

Antichaos and Adaptation

Biological evolution may have been shaped by more than just natural selection. Computer models suggest that certain complex systems tend toward self-organization

by Stuart A. Kauffman

Mathematical discoveries are inviting changes in biologists' thinking about the origins of order in evolution. All living things are highly ordered systems: they have intricate structures that are maintained and even duplicated through a precise ballet of chemical and behavioral activities. Since Darwin, biologists have seen natural selection as virtually the sole source of that order.

But Darwin could not have suspected the existence of self-organization, a recently discovered, innate property of some complex systems. It is possible that biological order reflects in part a spontaneous order on which selection has acted. Selection has molded, but was not compelled to invent, the native coherence of ontogeny, or biological development. Indeed, the capacity to evolve and adapt may itself be an achievement of evolution.

The studies supporting these conclusions remain tentative and incomplete. Nevertheless, on the basis of mathematical models for biological systems that exhibit self-organization, one can make predictions that are consistent with the observed properties of organisms. We may have begun to understand evolution as the marriage of selection and self-organization.

To understand how self-organization

can be a force in evolution, a brief overview of complex systems is necessary. During the past two decades, there has been an explosion of interest in such systems throughout the natural and social sciences. The efforts are still so new that there is not yet even a generally accepted, comprehensive definition of complexity.

Yet certain properties of complex systems are becoming clear. One phenomenon found in some cases has already caught the popular imagination: the randomizing force of deterministic "chaos." Because of chaos, dynamic, nonlinear systems that are orderly at first may become completely disorganized over time. Initial conditions that are very much alike may have markedly different outcomes. Chaos in the weather is exemplified by the so-called butterfly effect: the idea that a butterfly fluttering in Rio de Janeiro can change the weather in Chicago.

Chaos, fascinating as it is, is only part of the behavior of complex systems. There is also a counterintuitive phenomenon that might be called antichaos: some very disordered systems spontaneously "crystallize" into a high degree of order. Antichaos, I believe, plays an important part in biological development and evolution.

The discovery of antichaos in biology began more than 20 years ago with my efforts to understand mathematically how a fertilized egg differentiates into multitudes of cell types. Since then, mathematicians, computer scientists and solid state physicists, among

them my many colleagues at the Santa Fe Institute in New Mexico, have made substantial progress.

Biology is filled with complex systems: the thousands of genes regulating one another within a cell; the network of cells and molecules mediating the immune response; the billions of neurons in the neural networks underlying behavior and learning; the ecosystem webs replete with coevolving species. Of these, the self-regulating network of a genome (the complete set of genes in an organism)

STUART A. KAUFFMAN has been thinking about self-organization in living things since 1965. After studying at Dartmouth College and at the University of Oxford, he went on to receive his medical degree from the University of California, San Francisco, in 1968. Since then, he has held positions at the Massachusetts Institute of Technology, the University of Chicago and the National Cancer Institute. Kauffman is currently professor of biochemistry and biophysics at the University of Pennsylvania School of Medicine and external professor at the Santa Fe Institute in New Mexico. In 1987 he received a John D. and Catherine T. MacArthur Fellowship.

LOOPS OF DNA extruded by this bacterium contain thousands of genes. The genes act as a self-regulating network, turning one another on and off. Even more complex genetic circuits occur in higher cells. Computational models are now hinting at how such complex systems can spontaneously organize themselves to exhibit stable cycles of gene activity—an essential feature of all life.



offers a good example of how anti-chaos may govern development.

The genome of a higher organism such as a human being encodes the information for making about 100,000 different proteins. One of the central dogmas of developmental biology is that liver cells, neurons and other cell types differ because varied genes are active in them. Yet it is now also clear that all the cells in an organism contain roughly the same genetic instructions. Cell types differ because they have dissimilar patterns of genetic activity, not because they have different genes.

