NATURE'S ALGORITHMS

Brian Hayes

We have it on good authority (Adams 1979) that the earth is a gigantic computer, built to calculate the answer to some ultimate question (or vice versa). When you look around at the world, the idea seems plausible enough. There are certainly lots of things going on that look like computational processes. Ice feathers condense on a cold windowpane; lightning traces a jagged path across the sky; rivers carve the landscape into a network of bifurcating channels. In all of these phenomena there seems to be an algorithm at work, a set of rules telling the water molecules where to freeze next or telling the lightning bolt when to zig and when to zag. Biological systems offer still more examples of algorithmically generated patterns—the leopard's spots; the branching structures of tree limbs, corals and antlers; the fairy rings of fungi. All this computing in the great outdoors suggests an intriguing possibility: If nature has an algorithm, perhaps it can be adapted to a lesser computer, one that *we* know how to program.

In the past decade there has been substantial progress in finding algorithms for various kinds of growth, aggregation and deposition. In most of these algorithms a particle travels through a medium until it comes into contact with another particle or cluster of particles; then the roving particle sticks fast, becoming a member of the cluster. The process is repeated thousands of times, building up a connected aggregate whose geometry—dense or wispy, compact or ramified—depends on the motions that brought the particles together. The procedure serves as a model of many physical and biological processes, including crystallization, the condensation of colloids and polymers, the deposition of ions and molecules during the fabrication of integrated circuits, the "viscous fingering" of interpenetrating fluids, the breakdown of dielectrics, cracking and fracturing in solids, and the growth of tumors and bacterial colonies. The same methods might even be applied to the accretion of the solar nebula and to the filaments and sheets that make up the largest visible structures in the universe, not to mention the dust bunnies in the corner I haven't swept.

Brian Hayes, a former editor of American Scientist, *has been writing about computing for 10 years. Address: 211 Dacian Avenue, Durham, NC 27701. Internet: hayes@concert.net.*

Experiments and mathematical analysis have had roles to play in studies of aggregation, but the main tool has been computer simulation. Some of the simulations do not require outrageous amounts of computing power or great programming ingenuity, and so they are open to the enthusiastic amateur. The illustrations that accompany this article are souvenirs of my own recent adventures. They contribute nothing new to the world's knowledge of aggregation algorithms, but they have taught me much. The main lesson is one that I did not anticipate: Even as I have seen how successfully a simple program can mimic the effects of nature's algorithms, I have been made more aware of the differences between physics and computation. What comes naturally in nature is sometimes very unnatural for man and machine.

Particle Ballistics

My experiments have focused on deposition rather than aggregation *per se*; that is, I have considered particles falling onto a solid substrate rather than forming free-floating clusters. The deposition problem is somewhat easier to deal with in terms of both program complexity and computing time. As a further simplification I have looked only at two-dimensional systems, and I have imposed a lattice structure on the simulated space, so that the particles move discontinuously on a rectilinear grid, like pieces on a chessboard.

The simplest simulation carried out with this apparatus resembles a gentle snowfall. Particles are released one at a time at the top of the lattice and fall straight down until they touch another particle or else come to rest on the substrate at the bottom. The horizontal coordinate of each successive particle is chosen randomly, with a uniform distribution over all the columns. A particle is considered to be touching the substrate or the existing cluster if any of its four nearest-neighbor sites (north, south, east or west) are occupied. Once a particle has been deposited, it never moves again.

Figure 1 is an example of the pattern created when 40,000 particles are deposited in this way. The color coding indicates the sequence in which the particles were laid down. The earliest particles are violet, then later ones are red and orange, and the final ones are bright yellow. The structure is spongy, with many small voids. The bot-

