SPIN GLASS I: A SCALING LAW RESCUED

Philip W. Anderson

The history of spin glass may be the best example I know of the dictum that a real scientific mystery is worth pursuing to the ends of the Earth for its own sake, independently of any obvious practical importance or intellectual glamour. If a phenomenon seems likely to contradict the fundamental principles you thought you understood (that's what I mean by a *real* mystery—so long as you also believe the experiments!) you should stick with the phenomenon. This fundamental dictum of good science is increasingly neglected by our masters who provide the money; spin glass is even less popular with them than superconducting materials were before 1987. The pursuit of the spin glass mystery led, *inter alia* and aside from all the good solid-state physics that resulted, to new algorithms for computer optimization, a new statistical mechanics, a new view of protein structure, a new view of evolution and new ideas in neuroscience.

But it all started with some very simple physics. Measuring magnetic resonance and magnetic properties of dilute magnetic ions in insulators (such as Mn in ZnO) to probe magnetic interactions was a very useful game played in the 1950s and early 1960s by John Owen, among others. At about the same time, he and the Berkeley magnetic resonance group of Arthur Kip, Walter Knight, Charles Kittel and others tried diluting manganese into the nonmagnetic metal copper to test the proposed Ruderman–Kittel–Kasuya–Yosida interaction between spins via free electrons, and to see whether the free electrons' exchange interaction with the magnetic ions affected their resonance behavior—that is, to see if there was an electron magnetic resonance equivalent to the nmr Knight shift and the Korringa free-electron relaxation of nuclei.

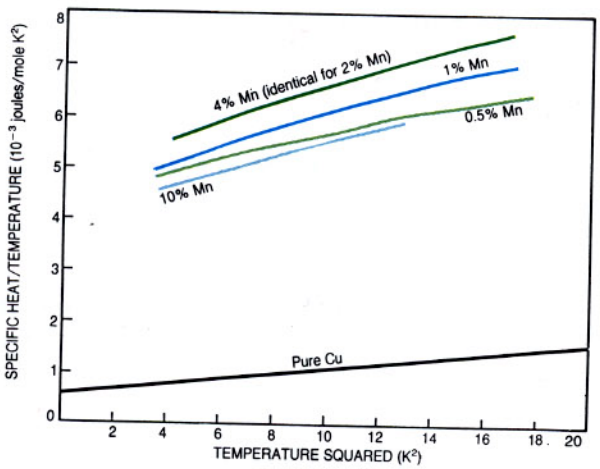

Specific heat of Cu–Mn alloys at low temperatures is independent of the manganese concentration, and it varies linearly with temperature when the contribution due to pure copper is subtracted. In the range of manganese concentrations shown here, the low-temperature properties of the alloy are dominated by the interaction between magnetic moments on manganese atoms, which leads to the spin glass behavior. (Adapted from J. E. Zimmerman, F. E. Hoare, *J. Phys. Chem. Solids* **17**, 52, 1960.)

This simple set of experiments opened up at least two cans of worms. The physics of the effect of the free electrons on the individual Mn ions is, of course, closely related to the Kondo problem and the Anderson model, but this is not our worm-receptacle of choice. We are interested in what they saw when the samples contained 0.1–10% manganese and the behavior was dominated by the ion–ion interactions.

The magnetic resonance effects were not very instructive at the time, involving complicated physics that was actually one of the last areas to be clarified. I will discuss this physics in a later column. What first attracted attention were simple thermodynamic measurements such as those of magnetic susceptibility and specific heat.

The extra specific heat at low temperature was perfectly linear in T, with a slope about five times that of the background copper, independently of concentration, but then the extra specific heat peaked at a concentration-dependent temperature T_0 and dropped down at a rate one could imagine to be approximately $1/T^2$. (See the figure above.) The susceptibility χ rose from a constant value to an (apparently) broad peak near T_0, then fell off at a good approximation to the Curie law $\chi = C/T$ (C is the Curie constant) with the right Bohr magneton number for Mn^{++} with its $g = 2$, $S = {}^5/_2$ ground state.

After a bit of to-ing and fro-ing, a

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

group at Grenoble showed that for these "classic spin glasses" (as they much later came to be known) all of these data fit onto a single universal curve with the temperature, energy and entropy scales all linearly proportional to concentration. We know now that this is a good confirmation of the extraordinarily simple underlying physics—that the dominant term in the Hamiltonian for the magnetic moments of solute atoms is the RKKY exchange integrals

$$J(r_{ij}) = J_0 \frac{1}{r^3} \cos(2k_F r_{ij})$$

where k_F is the Fermi momentum. These RKKY exchange integrals fall off slowly with distance and alternate in sign.

In these random, dilute solutions this behavior has the effect of randomizing the exchange integrals (since $1/(2k_F)$ is less than an interatomic distance) and each spin sees quite a few neighbors with each of which it has an exchange interaction that is random in both sign and magnitude. The probability distribution of these exchanges has a scale proportional to $\langle 1/r^3 \rangle_{avg}$, which is in turn equal to the density of solute atoms. The form of the probability distribution, however, is independent of the density except for the scale.