A genome acts like a complex parallel-processing computer, or network, in which genes regulate one another's activity either directly or through their products. The coordinated behavior of this system underlies cellular differentiation. Understanding the logic and structure of the genomic regulatory system has therefore become a central task of molecular biology.

Mathematical models can help researchers understand the features of such complex parallel-processing systems. Every complex system has what can be called local features: these characteristics describe how individual ele-

ments in the system are connected and how they may influence one another. For example, in a genome the elements are genes. The activity of any one gene is directly regulated by fairly few other genes or gene products, and certain rules govern their interactions.

Given any set of local features, one may construct a large ensemble, or class, of all the different complex systems consistent with them. A new kind of statistical mechanics can identify the average features of all the different systems in the ensemble. (Traditional statistical mechanics, in contrast, averages over all the possible states of a single system.) Individual systems in the ensemble might be very different; nonetheless, the statistically typical behaviors and structures are the best hypothesis for predicting the properties of any one system.


The approach begins by idealizing the behavior of each element in the system—each gene, in the case of the genome—as a simple binary (on or off) variable. To study the behavior of thousands of elements when they are coupled together, I used a class of systems called random Boolean networks. These systems are named after George

Boole, the English inventor of an algebraic approach to mathematical logic.

In a Boolean network, each variable is regulated by others that serve as inputs. The dynamic behavior of each variable—that is, whether it will be on or off at the next moment—is governed by a logical switching rule called a Boolean function. The function specifies the activity of a variable in response to all the possible combinations of activities in the input variables. One such rule is the Boolean OR function, which says that a variable will be active if any of its input variables is active. Alternatively, the AND function declares that a variable will become active only if all its inputs are currently active.

One can calculate how many Boolean functions could conceivably apply to any binary element in a network. If a binary element has K inputs, then there are 2^K possible combinations of inputs it could receive. For each combination, either an active or inactive result must be specified. Therefore, there are 2 to the 2^K power possible Boolean switching rules for that element.

The mathematically idealized ver-



sions of biological systems I shall discuss are called autonomous random Boolean *NK* networks. They consist of *N* elements linked by *K* inputs per element; they are autonomous because none of the inputs comes from outside the system. Inputs and one of the possible Boolean functions are assigned at random to each element. By assigning values to *N* and *K*, one can define an ensemble of networks with the same local features. A random network is one sampled at random from this ensemble.

Each combination of binary element activities constitutes one network "state." In each state, all the elements assess the values of their regulatory inputs at that moment. At the next clocked moment, the elements turn on or off in accordance with their individual functions. (Because all the elements act simultaneously, the system is also said to be synchronous.) A system pass-

es from one unique state to another. The succession of states is called the trajectory of the network.

A critical feature of random Boolean networks is that they have a finite number of states. A system must therefore eventually reenter a state that it has previously encountered. Because its behavior is determined precisely, the system proceeds to the same successor state as it did before. It will consequently cycle repeatedly through the same states.

Such state cycles are called the dynamic attractors of the network: once a network's trajectory carries it onto a state cycle, it stays there. The set of states that flow into a cycle or that lie on it constitutes the "basin of attraction" of the state cycle. Every network must have at least one state cycle; it may have more.

Left to itself, a network will eventual-

ly settle into one of its state cycle attractors and remain there. Yet if the network is perturbed in some way, its trajectory may change. Two types of perturbation are worth discussing here: minimal perturbations and structural perturbations.

A minimal perturbation is a transient flipping of a binary element to its opposite state of activity. If such a change does not move a network outside its original basin of attraction, the network will eventually return to its original state cycle. But if the change pushes the network into a different basin of attraction, the trajectory of the network will change: it will flow into a new state cycle and a new recurrent pattern of network behavior.

