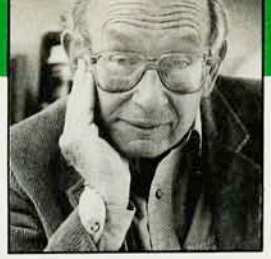

SPIN GLASS VI: SPIN GLASS AS CORNUCOPIA

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.

It's about TIME . . .

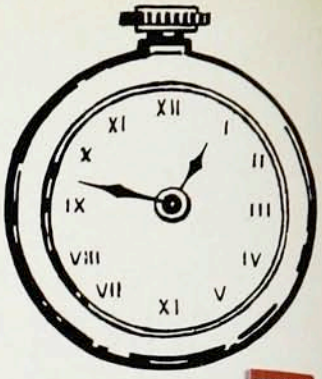

CAMAC Control of Constant-Fraction Timing

The **ALL NEW** CAMAC CCF8200 from EG&G ORTEC overcomes the timing limitations of leading-edge discriminators and offers UNIQUE benefits . . .

- Eight CFD channels with computer control in a CAMAC module
- An efficient solution for multi-detector experiments with Ge, Si, or scintillation detectors
- Individually programmable threshold settings
- Programmable output widths with programmable output masks
- Automatic walk adjustment
- Internally-selected shaping delay that eliminates bulky external cabling

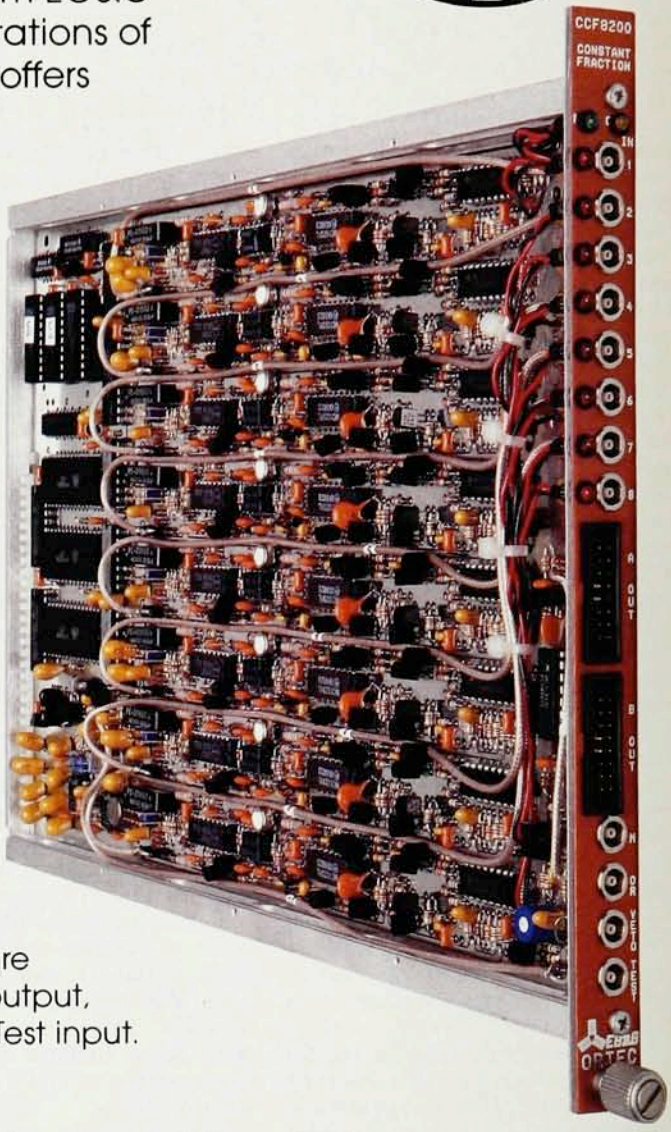

For further convenience in multiple-detector experiments, there are two sets of ECL outputs, a Multiplicity output, an OR output, an Inhibit input, and a Test input.

It's what you've been waiting for!

Save time by calling the **HOTLINE, 800-251-9750,** in the U.S.A. or contact your area EG&G ORTEC representative.

EG&G ORTEC

100 Midland Road, Oak Ridge, TN 37831-0895 U.S.A. • (615) 482-4411 • Telex 6843140 EGGOKRE • Fax (615) 483-0396

CANADA	FRANCE	WEST GERMANY	ITALY	JAPAN	THE NETHERLANDS	UNITED KINGDOM
(416) 475-8420	(01) 60.77.93.66	(089) 926920	(02) 7610267	(03) 638-1506	(030) 887520	(0734) 773003

Circle number 8 on Reader Service Card

"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. ■