

# The World in a Spin

Brian Hayes

A reprint from

## American Scientist

the magazine of Sigma Xi, the Scientific Research Society

Volume 88, Number 5  
September–October, 2000  
pages 384–388

This reprint is provided for personal and noncommercial use. For any other use, please send a request to Permissions, *American Scientist*, P.O. Box 13975, Research Triangle Park, NC, 27709, U.S.A., or by electronic mail to [perms@amsci.org](mailto:perms@amsci.org). Entire contents © 2000 Brian Hayes.

# The World in a Spin

Brian Hayes

**P**ierre Simon de Laplace had a plan for understanding everything. To this celestial mechanic, it looked simple: The particles of matter produce forces, and those forces in turn move the particles. So if we could just measure all the forces and motions at any one instant, we could calculate the entire history of the universe—past, present and future.

Two centuries of progress in the sciences have not fulfilled Laplace's vision; on the contrary, quantum mechanics, and lately chaos theory, have undermined faith in his program. But let's pretend. If we study a computational model of the universe rather than the real thing, we really can track all the forces and motions. The laws of physics can be kept as simple as we please, since we invent and enforce them. In this toy universe, we can banish all quantum uncertainties, and trace every last detail of every microscopic event. Yet even in such an open and transparent world, total knowledge is still elusive. Although we can follow the individual particles, we have trouble seeing how they act in the aggregate. For example, we may well fail to predict basic thermodynamic phenomena such as boiling and freezing. We could know the whereabouts of every molecule of water in an artificial ocean, but not know whether the stuff is solid or liquid or vapor.

The prototypical system for exploring issues of this kind is called the Ising model. It is a model of matter pared down to its barest essentials—just about the simplest imaginable system in which large numbers of particles might be expected to produce some kind of cooperative behavior. If Laplace's plan can be made to work anywhere, it should succeed here. But the Ising model has proved a difficult challenge, even when attacked with some heavy-duty mathematics and computer science. Indeed, the most important version of the model remains without an exact solution.

## A Model Magnet

The Ising model was invented in 1920 by Wilhelm Lenz, who proposed it as a simplified version of a ferromagnet—the kind of magnet that holds your grocery list to the refrigerator door. A

few years later Lenz's student Ernst Ising chose the model as the subject of his doctoral dissertation at the University of Hamburg.

The elements of the model are called spins, although the concept of rotation never enters the picture. You can think of a spin as nothing more than an arrow pointing either up or down (but in no other direction). The spins are arranged in a grid or lattice pattern. Spins at neighboring sites prefer to point the same way; in other words, the energy is lower when adjacent spins are parallel and higher when they are antiparallel. Except for these nearest-neighbor preferences, the spins don't interact at all. Thermal fluctuations tend to randomize the spins. Finally, an external magnetic field may impose a bias on the spin directions.

The Ising model is a crude cartoon of a ferromagnet, but it does capture the main features of the real thing. The Ising spins correspond to spinning electrons in iron atoms; the lattice represents the crystal structure; the nearest-neighbor interaction mimics the overlap of wave functions in adjacent iron atoms. The one element of the model that has no obvious counterpart in real physics is the requirement that spins take on only two possible orientations.

Is the model magnet magnetic? Do the spins line up in parallel the way they do in a real ferromagnet? It's easy to guess the answer at the extremes of the temperature range. At infinite temperature, thermal fluctuations overwhelm the nearest-neighbor interactions; each spin continually makes random flips, so that the average magnetization is zero. At the other end of the thermometer, thermal fluctuations disappear altogether at absolute zero, and the system falls into a state of minimum energy, with all the spins either *up* or *down*. (In the absence of an external field, the two choices are equally likely.)

Suppose we steadily reduce the temperature of an Ising model from infinity down to zero. Totally random spins must somehow become totally ordered, but is the change smooth and gradual, or does magnetization set in abruptly at some specific temperature? This question is an important one in statistical physics, where discontinuities—phase transitions and critical points—are major landmarks. For real ferromagnets, experiments give the answer. As iron cools from the melt, the magnetization remains zero until it suddenly

---

*Brian Hayes is a former editor of American Scientist. Address: 211 Dacian Avenue, Durham, NC 27701. Internet: bhayes@amsci.org*

leaps up at a temperature called the Curie point (about 1,040 kelvins). If the Ising model has anything to say about ferromagnetism, it should have a comparable discontinuity.