The stability of attractors subjected to minimal perturbations can differ. Some can recover from any single perturbation, others from only a few, whereas still others are destabilized by any perturbation. Flipping the activity of just one element may unleash an avalanche of changes in the patterns that would otherwise have occurred. The changes are "damage," and they may propagate to varying extents throughout a network [see "Self-Organized Criticality," by Per Bak and Kan Chen; *SCIENTIFIC AMERICAN*, January].

A structural perturbation is a permanent mutation in the connections or in the Boolean functions of a network. Such perturbations would include exchanging the inputs of two elements or switching an element's OR function to an AND function. Like minimal perturbations, structural perturbations can cause damage, and networks may vary in their stability against them.

As the parameters describing a complex Boolean system change, the system's behavior alters, too: a system can change from chaotic behavior to ordered behavior. A type of system that is perhaps surprisingly easy to understand is one in which the number of inputs to each element equals the total number of elements—in other words, everything is connected to everything else. (Such systems are called $K = N$ networks.) Because a random $K = N$ network is maximally disordered, the successor to each state is a completely random choice. The network behaves chaotically.


One sign of the disorder in $K = N$ systems is that as the number of elements increases, the length of the state cycles grows exponentially. For example, a $K = N$ network consisting of 200 elements can have 2^{200} (about 10^{60}) different states. The average length of a state

Boolean Functions and State Cycles

Boolean functions are logical rules that describe how binary (on or off) elements in networks will respond to combinations of signals from other elements. By applying Boolean logic to a network, one can predict the system's behavior.


In this simple network (a), there are three elements, each of which receives signals from the others. Element *A* obeys the Boolean AND function: it becomes active only if both elements *B* and *C* were previously active. Elements *B* and *C* obey OR functions: each one becomes active if either of the other elements was active. By listing every possible initial combination of states in the network, one can deduce from the Boolean functions what all the successor states will be (b).

The long-term behavior of the system is remarkably simple. Although the network can initially have any of eight different states, it will eventually settle into one of only three state cycles—each one a recurrent pattern of change (c). If all the elements are off initially, the network never changes. If only element *B* or element *C* is active, the system will cycle back and forth between the two states. Any other network state inevitably leads to all the elements' becoming active and staying that way.



b

| INITIAL STATES OF ELEMENTS | | | SUCCESSOR STATES OF ELEMENTS | | |
|----------------------------|---|---|------------------------------|---|---|
| A | B | C | A | B | C |
| 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 1 | 0 | 1 | 0 |
| 0 | 1 | 0 | 0 | 0 | 1 |
| 1 | 0 | 0 | 0 | 1 | 1 |
| 1 | 0 | 1 | 0 | 1 | 1 |
| 1 | 1 | 0 | 0 | 1 | 1 |
| 0 | 1 | 1 | 1 | 1 | 1 |
| 1 | 1 | 1 | 1 | 1 | 1 |



cycle in the network is roughly the square root of that number, about 10^{30} states. Even if each state transition took only one microsecond, it would take billions of times longer than the age of the universe for the network to traverse its attractor completely.

$K=N$ networks also exhibit maximum sensitivity to initial conditions. Because the successor to any state is essentially random, almost any perturbation that flips one element would sharply change the network's subsequent trajectory. Thus, minimal changes typically cause extensive damage—alterations in the activity patterns—almost immediately. Because the systems show extreme sensitivity to their initial conditions and because their state cycles increase in length exponentially, I characterize them as chaotic.


Despite these chaotic behaviors, however, $K=N$ systems do show one startling sign of order: the number of possible state cycles (and basins of attraction) is very small. The expected number of state cycles equals the number of elements divided by the logarithmic constant e . A system with 200 elements and 2^{200} states, for example, would have only about 74 different patterns of behavior.

Moreover, about two thirds of all the possible states fall within the basins of only a few attractors—sometimes of just one. Most attractors claim relatively few states. The stability of an attractor is proportional to its basin size, which is the number of states on trajectories that drain into the attractor. Big attractors are stable to many perturbations, and small ones are generally unstable.

