

A reprint from

American Scientist

the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request Brian Hayes by electronic mail to bhayes@amsci.org.

Flights of Fancy

Brian Hayes

A THOUSAND STARLINGS rose in unison from trees along a riverbank. The ascending cloud of birds took the form of a teardrop, then transformed itself into a butterfly, then a twisting vortex narrowing to a sinuous, quivering rope of birds stretched across the twilight sky. The flock had all the synchronized precision of a marching band, but none of the rigid, rank-and-file geometry. Instead the movements were smooth, fluid, organic, as if the flock were a single organism rather than a collection of individuals. The show went on for 10 minutes, then the birds swooped low over my head with a breathy rush of wing beats and returned to the same row of trees—only to rise again moments later for another performance.

The graceful aerial displays of starlings and other flocking birds have long inspired admiration and wonder. Lately they have also inspired serious work in mathematics, computer science, physics and biology. A theoretical framework for explaining the behavior of tightly clustered flocks emerged in the 1980s. The key idea, which came from computer simulations, is that purely local interactions between nearby birds are enough to hold the group together. Similar mechanisms are thought to operate in schools of fish, herds of grazing animals, swarms of insects and even crowds of people.

More recent work builds on this foundation. In particular, an ambitious program of stereoscopic photography and digital image processing has made it possible to reconstruct the positions and velocities of individual birds with-

*How birds
(and bird-watchers)
compute the behavior
of a flock on the wing*

in large flocks. The new data suggest some refinements to the algorithms that are thought to govern the birds' movements. The three-dimensional reconstructions have also brought some surprises about the overall shape and structure of the flocks, and show that the birds' movements are even more delicately coordinated than casual observation would suggest.

Avian Democracy

For early students of bird flocks, the most pressing questions had to do with decision making and communication. When hundreds of birds all take to the air at the same instant, how do they synchronize their liftoff? When the flock suddenly veers left or right, who decides which way to turn?

Human experience suggests one solution: Appoint a leader—a conductor, a choreographer, a dictator. Leadership might be the social role of a specific individual. (Tennyson wrote of “the many-winter’d crow” that leads its flock home.) Or leadership might shift from moment to moment, as a flock on the wing follows whichever bird is currently at the front of the formation.

The trouble is, biologists have found no clear evidence of either kind of leadership in flocks, and the whole idea has some deep conceptual difficulties. How do cues from the leader reach the rest of the birds? In large and dense groups, both sight and sound would be unreliable. And there’s another issue: A flock

attacked by a falcon responds instantly with evasive maneuvers; it’s implausible that the threatened birds await instructions from a leader.

A more democratic vision of flock governance was proposed early in the 20th century by Edmund Selous, an intrepid English bird-watcher who kept careful notes on flocking behavior for more than 30 years. Selous rejected the notion of designated leaders as “well-nigh unthinkable—it is *too* ridiculous”; his alternative, however, prompts the same kind of dismissive response. Selous found the birds’ synchronized movements so uncanny that he could explain them only as a product of “thought-transference” or “simultaneous collective thinking.” In other words, the birds are telepathic.

Selous was led to this extraordinary hypothesis from his assumption that flocks make decisions by some form of global consensus. The decision-making protocol is almost as hard to fathom as the spooky communication medium. Does each bird transmit its intention or preference to every other bird? From an algorithmic point of view, this process would be highly demanding: The N members of a flock would have to exchange almost N^2 messages. Thus the agility of the flock would depend on N , as larger flocks would necessarily take longer to make up their collective mind.

Boidland

In retrospect, neither democracy nor dictatorship was a very good point of reference for understanding bird society. Both theories were swept aside dramatically in the 1980s.

