveling the tangled web of causes of any outward form requires a tracking system almost as complex as the body itself. Biological Lamarckian evolution is hampered by a strict mathematical law: that it is supremely easy to multiply prime factors together, but supremely hard to derive the prime factors out of the result. The best encryption schemes work on this same asymmetrical difficulty. Biological Lamarckism probably hasn't happened because it requires an improbable biological decryption scheme.

But computational entities don't require bodies. In computer evolution (as in Tom Ray's electric-powered evolution machine) the computer code doubles as both gene and body. Thus, the dilemma of deriving a genotype from the phenotype is moot. (The restriction of monolithic representation is not all that artificial. Life on Earth must have passed through this stage, and perhaps any spontaneously organizing vivisystem must begin with a genotype that is restricted to its phenotype, as simple self-replicating molecules would be.)

In artificial computer worlds, Lamarckian evolution works. Ackley and Littman implemented a Lamarckian system on a parallel computer with 16,000 processors. Each processor held a subpopulation of 64 individuals, for a grand total of approximately one million individuals. To simulate the dual information lines of body and gene, the system made a copy of the gene for each individual and called the copy the "body." Each body was a slightly different bit of code trying to solve the same problem as its million siblings.

The Bellcore scientists set up two runs. In the Darwinian run, the body code would mutate over time. By chance a lucky guy might become code that provides a better solution, so the system chooses it to mate and replicate. But in Darwinism when it mates, it must use its original "gene" copy of the code -- the code it inherited, not the improved body code it acquired during its lifetime. This is the biological way; when the blacksmith mates, he uses the code for the body he inherited, not the body he acquired.

In the Lamarckian run, by contrast, when the lucky guy with the improved body code is chosen to mate, it can use the improved code acquired during its lifetime as the basis for its mating. It is as if a blacksmith could pass on his massive arms to his offspring.

Comparing the two systems, Ackley and Littman found that, at least for the complicated problems they looked at, the Lamarckian system discovered solutions almost twice as good as the Darwinian method. The smartest Lamarckian individual was far smarter than the smartest Darwinian one. The thing about Lamarckian evolution, says Ackley, is that it "very quickly squeezes out the idiots" in a population. Ackley once bellowed to a roomful of scientists, "Lamarck just blows the doors off of Darwin!"

In a mathematical sense, Lamarckian evolution injects a bit of learning into the soup. Learning is defined as adaptation within an individual's lifetime. In classical Darwinian evolution, individual learning doesn't count for much. But Lamarckian evolution permits information acquired during a lifetime (including how to build muscles or solve equations) to be incorporated into the long-term, dumb learning that takes place over evolution. Lamarckian evolution produces smarter answers because it is a smarter type of search.

The superiority of Lamarckism surprised Ackley because he felt that nature did things so well: "From a computer science viewpoint it seems really stupid that nature is Darwinian and not Lamarckian. But nature is stuck on chemicals. We're not." It got him thinking about other types of evolution and search methods that might be more useful if you weren't restricted to operating on molecules.

A group of researchers in Milan, Italy, have come up with a few new varieties of evolution and learning. Their methods fill a few holes in Ackley's proposed "space of all possible types of computation." Because they were inspired by the collective behavior of ant colonies, the Milan group call their searches "Ant Algorithms."

Ants have distributed parallel systems all figured out. Ants are the history of social organization and the future of computers. A colony may contain a million workers and hundreds of queens, and the entire mass of them can build a city while only dimly aware of one another. Ants can swarm over a field and find the choicest food in it as if the swarm were a large compound eye. They weave vegetation together in coordinated parallel rows, and collectively keep their nest at a steady temperature, although not a single ant has ever lived who knows how to regulate temperature.

An army of ants too dumb to measure and too blind to see far can rapidly find the shortest route across a very rugged landscape. This calculation perfectly mirrors the evolutionary search: dumb, blind, simultaneous agents trying to optimize a path on a computationally rugged landscape. Ants are a parallel processing machine.

Real ants communicate with each other by a chemical system called pheromones. Ants apply pheromones on each other and on their environment. These aromatic smells dissipate over time. The odors can also be relayed by a chain of ants picking up a scent and remanufacturing it to pass on to others. Pheromones can be thought of as information broadcasted or communicated within the ant system.

The Milan group (Alberto Colorni, Marco Dorigo, and Vittorio Maniezzo) constructed formulas modeled on ant logic. Their virtual ants ("vants") were dumb processors in a giant community operating in parallel. Each vant had a meager memory, and could communicate locally. Yet the rewards of doing well were shared by others in a kind of distributed computation.

The Italians tested their ant machine on a standard benchmark, the traveling salesman problem. The riddle was: what is the shortest route between a large number of cities, if you can only visit each city once? Each virtual ant in the colony would set out rambling from city to city leaving a trail of pheromones. The shorter the path between cities, the less the pheromone evaporated. The stronger the pheromone signal, the more other ants followed that route. Shorter paths were thus self-reinforcing. Run for 5,000 rounds or so, the ant group-mind would evolve a fairly optimal global route.

The Milan group played with variations. Did it make any difference if the vants all started at one city or were uniformly distributed? (Distributed was better.) Did it make any difference how many vants one ran concurrently? (More was better until you hit the ratio of one ant for every city, when the advantage peaked.) By varying parameters, the group came up with a number of computational ant searches.

Ant algorithms are a type of Lamarckian search. When one ant stumbles upon a short route, that information is indirectly broadcast to the other vants by the trail's pheromone strength. In this way learning in one ant's lifetime is indirectly incorporated into the whole colony's inheritance of information. Individual ants effectively broadcast what they have learned into their hive. Broadcasting, like cultural teaching, is a part of Lamarckian search. Ackley: "There are ways to exchange information other than sex. Like the evening news."

The cleverness of the ants, both real and virtual, is that the amount of information invested into "broadcasting" is very small, done very locally, and is very weak. The notion of introducing weak broadcasting into evolution is quite appealing. If there is any Lamarckism in earthly biology it is buried deep. But there remains a universe full of strange types of potential computation that might employ various modes of Lamarckian broadcasting. I know of programmers fooling around with algorithms to mimic "memetic" evolution -- the flow of ideas (memes) from one mind to another, trying to capture the essence and power of cultural evolution. Out of all the possible ways to connect the nodes in distributed computers, only a very few, such as the ant algorithms, have even been examined.

As late as 1990, parallel computers were derided by experts as controversial, specialized, and belonging the lunatic fringe. They were untidy and hard to program. The lunatic fringe disagreed. In 1989, Danny Hillis boldly made a widely publicized bet with a leading computer expert that as early as 1995, more bits per month would be processed by parallel machines than by serial machines. He is looking right. As serial computers audibly groaned under the burden of pushing complex jobs through the tiny funnel of von Neumann's serial processor, a change in expert opinion suddenly swept through the computer industry. Peter Denning signaled the new perspective when he wrote in a paper published by Science ("Highly Parallel Computation," November 30, 1990), "Highly parallel computing architectures are the only means to achieve the computational rates demanded by advanced scientific problems." John Koza of Stanford's Computer Science Department says flatly, "Parallel computers are the future of computing. Period."

But parallel computers remain hard to manage. Parallel software is a tangled web of horizontal, simultaneous causes. You can't check such nonlinearity for flaws since it's all hidden corners. There is no narrative to step through. The code has the integrity of a water balloon, yielding in one spot as another bulges. Parallel computers can easily be built but can't be easily programmed.

Parallel computers embody the challenge of all distributed swarm systems, including phone networks, military systems, the planetary 24-hour financial web, and large computer networks. Their complexity is taxing our ability to steer them. "The complexity of programming a massively parallel machine is probably beyond us," Tom Ray told me. "I don't think we'll ever be able to write software that fully uses the capacity of parallelism."

Little dumb creatures in parallel that can "write" better software than humans can suggests to Ray a solution for our desire for parallel software. "Look," he says, "ecological interactions are just parallel optimization techniques. A multicellular organism essentially runs massively parallel code of an astronomical scale. Evolution can 'think' of parallel programming ways that would take us forever to think of. If we can evolve software, we'll be way ahead." When it comes to distributed network kinds of things, Rays says, "Evolution is the natural way to program."

The natural way to program! That's an ego-deflating lesson. Humans should stick to what they do best: small, elegant, minimal systems that are fast and deep. Let natural evolution (artificially injected) do the messy big work.

Danny Hillis has come to the same conclusion. He is serious when he says he wants his Connection Machine to evolve commercial software. "We want these systems to solve a problem we don't know how to solve, but merely know how to state." One such problem is creating multimillion-line programs to fly airplanes. Hillis proposes setting up a swarm system which would try to evolve better software to steer a plane, while tiny parasitic programs would try to crash it. As his experiments have shown, parasites encourage a faster convergence to an error-free, robust software navigation program. Hillis: "Rather than spending uncountable hours designing code, doing error-checking, and so on, we'd like to spend more time making better parasites!"

Even when technicians do succeed in engineering an immense program such as navigation software, testing it thoroughly is becoming impossible. But things grown, not made, are different. "This kind of software would be built in an environment full of thousands of full-time adversaries who specialize in finding out what's wrong with it," Hillis says, thinking of his parasites. "Whatever survives them has been tested ruthlessly." In addition to its ability to create things that we can't make, evolution adds this: it can also make them more flawless than we can. "I would rather fly on a plane running software evolved by a program like this, than fly on a plane running software I wrote myself," says Hillis, programmer extraordinaire.

The call-routing program of long-distance phone companies tallies up to about 2 million lines of code. Three faulty lines in those 2 million caused the rash of national telephone system outages in the summer of 1990. And 2 million lines is no longer large. The combat computers aboard the Navy's Seawolf submarine contain 3.6 million lines of code. "NT," the new workstation computer operating system released by Microsoft in 1993, required 4 million lines of code. One-hundred-million-line programs are not far away.

When computer programs swell to billions of lines of code, just keeping them up and "alive" will become a major chore. Too much of the economy and too many people's lives will depend on billion-line programs to let them go down for even an instant. David Ackley thinks that reliability and up-time will become the primary chore of the software itself. "I claim that for a really complex program sheer survival is going to consume more of its resources." Right now only a small portion of a large program is dedicated to maintenance, error correction, and hygiene. "In the future," predicts Ackley, "99 percent of raw computer cycles are going to be spent on the beast watching itself to keep it going. Only that remaining 1 percent is going to be used for user tasks -- telephone switching or whatever. Because the beast can't do the user tasks unless it survives."

As software gets bigger, survival becomes critical yet increasingly difficult. Survival in the everyday world of daily use means flexibility and evolvability. And it demands more work to pull off. A program survives only if it constantly analyzes its status, adjusts its code to new demands, cleanses itself, ceaselessly dissects anomalous circumstances, and always adapts and evolves. Computation must seethe and behave as if it is alive. Ackley calls it "software biology" or "living computation." Engineers, even on 24-hour beepers, can't keep billion-line code alive. Artificial evolution may be the only way to keep software on its toes, looking lively.

Artificial evolution is the end of engineering's hegemony. Evolution will take us beyond our ability to plan. Evolution will craft things we can't. Evolution will make them more flawless than we can. And evolution will maintain them as we can't.

But the price of evolution is the title of this book. Tom Ray explains: "Part of the problem in an evolving system is that we give up some control."

Nobody will understand the evolved aviation software that will fly Danny Hillis. It will be an indecipherable spaghetti of 5 million strands of nonsense -- of which perhaps only 2 million are really needed. But it will work flawlessly.

No human will be able to troubleshoot the living software running Ackley's evolved telephone system. The lines of program are buried in an uncharted web of small machines, in an incomprehensible pattern. But, when it falters, it will heal itself.

No one will control the destination of Tom Ray's soup of critters. They are brilliant in devising tricks, but there is no telling them what trick to work on next. Only evolution can handle the complexities we are creating, but evolution escapes our total command.

At Xerox PARC, Ralph Merkle is engineering very small molecules that can replicate. Because these replicators dwell in the microscopic scale of nanometers (smaller than bacteria) their construction techniques are called nanotechnology. At some point in the very near future the engineering skills of nanotechnology and the engineering skills of biotechnology converge; they are both treating molecules as machines. Think of nanotechnology as bioengineering for dry life. Nanotechnology has the same potential for artificial evolution as biological molecules. Merkle told me, "I don't want nanotechnology to evolve. I want to keep it in a vat, constrained by international law. The most dangerous thing that could happen to nanotechnology is sex. Yes, I think there should be international regulations against sex for nanotechnology. As soon as you have sex, you have evolution, and as soon as you have evolution, you have trouble."