Those chaotic behavioral and structural features are not unique to $K=N$ networks. They persist as K (the number of inputs per element) decreases to about three. When K drops to two, however, the properties of random Boolean networks change abruptly: the networks exhibit unexpected, spontaneous collective order.

In $K=2$ networks, both the number and expected lengths of alternative state cycles fall to only about the square root of the number of elements. The state cycles of $K=2$ systems remain stable in the face of almost all minimal perturbations, and structural perturbations alter their dynamic behavior only slightly. (Networks with only a single input per element constitute a special ordered class. Their structure degenerates into isolated feedback loops that do not interact.)

It has been more than 20 years since I discovered those features of random networks, and they still surprise me. If



“FROZEN” ELEMENTS incapable of changing state can sometimes arise in a system. In this small sample network, all the elements are ruled by Boolean OR functions, and all are initially off. Changes cascade through the system after one element is turned on (*black*). Because of the configuration of the network and the Boolean functions involved, some elements (*blue*) freeze into the on state. Thereafter they will return to that state even if they or one of their inputs is altered.

one were to examine a network of 100,000 elements, each receiving two inputs, its wiring diagram would be a wildly complex scramble. The system could assume as many as $2^{100,000}$ (about $10^{30,000}$) different states. Yet order would emerge spontaneously: the system would settle into one of but 370 or so different state cycles. At a microsecond per transition, that $K=2$ network would traverse its tiny state-cycle attractor in only 370 microseconds—quite a bit less than the billions of times the age of the universe that the chaotic $K=N$ network requires.

In the ordered regime of networks with two or fewer inputs per element, there is little sensitivity to initial conditions: the butterfly sleeps. In the chaotic regime, networks diverge after beginning in very similar states, but in the ordered regime, similar states tend to converge on the same successor states fairly soon.

Consequently, in random networks with only two inputs per element, each attractor is stable to most minimal perturbations. Similarly, most mutations in such networks alter the attractors only slightly. The ordered network regime is therefore characterized by a homeostatic quality: networks typically return to their original attractors after perturbations. And homeostasis, as I shall discuss presently, is a property of all living things.


Why do random networks with two inputs per element exhibit such profound order? The basic answer seems to be that they develop a frozen core, or a connected mesh of elements that are effectively locked into either an ac-

tive or inactive state. The frozen core creates interlinked walls of constancy that “percolate” or grow across the entire system. As a result, the system is partitioned into an unchanging frozen core and islands of changing elements. These islands are functionally isolated: changes in the activities of one island cannot propagate through the frozen core to other islands. The system as a whole becomes orderly because changes in its behavior must remain small and local. Low connectivity is therefore a sufficient condition for orderly behavior to arise in disordered switching systems.

It is not a necessary condition, however. In networks of high connectivity, order will also arise if certain biases exist in the Boolean switching rules. Some Boolean functions turn elements on more often than off or vice versa. An OR function for two inputs, for example, will turn an element on in response to three out of the four possible combinations of binary signals.

A number of solid state physicists, including Deitrich Stauffer of the University of Köln and Bernard Derrida and Gerard Weisbuch of the Ecole Normale Supérieure in Paris, have studied the effects of biased functions. They have found that if the degree of bias exceeds a critical value, then “homogeneity clusters” of elements that have frozen values link with one another and percolate across the network. The dynamic behavior of the network becomes a web of frozen elements and functionally isolated islands of changeable elements.

That order, of course, is much the same as I have described for networks



PHASE CHANGES between “solid” and “gaseous” states can occur in self-regulating networks, depending on their local characteristics. If the elements’ Boolean functions are biased or if each element has only two inputs ($K = 2$), then a network in which all the elements can initially vary will eventually become stable and hence solid. Such ordered systems consist of a large web of frozen elements (blue) and isolated islands of variable elements (red). If the functions are unbiased or the interconnectedness of elements is high ($K > 3$), the system becomes a gas and behaves chaotically. Only small islands of elements will be frozen.

with low connectivity. Transient reversals in the activity of a single element typically cannot propagate beyond the confines of an isolated island and therefore cannot cause much damage. In contrast, if the level of bias is well below the critical value—as it is in chaotically active systems—then a web of oscillating elements spreads across the system, leaving only small islands of frozen elements. Minimal perturbations in those systems cause avalanches of damage that can alter the behavior of most of the unfrozen elements.