The key insight came not from an ornithologist but from an expert on computer graphics and animation. Craig W. Reynolds, now of Sony Computer Entertainment, wrote a program in which simulated birds—he called them “boids”—gave a convincing imi-

Brian Hayes is senior writer for American Scientist. Additional material related to the “Computing Science” column appears at <http://bit-player.org>. Address: 11 Chandler St. #2, Somerville, MA 02144. E-mail: brian@bit-player.org



Many thousands of starlings gather on winter evenings in Rome, where they make soaring display flights before roosting for the night in trees near the main railway terminal. A project called STARFLAG has employed stereophotography and computer analysis of images to determine the three-dimensional positions and velocities of birds in starling flocks. (The flocks successfully analyzed have up to about 4,000 members; some of the groups visible in this photograph are probably even larger.) The newly acquired 3D data suggest some revisions and refinements to algorithmic models that seek to explain how the birds coordinate and synchronize their movements. (Image reproduced courtesy of the STARFLAG project, INFN-CNR [Istituto Nazionale per la Fisica della Materia, Consiglio Nazionale delle Ricerche]).

tation of flocking behavior. And the boids did it without any need for leaders or thought transference. Each boid attends to a few near neighbors, ignoring the rest of the flock, and obeys three simple rules:

1. *Collision avoidance*: Try not to run into neighboring boids.
2. *Velocity matching*: Try to fly at the same speed and in the same direction as neighboring boids.
3. *Cohesion*: Try to stay close to neighboring boids.

The rules are to be applied in the order given. That is, collision avoid-


ance gets the highest priority; if no collisions are imminent, then the boid adjusts its speed and direction to match the average of its neighbors' velocities; if the velocities are already in good agreement, the boid moves a little closer to its flockmates. (Note that the third imperative is in mild conflict with the other two.)

When I first read about the boid model, I was astonished that it worked. A stable state in which all the boids fly on parallel trajectories seemed like a possible outcome of the rules, but just barely so. A far more likely result, I thought, was a fragmented collection

of many small flocks, each roughly the size of a boid neighborhood.

I changed my mind when I saw the boids in action and got a chance to play with implementations of the model. The flocking of boids is not a perfect simulation of bird behavior, but the resemblance is unmistakable. And the flocks are not only stable but robust. If you start the simulation with boids moving randomly, they spontaneously assemble into a flock. If an obstacle causes the flock to split, the factions eventually rejoin.

What accounts for the balletic coordination of a boid flock? Although a



Three-dimensional reconstruction of a flock entails solving a difficult matching problem: Given two photographs of a flock taken at the same instant from somewhat different points of view, how do you match the featureless blob that represents a bird in one image with the corresponding blob in the other image? The red boxes in the photos above indicate five such matches; the STARFLAG group developed image-analysis tools that automate the matching process. In the four panels below the photographs, the coordinates of birds are represented as green dots, providing a model of the flock that can be reoriented in 3D space and viewed from various directions. The two views at the bottom reveal something that is not obvious from the photos: The flock has a flattened shape. (Illustration reproduced from Ballerini et al., *Proceedings of the National Academy of Sciences of the U.S.A.* 105:1232–1237.)

boid interacts *directly* only with its immediate neighbors, those neighbors interact with their own neighbors, and thus a network of indirect connections binds the whole flock together. The surprise is that these tenuous

links are enough to do the job. When a boid at the north end of the flock starts turning, a boid at the south end gets the signal only through several intermediaries. Admittedly, *many* such chains of intermediaries connect

the two flocks. Even so, it's unexpected that distant events can exert such a powerful influence. It's almost like thought-transference.

Emergence Emerges

The Reynolds boid model was part of a revolutionary change in the prevailing style of thinking about collective behavior in animals. The model was bottom-up rather than top-down, local rather than global, algorithmic rather than equational. Instead of starting with the overall form of the flock and asking how the birds could organize themselves to produce it, the model merely specified how individual birds interact with one another; the large-scale structure of the flock is an "emergent property."

Reynolds was not the only one to discover the charms of such models. Frank Heppner of the University of Rhode Island had been studying bird behavior since the 1960s, including searches for leaders in flocks. In the late 1980s he began independently formulating leaderless models based on the idea of a force law acting between birds (with short-range repulsion and long-range attraction). A few years later Tamás Vicsek of Eötvös Loránd University in Budapest and several colleagues published another model, stripping away ornithological detail in a search for the simplest system that would exhibit interesting collective behavior. In essence, their model adopted just the second of Reynolds's rules (velocity matching) along with some random noise. The Vicsek model has been applied to everything from birds to bacteria.