The trouble of evolution is not entirely out of our control; surrendering some control is simply a tradeoff we make when we employ it. The things we are proud of in engineering -- precision, predictability, exactness, and correctness -- are diluted when evolution is introduced.

These have to be diluted because survivability in a world of accidents, unforeseen circumstances, shifting environments -- in short, the real world-demands a fuzzier, looser, more adaptable, less precise stance. Life is not controlled. Vivisystems are not predictable. Living creatures are not exact. "'Correct' will go by the board," Ackley says of complex programs. "'Correct' is a property of small systems. In the presence of great change, 'correct' will be replaced by 'survivability'."

When the phone system is run by adaptable, evolved software, there will be no correct way to run it. Ackley continues: "To say that a system is 'correct' in the future will sound like bureaucratic double-talk. What people are going to judge a system on is the ingenuity of its response, and how well it can respond to the unexpected." We will trade correctness for flexibility and durability. We will trade a clean corpse for messy life. Ackley: "It will be to your advantage to have an out-of-control, but responsive, monster spend 1 percent of itself on your problem, than to have a dedicated little correct ant of a program that hasn't got a clue about what in the world is going on."

A student at one of Stuart Kauffman's lectures once asked him, "How do you evolve for things you don't want? I see how you can get a system to evolve what you want; but how can you be sure it won't create what you don't want?" Good question, kid. We can define what we want narrowly enough to breed for it. But we often don't even know what we don't want. Or if we do, the list of things that are unacceptable is so long as to be impractical. How can we select out disadvantageous side effects?

"You can't." Kauffman replied bluntly.

That's the evolutionary deal. We trade power for control. For control junkies like us, this is a devil's bargain.

Give up control, and we'll artificially evolve new worlds and undreamed-of richness. Let go, and it will blossom.

Have we ever resisted temptation before?

An open universe

A swarm of honeybees absconds from the hive and then dangles in a cluster from a tree branch. If a nearby beekeeper is lucky, the swarm settles on a branch that is easy to reach. The bees, gorged with honey and no longer protecting their brood, are as docile as ladybugs.

I've found a swarm or two in my time hung no higher than my head, and I've moved them into an empty hive box for my own. The way you move 10,000 bees from a tree branch into a box is one of life's magic shows.

If there are neighbors watching you can impress them. You lay a white sheet or large piece of cardboard on the ground directly under the buzzing cluster of bees. You then slide the bottom entrance lip of an empty hive under one edge of the sheet so that the cloth or cardboard forms a gigantic ramp into the hive's opening. You pause dramatically, and then you give the branch a single vigorous shake.

The bees fall out of the tree in a single clump and spill onto the sheet like churning black molasses. Thousands of bees crawl over each other in a chaotic buzzing mass. Then slowly, you begin to notice something. The bees align themselves toward the hive opening and march into the entrance as if they were tiny robots under one command. And they are. If you bend down to the sheet and put your nose near the pool of crawling bees, you can smell a perfume like roses. You can see that the bees are hunched over and fanning their wings furiously as they walk. They are emitting the rose smell from a gland in their rear ends and fanning the scent back to the troops behind them. The scent says, "The queen is here. Follow me." The second follows the first and the third the second and five minutes later the sheet is almost empty as the last of the swarm sucks itself into the box.

The first life on Earth could not put on that show. It was not a matter of lacking the right variation. There simply was no room in all of the possibilities accorded by its initial genes for such a wild act. To use the smell of a rose to coordinate 10,000 flying beings into a purposeful crawling beast was beyond early life's reach. Not only had early life not yet created the space -- worker bee, queen relationship, honey from flowers, tree, hive, pheromones -- -- in which to stage the show, it had not created the tools to make the space.

Nature dispenses breathtaking diversity because its charter is open ended. Life did not confine itself to producing its dazzling variety within the limited space of the few genes it first made. On the contrary, one of the first things life discovered was how to create new genes, more genes, variable genes, and a bigger genetic library.

It is one of the hallmarks of life that it continues to enlarge the space of its own being. Nature is an ever-expanding library of possibilities. It is an open universe. At the same time that life turns up the most improbable books from the Library shelves, it is adding new wings to the collection, making room for more of its improbable texts.

We don't know how life crossed the threshold from fixed gene space to variable gene space. Perhaps it was one particular gene's duty to determine the total number of genes in the chromosome. Then by mutating that one gene, the sum of genes in the string would increase or decrease. Or the size of the genome might have been indirectly determined by more than one gene. Or, more likely, genome size is determined by the structure of the genetic system itself.

Tom Ray showed that in his world of self-replicators, variable genome length emerged instantaneously. His creatures determined their own genome (and thus the size of their possible libraries) in a range from his unexpectedly short "22" to one creature that was 23,000 bytes long.

The consequence of an open genome is open evolution. A system which predetermines what each gene must do or how many genes there are can only evolve to predetermined boundaries. The first systems of Dawkins, Latham, Sims and the Russian El-Fish programmers were grounded by this limitation. They may generate all possible pictures of a given size and depth, but not all possible art. A system that does not predetermine the role or number of genes can shoot the moon. This is why Tom Ray's critters stir such excitement. In theory, his world, run long enough, could evolve anything in the ultimate Library.

There is more than one way to organize an open genome. In 1990, Karl Sims took advantage of the supercomputing power of the CM2 to devise a new type of artificial world formed by genes of unfixed length, a world much improved over his botanical-picture world. Sims accomplished this trick by creating a genome composed of small equations rather than of long strings of digits. His original library of fixed genes each controlled one visual parameter of a plant; his second library held equations of variable and open-ended length which drew curves, colors, forms and shapes.

Sims's equation -- genes were small self-contained logical units of a computer language (LISP). Each module was an arithmetical command such as add, subtract, multiply, cosine, sine. Sims called these units "primitives" -- a logical alphabet. If you have a suitable primitive alphabet you can build all possible equations, just as with the appropriately diverse alphabet of sounds you could build all spoken sentences. Add, multiply, cosine, etc., can be combined to generate any mathematical equation we can think of. Since any shape can be described by an equation, this primitive alphabet can make any picture. Adding to the complexity of the equation will subtly enlarge the complexity of the resulting image.

There was a serendipitous second advantage to working with a library of equations. In Sims's original world (and in Tom Ray's Tierra and Danny Hillis's coevolutionary parasites), organisms were strings of digits that randomly flipped a digit, just as books in the Borgian Library altered by one letter at a time. In Sims's improved universe, organisms were strings of logical units that randomly flipped a unit. This would be like a Borgian Library where words, not letters, were flipped. Every word in every book was correctly spelled, so every page in every book had a more sensible pattern. But whereas the soup for a Borgian Library based on words would necessitate tens of thousands of words in the pot to begin with, Sims could make all possible equations starting with a soup of only a dozen or so mathematical primitives.

Yet, the most revolutionary advantage to evolving logic units rather than digital bits was that it immediately moved the system onto the road toward an opened-ended universe. Logic units are functions themselves and not mere values for functions, as digital bits are. By adding or swapping a logical primitive here or there, the entire functionality of the program shifts or enlarges. New kinds of functions and new kinds of things will emerge in such a system.

That's what Sims found. Entirely new kinds of pictures evolved by his equations and painted themselves onto the computer monitor. The first thing that struck him was how rich the space was. By restricting the primitives to logical parts, Sims's LISP alphabet ensured that most equations drew some pattern. Instead of being full of muddy gray patterns, there were astounding sights almost wherever he went. Just dipping in at random landed him in the middle of "art." The first screen was full of wild red and blue zigzags. The next screen pulsated with yellow hovering orbs. The next generation yielded yellow orbs with a misty horizon, the next, sharpened waves with a horizon of blue. And the next, circular smudges of pastel yellow color reminiscent of buttercups. Almost every turn reeled in a marvelously inventive scene. In an hour, thousands of stunning pictures were roused out of their hiding places and displayed to the living for the first and last time. It was like watching over the shoulder of the world's greatest painter as he sketched without ever repeating a theme or pattern.

While Sims selected one picture, bred variations of it, and then selected another, he was not only evolving pictures. Underneath it all, Sims was evolving logic. A relatively small logic equation drew an eye-boggling complex picture. At one point Sims's system evolved the following eight lines of logic code:

(cos (round (atan (log (invert y) (+ (bump (+ (round x y) y) #(0.46 0.82 0.65) 0.02 #(0.1 0.06 0.1) #(0.99 0.06 0.41) 1.47 8.7 3.7) (color-grad (round (+ y y) (log (invert x) (+ (invert y) (round (+ y x) (bump (warped-ifs (round y y) y 0.08 0.06 7.4 1.65 6.1 0.54 3.1 0.26 0.73 15.8 5.7 8.9 0.49 7.2 15.6 0.98) #(0.46 0.82 0.65) 0.02 #(0.1 0.06 0.1) #(0.99 0.06 0.41) 0.83 8.7 2.6))))) 3.1 6.8 #(0.95 0.7 0.59) 0.57))) #(0.17 0.08 0.75) 0.37) (vector y 0.09 (cos (round y y)))))

When fleshed out on Sims's color monitor, the equation painted what seems to be two sheets of icicles backlit by an arctic sunset. It's an arresting image. The ice is molded in great detail and translucent, the horizon in the background abstract and serene. It could have been painted by a weekend artist. As Sims points out, "This equation was evolved from scratch in only a few minutes -- probably much faster than it could be designed."

But Sims is at a total loss to explain the logic of the equation and why it produces a picture of ice. It looks as cryptic and muddled to him as to you. The equation's convoluted reason is beyond quick mathematical understanding.

The bombastic notion of evolving logic programs has been taken up in earnest by John Koza, a professor of computer science at Stanford. Koza was one of John Holland's students who brought knowledge of Holland's genetic algorithms out of the dark ages of the '60s and '70s into the renaissance of parallelism of the late '80s.

Rather than merely explore the space of possible equations, as Sims the artist did, Koza wanted to breed the best equation to solve a particular problem. One could imagine (as a somewhat silly example) that in the space of possible pictures there might be one that would induce cows gazing at it to produce more milk. Koza's method can evolve the equations that would draw that particular picture. In this farfetched idea, Koza would keep rewarding the equations which drew a picture that even minutely increased milk production until there was no further increase. For his actual experiments, though, Koza choose more practical tests, such as finding an equation that could steer a moving robot.

But in a sense his searches were similar to those of Sims and the others. He hunted in the Borgian Library of possible computer programs -- not on an aimless mission to see what was there, but to find the best equation for a particular practical problem. Koza wrote in Genetic Programming, "I claim that the process of solving these problems can be reformulated as a search for a highly fit individual computer program in the space of possible computer programs."

For the same reason computer experts said Ray's scheme of computer evolution couldn't work, Koza's desire to "find" equations by breeding them bucked convention. Everyone "knew" that logic programs were brittle and unforgiving of the slightest alteration. In computer science theory, programs had two pure states: (1) flawlessly working; or (2) modified and bombed. The third state -- modified at random yet working -- was not in the cards. Slight modifications were known as bugs, and people paid a lot of money to keep them out. If progressive modification and improvement (evolution) of computer equations was at all possible, the experts thought, it must be so only in a few precious areas or specialized types of programs.

The surprise of artificial evolution has been that conventional wisdom was so wrong. Sims, Ray, and Koza have wonderful evidence that logical programs can evolve by progressive modifications.

Koza's method was based on the intuitive hunch that if two mathematical equations are somewhat effective in solving a problem, then some parts of them are valuable. And if the valuable parts from both are recombined into a new program, the result might be more effective than either parent. Koza randomly recombined, in thousands of combinations, parts of two parents, banking on the probabilistic likelihood that one of those random recombinations would include the optimal arrangement of valuable parts to better solve the problem.

There are many similarities between Koza's method and Sims's. Koza's soup, too, was a mixture of about a dozen arithmetical primitives, such as add, multiply, cosine, rendered in the computer language LISP. The units were strung together at random to form logical "trees," a hierarchical organization somewhat like a computer flow chart. Koza's system created 500 to 10,000 different individual logic trees as the breeding population. The soup usually converged upon a decent offspring in about 50 generations.