Ising's dissertation examined this question for a one-dimensional version of the model—a line of spins. His result was a disappointment. There was no phase transition at any temperature above zero. It's not hard to understand why. Consider a line of 10,000 spins, all pointing *up*. This configuration is clearly a minimum-energy state, since all nearest-neighbor pairs are parallel. Now reverse each of the 5,000 spins lying to the right of the midpoint. The overall magnetization falls to zero (since there are now equal numbers of *up* and *down* spins), and yet the change in the energy is minuscule: Only one pair of neighbors is pointing in opposite directions, while 9,998 pairs remain parallel. In a one-dimensional system, it seems, the coupling between spins is just too tenuous to overcome even the slightest thermal agitation.

Ising reportedly believed that his negative result would hold in higher dimensions as well. In this conjecture he was wrong. But before going on to recount the further history of the Ising model, I want to mention the further history of Ising. After receiving his doctorate, he taught physics in German public high schools, but as a Jew he was dismissed when Hitler came to power in 1933. He then taught at a Jewish boarding school in Potsdam, until that was destroyed in the Kristallnacht pogrom of 1938. Ising and his wife fled Germany, but they got only as far as Luxembourg before the war overtook them. There they managed to survive the occupation, and finally reached the United States in 1947. Ising taught physics and mathematics in Minot, North Dakota, and then for almost 30 years more at Bradley University in Peoria, Illinois. He died two years ago at age 98.

### The Spin Cycle

The Ising model was rescued from obscurity in the 1930s, when Rudolf Peierls perceived that a two-dimensional array of spins might admit more interesting behavior than the one-dimensional system. Again a simple qualitative argument suggests the reason. Take 10,000 spins, all *up*, and arrange them in a square array. When you flip half of the spins to abolish the magnetization, at least 100 pairs of neighbors must be antiparallel. The energy penalty associated with these opposed spins is larger than it was in the one-dimensional case—perhaps large enough to maintain magnetization in the presence of thermal disruption.

This informal argument is encouraging, and Peierls gave a stronger version, but if you want to understand the Ising model mathematically, you need a way to calculate the probability of any configuration of the spins at any temperature. To see what this entails, it's helpful to work through a miniature example—a two-dimensional Ising array with just four spins arranged in a square. Since each spin has two possible values, the sys-

tem can take on any of 16 configurations, or states. In each state the magnetization is the number of *up* spins minus the number of *down* spins. Likewise the energy is the number of antiparallel neighbors minus the number of parallel neighbors. Calculating these properties calls for nothing but the simplest arithmetic. But what we want to know is the *probability* of each state at a given temperature, which is harder to determine.

A first step toward calculating probabilities is an exponential function called the Boltzmann weight, defined as  $e^{-H/T}$ , where  $H$  is the energy and  $T$  is the temperature. (If you want to measure  $H$  and  $T$  in joules and kelvins, a constant is need-



Figure 1. Two-dimensional Ising model is a grid of *up* and *down* spins, represented here by contrasting colors. In this example, temperature varies smoothly from top to bottom, producing nearly random spins in warm areas and large magnetized domains in cooler regions.

ed to make the units come out right, but in an abstract model these formalities can be ignored.) The formula for the Boltzmann weight reveals in a general way how the probability varies with temperature and energy. If the temperature is extremely high, then  $e^{-H/T}$  has a value close to 1 no matter what the energy, and all configurations are about equally likely. At low temperature, on the other hand, small differences in energy produce extreme changes in  $e^{-H/T}$ , so that only the states of lowest energy are likely to be seen.

The Boltzmann weight is proportional to a state's probability, but to get at the probability itself we need to know something more. The weights have to be *normalized*, so that the probabilities of all possible states add up to 1. In other words, we have to divide the Boltzmann weight for a single state by the sum of the weights for all possible states. This sum is called the partition function, and it plays a crucial role in the Ising model and in other thermodynamic systems.