A 1987 postscript to this story was written by Michael Stephen and Elihu Abrahams. The long-range $\cos(2k_F r)/r^3$ form occurs because of the sharp drop in occupation of free-electron states at k_F. Pierre-Gilles de Gennes had proposed, long before, that scattering of the free electrons by impurities reduced the range of this effect by a factor $\exp(-r/l)$, where l is the mean free path of the electrons, by dephasing the electronic wavefunctions at the Fermi surface; this unfortunately would ruin the scaling law for the probability distribution of the interaction. Until recently this argument was universally accepted, but it turns out not to be true, as Stephen and Abrahams have shown: The distribution is almost unchanged by scattering, because scattering merely shifts the phases of the wavefunctions near the Fermi surface without changing their relative phases in any given sample. This is a vital key to the properties of real spin glasses, and in addition restores the scaling laws.

But, as you'll see in my subsequent columns, aside from the scaling law (which has only just been rescued) none of the features of these measurements were to be understood for another 20 years or more, and the real physics is still a problem. ■

SPIN GLASS II: IS THERE A PHASE TRANSITION?

Philip W. Anderson

In the late 1950s and early 1960s Jim Kouvel, then at GE, and Paul Beck's group at the University of Illinois spent a lot of time exploring a phenomenon Paul called mictomagnetism. This phenomenon took place in dilute solutions of Mn atoms in Cu (and of other magnetic atoms in other nonmagnetic metals); I discussed these materials in my column in the January issue (page 9). These solutions, as I remarked, seemed to have small linear magnetic susceptibilities of typically paramagnetic magnitude (a few times 10^{-4} in dimensionless units). But Kouvel and Beck showed that the solutions exhibit, at a tiny scale and at very low temperatures, and in addition to the linear susceptibility, many of the phenomena typical of ferromagnetism: hysteresis, remanence and so on. In some ways these solutions are *more* hysteretic than ferromagnets, in that they can remember the sign and direction of the field they were cooled in, even when one applies an opposing field large enough to polarize them in the opposite direction.

Meanwhile Bernd Matthias and the rest of us at Bell Labs were very interested in the possibility that magnetism and superconductivity might coexist. Within the BCS theory the two should be quite incompatible, and in many cases they are; but (some 30 years too soon!) Bernd was determined to show that in at least some cases there would be a close relationship. In some of his dilute solutions of magnetic ions in superconductors (like Gd in $CeRu_2$) he noted the presence of the vague susceptibility peaks and remanences characteristic of spin glasses, and so he said: "Aha! Ferromagnetism and superconductivity are *not* incompatible!"

I always tried to listen more carefully to what Bernd's *results* said than what *he* said, since he had little regard for fine distinctions in statistical physics (like that between ferro- and antiferromagnetism, for instance, or between these and some vague bump in the susceptibility), but this is a case where he got to me. I was so certain that the transitions he was talking about were not true ferro- or antiferromagnetism that I failed to note what he had noted, that the transitions seemed remarkably sharp. I was particularly certain that a magnetic transition would involve a significant change in entropy and hence would certainly dominate the tiny energies and entropies of the superconducting state. (This was almost a decade before 1972, when Michael Kosterlitz and David Thouless, following my work on peculiar one-dimensional models, first showed that a phase transition could show no specific heat singularity at all.) Yet these bumps didn't seem to disturb the superconducting transition very much, which I felt meant that they were not phase transitions.

It is a bit ironic that only two or three years later, in 1965, an obscure journal called *Physics*, edited by none other than Bernd and myself, published the first evidence that there really was a spin glass transition, without either of us (or possibly even John C. Wheatley, the author) noticing. Wheatley was interested in testing his then new SQUID magnetometers in an interesting system and chose these same dilute solutions of Mn in Cu. His susceptibilities (measured, perforce, in a tiny magnetic field) followed a very precise Curie law C/T for each solution down to a temperature T_c, which was very exactly proportional to concentration, and then, as abruptly as he could measure, stopped changing with T and became constant. (Note that unlike the older measurements, Wheatley's did not exhibit a peak, because he cooled in a fixed magnetic field; a constant value of the susceptibility is characteristic of spin glasses when they are cooled in such a field.)

It was not until 1970 that the key measurements that woke the rest of us up to this peculiar transition were made—by Vincent D. Canella, John A. Mydosh and Joseph I. Budnick. This group measured ac magnetic susceptibilities with sensitive, but more conventional, methods, and discovered that the key variable is the magnitude of the measuring field. At 1000 gauss, there is only the conventional vague hump; at 1 gauss, a sharp, cusp-like peak appears whose width is less than 1% of T_c. Yet 1 gauss is 10^{-5} the magnitude of the internal field, since T_c is approximately 10 K. This tremendous nonlinearity is the appropriate characterization of the transition; later measurements, by P. Monod and Hélène Bouchiat, for instance, showed that the nonlinear susceptibility $\partial^2\chi/\partial H^2$ diverges as $(T-T_c)^{-P}$, where P is greater than 1. Thus, experimentally there is no doubt that the transition exists and is an equilibrium transition, since the nonlinearity can be measured above the transition point, where no one doubts that equilibrium is established in the system—after all, its natural relaxation frequency should be about 10^{12}–10^{13} sec^{-1}. Nonetheless, no measurement has ever revealed a specific heat singularity at $H=0$. As we shall see in the next column, the theoretical acceptance of a true phase transition, as well as an understanding of its nature, was much slower to come; and the most striking feature, the nonlinearity, is yet to be calculated, even roughly. ■

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

SPIN GLASS III: THEORY RAISES ITS HEAD

Philip W. Anderson

Sam Edwards was finishing out his term as head of the SRC, the British equivalent of the NSF, during the winter of 1974–75, after being appointed to a professorship at Cambridge. (He is now Sir Sam, the Cavendish Professor, successor in that chair to James Clerk Maxwell and four or five Nobel laureates, including Sir Nevill Mott.) Being Sam, he was unfazed by the full-time SRC job and needed research to do on the train back and forth to London, so he dropped in every Saturday at the Cavendish Laboratory for coffee and a chat with me and the theory group. I made a point that year of being there on Saturdays as well as during the week, and we did a lot of talking about localization and the "Fermi glass" (that is, the problem of electrons frozen in place by localization and interactions), the theory of liquids and the glass transition, and other problems of mutual interest.