Soon there were specialized programming languages and software systems that made it easy (and fun!) to explore such "agent-based models." Mitchel Resnick of MIT developed one of the first of these languages, called StarLogo (which I wrote about in an earlier column, "E Pluribus Unum," January–February 1999). Several other languages have appeared since then, and enthusiasm for agent-based modeling has spread from biology to the social sciences, finance and other areas. Moreover, the boids have come home to roost: The computational methods devised to help understand the behavior of animals in groups have come back as biologically inspired algorithms ("swarm intelligence," "ant-colony optimization") for hard computational problems.

Air Truth

The best computer simulations of bird flocks look wonderfully realistic—Reynolds won an Oscar for some of his flock animations—but do they tell us anything about real birds? Until recently it was hard to answer this question because so little was known about the trajectories of birds in large flocks. There were plenty of films and videos, but tracking individuals in three dimensions was not feasible.

The situation has changed with new data-gathering efforts, the largest being the STARFLAG project, an interdisciplinary collaboration extending across Italy, France, Hungary and the Netherlands. A STARFLAG group at the University of Rome, led by Andrea Cavagna and Giorgio Parisi, has succeeded in tracing the movements of individuals in starling flocks with up to 4,000 members. The Cavagna group did not have to go far to make their observations. Every winter, vast numbers of starlings gather just a few blocks away from the university at the Termini, Rome's main railway station. Cavagna and his colleagues set up digital cameras on the roof of a nearby building, the Palazzo Massimo.

Because video recorders lack sufficient resolution, the Rome group used cameras meant for still photography and made rapid-fire sequences of exposures. Pairs of cameras were mounted 25 meters apart so that the stereoscopic disparity would provide depth cues, as in binocular vision. Each camera could record five frames per second, which is not quite fast enough for following the movement of starlings. Therefore two pairs of cameras were triggered alternately to give an effective rate of 10 frames per second. Two more cameras (making a total of six) aided in the 3D reconstruction.

Birds and Blobs

Making the photos was only the first challenge; making sense of them was harder. Even with high-resolution images, a starling 100 meters from the camera is little more than a dark blob, with no individually distinguishable features. Given a pair of images that show the same configuration of birds from different points of view, how do you match up blobs that correspond to the same bird? Then how do you solve the similar problem of tracking a bird through successive frames? A further complication is that blobs often overlap.

To solve these problems the group built a hierarchy of image-analysis tools. The first step removes background features such as clouds, isolating the dark blobs that represent birds. Then a "blob-splitting" process breaks up overlapping bird images. In this way the photograph is reduced to an array of point coordinates in the two-dimensional plane of the image.

The leap into the third dimension requires combining information from multiple cameras. The basic geometric principle is well-known: An object photographed from two different points of view appears at different positions in the two images. This stereoscopic disparity provides information about the object's range, or distance from the cameras. The complication is that the photographs include a few thousand objects that all look alike. Before calculating ranges, it's necessary to figure out which blobs in image *A* go with which blobs in image *B*.

The STARFLAG group developed a multistage process for solving this matching problem. First a pattern-recognition algorithm searches for "constellations" of points in image *A* whose arrangement is distinctive enough that the same points can also be identified in *B*, even though the constellation will be distorted somewhat by the shift in point of view. Finding about 50 such matched pairs provides enough information to approximate the geometric transformation that maps

any point in *A* to the corresponding pixel position in *B*. The approximation is refined by exploiting images from a third camera, *C*, mounted close to *A*. The short baseline and small optical disparity make the matching problem easier for *A*–*C* image pairs. And once the *A*–*C* matches are found, they can guide the search for more *A*–*B* matches. On average, the algorithm assigned 3D positions to 88 percent of the birds in the images analyzed.