For the four-spin array, it's no great challenge to compute the partition function. Simply calculate  $e^{-H/T}$  for each of the 16 states, and add the results. For example, at a temperature of 2, the sum of the Boltzmann weights is 27.05. At the same temperature the weight of the specific state that has all four spins *up* is about 7.39; thus the probability of this state is 0.27. If you observe a four-spin Ising system at  $T=2$  long enough, you should see it with all spins *up* a little more than a quarter of the time.

Notice that in order to find the probability of a single state, we have to compute the Boltzmann weights of *all* the states. For a system of four spins, this calculation is easy, but for 40 spins the trillion possible states would challenge the fastest computers, and for 400 spins a brute-force enumeration is unthinkable. Yet it's the larger systems that

matter most. Indeed, what we really want to know is the partition function in the thermodynamic limit—as the number of spins tends to infinity.

Calculating the partition function of the two-dimensional model is hard but not impossible. The problem was solved—to much surprise and acclaim—in 1944 by Lars Onsager, a chemist at Yale University. Of course Onsager's method was not direct enumeration; he got his result through a virtuoso performance in operator algebra, which I'll not attempt to explain since my own grasp of it is tenuous at best. (There is a thorough exposition in Martin H. Krieger's book *Constitutions of Matter*, which also reprints two of Onsager's papers.) The final product was an exact expression for the partition function in the thermodynamic limit.


From the partition function flow all the macroscopic, observable properties of the model. In particular, Onsager's theory describes the onset of magnetization at the critical temperature  $T_C$ . The magnetization is equal to  $(T_C - T)^{1/8}$ ; note that this quantity has two values at any temperature below  $T_C$  but becomes undefined (within the field of real numbers) above  $T_C$ . In other words, the magnetization vanishes at the critical temperature.

### What a Difference a D Makes

Once Onsager had shown the way, the two-dimensional Ising model was solved again by several more methods. Solutions were also discovered for certain other two-dimensional models, which share a basic conceptual framework with the Ising model but differ in details of the lattice or the spin-spin interaction. But, significantly, all of these solutions are confined to flatland. In the three-dimensional world, not one Ising-like model has been solved exactly. This is a bit disappointing for three-dimensional physicists who would like to understand three-dimensional matter.

What is the essential difference between a 2D and a 3D model? Obviously, one is flat and the other isn't, but there's more to it than that. Geometry *per se* never enters the Ising model; distances and angles are simply not part of the problem description. All that matters is the topology—the pattern of connections between nearest-neighbor sites. If you were to build a three-dimensional cubic framework out of infinitely stretchy rubber, you could flatten this lattice by sliding layers of sites apart and then smashing them down into a plane. Now the 3D lattice would be geometrically two-dimensional, and yet it would still differ topologically from a true 2D lattice. The difference is that many of the bonds would be nonplanar—they would cross over one another—whereas a 2D lattice can always be drawn without crossings. This topological distinction seems to be at the root of the difficulty of 3D Ising models.

Tullio Regge of the Politecnico di Torino and Riccardo Zecchina of the Abdus Salam International Centre for Theoretical Physics in Trieste have recently looked at what happens to the partition function when you start with a two-dimen-



**Figure 2.** Four-spin Ising model has just 16 states. Magnetization is defined as the number of red spins minus the number of green, and energy is the number of unlike neighbors minus the number of like neighbors.

sional lattice and add nonplanar bonds one by one. In this way they explore the territory between 2D and 3D. Each added nonplanar bond doubles the labor needed to solve the model—an exponential increase. Whether or not this approach leads to a practical method for calculating the 3D partition function, it shows more clearly why that task is so hard.

Curiously, adding still more dimensions beyond three makes the Ising model simpler again rather than harder. A four-dimensional lattice is so densely connected that each spin responds to the averaged magnetic field of all the other spins, and analysis is easy. Thus it seems there's something special about three dimensions; perhaps the world we live in was created explicitly as a vexation to Ising modelers.

I do not want to leave the impression that nothing at all is known about the 3D Ising model. An exact mathematical solution is still lacking, but the region near the critical point has been explored by various methods of approximation. One approach, called series expansion, would have been familiar to Laplace. The idea is to start with the known solutions at very high and very low temperature and extrapolate into the more problematic region between. Another approximation method has a name that makes it sound like an organization of diplomats or economists: the renormalization group. The simplest version of this algorithm gathers sets of spins into blocks, replaces each block with a single new spin, and finally adjusts the couplings between spins to compensate for the coarsening of the lattice.