Variety was forced by sexually swapping branches from one tree to the next. Sometimes a long branch was grafted, other times a mere twig or terminal "leaf." Each branch could be thought of as an intact subroutine of logic made of smaller branches. In this way, bits of equation (a branch), or a little routine that worked and was valuable, had a chance of being preserved or even passed around.

All manner of squirrely problems can be solved by evolving equations. A typical riddle which Koza subjected to this cure was how to balance a broom on a skateboard. The skateboard must be moved back and forth by a motor to keep the inverted broom pivoted upright in the board's center. The motor-control calculations are horrendous, but not very different from the control circuits needed for maneuvering robot arms. Koza found he could evolve a program to achieve this control.

Other problems he tested evolutionary equations against included: strategies for navigating a maze; rules for solving quadratic equations; methods to optimize the shortest route connecting many cities (also known as traveling salesman problem); strategies for winning a simple game like tic-tac-toe. In each case, Koza's system sought a formula for the test problem rather than a specific answer for a specific instance of the test. The more varied instances a sound formula was tested against, the better the formula became with each generation.

While equation breeding yields solutions that work, they are usually the ugliest ones you could imagine. When Koza began to inspect the insides of his highly evolved prizes, he had the same shock that Sims and Ray did: the solutions were a mess! Evolution went the long way around. Or it burrowed through the problem by some circuitous loophole of logic. Evolution was chock-full of redundancy. It was inelegant. Rather than remove an erroneous section, evolution would just add a countercorrecting section, or reroute the main event around the bad sector. The final formula had the appearance of being some miraculous Rube Goldberg collection of items that by some happy accident worked. And that's exactly what it was, of course.

Take as an example a problem Koza once threw at his evolution machine. It was a graph of two intertwining spirals. A rough approximation would be the dual spirals in pinwheel. Koza's evolutionary equation machine had to evolve the best equation capable of determining on which of the two intertwined spiral lines each of about 200 data points lay.

Koza loaded his soup with 10,000 randomly generated computer formulas. He let them breed, as his machine selected the equations that came closest to getting the right formula. While Koza slept, the program trees swapped branches, occasionally birthing a program that worked better. He ran the machine while he was on vacation. When he returned, the system had evolved an answer that perfectly categorized the twin spirals.

This was the future of software programming! Define a problem and the machine will find a solution while the engineers play golf. But the solution Koza's machine found tells us a lot about the handiwork of evolution. Here's the equation it came up with:

(SIN (IFLTE (IFLTE (+ Y Y) (+ X Y) (- X Y) (+ Y Y)) (* X X) (SIN (IFLTE (% Y Y) (% (SIN (SIN (% Y 0.30400002))) X) (% Y 0.30400002) (IFLTE (IFLTE (% (SIN (% (% Y (+ X Y)) 0.30400002)) (+ X Y)) (% X 0.10399997) (- X Y) (* (+ -0.12499994 -0.15999997) (- X Y))) 0.30400002 (SIN (SIN (IFLTE (% (SIN (% (% Y 0.30400002) 0.30400002)) (+ X Y)) (% (SIN Y) Y) (SIN (SIN (SIN (% (SIN X) (+ -0.12499994 -0.15999997))))) (% (+ (+ X Y) (+ Y Y)) 0.30400002)))) (+ (+ X Y) (+ Y Y))))) (SIN (IFLTE (IFLTE Y (+ X Y) (- X Y) (+ Y Y)) (* X X) (SIN (IFLTE (% Y Y) (% (SIN (SIN (% Y 0.30400002))) X) (% Y 0.30400002) (SIN (SIN (IFLTE (IFLTE (SIN (% (SIN X) (+ -0.12499994 -0.15999997))) (% X -0.10399997) (- X Y) (+ X Y)) (SIN (% (SIN X) (+ -0.12499994 -0.15999997))) (SIN (SIN (% (SIN X) (+ -0.12499994 -0.15999997)))) (+ (+ X Y) (+ Y Y))))))) (% Y 0.30400002))))).

Not only is it ugly, it's incomprehensible. Even for a mathematician or computer programmer, this evolved formula is a tar baby in the briar patch. Tom Ray says evolution writes code that only an intoxicated human programmer would write, but it may be more accurate to say evolution generates code that only an alien would write; it is decidedly inhuman. Backtracking through the evolving ancestors of the equation, Koza eventually traced the manner in which the program tackled the problem. By sheer persistence and by hook and crook it found a laborious roundabout way to its own answer. But it worked.

The answer evolution discovered seems strange because almost any high school algebra student could write a very elegant equation in a single line that described the two spirals.

There was no evolutionary pressure in Koza's world toward simple solutions. His experiment could not have found that distilled equation because it wasn't structured to do so. Koza tried applying parsimony in other runs but found that parsimony added to the beginning of a run dampened the efficiency of the solutions. He'd find simple but mediocre to poor solutions. He has some evidence that adding parsimony at the end of evolutionary procedure -- that is, first let the system find a solution that kind of works and then start paring it down -- is a better way to evolve succinct equations.

But Koza passionately believes parsimony is highly overrated. It is, he says, a mere "human esthetic." Nature isn't particularly parsimonious. For instance, David Stork, then a scientist at Stanford, analyzed the neural circuits in the muscles of a crayfish tail. The network triggers a curious backflip when the crayfish wants to escape. To humans the circuit looks baroquely complex and could be simplified easily with the quick removal of a couple of superfluous loops. But the mess works. Nature does not simplify simply to be elegant.

Humans seek a simple formula such as Newton's f=ma, Koza suggests, because it reflects our innate faith that at bottom there is elegant order in the universe. More importantly, simplicity is a human convenience. The heartwarming beauty we perceive in f=ma is reinforced by the cold fact that it is a much easier formula to use than Koza's spiral monster. In the days before computers and calculators, a simple equation was more useful because it was easier to compute without errors. Complicated formulas were a grind and treacherous. But, within a certain range, neither nature nor parallel computers are troubled by convoluted logic. The extra steps we find ugly and stupefying, they do perfectly in tedious exactitude.

The great irony puzzling cognitive scientists is why human consciousness is so unable to think in parallel, despite the fact that the brain runs as a parallel machine. We have an almost uncanny blind spot in our intellect. We cannot innately grasp concepts in probability, horizontal causality, and simultaneous logic. We simply don't think like that. Instead our minds retreat to the serial narrative -- the linear story. That's why the first computers were programmed in von Neumann's serial design: because that's how humans think.

And this, again, is why parallel computers must be evolved rather than designed: because we are simpletons when it comes to thinking in parallel. Computers and evolution do parallel; consciousness does serial. In a very provocative essay in the Winter 1992 Daedalus, James Bailey, director of marketing at Thinking Machines, wrote of the wonderful boomeranging influence that parallel computers have on our thinking. Entitled "First We Reshape Our Computers. Then Our Computers Reshape Us," Bailey argues that parallel computers are opening up new territories in our intellectual landscape. New styles of computer logic in turn force new questions and new perspectives from us. "Perhaps," Bailey suggests, "whole new forms of reckoning exist, forms that only make sense in parallel." Thinking like evolution may open up new doors in the universe.

John Koza sees the ability of evolution to work on both ill-defined and parallel problems as another of its inimitable advantages. The problem with teaching computers how to learn to solve problems is that so far we have wound up explicitly reprogramming them for every new problem we come across. How can computers be designed to do what needs to be done, without being told in every instance what to do and how to do it?

Evolution, says Koza, is the answer. Evolution allows a computer's software to solve a problem to which the scope, kind, or range of the answer(s) may not be evident at all, as is usually the case in the real world. Problem: A banana hangs in a tree; what is the routine to get it? Most computer learning to date cannot solve that problem unless we explicitly clue the program in to certain narrow parameters such as: how many ladders are nearby? Any long poles?

Having defined the boundaries of the answer, we are half answering the question. If we don't tell it what rocks are near, we know we won't get the answer "throw a rock at it." Whereas in evolution, we might. More probably, evolution would hand us answers we could have never expected: use stilts; learn to jump high, employ birds to help you; wait until after storms; make children and have them stand on your head. Evolution did not narrowly require that insects fly or swim, only that they somehow move quick enough to escape predators or catch prey. The open problem of escape led to the narrow answers of water striders tiptoeing on water or grasshoppers springing in leaps.

Every worker dabbling in artificial evolution has been struck by the ease with which evolution produces the improbable. "Evolution doesn't care about what makes sense; it cares about what works," says Tom Ray.

The nature of life is to delight in all possible loopholes. It will break any rule it comes up with. Take these biological jaw-droppers: a female fish that is fertilized by her male mate who lives inside her, organisms that shrink as they grow, plants that never die. Biological life is a curiosity shop whose shelves never empty. Indeed the catalog of natural oddities is almost as long as the list of all creatures; every creature is in some way hacking a living by reinterpreting the rules.

The catalog of human inventions is far less diverse. Most machines are cut to fit a specific task. They, by our old definition, follow our rules. Yet if we imagine an ideal machine, a machine of our dreams, it would adapt, and -- better yet -- evolve.

Adaptation is the act of bending a structure to fit a new hole. Evolution, on the other hand, is a deeper change that reshapes the architecture of the structure itself -- how it can bend -- often producing new holes for others. If we predefine the organizational structure of a machine, we predefine what problems it can solve. The ideal machine is a general problem solver, one that has an open-ended list of things it can do. That means it must have an open-ended structure, too. Koza writes, "The size, shape, and structural complexity [of a solution] should be part of the answer produced by a problem solving technique -- not part of the question." In recognizing that a system itself sets the answers the system can make, what we ultimately want, then, is a way to generate machines that do not possess a predefined architecture. We want a machine that is constantly remaking itself.

Those interested in kindling artificial intelligence, of course, say "amen." Being able to come up with a solution without being unduly prompted to where the solution might exist -- lateral thinking it's called in humans -- is almost the definition of human intelligence.

The only machine we know of that can reshape its internal connections is the living gray tissue we call the brain. The only machine that would generate its own structure that we can presently even imagine manufacturing would be a software program that could reprogram itself. The evolving equations of Sims and Koza are the first step toward a self-reprogramming machine. An equation that can breed other equations is the basic soil for this kind of life. Equations that breed other equations are an open-ended universe. Any possible equation could arise, including self-replicating equations and formulas that loop back in a Uroborus bite to support themselves. This kind of recursive program, which reaches into itself and rewrites its own rules, unleashes the most magnificent power of all: the creation of perpetual novelty.

"Perpetual novelty" is John Holland's phrase. He has been crafting means of artificial evolution for years. What he is really working on, he says, is a new mathematics of perpetual novelty. Tools to create neverending newness.

Karl Sims told me, "Evolution is a very practical tool. It's a way of exploring new things you wouldn't have thought about. It's a way of refining things. And it's a way of exploring procedures without having to understand them. If computers are fast enough they can do all these things."

Exploring beyond the reach of our own understanding and refining what we have are gifts that directed, supervised, optimizing evolution can bring us. "But evolution," says Tom Ray, "is not just about optimization. We know that evolution can go beyond optimization and create new things to optimize." When a system can create new things to optimize we have a perpetual novelty tool and open-ended evolution.

Both Sims's selection of images and Koza's selection of software via the breeding of logic are examples of what biologists call breeding or artificial selection. The criteria for "fit" -- for what is selected -- is chosen by the breeder and is thus an artifact, or artificial. To get perpetual novelty -- to find things we don't anticipate -- we must let the system itself define the criteria for what it selects. This is what Darwin meant by "natural selection." The selection criteria was done by nature of the system; it arose naturally. Open-ended artificial evolution also requires natural selection, or if you will, artificial natural selection. The traits of selection should emerge naturally from the artificial world itself.

Tom Ray has installed the tool of artificial natural selection by letting his world determine its own fitness selection. Therefore his world is theoretically capable of evolving completely new things. But Ray did "cheat" a little to get going. He could not wait for his world to evolve self-replication on its own. So he introduced a self-replicating organism from the beginning, and once introduced, replication never vanished. In Ray's metaphor, he jump-started life as a single-celled organism, and then watched a "Cambrian explosion" of new organisms. But he isn't apologetic. "I'm just trying to get evolution and I don't really care how I get it. If I need to tweak my world's physics and chemistry to the point where they can support rich, open-ended evolution, I'm going to be happy. It doesn't make me feel guilty that I had to manipulate them to get it there. If I can engineer a world to the threshold of the Cambrian explosion and let it boil over the edge on its own, that will be truly impressive. The fact that I had to engineer it to get there will be trivial compared to what comes out of it."