Tracking a bird through time, and thereby learning its velocity as well as its position, entails matching across many successive frames rather than just the two images of a stereo pair. This is a challenging task, but the precise coordination of starling flocks helps to make it feasible. Because at any instant most of the birds are flying at the same speed and in the same direction, the average velocity vector gives a very good clue about where to look for a bird in the next frame.

Flattened Flocks

The findings of the STARFLAG program support the existing theoretical framework—the basic idea that flocks are held together by local interactions—but the results have also brought some surprises. For example, it turns out that the overall shape of starling flocks is not what it appears to be. To a casual observer, most flocks look globular; they appear to be deformed but fully rounded spheres or cylinders. But the



Analysis of starling velocities reveals a pattern of correlations extending across the entire flock. At left are vectors representing bird velocities; at right the flock's mean velocity has been subtracted from each vector, leaving fluctuations about the mean. (The fluctuations are actually much smaller than the velocities; the vectors have been rescaled for clarity.) The fluctuation vectors necessarily sum to zero, since the net motion of the flock has been subtracted; thus in some regions the prevailing direction is up, elsewhere it is down, and so on. The scale of these regions is comparable to the size of the flock, which indicates that velocities are correlated throughout the group. (Reproduced from Cavagna et al., *Proceedings of the National Academy of Sciences of the U.S.A.* 107:11865–11870.)

3D reconstructions show that the typical flock is quasi-two-dimensional: It is extended along two dimensions but squashed to a fairly thin layer along the third dimension. The average aspect ratio of length to width to thickness is 6:3:1. In general, the shortest axis is vertical, so that the flock is spread out in a thin horizontal sheet.

Cavagna and his colleagues suggest a physical or physiological reason for this flight formation: Vertical movements have a higher energy cost than lateral ones, and so birds favor flight at constant altitude. What's harder to explain is why most observers of flocks have a very different impression of their shapes and motions. The suggested explanation is a perceptual effect: Looking up at a steep angle makes it hard to distinguish between vertical motion of the flock and a component of horizontal motion directed toward or away from the observer.

Another curious finding is that starling flocks have a dense outer rind and a mushy interior. The density of birds is greatest at the boundary of the flock, and it declines steadily toward the core. This is the opposite of the distribution that evolves when free-moving objects are bound together by a long-range force such as gravity; globular clusters of stars, for example, are densest in the middle and sparse at the edges. The Rome investigators discuss some ideas about why the inside-out density gradient might be useful as a defense against predators, but they do not propose a mechanism for generating the gradient. What modification of the Reynolds boid rules would account for this observation?

The Rome group *does* introduce another, more fundamental revision of the Reynolds model. In the original simulation, a bird's neighborhood is defined geometrically: The bird keeps track of all other birds within a certain radius of its own position. Cavagna and his colleagues propose a "topological" alternative: A bird interacts with a fixed number of nearest neighbors, regardless of their geometric distance. The number of birds comprising the neighborhood is probably six or seven.

What led to this proposal was the discovery that flocks differ substantially in overall density—by a factor of two or more. If interactions were limited to a fixed radius, then birds in dense flocks would effectively have more neighbors than those in looser

groupings. But the observed behavior of the flocks is identical across the range of densities.

Scale-Free Flocking

Of all the findings announced so far, the most intriguing have come from a study of correlations in velocity. The existence of strong correlations is hardly a surprise. After all, if the birds weren't all flying in roughly the same direction, the flock would disintegrate. But the result is stronger than this. Subtracting the flock's mean velocity from each bird's individual velocity vector leaves a field of residual fluctuations—new, smaller vectors that represent the discrepancy between each bird's motion and the average. It turns out that these vectors too are strongly correlated.

The statistic called the correlation length is the distance beyond which one bird's influence on another fades away, so that the two birds fly independently. Remarkably, the correlation length in starling flocks grows along with the flock itself, so that even birds at opposite ends remain connected. The correlations are said to be "scale-free."

