

## Out of Control

The New Biology of Machines, Social Systems, and the Economic World



Kevin Kelly

Out of Control is a summary of what we know about self-sustaining systems, both living ones such as a tropical wetland, or an artificial one, such as a computer simulation of our planet. The last chapter of the book, "The Nine Laws of God," is a distillation of the nine common principles that all life-like systems share. The major themes of the book are:

- As we make our machines and institutions more complex, we have to make them more biological in order to manage them.
- The most potent force in technology will be artificial evolution. We are already evolving software and drugs instead of engineering them.
- Organic life is the ultimate technology, and all technology will improve towards biology.
- The main thing computers are good for is creating little worlds so that we can try out the Great Questions. Online communities let us ask the question "what is a democracy; what do you need for it?" by trying to wire a democracy up, and re-wire it if it doesn't work. Virtual reality lets us ask "what is reality?" by trying to synthesize it. And computers give us room to ask "what is life?" by providing a universe in which to create computer viruses and artificial creatures of increasing complexity. Philosophers sitting in academies used to ask the Great Questions; now they are asked by experimentalists creating worlds.
- As we shape technology, it shapes us. We are connecting everything to everything, and so our entire culture is migrating to a "network culture" and a new network economics.
- In order to harvest the power of organic machines, we have to instill in them guidelines and self-governance, and relinquish some of our total control.

## Chapter 2: HIVE MIND

Bees do it: distributed governance  
The collective intelligence of a mob  
Asymmetrical invisible hands  
Decentralized remembering as an act of perception  
More is more than more, it's different  
Advantages and disadvantages of swarms  
The network is the icon of the 21st century

### Bees do it: distributed governance

**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.

### [The collective intelligence of a mob](#)

**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! REEEEEED!" 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.

### Asymmetrical invisible hands

**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,  $2 + 2 = \text{apples}$ . "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.

#### Decentralized remembering as an act of perception

**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 *džj^ 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.

### [More is more than more, it's different](#)

**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 H<sub>2</sub>O, 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.

#### Advantages and disadvantages of swarms

**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.

### The network is the icon of the 21st century

**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, it's 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 responsiveness -- 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.