Ray decided that getting artificial open-ended evolution up and running was enough of a challenge that he didn't need to evolve it to that stage. He would engineer his system until it could evolve on its own. As Karl Sims said, evolution is a tool. It can be combined with engineering. Ray used artificial natural selection after months of engineering. But it can go both ways. Other workers will engineer a result after months of evolution.

As a tool, evolution is good for three things:

- How to get somewhere you want but can't find the route to.
- How to get to somewhere you can't imagine.
- How to open up entirely new places to get to.

The third use is the door to an open universe. It is unsupervised, undirected evolution. It is Holland's ever-expanding perpetual novelty machine, the thing that creates itself.

Amateur gods such as Ray, Sims, and Dawkins have all expressed their astonishment at the way evolution seems to amplify the fixed space they thought they had launched. "It's a lot bigger than I thought" is the common refrain. I had a similar overwhelming impression when I stepped and jumped (literally) through the picture space of Karl Sims's evolutionary exhibit. Each new picture I found (or it found for me) was gloriously colored, unexpectedly complex, and stunningly different from anything I had ever seen before. Each new image seemed to enlarge the universe of possible pictures. I realized that my idea of a picture had previously been defined by pictures made by humans, or perhaps by biological nature. But in Sims's world an equally vast number of breathtaking vistas that were neither human-made nor biologically made -- but equally rich -- were waiting to be unwrapped.

Evolution was expanding my notions of possibilities. Life's biological system is very much like this. Bits of DNA are functional units -- logical evolvers that expand the space of possibilities. DNA directly parallels the operation of Sims's and Koza's logical units. (Or should we say their logical units parallel DNA?) A handful of units can be mixed and matched to code for any one of an astronomical number of possible proteins. The proteins produced by this small functional alphabet serve as tissue, disease, medicines, flavors, signals, and the bulk infrastructure of life.

Biological evolution is the open-ended evolution of DNA units breeding new DNA units in a library that is ever-expanding and without known boundaries.

Gerald Joyce, the molecular breeder, says he is happily into "evolving molecules for fun and profit." But his real dream is to hatch an alternative open-ended evolution scheme. He told me, "My interest is to see if we can set in motion, under our own control, the process of self-organization." The test case Joyce and colleagues are working on is to try to get a simple ribozyme to evolve the ability to replicate itself -- that very crucial step that Tom Ray skipped over. "The explicit goal is to set an evolving system in motion. We want molecules to learn how to make copies of themselves by themselves. Then it would be autonomous evolution instead of directed evolution."

Right now autonomous and self-sustained evolution is a mere dream for biochemists. No one has yet coerced an evolutionary system to take an "evolutionary step," one that develops a chemical process that heretofore didn't exist. To date, biochemists have only evolved new molecules which resolve problems they already knew how to solve. "True evolution is about going somewhere novel, not just reeling in interesting variants," says Joyce.

A working, autonomous, evolving, molecular system would be an incredibly powerful tool. It would be an open-ended system that could create all possible biologies. "It would be biology's triumph!" Joyce exclaims, equivalent, he believes, to the impact of "finding another life form in the universe that was happy to share samples with us."

But Joyce is a scientist and does not want to let his enthusiasm run over the edge: "We're not saying we are going to make life and it's going to develop its own civilization. That's goofy. We're saying we are going to make an artificial life form that is going to do slightly different chemistry than it does now. That's not goofy. That's realistic."

But Chris Langton doesn't find the prospect of artificial life creating its own civilization so goofy. Langton has gotten a lot of press for being the maverick who launched the fashionable field of artificial life. He has a good story, worth retelling very briefly because his own journey recapitulates the awakening of human-made, open-ended evolution.

Several years ago Langton and I attended a week-long science conference in Tucson, and to clear our heads, we played hooky for an afternoon. I had an invitation to visit the unfinished Biosphere 2 project an hour away, and so as we cruised the black ribbon of asphalt that winds through the basins of southern Arizona, Langton told me his life story.

At the time, Langton worked at the Los Alamos National Laboratory as a computer scientist. The entire town and lab of Los Alamos were originally built to invent the ultimate weapon. So I was surprised to hear Langton begin his story by saying he was a conscientious objector during the Vietnam War.

As a CO, Langton scored a chance to do alternative service as a hospital orderly at Boston's Massachusetts General Hospital. He was assigned the undesirable chore of transporting corpses from the hospital basement to the morgue basement. On the first week of the job, Langton and his partner loaded a corpse onto a gurney and pushed it through the dank, underground corridor connecting the two buildings. They needed to push it over a small concrete bridge under the only light in the tunnel, and as the gurney hit the bump, the corpse belched, sat upright, and started to slide off its perch! Chris spun around to grab his partner, but he saw only the distant doors flapping behind his coworker. Dead things could behave as if they were alive! Life was behavior; that was the first lesson.

Langton told his boss he couldn't go back to that job. Could he do something else? "Can you program computers?" he was asked. "Sure."

He got a job programming early-model computers. Sometimes he would let a silly game run on the unused computers at night. The game was called Life, devised by John Conway, and written for the mainframe by an early hacker named Bill Gosper. The game was a very simple code that would generate an infinite variety of forms, in patterns reminiscent of biological cells growing, replicating, and propagating on an agar plate. Langton remembered working alone late one night and suddenly feeling the presence of someone, something alive in the room, staring at him. He looked up and on the screen of Life he saw an amazing pattern of self-replicating cells. A few minutes later he felt the presence again. He looked up again and saw that the pattern had died. He suddenly felt that the pattern had been alive -- alive and as real as mold on an agar plate -- but on a computer screen instead. The bombastic idea that perhaps a computer program could capture life sprouted in Langton's mind.

He started fooling around with the game, probing it, wondering if it was possible to design a game like Life that would be open ended -- so that things would start to evolve on their own. He honed his programmer skills. On the job Langton was given the task of transferring a program from an out-of-date mainframe computer to a very different newer one. In order to do this, the trick was to abstract the operation of the hardware of the old computer and put it into the software of the newer one -- to extract the essential behavior of the hardware and cast it in intangible symbols. This way, old programs running on the new machine would be running in a virtual old computer emulated in software in the new computer. Langton said, "This was a first-hand experience of moving a process from one medium to another. The hardware didn't matter. You could run it on any hardware. What mattered was capturing the essential processes." It made him wonder if life could be taken from carbon and put into silicon.

After his service stint Langton spent his summers hang-gliding. He and a friend got a job hang-gliding over Grandfather Mountain in North Carolina for $25 per day as an airborne tourist attraction. They stayed aloft for hours at a time in 40-mile-per-hour winds. Swiped by a freak gust one day, Langton crashed from the sky. He hit the ground in a fetus position and broke 35 bones, including all the bones in his head except his skull. Although he smashed his knees through his face, he was alive. He spent the next six months on his back, half-conscious.

As he recovered from his massive concussions, Langton felt he was watching his brain "reboot," just as computers that are turned off have to rebuild their operating system when turned back on. One by one certain deep functions of his mind reappeared. In an epiphany of sorts, Langton remembers the moment when his sense of proprioception -- the sense of being centered in a body -- returned. He was suddenly struck with a "deep emotional gut feeling" of his own self becoming integrated, as if his machine had completed its reboot and was now waiting for an application. "I had a personal experience of what growing a mind feels like," he told me. Just as he had seen life in a computer, he now had a visceral appreciation of his own life being in a machine. Surely, life must be independent of its matrix? Couldn't life in both his body and his computer be the same?

Wouldn't it be great, he thought, if he could get something alive with evolution going in a computer! He thought he would start with human culture. That seemed an easier simulation to start with than simulated cells and DNA. As a senior at the University of Arizona, Langton wrote a paper on "The Evolution of Culture." He wanted his anthropology, physics, and computer science professors to let him design a degree around building a computer to run artificial evolution, but they discouraged him. On his own he bought an Apple II and wrote his first artificial world. He couldn't get self-reproduction or natural selection, but he did discover the literature of cellular automata -- of which the Game of Life, it turned out, was only one example.

And he came across John von Neumann's proofs of artificial self-replication from the 1940s. Von Neumann had come up with a landmark formula that would self-replicate. But the program was unwieldy, inelegantly large and clumsy. Langton spent months of long nights coding his Apple II (a handy advantage that von Neumann didn't have; he did his with pencil on paper). Eventually guided only by his dream to create life in silicon, Langton came up with the smallest self-replicating machine then known to anyone. On the computer screen the self-replicator looked like a small blue Q. Langton was able to pack into its loop of only 94 symbols a complete representation of the loop, instructions on how to reproduce, and the trick of throwing off another just like itself. He was delirious. If he could engineer such a simple replicator, how many of life's other essential processes could he also mimic? Indeed, what were life's other essential processes?

A thorough search of the existing literature showed that very little science had been written on such a simple question, and what little there was, was scattered here and there in hundreds of tiny corners. Emboldened by his new research position at the Los Alamos Labs, in 1987 Langton staked his career on gathering an "Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems," -- the first conference on what Langton was now calling Artificial Life. In his search for any and all systems that exhibit the behavior of living systems, Langton opened the workshop to chemists, biologists, computer scientists, mathematicians, material scientists, philosophers, roboticists, and computer animators. I was one of the few journalists attending.

At the workshop Langton began with his quest for a definition of life. Existing ones seemed inadequate. As more research was started over the years following the first conference, physicist Doyne Farmer proposed a list of traits that defined life. Life, he said, has:

- Patterns in space and time
- Self-reproduction
- Information storage of its self-representation (genes)
- Metabolism, to keep the pattern persisting
- Functional interactions -- it does stuff
- Interdependence of parts, or the ability to die
- Stability under perturbations
- Ability to evolve.

The list provokes. For although we do not consider computer viruses alive, computer viruses satisfy most of the qualifications above. They are a pattern that reproduce; they include a copy of their own representation; they capture computer metabolistic (CPU) cycles; they can die; and they can evolve. We could say that computer viruses are the first examples of emergent artificial life.

On the other hand, we all know of a few things whose aliveness we don't doubt yet are exceptions to this list. A mule can not self-reproduce, and a herpes virus has no metabolism. Langton's success in creating a self-reproducing entity made him skeptical of arriving at a consensus: "Every time we succeed in synthetically satisfying the definition of life, the definition is lengthened or changed. For instance if we take Gerald Joyce's definition of life -- a self-sustaining chemical system capable of undergoing Darwinian evolution -- I believe that by the year 2000 one lab somewhere in the world will make a system satisfying this definition. But then biologists will merely redefine life."

Langton had better luck defining artificial life. Artificial life, or "a-life" in short hand, is, he said, "the attempt to abstract the logic of life in different material forms." His thesis was that life is a process -- a behavior that is not bound to a specific material manifestation. What counts about life is not the stuff it is made of, but what it does. Life is a verb not a noun. Farmer's list of qualifications for life represent actions and behaviors. It is not hard for computer scientists to think of the list of life's qualities as varieties of processing. Steen Rasmussen, a colleague of Langton who was also interested in artificial life, once dropped a pencil onto the desk and sighed, "In the West we think a pencil is more real than its motion."

If the pencil's motion is the essence -- the real part -- then "artificial" is a deceptive word. At the first Artificial Life Conference, when Craig Reynolds showed how he was able to use three simple rules to get dozens of computer-animated birds to flock in the computer autonomously, everyone could see that the flocking was real. Here were artificial birds really flocking. Langton summarized the lesson: "The most important thing to remember about a-life is that the part that is artificial is not the life, but the materials. Real things happen. We observe real phenomena. It is real life in an artificial medium."

Biology -- the study of life's general principles -- is undergoing an upheaval. Langton says biology faces "the fundamental obstacle that it is impossible to derive general principles from single examples." Since we have only a single collective example of life on Earth, it is pointless to try to distinguish its essential and universal properties from those incidental properties due to life's common descent on the planet. For instance, how much of what we think life is, is due to its being based on carbon chains? We can't know without at least a second example of life not based on carbon chains. To derive general principles and theories of life -- that is, to identify properties that would be shared by any vivisystem or any life -- Langton argues that "we need an ensemble of instances to generalize over. Since it is quite unlikely that alien life-forms will present themselves to us for study in the near future, our only option is to try to create alternative life-forms ourselves." This is Langton's mission -- to create an alternative life, or maybe even several alternative "lifes," as a basis for a true biology, a true logic of Bios. Since these other lifes are artifacts of humans rather than nature, we call them artificial life; but they are as real as we are.