One of these problems was that of dilute magnetic alloys, which seems to have acquired the name "spin glass" in a 1970 paper I wrote with Wai-Chao Kok (now at Singapore University) for a 65th-birthday festschrift for Mott in the *Materials Research Bulletin*. (See my columns in the January and March issues of PHYSICS TODAY.) I described to Sam the old mystery of continuous, disordered freezing in these alloys, and the new mystery of the sharp cusps and nonlinear behavior that John A. Mydosh had reported. Sam's ears pricked up. He had a notebook full of methods he had been trying on gelation, the glass transition and various polymer questions, but had been frustrated because these are not clean, well-posed problems. (He had also tried the methods on localization; and later, in the hands of Franz Wegner and Shinobu Hikami, they did work reasonably well on that problem—but some extra frills were needed.) I was convinced that the random Heisenberg Hamiltonian

$$H = \sum_{(ij)} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j$$

where $\mathbf{S}_i$ is a classical spin vector at site i and J_{ij} is the interaction between the spins on sites i and j, was almost certainly the proper statement of the spin glass problem. Sam was overjoyed when he learned this, because here was a nice clean problem to work his methods on.

The methodology of the resulting, justifiably famous paper was almost entirely his. (That's why I can make such an immodest statement.) But the basic physical concept we worked out together. We decided that the thing to do was to ignore the spatial ordering, that is, to neglect the long-range ordering of spins in space, if any, and instead to look for long-range order in *time*. Richard Palmer later named this concept "nonergodicity" because it means, when present, that the system does not explore all possible states in the course of time. As the measure of long-range order in time, we introduced q, which is the average correlation between a spin $\mathbf{S}_i$ measured at one time and the same spin measured a macroscopic time t later. The equations are

$$q(t) = \langle \mathbf{S}_i(0) \cdot \mathbf{S}_i(t) \rangle_{\text{ave over } i}$$

$$q = \lim_{t \to \infty} q(t)$$

First we did a little physical calculation of the transition temperature below which q became nonzero. One separates out one of the sites, say i, and assumes that all the neighboring sites have a finite value q_0 of the "spin glass order parameter" q. Then one calculates the correlation enforced on site i by the effective fields due to the other sites, and finally one averages it over all sites i in the sample. For self-consistency, the order parameter thus calculated must have the value q_0 initially assumed for sites that are neighbors of i. The order parameter we calculated,

$$q_0 = \langle \langle \mathbf{S}_i(0) \cdot \mathbf{S}_i(\infty) \rangle_{\text{ave}} \rangle_{\text{ave over } i}$$

had a nonzero solution below a certain T_c. (The subscript "ave" stands for the statistical mechanics average over thermal fluctuations.)

Much more devious is Sam's so-called replica method of calculating thermodynamic properties. To do thermodynamics properly, one must average extensive quantities such as free energy and entropy. These are all derivable from the logarithm of the partition function

$$Z = \mathrm{Tr} \exp(e^{-\beta H})$$

which grows exponentially with the size of the system. It is dangerous—in fact wrong—to average the partition function in random systems like the spin glass. This is because the partition function fluctuates too much: Special configurations, such as regions where all the J_{ij}'s are accidentally positive, will dominate its average value.

The key point of principle that makes studies of the spin glass and similar systems difficult problems is this: They are "quenched" random systems, with the values of J_{ij} fixed for all time by the conditions of preparation of the sample. But we want to average over macroscopic samples, in which many different configurations of J_{ij}'s occur, in such a way that the average represents the behavior of a *typical* system and is the proper "extensive" or "intensive" thermodynamic quantity that varies sensibly with the size of the system. Sam recalled the obvious identity

$$\ln Z = \lim_{m \to 0} \frac{(Z)^m - 1}{m}$$

The necessary average is then that of the m-th power of Z, not of its log; but when m is small this is no easier.

Now we do an outrageous thing: We note that the average of Z^m for $m = 1,2,3,4,\ldots$ is calculable because it is the average of an exponential containing J. For m an integer,

$$\langle Z^m \rangle = \mathrm{Tr}_{(\mathbf{S}^\alpha)} \int \exp\left(-\beta \sum_{\alpha=1}^{m} \sum_{ij} J_{ij} \mathbf{S}_i^\alpha \cdot \mathbf{S}_j^\alpha \right) \times P(J_{ij}) \mathrm{d}(J_{ij})$$

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

which is the average over m identical "replicas" of the system. If

$$P(J_{ij}) \propto \exp(-J_{ij}^2/2J^2)$$

this integration is easily done. It gives rise to the following type of statistical problem:

$$\langle Z^m \rangle = \operatorname*{Tr}_{(S^\alpha)} \exp\left[\frac{\beta^2 J^2}{2}\left(\sum_{\alpha=1}^{m} \mathbf{S}_i^\alpha \cdot \mathbf{S}_j^\alpha\right)^2\right]$$