Still another important technique for Ising studies is the Monte Carlo method, which relies on a random process to approximate the probability distribution of the spin states. I shall say a little more about Monte Carlo methods below.

### The Blinking Checkerboard Catastrophe

Subjecting the Ising model to all this industrial-strength mathematics—from Onsager's algebra down through the renormalization group—has produced a wealth of useful answers, and yet it seems rather remote from Laplace's simple notion of understanding nature by following the microscopic events as they unfold. Isn't there a more direct way to learn what happens in the lattice? Can't we just program the system into a computer and let it evolve under its own internal rules?

As it happens, the first computer program I ever wrote was a naive attempt to do just that. It was a two-dimensional Ising model implemented in an early version of the spreadsheet Lotus 1-2-3. Each cell of the spreadsheet held an identical evaluation rule, which examined the four surrounding cells and calculated its own next state accordingly. As I gradually cooled the model, I expected to see a magnetized phase spread across the array. What I saw instead was perplexing. Although some areas were magnetized, other regions were taken over by a blinking checkerboard

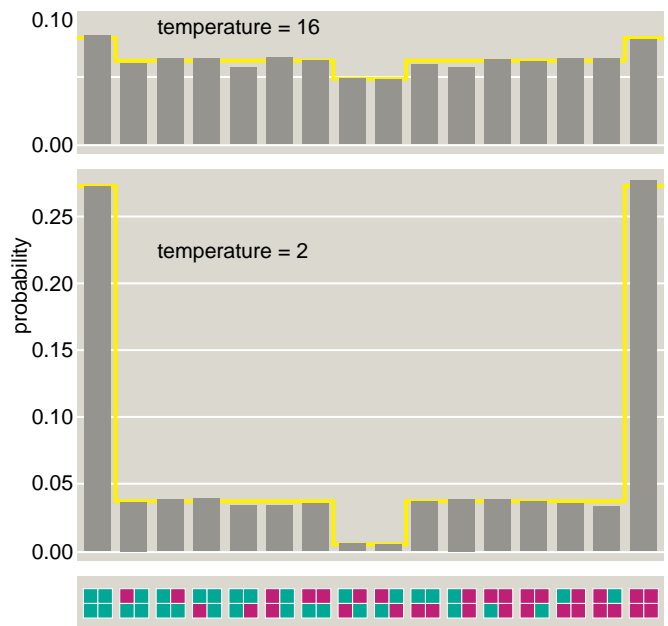


Figure 3. Ising probabilities vary with temperature, favoring parallel spins at  $T=2$  but having a flatter distribution at  $T=16$ . Yellow lines give calculated probabilities; bars are results of a single simulation.

pattern, with alternating up and down spins that flipped at every time step. This pattern should be the least favorable of all configurations. I suspected some elementary bug, such as a missing minus sign, but I wasn't able to track it down.

The problem was quickly diagnosed by more experienced friends: Tommaso Toffoli, Norman Margolus and Gérard Vichniac, who were then all at MIT. Without even looking at my code, they recognized the symptoms. The bug was in fact a fairly interesting one. Imagine you are an *up* spin somewhere in the middle of a blinking-checkerboard area. Your four neighbors are all *down*, and so you are strongly inclined to flip to the *down* state and join them. But each of those neighbors is surrounded by four *up* spins, and so they too flip. Thus all the spins reverse on each step, and the pattern is perpetuated.

Whether or not this bug appears depends on fine details of implementation, in particular on the sequence in which the spins examine their neighbors' states and make decisions about their own next state. My mentors suggested a remedy for my particular situation: The problem would go away if I updated the spreadsheet array in a checkerboard sequence, first all the black squares and then the white. It was an ingenious trick, and it worked like a charm—but what would Laplace think about such stratagems? Does a real ferromagnet require careful sequencing of operations to avoid falling victim to a subtle bug? I find it hard to believe that alternate iron atoms politely take turns updating their spin states. And if nature can get along without such artifice, why should it be needed in a computer simulation of nature?