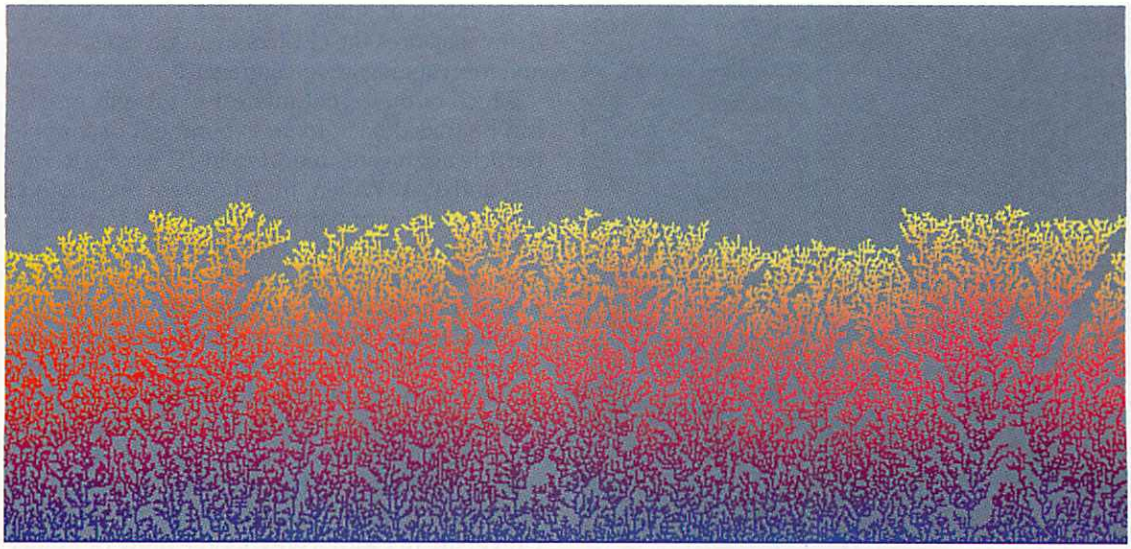

Figure 1. Spongelike texture evolves when particles descend vertically and stick as soon as they touch another particle.

tommost strata are slightly denser than average, probably because of the solid substrate under them, but beyond that boundary layer the texture appears to be statistically uniform.

The smooth vertical descent of particles in this simulation brings to mind a number of natural phenomena, such as the settling of sediments in still water. It also resembles an important industrial process: molecular-beam epitaxy (MBE), a kind of molecular-scale spray-painting in which thin layers of material are laid down on semiconductor surfaces. Because MBE is done in a vacuum chamber, the molecules in the beam are not scattered by collisions with air molecules; they follow parallel ballistic trajectories, like the particles in the simulation. At low temperatures and high deposition rates MBE films have the spongy texture seen in the model; creating smooth films requires higher temperatures and slower deposition, so that the molecules have a chance to rearrange themselves on the surface before sticking permanently.

In the program implementing this algorithm, the lattice is represented by a two-dimensional array whose elements are specified by x and y coordinates. Initially, all lattice sites are empty except in the bottom row (the row where $y = 0$), where the sites are marked as filled in order to create an impermeable substrate. The program is organized as a loop within a loop. The outer loop selects a random x coordinate and places a particle in the cell at the top of the corresponding column. Then control passes to the inner loop, which has two phases. First it examines the particle's neighborhood; if any adjacent sites are occupied, the particle's present position is also marked as occupied. Otherwise, in the second phase, the particle is moved one square south, and the inner loop repeats. The inner loop is guaranteed to terminate, since the downward-moving particle must eventually reach the substrate at $y = 0$ (if nothing else stops it first). When a particle has been deposited, the outer loop rssumes, selecting a new random x coordinate.

One technical point in the program requires close attention: boundary conditions. The lower boundary of the lattice is supplied by the substrate, and the upper boundary is of no consequence because nothing ever crosses it, but the two sides of the lattice are more problematic. If these boundaries remain open, particles deposited in the extreme left and right columns will have a different environment from all other particles—they will be missing one neighbor—and this anomaly will be reflected in the configuration of occupied cells. A solution is to make the boundary conditions periodic by joining the left edge to the right edge, as if the simulation space were wrapped around a cylinder. Thus if a particle moves out of the space by crossing the right-hand edge, it immediately re-enters from the left.

What happens if the trajectories are not vertical but oblique, as if the particles were driven by a steady wind? Figure 2 shows a deposit formed by particles on parallel ballistic paths inclined about 30 degrees below the horizontal, moving toward the right. The striated texture reminds me of a kelp bed swaying in the tide.