This is no longer a random problem, but a regular one. It is more difficult, however, because it is biquadratic in spins. It can be solved in mean-field theory—assuming that $\mathbf{S}_i^\alpha$ and $\mathbf{S}_j^\alpha$ are uncorrelated (correlation only gives terms of order m^2) but that $\mathbf{S}_i^\alpha$ and $\mathbf{S}_i^\beta$ *are* correlated. It can also occasionally even be solved by renormalization-group methods. But one has the awful problem of extrapolating from all positive integral values of m to small m. In principle it is not rigorous to take the limit as m goes to 0 when the function is known only for integral values of m. In practice, however, it turns out to be easy since one keeps only terms of order m. What is more, so far none of the real difficulties encountered in the spin glass theory seem to have come from failure of the mathematical extension to $m \to 0$! Recently Haim Sompolinsky (*Phys. Rev. B* **25**, 6860, 1982) and Miguel Virasoro (*Europhys. Lett.* 1, 77, 1986) have given us some ideas about why that is true.

In the mean-field solution the "Edwards–Anderson order parameter" q reappears in a new guise, as a replica–replica correlation function

$$q_{\alpha\beta} = \langle \mathbf{S}_i^\alpha \cdot \mathbf{S}_i^\beta \rangle$$

Thus in some real sense the different replicas represent very widely separated instants in time at which we choose to look at the same system.

Sure enough, the mean-field theory we worked out showed a nice sharp cusp in the susceptibility, in qualitative agreement with experiment, and weakly nonlinear behavior, qualitatively correct but too small. Unfortunately it also gives a cusp in the specific heat, which to this day has never been seen, and which is surely unphysical for real, finite-dimensional spin glasses. Nonetheless the result, giving a sharp freezing transition and describing a true nonergodicity, seemed sufficiently promising that we felt that the replica methodology was the doorway into the problem and that final solutions were just around the corner.

Little did we know! See next time, when I reveal the Negative-Entropy Catastrophe. ■

SPIN GLASS IV: GLIMMERINGS OF TROUBLE

Philip W. Anderson

In my last column (June, page 9), it seemed as though Sam Edwards's beautiful "replica" scheme had brought us to a highly satisfactory resolution of the old problem of magnetic systems with random exchange interactions—what we now call "spin glasses." (In the replica method, one calculates the partition function of n replicas α of the same random Hamiltonian, averages over the randomness and takes the logarithm by studying the formal limit as $n \to 0$: an indirect, shaky but often useful procedure.) In 1975 David Sherrington, who had been Sam's student and is now at Imperial College, London, tried applying the methods and ideas of the Edwards–Anderson paper to an especially simple model in which the "mean field" version should certainly be exact. In Sherrington's model, every spin in a macroscopic sample of N spins is connected by a random exchange integral $J_{ij} \times 1/\sqrt{N}$ to every other spin. This is precisely the kind of artificial system for which mean-field theory is exact in other magnetic models. Sherrington brought the model with him that summer on a visit to IBM (Yorktown Heights), where he worked with Scott Kirkpatrick. The model is now famous as the SK model. Their conclusion was that the EA method led to a solution that, while superficially plausible, was unequivocally *nonsense*—specifically, they showed that as the temperature approached zero the calculated entropy passed through zero and became negative. Since entropy is the log of an integer (the number of states at energy E), its acquiring a negative value is forbidden in statistical mechanics. The energy near $T = 0$ also seemed to be a little lower (by about 2 percent) in the SK solution than the best that Scott could achieve by simulating the model on a computer.

Naturally, everyone at first assumed that the replica method itself was at fault. In fact David Thouless, Richard Palmer and I set out to produce a solution directly, without the replica method. This so-called TAP theory (1977) adapted the ancient cavity-field method of Lars Onsager and Hans Bethe to include a local-field correction for the response of all the spins affected by the fluctuations of a given spin. This correction, which is absolutely negligible (of order $1/N$) in the corresponding long-range ferromagnet, is finite here, changing T_c by a factor of 2, for instance. But we could "prove" that all further corrections *were* negligible. The results agreed with the SK findings near and above T_c, where in fact we now know that both solutions are right, but they deviated subtly below T_c. One important difference was that we got rid of the negative entropy of the SK solution.

Again we thought we had the answer, and again we were to be disappointed, though the problem surfaced in more subtle ways this time. Central to the TAP solution is a mean-field equation,

$$m_i = \tanh \frac{h_i}{2k_B T}$$

where

$$h_i = \sum_j J_{ij} m_j - (\text{local field correction})$$

Here m_i is the mean magnetization at site i. Unlike in the simple case of a ferromagnet, where similar equations are encountered, looking for a nontrivial solution for the magnetization m_i in this equation involves an infinite random-matrix problem at every T. Near T_c, this problem appeared to be just expressible in terms of the known statistical properties of the eigenvalues of the random matrix J, and near zero it depends only on properties of "the" solution at $T = 0$. The former case we solved in lowest order, and it seemed to look OK; and for the latter

"WHAT YOU HAVE DONE, GRUNDIG, IS HELP CONTRIBUTE TO A DIS-UNIFIED THEORY."