The nature of this ambitious challenge initially sets the science of artificial life apart from the science of biology. Biology seeks to understand the living by taking it apart and reducing it to it pieces. Artificial life, on the other hand, has nothing to dissect, so it can only make progress by putting the living together and assembling it from pieces. Rather than analyze life, synthesize it. For this reason, Langton says, "Artificial life amounts to the practice of synthetic biology."

Artificial life acknowledges new lifes and a new definition of life. "New" life is an old force that organizes matter and energy in new ways. Our ancient ancestors were often generous in deeming things alive. But in the age of science, we make a careful distinction. We call creatures and green plants alive, but when we call an institution such as the post office an "organism," we say it is lifelike or "as if it were alive."

We (and by this I mean scientists first) are beginning to see that those organizations once called metaphorically alive are truly alive, but animated by a life of a larger scope and wider definition. I call this greater life "hyperlife." Hyperlife is a particular type of vivisystem endowed with integrity, robustness, and cohesiveness -- a strong vivisystem rather than a lax one. A rain forest and a periwinkle, an electronic network and a servomechanism, SimCity and New York City, all possess degrees of hyperlife. Hyperlife is my word for that class of life that includes both the AIDS virus and the Michelangelo computer virus.

Biological life is only one species of hyperlife. A telephone network is another species. A bullfrog is chock-full of hyperlife. The Biosphere 2 project in Arizona swarms with hyperlife, as do Tierra, and Terminator 2. Someday hyperlife will blossom in automobiles, buildings, TVs, and test tubes.

This is not to say that organic life and machine life are identical; they are not. Water striders will forever retain certain characteristics unique to carbon-based life. But organic and artificial life share a set of characteristics that we have only begun to discern. And of course there easily may be other types of hyperlife to come that we can't describe yet. One can imagine various possibilities of life -- weird hybrids bred from both biological and synthetic lines, the half-animal/half-machine cyborgs of old science fiction -- that may have emergent properties of hyperlife not found in either parent.

Man's every attempt to create life is a probe into the space of possible hyperlifes. This space includes all endeavors to re-create the origins of life on Earth. But the challenge goes way beyond that. The goal of artificial life is not to merely describe the space of "life-as-we-know-it." The quest that fires up Langton is the hope of mapping the space of all possible lifes, a quest that moves us into the far, far vaster realm of "life-as-it-could-be." Hyperlife is that library which contains all things alive, all vivisystems, all slivers of life, anything bucking the second law of thermodynamics, all future and all past arrangements of matter capable of open-ended evolution, and all examples of a type of something marvelous we can't really define yet.

The only way to explore this terra incognita is to build many examples and see if they fit in the space. As Langton wrote in his introduction to the proceedings of the Second Artificial Life conference, "If biologists could 're-wind the tape' of evolution and start it over, again and again, from different initial conditions, or under different regimes of external perturbations along the way, they would have a full ensemble of evolutionary pathways to generalize over." Keep starting from zero, alter the rules a bit and then build an example of artificial life. Do it dozens of times. Each instance of synthetic life is added to the example of Earth-bound organic life to form the complete ensemble of hyperlife.

Since life is a property of form, and not matter, the more materials we can transplant living behaviors into, the more examples of "life-as-it-could-be" we can accumulate. Therefore the field of artificial life is broad and eclectic in considering all avenues to complexity. A typical gathering of a-life researchers includes biochemists, computer wizards, game designers, animators, physicists, math nerds, and robot hobbyists. The hidden agenda is to hack the definition of life.

One evening after a late-night lecture session at the First Artificial Life Conference, while some of us watched the stars in the desert night sky, mathematician Rudy Rucker came up with the most expansive motivation for artificial life I've heard: "Right now an ordinary computer program may be a thousand lines and take a few minutes to run. Artificial life is about finding a computer code that is only a few lines long and that takes a thousand years to run."

That seems about right. We want the same in our robots: Design them for a few years and then have them run for centuries, perhaps even manufacturing their replacements. That's what an acorn is too -- a few lines of code that run out as a 180-year-old tree.

The conference-goers felt the important thing about artificial life was that it not only was redefining biology and life, but it was also redefining the concept of both artificial and real. It was radically enlarging the realm of what seemed important -- that is, the realm of life and reality. Unlike the "publish or perish" mode of academic professionalism of yesteryear, most of the artificial life experimenters -- even the mathematicians -- espoused the emerging new academic creed of "demo or die." The only way to make a dent in artificial and hyperlife was to get a working example up and running. Explaining how he got started in life-as-it-could-be, Ken Karakotsios, a former Apple employee, recalled, "Every time I met a computer I tried to program the Game of Life into it." This eventually led to a remarkable Macintosh a-life program called SimLife. In SimLife you create a hyperlife world and set loose little creatures into it to coevolve into a complexifying artificial ecology. Now Karakotsios seeks to write the biggest and best game of life, an ultimate living program: "You know, the universe is the only thing big enough to run the ultimate game of life. The only problem with the universe as a platform, though, is that it is currently running someone else's program."

Larry Yaeger, a current Apple employee, once handed me his business card. It ran: "Larry Yaeger, Microcosmic God." Yaeger created Polyworld, a sophisticated computer world with organisms in the shape of polygons. The polys fly around by the hundreds, mating, breeding, consuming resources, learning (a power God Yaeger gave them), adapting, and evolving. Yaeger was exploring the space of possible life. What would appear? "At first," said Yaeger, "I did not charge the parents an energy cost when offspring was born. They could have offspring for free. But I kept getting this particular species, these indolent cannibals, who liked to hang around the corner in the vicinity of their parents and children and do nothing, never leave. All they would do was mate with each other, fight with each other, and eat each other. Hey, why work when you can eat your kids!" Life of some hyper-type had appeared.

"A central motivation for the study of artificial life is to extend biology to a broader class of life forms than those currently present on the earth," writes Doyne Farmer, understating the sheer, great fun artificial life gods are having.

But Farmer is onto something. Artificial life is unique among other human endeavors for yet another reason. Gods such as Yaeger are extending the class of life because life-as-it-could-be is a territory we can only study by first creating it. We must manufacture hyperlife to explore it, and to explore it we must manufacture it.

But as we busily create ensembles of new forms of hyperlife, an uneasy thought creeps into our minds. Life is using us. Organic carbon-based life is merely the first, earliest form of hyperlife to evolve into matter. Life has conquered carbon. But now under the guise of pond weed and kingfisher, life seethes to break out into crystal, into wires, into biochemical gels, and into hybrid patches of nerve and silicon. If we look at where life is headed, we have to agree with developmental biologist Lewis Held when he said, "Embryonic cells are just robots in disguise." In his report for the proceedings of Second Artificial Life Conference Tom Ray wrote, "Virtual life is out there, waiting for us to create environments for it to evolve into." Langton told Steven Levy, reporting in Artificial Life, "There are these other forms of life, artificial ones, that want to come into existence. And they are using me as a vehicle for its reproduction and its implementation."

Life -- the hyperlife -- wants to explore all possible biologies and all possible evaluations, but it uses us to create them because to create them is the only way to explore or complete them. Humanity is thus, depending on how you look at it, a mere passing station on hyperlife's gallop through space, or the critical gateway to the open-ended universe.

"With the advent of artificial life, we may be the first species to create its own successors," Doyne Farmer wrote in his manifesto, Artificial Life: The Coming Evolution. "What will these successors be like? If we fail in our task as creators, they may indeed be cold and malevolent. However, if we succeed, they may be glorious, enlightened creatures that far surpass us in their intelligence and wisdom." Their intelligence might be "inconceivable to lower forms of life such as us." We have always been anxious about being gods. If through us, hyperlife should find spaces where it evolves creatures that amuse and help us, we feel proud. But if superior successors should ascend through our efforts, we feel fear.

Chris Langton's office sat catty-corner to the atomic museum in Los Alamos, a reminder of the power we have to destroy. That power stirred Langton. "By the middle of this century, mankind had acquired the power to extinguish life," he wrote in one of his academic papers. "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 responsibilities on our shoulders."

Here and there we create space for other varieties of life to emerge. Juvenile delinquent hackers launch potent computer viruses. Japanese industrialists weld together smart painting robots. Hollywood directors create virtual dinosaurs. Biochemists squeeze self-evolving molecules into tiny plastic test tubes. Someday, we will create an open-ended world that can keep going, and keep creating perpetual novelty. When we do we will have created another living vector in the life space.

When Danny Hillis says he wants to make a computer that would be proud of him, he isn't kidding. What could be more human than to give life? I think I know: to give life and freedom. To give open-ended life. To say, here's your life and the car keys. Then you let it do what we are doing -- making it all up as we go along. Tom Ray once told me, "I don't want to download life into computers. I want to upload computers into life."

The butterfly sleeps

Some ideas are reeled into our mind wrapped up in facts; and some ideas burst upon us naked without the slightest evidence they could be true but with all the conviction they are. The ideas of the latter sort are the more difficult to displace.

The idea of antichaos -- order for free -- came in a vision of the unverifiable sort.

The idea was dealt to Stuart Kauffman, an undergraduate medical student at Dartmouth College some thirty years ago. As Kauffman remembers it, he was standing in front of a bookstore window daydreaming about the design of a chromosome. Kauffman was a sturdy guy with curly hair, easy smile, and no time to read. As he stared in the window, he imagined a book, a book with his name on it in the author's slot, a book that he would write in the future.

In his vision the pages of the book were filled with a web of arrows connecting other arrows, weaving in and out of a living tangle. It was the icon of the Net. But the mess was not without order. The tangle sparked mysterious, almost cabalistic, "currents of meanings" along the threads. Kauffman discerned an image emerging out of the links in a "subterranean way," just as recognition of a face springs from the crazy disjointed surfaces in a cubist painting.

As a medical student studying cell development, Kauffman saw the intertwined lines in his fantasy as the interconnections between genes. Out of that random mess, Kauffman suddenly felt sure, would come inadvertent order -- the architecture of an organism. Out of chaos would come order for no reason: order for free. The complexity of points and arrows seemed to be generating a spontaneous order. To Kauffman the depiction was intimately familiar; it felt like home. His task would be to explain and prove it. "I don't know why this question, this ill-lit path," he says, but it has become a "deeply felt, deeply held image."

Kauffman pursued his vision by taking up academic research in cell development. As many other developmental biologists had, he studied Drosophila, the famous fruit fly, as it progressed from fertilized egg to adult. How did the original lone egg cell of any creature manage to divide and specialize first into two, then four, then eight new kinds of cells? In a mammal the original egg cell would propagate an intestinal cell line, a brain cell line, a hair cell line; yet each substantially specialized line of cells presumably ran the same operating software. After a relatively few generations of division, one cell type could split into all the variety and bulk of an elephant or oak. A human embryo egg needed to divide only 50 times to produce the trillions of cells that form a baby.

What invisible hand controlled the fate of each cell, as it traveled along a career path forking 50 times, guiding it from general egg to hundreds of kinds of specialized cells? Since each cell was supposedly driven by identical genes (or were they actually different?), how could cells possibly become different? What controlled the genes?

Françoise Jacob and Jacques Monod discovered a major clue in 1961 when they encountered and described the regulatory gene. The regulatory gene's function was stunning: to turn other genes on. In one breath it blew away all hopes of immediately understanding DNA and life. The regulatory gene set into motion the quintessential cybernetic dialogue: What controls genes? Other genes! And what controls those genes? Other genes! And what...

That spiraling, darkly modern duet reminded Kauffman of his home image. Some genes controlling other genes which in turn might control still others was the same tangled web of arrows of influence pointing in every direction in his vision book.

Jacob and Monod's regulatory genes reflected a spaghetti-like vision of governance -- a decentralized network of genes steering the cellular network to its own destiny. Kauffman was excited. His picture of "order for free" suggested to him a fairly far-out idea: that some of the differentiation (order) each egg underwent was inevitable, no matter what genes you started out with!

He could think of a test for this notion. Replace all the genes in the fruitfly with random genes. His bet: you would not get Drosophila, but you would get the same order of monsters and freak mutations Drosophila produced in the natural course of things. "The question I asked myself," Kauffman recalls, "was the following. If you just hooked up genes at random, would you get anything that looked useful?" His intuitive hunch was that simply because of distributed bottom-up control and everything-is-connected-to-everything type of cell management, certain classes of patterns would be inevitable. Inevitable! Now here was a germ of heresy. Something to devote one's years to!