The feathery pattern in Figure 3 was generated by particles traveling on headings that varied randomly from just south of due east to just south of due west. Note that the most recently deposited particles—the bright yellow ones—are concentrated on the extremities of the tallest plumes. This distribution is a clue to the mechanism of pattern formation, which appears to be a kind of competitive growth. Once a cluster emerges above its surroundings, it casts a shadow, capturing particles that speed its own growth at the expense of its neighbors'.

Taking a Random Walk

The common thread running through the three simulations seen so far is that particles always move in a straight line. Relaxing that constraint leads to a new set of models with distinctly differ-

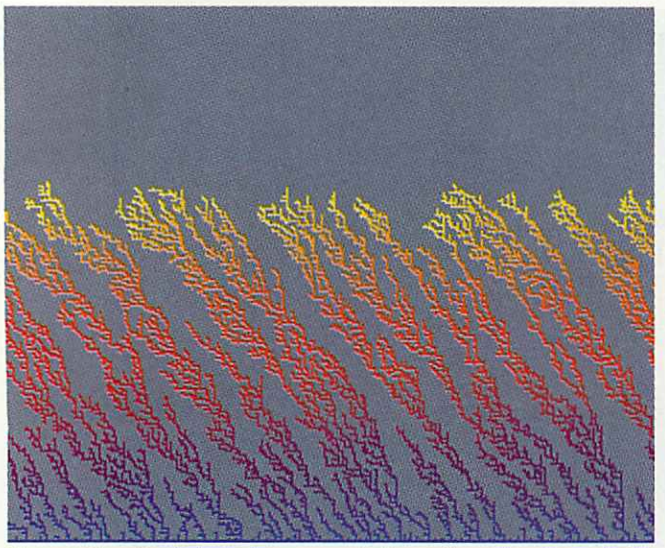

Figure 2. Swaying kelp beds were created by paths with a slope of –1/2.

ent growth habits. Figure 4 is a deposition pattern created by particles diffusing randomly through the simulation space. At each time step a particle moves one lattice unit in any of the four cardinal directions, with the direction being determined randomly and independently each time. As in the ballistic models, the random walk ends whenever the particle comes in contact with the substrate or with another particle that is already part of the deposit.

This process of "diffusion-limited deposition" is an appealing model for many real-world phenom-

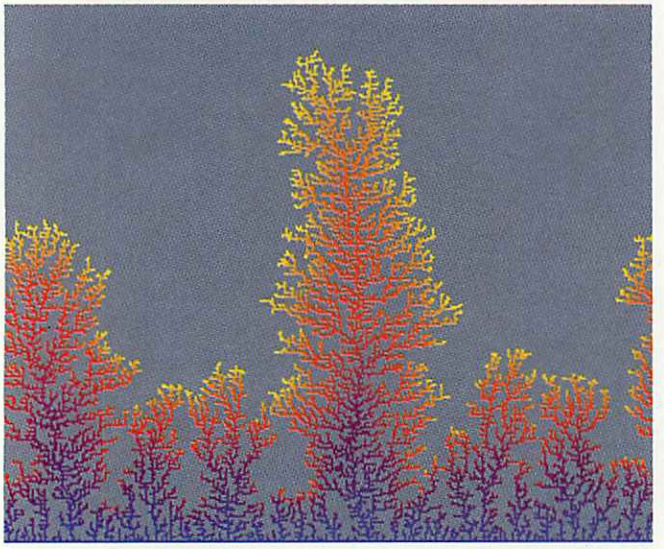

Figure 3. Feathery patterns grow from random ballistic trajectories.

ena. An example is electrochemical deposition, where ions diffusing through a liquid in response to an electric field are ultimately deposited on one of the electrodes. Microscopic examination of such deposits reveals treelike structures much like those seen in the simulations. The closely related process of diffusion-limited aggregation (which creates free clusters instead of deposits on a substrate) simulates a variety of other physical systems where particles are in random motion, such as the condensation of gold black from colloidal gold and the coagulation of smoke particles to form soot. In many instances the clusters themselves can move and form aggregates of aggregates.