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

case, Richard set out to calculate "typical" ground state solutions (that is, solutions at $T=0$). The limiting form of the above set of random equations at $T=0$ is

$$m_i = \text{sign}(h_i)$$

$$h_i = \sum_j J_{ij} m_j$$

("Sign" just means that m points in the same direction as h and has the maximum possible magnitude.) Both Richard and Scott had been trying to solve these equations numerically for some time.

Both of them gradually came to the same paradoxical result: They could find no "the" solution to this set of equations. Instead, they found many, many solutions of nearly identical energy. They also noticed that it is very difficult, once one's computer has found one solution, to persuade it to move to another, even if the first has a much higher energy than the optimal one. Incidentally, both Richard and Scott found that the easiest way to find a new solution was to raise the temperature nearly to T_c and come back down again—a procedure that Scott called "simulated annealing."

This peculiar feature, enormously annoying at the time, was the beginning of one of the important discoveries of modern theoretical physics, a discovery comparable to that of chaos in its broad applicability to science. But we didn't quite understand that yet.

Because of this unusual feature, and also for other reasons—Thouless, for instance, was unhappy that our solution near T_c might not be quite stable—the TAP "solution" still did not satisfy. We also needed to know why the replica method had failed. Thouless and a student, Jairo de Almeida, soon discovered the rather unexpected reason. Below a certain line in magnetic field–temperature space, a solution with "replica symmetry"—that is, where every replica has the same correlation $q = q_{\alpha\beta}$ with every other replica—is dynamically unstable to "replica symmetry breaking." This implied that there was some new structure in $q_{\alpha\beta}$ that depended on α and β: According to the ideas that underlay the Edwards–Anderson paper, then, not every time you tried to compare one specimen of a system with another specimen of the *same* system would you get the same answer! Looking back, it seems obvious that this was closely related to the simulation problem, but it was a few years before we caught on to that. In my next column I'll try to explain the final resolution. ■

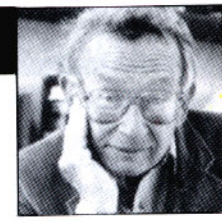

SPIN GLASS V: REAL POWER BROUGHT TO BEAR

Philip W. Anderson

Gérard Toulouse had always been interested in the spin glass problem. In 1977, subsequent to the work I discussed in my last column (September 1988, page 9), Gérard, then at the Ecole Normale Supérieure in Paris, began to discuss the problem with the powerful "Cargèse" group of theoretical physicists in Paris and Rome: Cyrano de Dominicis, Giorgio Parisi and Miguel Virasoro in particular, and later Bernard Derrida, Nick Sourlas and others. Gérard was the originator of the formal theory of "frustration" (I believe I introduced the *term* originally) as *the* important feature of the spin glass problem. Because the exchange bonds J between spins have random signs in most circuits (loops) of spins returning upon themselves, not all the spins can be made "happy"—hence the "frustration." In a square of four spins, for instance, only if an even number of the J's have the same sign can one satisfy everybody—that is, find a unique minimum-energy configuration of the four spins. Since the world is made up of systems of conflicting desires, from game strategies to a group of people choosing a menu, one begins to see that the spin glass is not that bad a model for many aspects of life.

It was Giorgio Parisi who developed the replica-symmetry-breaking scheme that solved the problems, such as negative entropy, that had been troubling us. You will remember from my column of June 1988 (page 9) that Sam Edwards had introduced the "replica" scheme of making m identical copies $1, \ldots, \alpha, \ldots, m$ of the same system and calculating the average of the product of all m partition functions. One then schematically takes the limit $m \to 0$ and uses the formula

$$\lim_{m \to 0} \frac{Z^m - 1}{m} = \ln Z$$

to calculate the free energy

$$F = -kT\langle \ln Z \rangle_{\text{ave}}$$

David Thouless and Jairo de Almeida later showed that not all pairs of copies gave the same average correlation

$$q_{\alpha\beta} = \langle S_i^\alpha S_i^\beta \rangle_{\text{ave over } i}$$

This finding was called "replica symmetry breaking." Giorgio was able to produce a form of $q_{\alpha\beta}$ that worked, in the sense that it was a self-consistent and stable solution of the equations in this replica formalism. It would not be wise for me to go through the complicated mathematical structure that evolved from this beautiful and unexpected solution. Instead I shall try to describe what was eventually understood about its implications—with contributions from Toulouse, Virasoro, and later Haim Sompolinsky, Peter Young, Richard Palmer and many others.

What Giorgio's solution means is that at any temperature below T_c there is no *unique* locally stable thermodynamic state that solves the "TAP" mean-field equations I described in my last column, but rather many such states, which resemble one another to different degrees. Each replica α corresponds to a different solution of the TAP equations; the solutions can be thought of as clusters of states in the N-dimensional configuration space of the N spins. The TAP equations are obtained when the thermodynamic average is restricted to these local clusters of states. The off-diagonal terms in the order parameter $q_{\alpha\beta}$, which represent the average overlap between states in the cluster belonging to replica (solution) α and those belonging to replica β, are a measure of the degree of resemblance between clusters α and β. The diagonal elements $q_{\alpha\alpha}$, which are all the same, are the average overlaps of