"I had a hard time in medical school," he continues, "because instead of studying anatomy I was scribbling all these notebooks with little model genomes." The way to prove this heresy, Kauffman cleverly decided, was not to fight nature in the lab, but to model it mathematically. Use computers as they became accessible. Unfortunately there was no body of math with the ability to track the horizontal causality of massive swarms. Kauffman began to invent his own. At the same time (about 1970) in about a half-dozen other fields of research, the mathematically inclined (such as John Holland) were coming up with procedures that allowed them to simulate the effects of a mob of interdependent nodes whose values simultaneously depend on each other.

This set of math techniques that Kauffman, Holland and others devised is still without a proper name, but I'll call it here "net math." Some of the techniques are known informally as parallel distributed processing, Boolean nets, neural nets, spin glasses, cellular automata, classifier systems, genetic algorithms, and swarm computation. Each flavor of net math incorporates the lateral causality of thousands of simultaneous interacting functions. And each type of net math attempts to coordinate massively concurrent events -- the kind of nonlinear happenings ubiquitous in the real world of living beings. Net math is in contradistinction to Newtonian math, a classical math so well suited to most physics problems that it had been seen as the only kind of math a careful scientist needed. Net math is almost impossible to use practically without computers.

The wide variety of swarm systems and net maths got Kauffman to wondering if this kind of weird swarm logic -- and the inevitable order he was sure it birthed -- were more universal than special. For instance, physicists working with magnetic material confronted a vexing problem. Ordinary ferromagnets -- the kind clinging to refrigerator doors and pivoting in compasses -- have particles that orient themselves with cultlike uniformity in the same direction, providing a strong magnetic field. Mildly magnetic "spin glasses," on the other hand, have wishy-washy particles that will magnetically "spin" in a direction that depends in part on which direction their neighbors spin. Their "choice" places more clout on the influence of nearby ones, but pays some attention to distant particles. Tracing the looping interdependent fields of this web produces the familiar tangle of circuits in Kauffman's home image. Spin glasses used a variety of net math to model the material's nonlinear behavior that was later found to work in other swarm models. Kauffman was certain genetic circuitry was similar in its architecture.

Unlike classical mathematics, net math exhibits nonintuitive traits. In general, small variations in input in an interacting swarm can produce huge variations in output. Effects are disproportional to causes -- the butterfly effect.

Even the simplest equations in which intermediate results flow back into them can produce such varied and unexpected turns that little can be deduced about the equations' character merely by studying them. The convoluted connections between parts are so hopelessly tangled, and the calculus describing them so awkward, that the only way to even guess what they might produce is to run the equations out, or in the parlance of computers, to "execute" the equations. The seed of a flower is similarly compressed. So tangled are the chemical pathways stored in it, that inspection of a unknown seed -- no matter how intelligent -- cannot predict the final form of the unpacked plant. The quickest route to describing a seed's output is therefore to sprout it.

Equations are sprouted on computers. Kauffman devised a mathematical model of a genetic system that could sprout on a modest computer. Each of the 10,000 genes in his simulated DNA is a teeny-weeny bit of code that can turn other genes either on or off. What the genes produced and how they were connected were assigned at random.

This was Kauffman's point: that the very topology of such complicated networks would produce order -- spontaneous order! -- no matter what the tasks of the genes.

While he worked on his simulated gene, Kauffman realized that he was constructing a generic model for any kind of swarm system. His program could model any bunch of agents that interact in a massive simultaneous field. They could be cells, genes, business firms, black boxes, or simple rules -- anything that registers input and generates output interpreted as input by a neighbor.

He took this swarm of actors and randomly hooked them up into an interacting network. Once they were connected he let them bounce off one another and recorded their behavior. He imagined each node in the network as a switch able to turn certain neighboring nodes off or on. The state of the neighbor nodes looped back to regulate the initial node. Eventually this gyrating mess of he-turns-her-who-turns-him-on settled down into a stable and measurable state. Kauffman again randomly rearranged the entire net's connections and let the nodes interact until they all settled down. He did that many times, until he had "explored" the space of possible random connections. This told him what the generic behavior of a net was, independent of its contents. An oversimplified analogous experiment would be to take ten thousand corporations and randomly link up the employees in each by telephone networks, and then measure the average effects of these networks, independent of what people said over them.

By running these generic interacting networks tens of thousands of times, Kauffman learned enough about them to paint a rough portrait of how such swarm systems behaved under specific circumstances. In particular, he wanted to know what kind of behavior a generic genome would create. He programmed thousands of randomly assembled genetic systems and then ran these ensembles on a computer -- genes turning off and on and influencing each other. He found they fell into "basins" of a few types of behaviors.

At a slow speed water trickles out of a garden hose in one uneven but consistent pattern. Turn up the tap, and it abruptly sprays out in a chaotic (but describable) torrent. Turn it up full blast, and it gushes out in a third way like a river. Carefully screw the tap to the precise line between one speed and a slower one, and the pattern refuses to stay on the edge but reverts to one state or the other, as if it were attracted to a side, any side. Just as a drop of rain falling on the ridge of a continental divide must eventually find its way down to either the Pacific Basin or the Atlantic Basin, roll down one side or the other it must.

Sooner or later the dynamics of the system would find its way to at least one "basin" that entrapped the shifting motions into a persistent pattern. In Kauffman's view a randomly assembled system would find its way to a stock pattern (a basin); thus, out of chaos, order for free emerges.

As he ran uncounted genetic simulations, Kauffman discovered a rough ratio (the square root) between the number of genes and the number of basins the genes in the system settled into. This proportion was the same as the number of genes in biological cells and the number of cell types (liver cells, blood cells, brain cells) those genes created, a ratio that is roughly constant in all living things.

Kauffman claims this universal ratio across many species suggests that the number of cell types in nature may derive from cellular architecture itself. The number of types of cells in your body, then, may have little to do with natural selection and more to do with the mathematics of complex gene interactions. How many other biological forms, Kauffman gleefully wonders, might also owe little to selection?

He had a hunch about a way to ask the question experimentally. But first he needed a method to cook up random ensembles of life. He decided to simulate the origin of life by generating all possible pools of prelife parts -- at least in simulation. He would let the virtual pool of parts interact randomly. If he could then show that out of this soup order inevitably emerged, he would have a case. The trick would be to allow molecules to converge into a lap game.

The lap game peaked in popularity a decade ago. It is a spectacular outdoor game that advertises the power of cooperation. The facilitator of the lap game takes a group of 25 or more people and has them stand fairly close together in a circle, so that each participant is staring at the back of the head of the person in front of him. Just picture a queue of people waiting in line for a movie and connect them in a tidy circle.

At the facilitator's command this circle of people bend their knees and sit on the spontaneously generated knee-lap of the person behind them. If done in unison, the ring of people lowering to sit are suddenly propped up on a self-supporting collective chair. If one person misses the lap behind him, the whole circling line crashes. The world's record for a stable lap game is several hundred people.

Auto-catalytic sets and the selfish Uroborus snake circle are much like lap games. Compound (or function) A makes compound (or function) B with the aid of compound (or function) C. But C itself is produced by A and D. And D is generated by E and C, and so on. Without the others none can be. Another way of saying this is to state that the only way for a particular compound or function to survive in the long run is for it to be a product of another compound or function. In this circular world all causes are results, just as all knees are laps. Contrary to common sense, all existences depend on the consensual existence of all others.

As the reality of the lap game proves, however, circular causality is not impossible. Tautology can hold up 200 pounds of flesh. It's real. Tautology is, in fact, an essential ingredient of stable systems.

Cognitive philosopher Douglas Hofstadter calls these paradoxical circuits "Strange Loops." As examples, Hofstadter points to the seemingly ever rising notes in a Bach canon, or the endlessly rising steps in an Escher staircase. He also includes as Strange Loops the famous paradox about Cretan liars who say they never lie, and Gödel's proof of unprovable mathematical axioms. Hofstadter writes in Gödel, Escher, Bach: "The 'Strange Loop' phenomenon occurs whenever, by moving upwards (or downwards) through the levels of some hierarchical system, we unexpectedly find ourselves right back where we started."

Life and evolution entail the necessary strange loop of circular causality -- of being tautological at a fundamental level. You can't get life and open-ended evolution unless you have a system that contains that essential logical inconsistency of circling causes. In complex adapting processes such as life, evolution, and consciousness, prime causes seem to shift, as if they were an optical illusion drawn by Escher. Part of the problem humans have in trying build systems as complicated as our own human biology is that in the past we have insisted on a degree of logical consistency, a sort of clockwork logic, that blocks the emergence of autonomous events. But as the mathematician Gödel showed, inconsistency is an inevitable trait of any self-sustaining system built up out of consistent parts.

Gödel's 1931 theorem demonstrates, among other things, that attempts to banish self-swallowing loopiness are fruitless, because, in Hofstadter's words, "it can be hard to figure out just where self-referencing is occurring." When examined at a "local" level every part seems legitimate; it is only when the lawful parts form a whole that the contradiction arises.

In 1991, a young Italian scientist, Walter Fontana, showed mathematically that a linear sequence of function A producing function B producing function C could be very easily circled around and closed in a cybernetic way into a self-generating loop, so that the last function was coproducer of the initial function. When Kauffman first encountered Fontana's work he was ecstatic with the beauty of it. "You have to fall in love with it! Functions mutually making one another. Out of all function space, they come gripping one another's arms in an embrace of creating!" Kauffman called such a autocatalytic set an "egg." He said, "An egg would be a set of rules having the property that the rules they pose are precisely the ones that create them. That's really not crazy at all."

To get an egg you start with a huge pool of different agents. They could be varieties of protein pieces or fragments of computer code. If you let them interact upon each other long enough, they will produce small loops of thing-producing-other things. Eventually, if given time and elbowroom the spreading network of these local loops in the system will crowd upon itself, until every producer in the circuit is a product of another, until every loop is incorporated into all the other loops in massively parallel interdependence. At this moment of "catalytic closure" the web of parts suddenly snaps into a stable game -- the system sits in its own lap, with its beginning resting on its end, and vice versa.

Life began in such a soup of "polymers acting on polymers to form new polymers," Kauffman claims. He demonstrated the theoretical feasibility of such a logic by running experiments of "symbol strings acting on symbol strings to form new symbol strings." His assumption was that he could equate protein fragments and computer code fragments as logical equivalents. When he ran networks of bits of code-which-produce-code as a model for proteins, he got autocatalytic systems that are circular in the sense of the lap game: they have no beginning, no center, and no end.

Life popped into existence as a complete whole much as a crystal suddenly appears in its final (though miniature) form in a supersaturated solution: not beginning as a vague half-crystal, not appearing as a half-materialized ghost, but wham, being all at once, just as a lap game circle suddenly emerges from a curving line of 200 people. "Life began whole and integrated, not disconnected and disorganized," writes Stuart Kauffman. "Life, in a deep sense, crystallized."

He goes on to say, "I hope to show that self-reproduction and homeostasis, basic features of organisms, are natural collective expressions of polymer chemistry. We can expect any sufficiently complex set of catalytic polymers to be collectively autocatalytic." Kauffman was creeping up on that notion of inevitability again. "If my model is correct then the routes to life in the universe are boulevards, rather than twisted back alleyways." In other words, given the chemistry we have, "life is inevitable."

"We've got to get used to dealing in billions of things!" Kauffman once told an audience of scientists. Huge multitudes of anything are different: the more polymers, the exponentially more possible interactions where one polymer can trigger the manufacture of yet another polymer. Therefore, at some point, a droplet loaded up with increasing diversity and numbers of polymers will reach a threshold where a certain number of polymers in the set will suddenly fall out into a spontaneous lap circle. They will form an auto-generated, self-sustaining, self-transforming network of chemical pathways. As long as energy flows in, the network hums, and the loop stands.

Codes, chemicals, or inventions can in the right circumstances produce new codes, chemicals, or inventions. It is clear this is the model of life. An organism produces new organisms which in turn create newer organisms. One small invention (the transistor) produces other inventions (the computer) which in turn permit yet other inventions (virtual reality). Kauffman wants to generalize this process mathematically to say that functions in general spawn newer functions which in turn birth yet other functions.