Much of the recent interest in growth models can be traced back to an experiment on smoke-particle aggregates reported in 1979 by S. R. Forrest and T. A. Witten, Jr., who were then both at the University of Michigan (Forrest and Witten 1979). Two years later Witten and L. M. Sander, also of Michigan, introduced a computer simulation of diffusion-limited aggregation (Witten and Sander 1981). Hundreds of others have contributed to the subsequent elaboration and analysis of various growth and aggregation models, but three investigators have been so prominent in these studies that their work needs to be mentioned individually. They are Paul Meakin of E. I. du Pont de Nemours and Company (now at the University of Oslo), Fereydoon Family of Emory University and Tamás Vicsek of Eötvös University in Hungary. In addition to publications describing their own work, Meakin, Family and Vicsek have written useful review articles (Family and Vicsek 1990, Meakin 1991a, 1991b). Vicsek's recent book (Vicsek 1992) is the most comprehensive and up-to-date survey of work in the field. *Dynamics of Fractal Surfaces* (Family and Vicsek 1991) describes recent studies of deposition.

The structures formed by diffusion-limited growth are of particular interest because of their mathematical properties: They are *fractals*, or objects of fractional dimensionality. One way of measuring dimensionality is to examine how mass increases as a function of size. For a one-dimensional object, such as a line, mass increases as the first power of the size, and for a two-dimensional object, such as a disk or a square, mass varies as the second power of size. Measurements of the 10 largest trees in Figure 4 reveal that their mass (as represented by the number of particles) is proportional to their height raised to a power of approximately 1.64. More careful studies with much larger samples indicate that the true average exponent is about 1.7; in other words, the trees are 1.7-dimensional objects.

Fractals have the interesting property of scale-invariance, or self-similarity: Each part reproduces the structure of the whole, so that the pattern looks the same no matter how much it is magnified or reduced. For finite objects made of finite subunits there must be limits to this invariance, but the trees of Figure 4 do appear to be geometrically similar over a broad range of sizes.

A straightforward program for diffusion-limited deposition is easy to describe. Start each particle somewhere above the highest occupied lattice site, then at each time step move the particle one unit to the north, south, east or west, choosing the direction at random. Stop when the particle arrives at a site neighboring an occupied site or the substrate. A theorem guarantees that the random walk will eventually reach a terminating site. Unfortunately, the theorem makes this

guarantee by stating that eventually the walk will visit every site on the lattice, and so the particle may wander arbitrarily far before returning to the vicinity of the substrate. If no measures are taken to speed the process, the guarantee of termination could well be a lifetime guarantee.

There is a strategy for addressing this problem. Whenever a particle wanders too far from the substrate, pick it up and put it down again closer to the growing cluster, at a randomly chosen x coordinate. The justification for this shortcut is that when the particle eventually returns on its own, it will have equal probability of appearing at any x coordinate. Even with this strategy, however, the algorithm is painfully slow; in my experiments depositing 10,000 particles took 30 hours. To make the program practical, the particle must be allowed to take longer strides when it is far from the cluster, instead of moving one square at a time. With this improvement, execution time for 10,000 particles came down to a few minutes.

Organic Growth

Another kind of growth model was explored long before experiments began with ballistic and diffusive aggregation. In 1960 Murray Eden of the Massachusetts Institute of Technology introduced a model of what might be called organic growth (Eden 1961). It is growth from within rather than by accretion from without, and it is meant to represent the proliferation of bacteria in a culture medium. Instead of shooting particles at a growing cluster, each site on the periphery of the cluster is given a chance to spawn a new occupied site. In the simplest of several variations, every site on the periphery has an equal chance of being selected as the next growth site. Figure 5 shows what develops when the Eden rule is applied to growth on a horizontal substrate: a compact mass with a mottled internal structure and a somewhat rough surface.

A close comparison of the Eden model and the various accretion models suggests they are not quite as different as they seem. An accretion model exhibits Eden growth if every peripheral site has an equal probability of being struck by an incoming particle. Likewise, the Eden mechanism could reproduce all the effects of a random-walk model if the peripheral sites could be assigned growth probabilities proportional to the expected flux of incoming particles at each site. Surprisingly, the calculation of these probabilities is entirely feasible; it turns out to be a matter of solving the equation for an electrostatic field (Mandelbrot and Evertsz 1990). In this calculation the substrate and the clusters attached to it are assumed to be perfect conductors at a potential of zero volts. Well above the tallest cluster, the horizontal line where diffusing particles are released is taken to be an electrode with a fixed potential of one volt. The potential field at each point in the intervening space corresponds to the probability of finding a randomly moving particle at that point. The flux of particles onto the substrate is

Figure 4. Particles on random walks yield a forest of fractal trees.

equal to the gradient of this field, or in other words to its spatial rate of change.