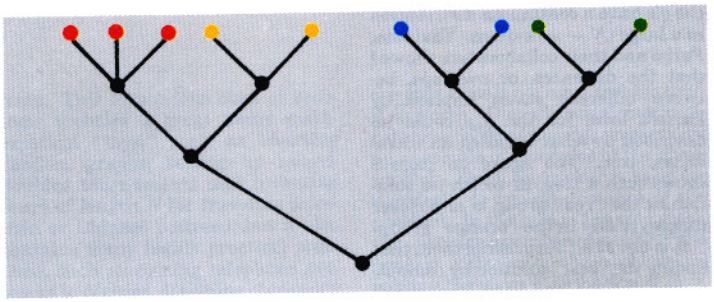

Ultrametric tree. This structure is a convenient way to represent the degree of resemblance between spin glass states (colored dots). The overlap between any pair of states of the same color is q_1; that between "red" and "orange" or between "blue" and "green" states is $q_2 < q_1$. Similarly, the overlap of any of the five states on the left with any of the four states on the right is $q_3 < q_2$. Thus the overlap between two states depends on how deep into the tree one has to go to find a node (black dot) that connects them. One may verify from the figure the amazing property that when any three states are picked at random at least two of the overlaps are equal.

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

states within a given replica, or cluster. There is a hierarchy of overlaps, with $q_{\alpha\alpha} \equiv q_0$ being the largest. The next in value is q_1, the overlap between the inequivalent groups of solutions α and β that are closest in the configuration space. Then there is a $q_2 < q_1$ and so on.

At T_c the solutions begin to separate, and the "distance" between them, measured by the deviation of their overlaps from unity, increases until, at $T = 0$, q_0 is 1 but the smallest $q_{\alpha\beta}$ may be nearly zero. Of course, the q's have a continuous distribution in a large ($N \to \infty$) system. Toulouse, Parisi and their collaborators showed that the distances, or overlaps, between different states implied by Parisi's form for the $q_{\alpha\beta}$ could be described by what is called an ultrametric tree. The figure on page 9 shows such a tree, in which no solution in the "red" group is any closer than q_2 to any in the "orange" group.

It is not at all surprising, then, that finding the "best" solution by computer simulations had been impossible: The solutions that separated at T_c became increasingly different as T was lowered. From thermodynamics and the extensive nature of the thermodynamic variables such as entropy and energy, one can show explicitly, as I did, that the only route from one set of solutions to another—through configuration space—passes over energy barriers whose height grows with N, the total number of spins. Thus if you try to get from one solution to another by flipping spins a few at a time, you must make flips that increase the energy by amounts of order N before you can ever get to one of the other solutions or, in particular, to the *best* one. Thus one can represent the solutions as deep valleys connected only by very high passes in a "rugged energy landscape" (to use Stu Kauffman's terms). This is a remarkable result—how truly remarkable and powerful we are only beginning to understand. It implies, among other things, a new thermodynamics—a thermodynamics of systems that are never in thermodynamic equilibrium. Richard Palmer and I called these systems "nonergodic." That one can nonetheless use statistical mechanical methods to get not only the quantitative solutions relevant to such systems but also the structure of the set of solutions is, to say the least, fantastic.

Next time I shall begin my discussion of the implications of this work in fields as far apart as computer science, biology and neuroscience, which normally have been quite outside the purview of physics. ■

SPIN GLASS VI: SPIN GLASS AS CORNUCOPIA

Philip W. Anderson

Some attentive readers will recall a remark I made in my fourth column (September 1988, page 9), to the effect that in the difficulties and annoying features encountered in the study of spin glasses, we were beginning to have an inkling of results that would turn out to be among the most important of modern theoretical physics. I shall now try to make that clear to you. I explained one of the key results last time (July, page 9): the discovery by Gérard Toulouse and his collaborators that there are many inequivalent solutions of the TAP theory of the SK long-range spin glass and that those solutions can be arranged in an "ultrametric tree" whose branches already begin dividing as T is lowered below T_c. To remind you what this jargon means: The TAP theory is the mean-field theory David Thouless, Richard Palmer and I constructed. That theory, we thought, would in principle be exact because fluctuations about it should be negligible in view of the many long-range interactions each spin has in the SK spin glass. "Ultrametric" is an ant's-eye view of a tree, in which the only way to get to another leaf is to climb all the way down to the common branch point and back up (see the illustration in my last column).

Scott Kirkpatrick made a second important connection. Scott observed that finding the lowest-energy state of the SK spin glass—in fact, of almost any spin glass—is a complex optimization problem equivalent to one of the classic examples of what computer theorists call the NP-complete problems. This mysterious class of problems includes a great many mathematical "toys," such as bisecting random graphs, setting up mixed-doubles tournaments and inventing tours of length N for traveling salesmen or Chinese postmen; but it also contains many highly practical problems, such as routing telephone networks to N cities, designing chips with N transistors, connecting N chips together, evolving the fittest animal with N genes and doing almost anything useful with N neurons. Large complex optimization problems are everywhere around us, and almost anything that can be learned about them is of immense importance.

An important branch of computer science is complexity theory, which classifies such large problems according to their "size" N. The size of a complex problem may be thought of as the number of bits necessary to state that problem. For instance, the size of the SK spin glass problem is $N(N-1)/2$, the number of J_{ij}'s. It is strongly conjectured that the number of steps it takes a computer to solve an NP-complete problem cannot be less than a number proportional to an exponential of a positive power of the size. For large N, then, it could take forever. This is clearly the reason why Scott, Richard and others had been unable to find a unique lowest-energy state.