"Five years ago," recalls Kauffman, "Brian Goodwin [an evolutionary biologist] and I were sitting in some World War I bunker in northern Italy during a rainstorm talking about autocatalytic sets. I had this profound sense then that there's a deep similarity between natural selection -- what Darwin told us -- and the wealth of nations -- what Adam Smith told us. Both have an invisible hand. But I didn't know how to proceed any further until I saw Walter Fontana's work with autocatalytic sets, which is gorgeous."

I mentioned to Kauffman the controversial idea that in any society with the proper strength of communication and information connection, democracy becomes inevitable. Where ideas are free to flow and generate new ideas, the political organization will eventually head toward democracy as an unavoidable self-organizing strong attractor. Kauffman agreed with the parallel: "When I was a sophomore in '58 or '59 I wrote a paper in philosophy that I labored over with much passion. I was trying to figure out why democracy worked. It's obvious that democracy doesn't work because it's the rule of the majority. Now, 33 years later, I see that democracy is a device that allows conflicting minorities to reach relative fluid compromises. It keeps subgroups from getting stuck on some locally good but globally inferior solution."

It is not difficult to imagine Kauffman's networks of Boolean logic and random genomes mirroring the workings of town halls and state capitals. By structuring miniconflicts and microrevolutions as a continuous process at the local level, large scale macro- and mega-revolutions are avoided, and the whole system is neither chaotic nor stagnant. Perpetual change is fought out in small towns, while the nation remains admirably stable -- thus creating a climate to keep the small towns in ceaseless compromise-seeking modes. That circular support is another lap game, and an indication that such systems are similar in dynamics to the self-supporting vivisystems.

"This is just intuitive," Kauffman cautions me, "but you can feel your way from Fontana's 'string-begets-string-begets-string' to 'invention-begets-invention-begets-invention' to cultural evolution and then to the wealth of nations." Kauffman makes no bones about the scale of his ambition: "I am looking for the self-consistent big picture that ties everything together, from the origin of life, as a self-organized system, to the emergence of spontaneous order in genomic regulatory systems, to the emergence of systems that are able to adapt, to nonequilibrium price formation which optimizes trade among organisms, to this unknown analog of the second law of thermodynamics. It is all one picture. I really feel it is. But the image I'm pushing on is this: Can we prove that a finite set of functions generates this infinite set of possibilities?"

Whew. I call that a "Kauffman machine." A small but well-chosen set of functions that connect into an auto-generating ring and produce an infinite jet of more complex functions. Nature is full of Kauffman machines. An egg cell producing the body of a whale is one. An evolution machine generating a flamingo over a billion years from a bacterial blob is another. Can we make an artificial Kauffman machine? This may more properly be called a von Neumann machine because von Neumann asked the same question in the early 1940s. He wondered, Can a machine make another machine more complex that itself? Whatever it is called, the question is the same: How does complexity build itself up?

"You can't ask the experimental question until, roughly speaking, the intellectual framework is in place. So the critical thing is asking important questions," Kauffman warned me. Often during our conversations, I'd catch Kauffman thinking aloud. He'd spin off wild speculations and then seize one and twirl it around to examine it from various directions. "How do you ask that question?" he asked himself rhetorically. His quest was for the Question of All Questions rather than the Answer of All Answers. "Once you've asked the question," he said, "there's a good chance of finding some sort of answer.

A Question Worth Asking. That's what Kauffman thought of his notion of self-organized order in evolutionary systems. Kauffman confided to me: "Somehow, each of us in our own heart is able to ask questions that we think are profound in the sense that the answer would be truly important. The enormous puzzle is why in the world any of us ask the questions that we do."

There were many times when I felt that Stuart Kauffman, medical doctor, philosopher, mathematician, theoretical biologist, and MacArthur Award recipient, was embarrassed by the wild question he had been dealt. "Order for free" flies in the face of a conservative science that has rejected every past theory of creative order hidden in the universe. It would probably reject his. While the rest of the contemporary scientific world sees butterflies of random chance sowing out-of-control, nonlinear effects in every facet of the universe, Kauffman asks if perhaps the butterflies of chaos sleep. He wakes the possibility of an overarching design dwelling within creation, quieting disorder and birthing an ordered stillness. It's a notion that for many sounds like mysticism. At the same time, the pursuit and framing of this single huge question is the quasar source of Kauffman's considerable pride and energy: "I would be lying if I didn't tell you that when I was 23 and started wondering how in the world a genome with 100,000 genes controls the emergence of different cell types, I felt that I had found something profound, I had found a profound question. And I still feel that way. I think God was very nice to me."

"If you write something about this," Kauffman says softly, "make sure you say that this is only something crazy that people are thinking about. But wouldn't it be wonderful if somehow there are laws that make laws that make laws, so that the universe is, in John Wheeler's words, something that is looking in at itself!? The universe posts its own rules and emerges out of a self-consistent thing. Maybe that's not impossible, this notion that quarks and gluons and atoms and elementary particles have invented the laws by which they transform one another."

Deep down Kauffman felt that his systems built themselves. In some way he hoped to discover, evolutionary systems controlled their own structure. From the first glimpse of his visionary network image, he had a hunch that in those connections lay the answer to evolution's self-governance. He was not content to show that order emerged spontaneously and inevitably. He also felt that control of that order also emerged spontaneously. To that end he charted thousands of runs of random ensembles in computer simulation to see which type of connections permitted a swarm to be most adaptable. "Adaptable" means the ability of system to adjust its internal links so that it fits its environment over time. Kauffman views an organism, a fruitfly say, as adjusting the network of its genes over time so that the result of the genetic network -- a fly body -- best fits its changing surroundings of food, shelter, and predators. The Question Worth Asking was: what controlled the evolvability of the system? Could the organism itself control its evolvability?

The prime variable Kauffman played with was the connectivity of the network. In a sparsely connected network, each node would on average only connect to one other node, or less. In a richly connected network, each node would link to ten or a hundred or a thousand or a million other nodes. In theory the limit to the number of connections per node is simply the total number of nodes, minus one. A million-headed network could have a million-minus-one connections at each node; every node is connected to every other node. To continue our rough analogy, every employee of GM could be directly linked to all 749,999 other employees of GM.

As Kauffman varied this connectivity parameter in his generic networks, he discovered something that would not surprise the CEO of GM. A system where few agents influenced other agents was not very adaptable. The soup of connections was too thin to transmit an innovation. The system would fail to evolve. As Kauffman increased the average number of links between nodes, the system became more resilient, "bouncing back" when perturbed. The system could maintain stability while the environment changed. It would evolve. The completely unexpected finding was that beyond a certain level of linking density, continued connectivity would only decrease the adaptability of the system as a whole.

Kauffman graphed this effect as a hill. The top of the hill was optimal flexibility to change. One low side of the hill was a sparsely connected system: flat-footed and stagnant. The other low side was an overly connected system: a frozen grid-lock of a thousand mutual pulls. So many conflicting influences came to bear on one node that whole sections of the system sank into rigid paralysis. Kauffman called this second extreme a "complexity catastrophe." Much to everyone's surprise, you could have too much connectivity. In the long run, an overly linked system was as debilitating as a mob of uncoordinated loners.

Somewhere in the middle was a peak of just-right connectivity that gave the network its maximal nimbleness. Kauffman found this measurable "Goldilocks'" point in his model networks. His colleagues had trouble believing his maximal value at first because it seemed counterintuitive at the time. The optimal connectivity for the distilled systems Kauffman studied was very low, "somewhere in the single digits." Large networks with thousands of members adapted best with less than ten connections per member. Some nets peaked at less than two connections on average per node! A massively parallel system did not need to be heavily connected in order to adapt. Minimal average connection, done widely, was enough.

Kauffman's second unexpected finding was that this low optimal value didn't seem to fluctuate much, no matter how many members comprised a specific network. In other words, as more members were added to the network, it didn't pay (in terms of systemwide adaptability) to increase the number of links to each node. To evolve most rapidly, add members but don't increase average link rates. This result confirmed what Craig Reynolds had found in his synthetic flocks: you could load a flock up with more and more members without having to reconfigure its structure.

Kauffman found that at the low end, with less than two connections per agent or organism, the whole system wasn't nimble enough to keep up with change. If the community of agents lacked sufficient internal communication, it could not solve a problem as a group. More exactly, they fell into isolated patches of cooperative feedback but didn't interact with each other.

At the ideal number of connections, the ideal amount of information flowed between agents, and the system as a whole found the optimal solutions consistently. If their environment was changing rapidly, this meant that the network remained stable -- persisting as a whole over time.

Kauffman's Law states that above a certain point, increasing the richness of connections between agents freezes adaptation. Nothing gets done because too many actions hinge on too many other contradictory actions. In the landscape metaphor, ultra-connectance produces ultra-ruggedness, making any move a likely fall off a peak of adaptation into a valley of nonadaptation. Another way of putting it, too many agents have a say in each other's work, and bureaucratic rigor mortis sets in.

Adaptability conks out into grid-lock. For a contemporary culture primed to the virtues of connecting up, this low ceiling of connectivity comes as unexpected news.

We postmodern communication addicts might want to pay attention to this. In our networked society we are pumping up both the total number of people connected (in 1993, the global network of networks was expanding at the rate of 15 percent additional users per month!), and the number of people and places to whom each member is connected. Faxes, phones, direct junk mail, and large cross-referenced data bases in business and government in effect increase the number of links between each person. Neither expansion particularly increases the adaptability of our system (society) as a whole.

Stuart Kauffman's simulations are as rigorous, original, and well- respected among scientists as any mathematical model can be. Maybe more so, because he is using a real (computer) network to model a hypothetical network, rather than the usual reverse of using a hypothetical to model the real. I grant, though, it is a bit of a stretch to apply the results of a pure mathematical abstraction to irregular arrangements of reality. Nothing could be more irregular than online networks, biological genetic networks, or international economic networks. But Stuart Kauffman is himself eager to extrapolate the behavior of his generic test-bed to real life. The grand comparison between complex real-world networks and his own mathematical simulations running in the heart of silicon is nothing less than Kauffman's holy grail. He says his models "smell like they are true." Swarmlike networks, he bets, all behave similarly on one level. Kauffman is fond of speculating that "IBM and E. coli both see the world in the same way."

I'm inclined to bet in his favor. We own the technology to connect everyone to everyone, but those of us who have tried living that way are finding that we are disconnecting to get anything done. We live in an age of accelerating connectivity; in essence we are steadily climbing Kauffman's hill. But we have little to stop us from going over the top and sliding into a descent of increasing connectivity but diminishing adaptability. Disconnection is a brake to hold the system from overconnection, to keep our cultural system poised on the edge of maximal evolvability.

The art of evolution is the art of managing dynamic complexity. Connecting things is not difficult; the art is finding ways for them to connect in an organized, indirect, and limited way.

From his experiments in artificial life in swarm models, Chris Langton, Kauffman's Santa Fe Institute colleague, derived an abstract quality (called the lambda parameter) that predicts the likelihood that a particular set of rules for a swarm will produce a "sweet spot" of interesting behavior. Systems built upon values outside this sweet spot tend to stall in two ways. They either repeat patterns in a crystalline fashion, or else space out into white noise. Those values within the range of the lambda sweet spot generate the longest runs of interesting behavior.

By tuning the lambda parameter Langton can tune a world so that evolution or learning can unroll most easily. Langton describes the threshold between a frozen repetitious state and a gaseous noise state as a "phase transition" -- the same term physicists use to describe the transition from liquid to gas or liquid to solid. The most startling result, though, is Langton's contention that as the lambda parameter approaches that phase transition -- the sweet spot of maximum adaptability -- it slows down. That is, the system tends to dwell on the edge instead of zooming through it. As it nears the place it can evolve the most from, it lingers. The image Langton likes to raise is that of a system surfing on an endless perfect wave in slow motion; the more perfect the ride, the slower time goes.

This critical slowing down at the "edge" could help explain why a precarious embryonic vivisystem could keep evolving. As a random system neared the phase transition, it would be "pulled in" to rest at that sweet spot where it would undergo evolution and would then seek to maintain that spot. This is the homeostatic feedback loop making a lap for itself. Except that since there is little "static" about the spot, the feedback loop might be better named "homeodynamic."

Stuart Kauffman also speaks of "tuning" the parameters of his simulated genetic networks to the "sweet spot." Out of all the uncountable ways to connect a million genes, or a million neurons, some relatively few setups are far more likely to encourage learning and adaptation throughout the network. Systems balanced to this evolutionary sweet spot learn fastest, adapt more readily, or evolve the easiest. If Langton and Kauffman are right, an evolving system will find that spot on its own.