Christopher Langton, a computer scientist at Los Alamos National Laboratory, has introduced an analogy that helps one think about the change between order and disorder in different ensembles of networks. He has related network behavior to the phases of matter: ordered networks are solid, chaotic networks are gaseous and networks in an intermediate state are liquid. (The analogy should not be interpreted too literally, of course: true liquids are a distinct phase of matter and not just a transitional regime between gases and solids.)

If the biases in an ordered network are lowered to a point near the critical value, it is possible to “melt” slightly

the frozen components. Interesting dynamic behaviors emerge at the edge of chaos. At that phase transition, both small and large unfrozen islands would exist. Minimal perturbations cause numerous small avalanches and a few large avalanches. Thus, sites within a network can communicate with one another—that is, affect one another’s behavior—according to a power law distribution: nearby sites communicate frequently via many small avalanches of damage; distant sites communicate less often through rare large avalanches.

These characteristics inspired Langton to suggest that parallel-processing networks poised at the edge of chaos might be capable of extremely complex computations. On the face of it, the idea is plausible. Highly chaotic networks would be so disordered that control of complex behaviors would be hard to maintain. Highly ordered networks are too frozen to coordinate complex behavior. But as frozen components melt, more complicated dynamics involving the complex coordination of activities throughout a network become feasible. The complexity that a network can coordinate peaks at the liquid transition between solid and gaseous states.

Systems poised in the liquid transi-

tion state may also have special relevance to evolution because they seem to have the optimal capacity for evolving. As Darwin taught, mutations and natural selection can improve a biological system through the accumulation of successive minor variants, just as tinkering can improve technology. Yet not all systems have the capacity to adapt and improve in that way. A complex program on a standard computer, for example, cannot readily evolve by random mutations: almost any change in its code would catastrophically alter the computation. The more compressed the code, the less capacity it has to evolve.

Networks on the boundary between order and chaos may have the flexibility to adapt rapidly and successfully through the accumulation of useful variations. In such poised systems, most mutations have small consequences because of the systems’ homeostatic nature. A few mutations, however, cause larger cascades of change. Poised systems will therefore typically adapt to a changing environment gradually, but if necessary, they can occasionally change rapidly. These properties are observed in organisms.

If parallel-processing Boolean networks poised between order and chaos can adapt most readily, then they may be the inevitable target of natural selection. The ability to take advantage of natural selection would be one of the first traits selected.

The hypothesis is bold, perhaps even beautiful, but is it true? Physicist Norman H. Packard of the University of Illinois at Champaign-Urbana may have been the first person to ask whether selection could drive parallel-processing Boolean networks to the edge of chaos. Sometimes at least the answer is yes. Packard found such evolution occurring in a population of simple Boolean networks called cellular automata, which had been selected for their ability to perform a specific simple computation.

Recently my colleague Sonke Johnsen of the University of Pennsylvania and I have found further evidence of evolution proceeding to the edge of chaos. We have begun studying the question by making Boolean networks play a variety of games with one another [see box on opposite page]. Our results, too, suggest that the transition between chaos and order may be an attractor for the evolutionary dynamics of networks performing a range of simple and complex tasks. All the network populations improved at playing the games faster than chance alone could

accomplish. The organization of the successful networks also evolved: their behaviors converged toward the boundary between order and chaos.

If these results hold up under further scrutiny, then the liquid transition between ordered and chaotic organizations may be the characteristic target of selection for systems able to coordinate complex tasks and adapt. By that reasoning, such poised systems should occur in biology.