Figure 6 shows the electric field surrounding a structure grown by diffusion-limited deposition. The skeleton of the deposit, at zero volts, is white, and the top row of the lattice, at one volt, is dark blue. The field is draped like a tissue of cobwebs over the gaps between clusters, showing how these areas are screened from further growth. The field gradient is highest at the exposed tips of the clusters, which are the regions

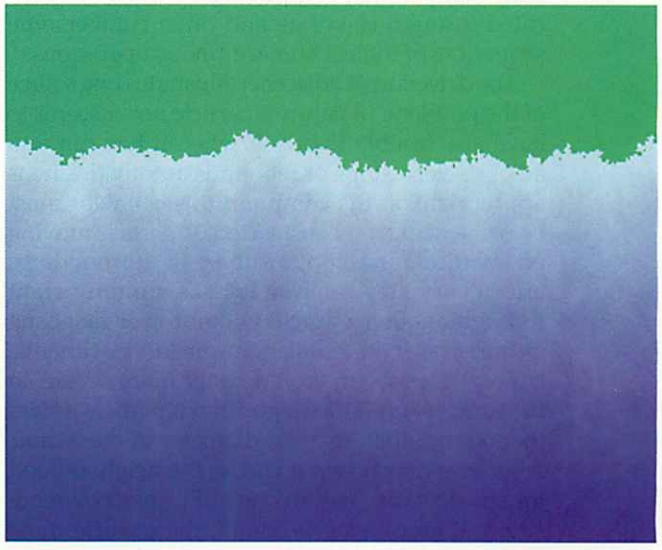

Figure 5. The Eden model resembles the growth of a bacterial colony.

of fastest growth. Lighter areas are depleted of particles and so will grow slowly.

The calculation of probability fields could in principle replace the entire methodology of launching and following particles on random trajectories, but it is not clear there is anything to be gained by such a change in technique. Solving the field equations takes far longer than a random walk. Indeed, the random walk can be viewed as a shortcut method of calculating the distribution of growth probabilities.

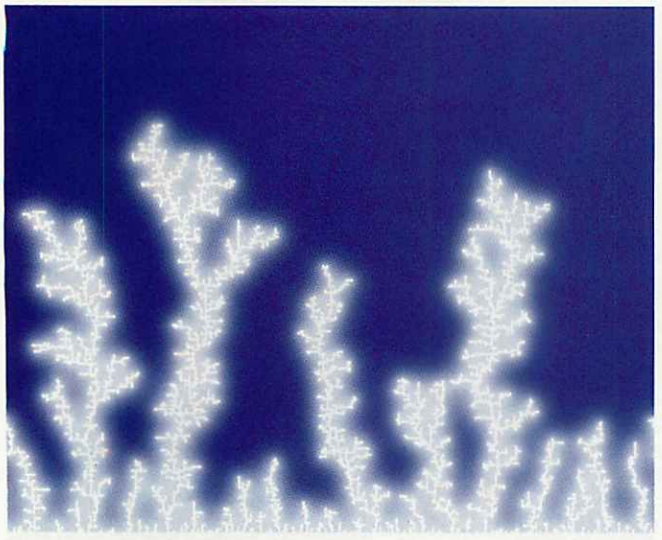

Figure 6. Field surrounding a cluster predicts its growth probability.

Bugs Bunny Dynamics

In this article I have given more attention to details of program implementation than is usual in accounts of these algorithms. I believe the details are important. Merely saying "particles execute a random walk until they reach a site adjacent to an occupied site" may give the misleading impression that random walks and the concept of adjacency are simple, built-in faculties of the computer. While it's true that in *nature's* computer they do seem to be primitive operations, in the man-made digital computer they must be fabricated through elaborate and often cumbersome sequences of logical and arithmetic operations.