Each instance of the dozens of known NP-complete problems can be converted to an instance of any of the other problems by an algorithm taking only N^p time steps—that is, the number of time steps is a polynomial function of the size of the problem. This suggests that a statistical mechanical "solution" of the spin glass problem might be of general interest for all NP-complete problems. But that is not the case, even if one assumes that the "polynomial" algorithm that maps other problems to the spin glass is not more trouble than it is worth. Our statistical mechanical solution gives *average* answers for an ensemble of examples of the given problem. Such an answer is valid for a generic, or typical, instance of the problem. In the case of the spin glass, the average number describes the generic instance of the problem involving the given distribution of J_{ij}'s. But the mapping algorithm might transform that generic instance into a special case or vice versa. This issue was perhaps somewhat clarified in an exchange between Eric Baum (Princeton), on the one hand, and Daniel Stein (University of Arizona), G. Baskaran (MATSCIENCE, Madras, India) and myself, on the other, about NP-complete problems with "golf course" energy landscapes—landscapes that are flat everywhere except one point! Furthermore, proofs of NP completeness in computer science often refer only to the worst possible case, and some NP-complete problems do not look very hard in generic terms. Finally, the computer scientist discusses—for obvious reasons—the problem of finding the *exact* answer for a particular case, not the average answer correct to order N for the generic case.

Nonetheless, specifying exactly the structure of the landscape of energy values as a function in the 2^N-dimensional space of spins tells us a very great deal about such problems. For instance, the existence of a transition temperature T_c tells us that below some value of energy per site E_c the space bifurcates into regions corresponding to different "solutions," and that as we go lower and lower in energy (or temperature) the space breaks up more and more. This gives us a clear reason why such a problem is "exponentially" hard: If we are in the wrong region, we have to cross an energy and entropy barrier of order N to get a better solution. This kind of

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

"freezing" phenomenon had been conjectured by computer scientists but never rigorously proved. To counter it, they had evolved a number of heuristic techniques for getting approximate solutions. We now know why this was necessary—namely, to get over the high barriers and sample the entire space of solutions.

Almost the first effect of the kind of thinking developed to understand spin glasses was to provide a *new* heuristic algorithm for the solution of complex optimization problems. That algorithm is called simulated annealing, and it was introduced by Scott and his colleague C. Daniel Gelatt Jr. Kirkpatrick and Gelatt proposed that one imitate the procedure the spin glassers had already been using, of "warming up" the problem above T_c and slowly cooling it back down, or "annealing" it. This could be done by regarding the "cost" for a given problem—say, the cost of connections on a chip—as a "Hamiltonian" function C of the positions to be varied. One plugs this Hamiltonian into a statistical mechanics simulator program, such as the well-known Metropolis algorithm. Then one chooses an appropriately scaled "temperature" T and minimizes $<e^{-C/T}>_{ave}$ for increasingly low temperatures. Simulated annealing, it turns out, is the most effective algorithm only for certain problems, but where it works it is very good indeed, and it is already in regular, profitable commercial use. The question of *why* simulated annealing works as well as it does was approached theoretically by Miguel Virasoro, who showed that, at least for the SK model, the lower the energy of a solution is, the larger is the entropy associated with it near T_c. That is, deeper valleys have bigger basins of attraction near T_c, and so one is more likely to start out in such a valley at T_c.

To me the key result here is the beautiful revelation of the structure of the randomly "rugged landscape" that underlies many complex optimization problems. Physics, however, has its own "nattering nabobs of negativism" (in the immortal phrase of William Safire), and they recently have been decrying the importance of the ultrametric structure, saying that it is a property of the SK model, not of physical spin glasses. Such criticism misses the point: Physical spin glasses and the SK model are only a jumping-off point for an amazing cornucopia of wide-ranging applications of the same kind of thinking. I will write about this in the next—and I hope the last—of these columns. ■

SPIN GLASS VII: SPIN GLASS AS PARADIGM

Philip W. Anderson

In my last column on spin glasses (September 1989, page 9) I tried to show you that the exact solutions of a particular spin glass problem, by Giorgio Parisi and Gérard Toulouse, gave us great insight into the theory of complex optimization problems, as well as an algorithm for solving some of them. One such problem, which has been exhaustively studied by methods of spin glass theory, is the graph partition problem. This is the question of how to divide an arbitrary graph into two pieces, cutting the fewest possible bonds. My student Yao-Tian Fu (now at Washington University, St. Louis) initiated the study of the graph partition problem by replica theory. This classic problem of complexity theory was difficult to solve for sparse graphs by those methods, but another of my students, Wuwell Liao, seems to have done it.

Even more interesting than these applications to complexity theory is the way apparently unrelated areas of science have been stimulated into parallel growth by the spin glass work. John Hopfield (Caltech), who was instrumental in bringing me to Princeton in 1975, became interested in models for neural networks and brain function about 1979–80. It was natural for him to realize that complex, interconnected systems of simple units could have the "rugged landscape," multistable properties of spin glasses. Using, very ingeniously, an *ad hoc* and apparently unrealistic assumption of symmetric coupling between neurons, John got the following results:

▷ For a given set of coupling synapses (interactions) J_{ij} between neurons (spins) i and j, the conventional McCulloch–Pitts model of neuronal interactions maps onto a "greedy" algorithm for finding the *local* ground state of a corresponding spin glass. ("Greedy" is the computer scientists' self-evident jargon term for jumping directly to the lowest local energy for each spin.)