How much order and chaos do the genomic systems of viruses, bacteria, plants and animals exhibit? Usually each gene is directly regulated by few other genes or molecules—perhaps no more than 10. The Boolean wiring diagram for the genome is therefore sparse, and the individual gene elements have few inputs. Furthermore, almost all known regulated genes are governed by a particular class of Boolean switching rules called canalizing functions. In canalizing functions, at least one input has a value that can by itself determine the activity of the regulated element. (The OR function is a typical canalizing function.)

Like low connectivity or biases in the Boolean rules, an abundance of canalizing functions in a network can create an extensive frozen core. Increasing the proportion of canalizing functions

used in a network can therefore drive the system toward a phase transition between chaos and order. Because genomic regulatory systems are sparsely connected and typically appear to be governed by canalizing functions, such networks are very likely to exhibit the traits of parallel-processing systems with frozen percolating elements: a modest number of small, stable attractors, the confinement of damage to small cascading avalanches and modest alterations in dynamics in response to mutations.

One interpretation of the meaning of antichaos in complex systems has particular relevance to biology: a cell type may correspond to an attractor in the genomic dynamics. A genome that contains 100,000 genes has the potential for at least $10^{30,000}$ patterns of gene expression. The genomic regulatory network orchestrates those possibilities into changing patterns of gene activity over time. But a stable cell type persists in expressing restricted sets of genes. The natural suggestion is that a cell type corresponds to a state-cycle attractor: it embodies a fairly stable cycle of expression in a specific set of genes.

Given that interpretation, the spontaneous order arising in networks with

low connectivity and canalizing Boolean functions sets up several predictions about real biological systems. First, each cell type should correspond to a very small number of gene expression patterns through which it cycles. One can therefore calculate how long such cell cycles should be.




After receiving an appropriate stimulus, a gene in a eukaryotic cell needs about one to 10 minutes to become active. The length of an attractor in a genome with 100,000 genes would be about 370 states. Consequently, a cell should run through all the gene expression patterns of its type in roughly 370 to 3,700 minutes. This figure approximates the correct range for real biological systems. As predicted, the length of cell cycles does seem to be proportional to roughly the square root of the amount of DNA in the cells of bacteria and higher organisms.


If a cell type is an attractor, it should be possible to predict how many cell types could appear in an organism. The number of attractors is about the square root of the number of elements in a network; therefore, the number of cell types should be approximately the square root of the number of genes. If we assume that the number of genes is proportional to the amount of DNA in a cell, then humans should

The Mismatch Game

A computer game demonstrates how natural selection can push random networks to a point near the edge of chaos. Pairs of networks receive high scores for having patterns in six of their elements, designated as "matchsites," that differ maximally from one another. After a comparison, each network takes the pattern in the matchsites of the other as an input, works on its response and makes another comparison [see below]. After 10 comparisons, each network finds a new opponent. The networks with the highest scores reproduce preferentially to form the next generation. Occasional mutations randomly alter the local characteristics of the networks. Over several generations, the networks generally improve at playing the game. Regardless of their initial conditions, the networks approach the boundary between order and chaos.

These changes can be observed by measuring how a network's trajectories diverge or converge at various times. D_T is the percentage of the elements that have different activities in two initial states of one network. D_{T+1} is that percentage measured for the successor states. When D_T is small in very chaotic ($K_i = 5$) networks, trajectories diverge ($D_{T+1} > D_T$). In very ordered ($K_i = 1$) networks, trajectories converge ($D_{T+1} < D_T$). After several generations of playing the Mismatch Game, chaotic and ordered networks moderate their behavior [see graphs at right].