The detection of adjacency illustrates one source of the problem. In nature a particle never seems to have any trouble knowing when it has run into an obstacle; it stops or rebounds automatically, as appropriate. In the computer this behavior must be painstakingly reconstructed. A particle moving on a lattice must stop to survey its surroundings after every step, or else it risks running right through another particle without ever detecting its presence. If the simulation were done without a lattice, the problem would be far more severe. In the simplest off-lattice algorithm a particle moves by no more than its own diameter in one stride. After every such step it checks the neighborhood for any obstacles within one radius of its own position. If an overlap is found, the particle must back up until it is just touching the obstacle. This method of blundering through things and then retreating gives correct results, but it hardly seems a natural representation of particle dynamics in the physical world. It is more like the cartoon world where Bugs Bunny doesn't begin falling until he notices that he has stepped off the cliff.

Simulations done with a lattice are more efficient than those without (largely because it is easier to detect collisions), but they are also one step further removed from reality. There is always the worry that the geometry of the lattice will "show through" in the geometry of the aggregates. Simulations using 10,000 or even 100,000 particles show no obvious signs of lattice influence, but this finding does not stand up to larger-scale testing. Through prodigious feats of computing Meakin has shown that clusters of several million particles on a square lattice form diamond or cross shapes aligned with the axes of the lattice; the shapes do not appear in simulations done on a triangular or hexagonal lattice, or with no lattice at all (Meakin 1986a).

In calling attention to the various difficulties of imitating nature I do not mean to question the legitimacy of computer simulation as a means of studying aggregation (or anything else). As a matter of fact, the difficulties may be where we stand to learn the most. A concept like adjacency seems so simple and self-evident that you are not likely to think clearly about it until you try living in a world where it does not exist. Simulations of fractal growth offer many such opportunities to learn things you thought you always knew.

Looking at the bizarre contortions needed to make a digital computer deal with continuous motion, one might conclude that present-day computational science must be in a feeble and primitive state. That may turn out to be a fair assessment, and yet it is only part of the story. The other part is that the continuum of space and time really is a deep mystery, as Zeno of Elea knew. We become most aware of how singular that continuum is only when we set out to construct a world of our own inside the computer—only when we formulate and make fully explicit our own laws of nature.

Bibliography

Adams, Douglas. 1979. *The Hitchhiker's Guide to the Galaxy*. New York: Pocket Books.

Eden, Murray. 1961. A two-dimensional growth process. In *Proceedings of the Fourth Berkeley Symposium on Mathematics, Statistics, and Probability*, edited by Jerzy Neyman, Vol. 4, pp. 223–229. Berkeley: University of California Press. [Also reprinted in (Family and Vicsek 1991).]

Family, Fereydoon, and David P. Landau (eds.). 1984. *Kinetics of Aggregation and Gelation*. Proceedings of an International Topical Conference. Amsterdam: North-Holland.

Family, Fereydoon, and Tamás Vicsek. 1990. Simulating fractal aggregation. *Computers in Physics* 4 (January–February): 44–49.

Family, Fereydoon, and Tamás Vicsek (eds.). 1991. *Dynamics of Fractal Surfaces*. Singapore: World Scientific.

Forrest, S. R., and T. A. Witten, Jr. 1979. Long-range correlations in smoke-particle aggregates. *Journal of Physics A* 12:L109–L117.

Mandelbrot, Benoit B., and Carl J. G. Evertsz. 1990. The potential distribution around growing fractal clusters. *Nature* 348 (8 Nov.):143–145.

Meakin, Paul. 1986a. Universality, nonuniversality, and the effects of anisotropy on diffusion-limited aggregation. *Physical Review A* 33:3371–3382.

Meakin, Paul. 1991a. Models for material failure and deformation. *Science* 252: 226–234.

Meakin, Paul. 1991b. Fractal aggregates in geophysics. *Reviews of Geophysics* 29:317–354.

Vicsek, Tamás. 1992. *Fractal Growth Phenomena*. Second Edition. Singapore: World Scientific.

Witten, T. A., Jr., and L. M. Sander. 1981. Diffusion-limited aggregation, a kinetic critical phenomenon. *Physical Review Letters* 47:1400–1403.