Langton discovered a clue as to how that may happen. He found that this spot teeters right on the edge of chaotic behavior. He says that systems that are most adaptive are so loose they are a hairsbreadth away from being out of control. Life, then, is a system that is neither stagnant with noncommunication nor grid-locked with too much communication. Rather life is a vivisystem tuned "to the edge of chaos" -- that lambda point where there is just enough information flow to make everything dangerous.

Rigid systems can always do better by loosening up a bit, and turbulent systems can always improve by getting themselves a little more organized. Mitch Waldrop explains Langton's notion in his book Complexity, thusly: if an adaptive system is not riding on the happy middle road, you would expect brute efficiency to push it toward that sweet spot. And if a system rests on the crest balanced between rigidity and chaos, then you'd expect its adaptive nature to pull it back onto the edge if it starts to drift away. "In other words," writes Waldrop, "you'd expect learning and evolution to make the edge of chaos stable." A self-reinforcing sweet spot. We might call it dynamically stable, since its home migrates. Lynn Margulis calls this fluxing, dynamically persistent state "homeorhesis" -- the honing in on a moving point. It is the same forever almost-falling that poises the chemical pathways of the Earth's biosphere in purposeful disequilibrium.

Kauffman takes up the theme by calling systems set up in the lambda value range "poised systems." They are poised on the edge between chaos and rigid order. Once you begin to look around, poised systems can be found throughout the universe, even outside of biology. Many cosmologists, such as John Barrow, believe the universe itself to be a poised system, precariously balanced on a string of remarkably delicate values (such as the strength of gravity, or the mass of an electron) that if varied by a fraction as insignificant as 0.000001 percent would have collapsed in its early genesis, or failed to condense matter. The list of these "coincidences" is so long they fill books. According to mathematical physicist Paul Davies, the coincidences "taken together...provide impressive evidence that life as we know it depends very sensitively on the form of the laws of physics, and on some seemingly fortuitous accidents in the actual values that nature has chosen for various particle masses, force strengths, and so on." In brief, the universe and life as we know are poised on the edge of chaos.

What if poised systems could tune themselves, instead of being tuned by creators? There would be tremendous evolutionary advantage in biology for a complex system that was auto-poised. It could evolve faster, learn more quickly, and adapt more readily. If evolution selects for a self-tuning function, Kauffman says, then "the capacity to evolve and adapt may itself be an achievement of evolution." Indeed, a self-tuning function would inevitably be selected for at higher levels of evolution. Kauffman proposes that gene systems do indeed tune themselves by regulating the number of links, size of genome, and so on, in their own systems for optimal flexibility.

Self-tuning may be the mysterious key to evolution that doesn't stop -- the holy grail of open-ended evolution. Chris Langton formally describes open-ended evolution as a system that succeeds in ceaselessly self-tuning itself to higher and higher levels of complexity, or in his imagery, a system that succeeds in gaining control over more and more parameters affecting its evolvability and staying balanced on the edge.

In Langton's and Kauffman's framework, nature begins as a pool of interacting polymers that catalyze themselves into new sets of interacting polymers in such a networked way that maximal evolution can occur. This evolution-rich environment produces cells that also learn to tune their internal connectivity to keep the system at optimal evolvability. Each step extends the stance at the edge of chaos, poised on the thin path of optimal flexibility, which pumps up its complexity. As long as the system rides this upwelling crest of evolvability, it surfs along.

What you want in artificial systems, Langton says, is something similar. The primary goal that any system seeks is survival. The secondary search is for the ideal parameters to keep the system tuned for maximal flexibility. But it is the third order search that is most exciting: the search for strategies and feedback mechanisms that will increasingly self-tune the system each step on the way. Kauffman's hypothesis is that if systems constructed to self-tune "can adapt most readily, then they may be the inevitable target of natural selection. The ability to take advantage of natural selection would be one of the first traits selected."

As Langton and colleagues explore the space of possible worlds searching for that sweet spot where life seems poised on the edge, I've heard them call themselves surfers on an endless summer, scouting for that slo-mo wave.

Rich Bageley, another Santa Fe Institute fellow, told me "What I'm looking for are things that I can almost predict, but not quite." He explained further that it was not regular but not chaotic either. Some almost-out-of-control and dangerous edge in between.

"Yeah," replied Langton who overheard our conversation. "Exactly. Just like ocean waves in the surf. They go thump, thump, thump, steady as a heartbeat. Then suddenly, WHUUUMP, an unexpected big one. That's what we are all looking for. That's the place we want to find."

Epilogue: Nine Laws of God

So how *do* you make something from nothing? From the frontiers of computer science, and the edges of biological research, and the odd corners of interdisciplinary experimentation, I have compiled Nine Laws of God governing the incubation of somethings from nothing. These nine laws are the organizing principles that can be found operating in systems as diverse as biological evolution and SimCity. Of course I am not suggesting that they are the only laws needed to make something from nothing; but out of the many observations accumulating in the science of complexity, these principles are the broadest, crispest, and most representative generalities. I believe that one can go pretty far as a god while sticking to these nine rules:

- Distribute being
- Control from the bottom up
- Cultivate increasing returns
- Grow by chunking
- Maximize the fringes
- Honor your errors
- Pursue no optima; have multiple goals
- Seek persistent disequilibrium
- Change changes itself.

Distribute being. The spirit of a beehive, the behavior of an economy, the thinking of a supercomputer, and the life in me are distributed over a multitude of smaller units (which themselves may be distributed). When the sum of the parts can add up to more than the parts, then that extra being (that something from nothing) is distributed among the parts. Whenever we find something from nothing, we find it arising from a field of many interacting smaller pieces. All the mysteries we find most interesting -- life, intelligence, evolution -- are found in the soil of large distributed systems.

Control from the bottom up. When everything is connected to everything in a distributed network, everything happens at once. When everything happens at once, wide and fast moving problems simply route around any central authority. Therefore overall governance must arise from the most humble interdependent acts done locally in parallel, and not from a central command. A mob can steer itself, and in the territory of rapid, massive, and heterogeneous change, only a mob can steer. To get something from nothing, control must rest at the bottom within simplicity.

Cultivate increasing returns. Each time you use an idea, a language, or a skill you strengthen it, reinforce it, and make it more likely to be used again. That's known as positive feedback or snowballing. Success breeds success. In the Gospels, this principle of social dynamics is known as "To those who have, more will be given." Anything which alters its environment to increase production of itself is playing the game of increasing returns. And all large, sustaining systems play the game. The law operates in economics, biology, computer science, and human psychology. Life on Earth alters Earth to beget more life. Confidence builds confidence. Order generates more order. Them that has, gets.

Grow by chunking. The only way to make a complex system that works is to begin with a simple system that works. Attempts to instantly install highly complex organization -- such as intelligence or a market economy -- without growing it, inevitably lead to failure. To assemble a prairie takes time -- even if you have all the pieces. Time is needed to let each part test itself against all the others. Complexity is created, then, by assembling it incrementally from simple modules that can operate independently.

Maximize the fringes. In heterogeneity is creation of the world. A uniform entity must adapt to the world by occasional earth-shattering revolutions, one of which is sure to kill it. A diverse heterogeneous entity, on the other hand, can adapt to the world in a thousand daily minirevolutions, staying in a state of permanent, but never fatal, churning. Diversity favors remote borders, the outskirts, hidden corners, moments of chaos, and isolated clusters. In economic, ecological, evolutionary, and institutional models, a healthy fringe speeds adaptation, increases resilience, and is almost always the source of innovations.

Honor your errors. A trick will only work for a while, until everyone else is doing it. To advance from the ordinary requires a new game, or a new territory. But the process of going outside the conventional method, game, or territory is indistinguishable from error. Even the most brilliant act of human genius, in the final analysis, is an act of trial and error. "To be an Error and to be Cast out is a part of God's Design," wrote the visionary poet William Blake. Error, whether random or deliberate, must become an integral part of any process of creation. Evolution can be thought of as systematic error management.

Pursue no optima; have multiple goals. Simple machines can be efficient, but complex adaptive machinery cannot be. A complicated structure has many masters and none of them can be served exclusively. Rather than strive for optimization of any function, a large system can only survive by "satisficing" (making "good enough") a multitude of functions. For instance, an adaptive system must trade off between exploiting a known path of success (optimizing a current strategy), or diverting resources to exploring new paths (thereby wasting energy trying less efficient methods). So vast are the mingled drives in any complex entity that it is impossible to unravel the actual causes of its survival. Survival is a many-pointed goal. Most living organisms are so many-pointed they are blunt variations that happen to work, rather than precise renditions of proteins, genes, and organs. In creating something from nothing, forget elegance; if it works, it's beautiful.

Seek persistent disequilibrium. Neither constancy nor relentless change will support a creation. A good creation, like good jazz, must balance the stable formula with frequent out-of-kilter notes. Equilibrium is death. Yet unless a system stabilizes to an equilibrium point, it is no better than an explosion and just as soon dead. A Nothing, then, is both equilibrium and disequilibrium. A Something is persistent disequilibrium -- a continuous state of surfing forever on the edge between never stopping but never falling. Homing in on that liquid threshold is the still mysterious holy grail of creation and the quest of all amateur gods.

Change changes itself. Change can be structured. This is what large complex systems do: they coordinate change. When extremely large systems are built up out of complicated systems, then each system begins to influence and ultimately change the organizations of other systems. That is, if the rules of the game are composed from the bottom up, then it is likely that interacting forces at the bottom level will alter the rules of the game as it progresses. Over time, the rules for change get changed themselves. Evolution -- as used in everyday speech -- is about how an entity is changed over time. Deeper evolution -- as it might be formally defined -- is about how the rules for changing entities over time change over time. To get the most out of nothing, you need to have self-changing rules.

These nine principles underpin the awesome workings of prairies, flamingoes, cedar forests, eyeballs, natural selection in geological time, and the unfolding of a baby elephant from a tiny seed of elephant sperm and egg.

These principles of bio-logic are now being implanted in computer chips, electronic communication networks, robot modules, pharmaceutical searches, software design, and corporate management, in order that these artificial systems may overcome their own complexity.

All complex things taken together form an unbroken continuum between the extremes of stark clockwork gears and ornate natural wilderness. The hallmark of the industrial age has been its exaltation of mechanical design. The hallmark of a neo-biological civilization is that it returns the designs of its creations toward the organic, again. But unlike earlier human societies that relied on found biological solutions -- herbal medicines, animal proteins, natural dyes, and the like -- neo-biological culture welds engineered technology and unrestrained nature until the two become indistinguishable, as unimaginable as that may first seem.

The intensely biological nature of the coming culture derives from five influences:

- Despite the increasing technization of our world, organic life -- both wild and domesticated -- will continue to be the prime infrastructure of human experience on the global scale.
- Machines will become more biological in character.
- Technological networks will make human culture even more ecological and evolutionary.
- Engineered biology and biotechnology will eclipse the importance of mechanical technology.
- Biological ways will be revered as ideal ways.

In the coming neo-biological era, all that we both rely on and fear will be more born than made. We now have computer viruses, neural networks, Biosphere 2, gene therapy, and smart cards -- all humanly constructed artifacts that bind mechanical and biological processes. Future bionic hybrids will be more confusing, more pervasive, and more powerful. I imagine there might be a world of mutating buildings, living silicon polymers, software programs evolving offline, adaptable cars, rooms stuffed with coevolutionary furniture, gnatbots for cleaning, manufactured biological viruses that cure your illnesses, neural jacks, cyborgian body parts, designer food crops, simulated personalities, and a vast ecology of computing devices in constant flux.

Yet as we unleash living forces into our created machines, we lose control of them. They acquire wildness and some of the surprises that the wild entails. This, then, is the dilemma all gods must accept: that they can no longer be completely sovereign over their finest creations.

The world of the made will soon be like the world of the born: autonomous, adaptable, and creative but, consequently, out of our control. I think that's a great bargain. Even without the control we must surrender, a neo-biological technology is far more rewarding than a world of clocks, gears, and predictable simplicity.

As complex as things are today, everything will be more complex tomorrow. The scientists and projects reported here have been concerned with harnessing the laws of design so that order can emerge from chaos, so that organized complexity can be kept from unraveling into unorganized complications, and so that something can be made from nothing.