What's unusual about this situation is not just that the correlation length is greater than the interaction distance. That happens in many systems. An example from physics is a ferromagnet, in which short-range interactions between atoms produce alignments over much larger distances. But the correlation length in the magnet is not scale-free; a ferromagnet of macroscopic size breaks up into many independent domains.

Why isn't there a similar limit on the correlation length in bird flocks? One possibility is that the flock operates at a "critical point," a set of circumstances where fluctuations extend to all possible scales of length. For a ferromagnet the critical point is a temperature known as the Curie point. In the case of bird flocks it's not clear what would correspond to the Curie point, or even what variable would play the role of temperature. The Cavagna group explains:

Scale-free correlations imply that the group is, in a strict sense, different from and more than the sum of its parts. The effective perception range of each individual is as large as the entire group and it becomes possible to transfer undamped information to all animals, no matter their distance, making the group respond as one.

In this passage I hear an echo of Selous's "collective thinking." Of course no telepathy is implied; the birds communicate by ordinary physical means. And yet the outcome seems all the more impressive—and perhaps even more mysterious—for that very reason.

Bibliography

- Bajec, Iztok Lebar, and Frank H. Heppner. 2009. Organized flight in birds. *Animal Behaviour* 78:777–789.
- Ballerini, Michele, Nicola Cabibbo, Raphael Candelier, Andrea Cavagna, Evaristo Cisbani, Irene Giardina, Alberto Orlandi, Giorgio Parisi, Andrea Procaccini, Massimiliano Viale and Vladimir Zdravkovic. 2008. Empirical investigation of starling flocks: A benchmark study in collective animal behaviour. *Animal Behaviour* 76:201–215.
- Ballerini, Michele, Nicola Cabibbo, Raphael Candelier, Andrea Cavagna, Evaristo Cisbani, Irene Giardina, Vivien Lecomte, Alberto Orlandi, Giorgio Parisi, Andrea Procaccini, Massimiliano Viale and Vladimir Zdravkovic. 2008. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences of the U.S.A.* 105:1232–1237.
- Cavagna, Andrea, Irene Giardina, Alberto Orlandi, Giorgio Parisi, Andrea Procaccini, Massimiliano Viale and Vladimir Zdravkovic. 2008. The STARFLAG handbook on collective animal behaviour. 1: Empirical methods. *Animal Behaviour* 76:217–236.
- Cavagna, Andrea, Irene Giardina, Alberto Orlandi, Giorgio Parisi and Andrea Procaccini. 2008. The STARFLAG handbook on collective animal behaviour. 2: Three-dimensional analysis. *Animal Behaviour* 76:237–248.
- Cavagna, Andrea, Alessio Cimorelli, Irene Giardina, Giorgio Parisi, Raffaele Santagati, Fabio Stefanini and Massimiliano Viale. 2010. Scale-free correlations in starling flocks. *Proceedings of the National Academy of Sciences of the U.S.A.* 107:11865–11870.
- Couzin, Iain. 2007. Collective minds. *Nature* 445:715.
- Heppner, Frank, and Ulf Grenander. 1990. A stochastic nonlinear model for coordinated bird flocks. In *The Ubiquity of Chaos* (ed. by Saul Krasner), pp. 233–238. Washington, D.C.: American Association for the Advancement of Science.
- Resnick, Mitchel. 1994. *Turtles, Termites and Traffic Jams: Explorations in Massively Parallel Microworlds*. Cambridge, Mass.: MIT Press.
- Reynolds, Craig W. 1987. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics* 21(4):25–34.
- Selous, Edmund. 1901. *Bird Watching*. London: J. M. Dent & Co.
- Selous, Edmund. 1931. *Thought-transference (or What?) in Birds*. London: Constable & Co.
- Vicsek, Tamás, András Czirók, Eshel Ben-Jacob, Inon Cohen and Ofer Shochet. 1995. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters* 75:1226–1229.
- Vicsek, Tamás, and Anna Zafiris. 2010 (preprint). Collective motion. <http://arxiv.org/abs/1010.5017>