NUMBER OF CELL TYPES in organisms seems to be related mathematically to the number of genes in the organism. In this diagram the number of genes is assumed to be proportional to the amount of DNA in a cell. If the gene regulatory systems are $K=2$ networks, then the number of attractors in a system is the square root of the number of genes. The actual number of cell types in various organisms appears to rise accordingly as the amount of DNA increases.

have about 100,000 genes and 370 cell types. By the most recent count, humans have about 254 distinct cell types, so that prediction is also in the right range.

Across many phyla, the number of cell types seems to increase with approximately the square root of the number of genes per cell (that is, with the number of genes raised to a fractional power that is roughly one half). Thus, bacteria have one or two cell types, sponges have perhaps from 12 to 15 and annelid worms have about 60.

Because not all DNA may have a function, the number of genes may not rise directly with the amount of DNA. The predicted number of cell types could therefore increase according to a fractional power greater than one half (the square root) but less than one. In fact, by conservative estimates, the number of cell types appears to increase at most as a linear function. Such a range of behavior is found in complex Boolean networks. In contrast, other simple mathematical models for genomic systems predict that the number of cell types would increase exponentially with the number of genes.

Another prediction refers to the stability of cell types. If a cell type is an attractor, then it cannot be altered by most perturbations: its stability is an emergent property of the gene regulatory system.

Differentiation, according to this model, would be a response to perturbations that carried a cell into the basin of attraction for another cell type. In a canalizing ensemble, however, each model cell can differentiate directly into only a few alternative cell types because each attractor is "near" only a few others. Consequently, ontological development from a fertilized egg should proceed by successive branching pathways of differentiation. In other words, once a cell has begun to differentiate along certain lines, it loses the choice of differentiating in other ways. As far as biologists know, cell differentiation in multicellular organisms has been fundamentally constrained and organized by successive branching pathways since the Cambrian period almost 600 million years ago.

In canalizing networks, order emerges because a large fraction of the binary elements falls into a stable, frozen

state. That stable core of elements is identical in almost all the attractors. Hence, all the cell types in an organism should express most of the same genes. Typically only a few percent of the genes should show different activities. Both claims hold true for biological systems.

The attractor model for cell types also predicts that the mutation of a single gene should usually have fairly limited effects. Avalanches of damage (or changed activity) caused by the mutation should not propagate to the vast majority of genes in the regulatory network. Changes in activity should be restricted to small, isolated islands of genes. These expectations are met by real genetic systems.

Moreover, the expected sizes of the unfrozen islands in the gene systems come close to predicting the sizes of such avalanches. For example, a hormone called ecdysone in the fruit fly *Drosophila* can unleash a cascade that changes the activity of about 150 genes out of at least 5,000. The expected size of avalanches in canalizing genomes with 5,000 elements or in those with low connectivity and a frozen core containing roughly 80 percent of the genes is about 160.

Taken as models of genomic systems, systems poised between order and chaos come close to fitting many features of cellular differentiation during ontogeny—features common to organisms that have been diverging evolutionarily for more than 600 million years. The parallels support the hypothesis that evolution has tuned adaptive gene regulatory systems to the ordered region and perhaps to near the boundary between order and chaos. If the hypotheses continue to hold up, biologists may have the beginnings of a comprehensive theory of genomic organization, behavior and capacity to evolve.

FURTHER READING

PHASE-TRANSITIONS IN TWO DIMENSIONAL KAUFFMAN CELLULAR AUTOMATA. B. Derrida and D. Stauffer in *Europhysics Letters*, Vol. 2, No. 10, pages 739-745; 1986.

RANDOM BOOLEAN NETWORKS: ANALOGY WITH PERCOLATION. D. Stauffer in *Philosophical Magazine B*, Vol. 56, No. 6, pages 901-916; 1987.

LECTURES IN THE SCIENCES OF COMPLEXITY. Edited by Daniel L. Stein. Addison-Wesley, 1989.

ORIGINS OF ORDER: SELF-ORGANIZATION AND SELECTION IN EVOLUTION. Stuart A. Kauffman. Oxford University Press (in press).