▷ Modifying the couplings, or choosing the J_{ij}, in such a way as to form the so-called Hebb synapses makes the neural network into a model for "content addressable" memory: a memory like our own, which can reproduce full detail from fragmentary information. A system of N neurons connected by $N(N-1)/2$ symmetric synapses can remember about $N/6$ N-bit messages in the form of locally stable "spin" (that is, neuron firing rate) configurations. (John made this conjecture about the capacity of a neural network on the basis of simulations, but it later turned out to agree with exact analytic theory.) Thus, in exchange for a capacity reduction of a factor of 3 relative to the information-theoretic maximum, one gets the content-addressable feature.

▷ Finally, several other brain functions, such as pattern recognition, could be modeled with the spin glass type of neural network.

Many of you may be aware of the gigantic growth of neural network science in recent years. In 1979, however, when I tried to whip up interest in John's ideas among computer scientists at Bell Labs, there was little response; and he, Alan Gelperin and John Connor got nearly equally short shrift among neuroscientists. Nowadays the neuroscientists and computer scientists like to point to prior claims for each component of John's achievements. I can hardly believe, however, that such further developments in neural networks as the revival of the perceptron would have occurred except as a response to John's beautiful demonstrations that, after all, *one* such system—the Hopfield neural network—does work and has a rigorous, mathematically respectable basis. In particular, John's work has generated a very healthy trend toward rigorous mathematical demonstrations of limits on capacity, accuracy and so on in neural networks and perceptron-like models, using the statistical mechanics methods provided by Toulouse, Haim Somopolinsky, Miguel Virasoro, John Hertz, Richard Palmer (who was John's associate at Princeton in 1975–78) and many others. It turns out that statistical mechanics can be applied to realistic, asymmetric networks as well, and that there is no real difference between the capabilities of symmetric and asymmetric networks.

I promised I would close my series on spin glasses with this column, so it must be descriptive, not detailed. But I must also mention how the "rugged landscape" of spin glasses relates to theories of biological evolution. In 1981 I visited John Hopfield at Caltech and helped with the course on "physics of information" that he, Richard Feynman and Carver Mead were giving. Stimulated by John's work on neural nets, I came back to Princeton with the realization that I could put my own rugged-landscape ideas into a theory of prebiotic evolution that Daniel Stein (now at the University of Arizona) and I were already working on. The genome of an organism can be thought of as a set of Ising spins—two for each base in the DNA because there are four types of bases and the Ising spin has two possible values. The fitness, or reproductive capacity, of the genome can be modeled by a frustrated, quenched random function of this list of spins, and the simplest random function that satisfies the requisite plus–minus symmetry is a spin glass Hamiltonian function. (The plus–minus symmetry is imposed by the complementarity of base pairing.) With a senior thesis student, Dan Rokhsar, Stein and I made a simple model of primitive evolutionary processes using this idea.

Related ideas, but without the statistical mechanics insights, had already occurred to Gérard Weisbuch at the Ecole Normale in Paris and to Stu Kauffman at the University of Pennsylvania. Nonetheless the extra understanding those insights provide has encouraged us, and especially Stu, to go on and attack all kinds of evolutionary—and other—problems with spin-glass-like random, rugged-landscape models. This approach has become an important part of the program at the Santa Fe Institute, of

Philip Anderson is a condensed matter theorist whose work has also had impact on field theory, astrophysics, computer science and biology. He is Joseph Henry Professor of Physics at Princeton University.

which Stu and I are members. Unfortunately I cannot discuss here the many other ramifications of this way of thinking. From my point of view, its attractiveness lies in that it allows us to explain simultaneously the contradictory aspects of variety and stability of certain special forms and patterns. Life exhibits these contradictory aspects: Among the countless possible mutations, many lead to stable species. In the language of spin glass theory, there are many "basins of attraction," but each is stable (or at least metastable). It is also clear, as Gérard Weisbuch first pointed out, that evolution in such a landscape will exhibit "punctuated equilibrium," or sudden changes from one deep maximum of fitness to another, a feature that has been emphasized recently as characteristic of much of evolution.

I realize that I never returned to "real" spin glasses, even though it was studies of the low-temperature properties of those dilute magnetic alloys that led to the theoretical ideas I have been discussing. There is a reason: In spite of much beautiful experimental, computational and theoretical work, a complete and consistent understanding of those materials is not yet at hand. Helène Bouchiat and Pierre Monod, among others such as Laurent Levy, have beautifully demonstrated that real spin glasses have divergent nonlinear magnetic susceptibility at T_{SG}, verifying that there is a real phase transition, albeit one without a visible specific heat singularity. Spectacular simulations carried out on special purpose machines by Peter Young (University of California, Santa Cruz) and Andrew Ogielski (AT&T Bell Labs) have also verified the existence of a phase transition in three dimensions. Daniel Fisher (Princeton) and David Huse (Bell Labs), among others, speculate, however, that real spin glasses are really *not* ultrametric or replica-symmetry breaking. The theory is still under development. Some of it was explored in the December 1988 PHYSICS TODAY special issue on disordered solids.

For further information on random landscapes and evolution, Stu Kauffman's forthcoming book *Origins of Order: Self-Organization and Selection in Evolution* (Oxford U. P., New York) is perhaps the best source. John Hopfield has written (with David W. Tank) a *Scientific American* article on his neural network ideas (December 1987, page 104). An article on neural networks by Haim Sompolinsky appeared in the December 1988 PHYSICS TODAY (page 70). ■