One answer to this question is that nature has a better computer than we do. In particular, the

checkerboard bug might be seen as an artifact of mapping a parallel process onto a sequential machine. If we ran the program on a computer with one processor per spin, the question of how to sequence the updates wouldn't arise at all; the spins would all be updated at once. But I'm not convinced that parallelism solves the problem; rather, it just converts a problem of sequencing into a problem of synchronization. Figuring out what's supposed to happen when all the spins interrogate their neighbors at exactly the same instant they are themselves being interrogated is an even deeper conundrum than the sequential case.

The Monte Carlo method evades problems of sequencing and synchronization by introducing randomness. Each iteration of the algorithm selects a single spin and looks at the effect of reversing it. If flipping the spin would lower the system's overall energy, the change is made. If the energy would be increased, the spin may or may not be flipped; the choice is made randomly, with a probability determined by the change in the Boltzmann weight. Repeating this process many times should produce a statistical sample of spin configurations whose frequencies are proportional to their Boltzmann weights.

Most practitioners look upon the Monte Carlo method as a tool for estimating the partition function, and by this criterion all that matters is that it produces correct answers. But the algorithm seems so physical and mechanistic that it's tempting to view the Monte Carlo process as a simulation of what might actually go on inside an Ising system. Does this view stand up to scrutiny? Laplace might raise an eyebrow at the presence of random numbers in the algorithm, but modern sensibilities seldom take offense at a little randomness. What I find more unsettling is the explicit use of Boltzmann weights to calculate probabilities. I don't want to have to imagine that individual atoms know how to evaluate  $e^{-H/T}$ . Ideally, the Boltzmann probability distribution would not be built into the model but would emerge from some simpler rule for strictly local interactions of spins.

Among all the computer implementations of the Ising model, my personal favorite is one devised in 1985 by Michael Creutz of Brookhaven National Laboratory. It is strictly deterministic—no random numbers are needed, except perhaps to set the initial conditions—and individual spins know only their energy, not their Boltzmann weight. The model works by giving each spin a kinetic energy as well as the energy associated with nearest-neighbor interactions. The kinetic term acts as a reserve, absorbing excess energy and giving it back when needed. On each update step, each spin is flipped if and only if the corresponding kinetic reserve can accommodate the change in the interaction energy. But even this model is marred by the blemish of the blinking-checkerboard bug, which Creutz avoids by imposing a checkerboard update sequence. Again:

Why should such strange and unphysical contortions be necessary?

In a sense, asking a computer program to show us what “actually” happens inside an Ising model is rather silly, because no one will ever build a physical Ising system. The model is an idealization. Everything about it is perfectly discrete and symmetrical. In such a world of mathematical exactitude, the blinking checkerboard may not be a bug at all; it may be a genuine state of the system, which we don't see in real ferromagnets only because there are no perfect ones. Vichniac showed that when Ising spins are allowed to wiggle even slightly, the checkerboard bug is eradicated. Breaking the spatial or the temporal symmetry would surely have the same effect. In other words, maybe what the model needs most is a little imperfection, but I don't think Laplace would be pleased with this thought.

### Acknowledgment

*This article was inspired and informed by a visit to the Scuola Internazionale Superiore di Studi Avanzati and the Abdus Salam International Centre for Theoretical Physics, both in Trieste. I am grateful for the advice and criticism of Riccardo Zecchina, Federico Ricci-Tersenghi, Matteo Marsili, Anna Delin and Johannes Berg, and for the hospitality of both institutions.*

### Bibliography

- Baxter, Rodney J. 1982. *Exactly Solved Models in Statistical Mechanics*. London, New York: Academic Press.
- Brush, Stephen G. 1967. History of the Lenz-Ising model. *Reviews of Modern Physics* 39:883–893.
- Cipra, Barry A. 1987. An introduction to the Ising model. *American Mathematical Monthly* 94:937–959.
- Creutz, Michael. 1986. Deterministic Ising dynamics. *Annals of Physics* 167:62–76.
- Ising, Ernst. 1925. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik* 31:253–258.
- Kobe, S. 1995. Ernst Ising—physicist and teacher. Actas: Noveno Taller Sur de Física del Sólido, Misión Boroa, Chile, 26–29 April 1995. <http://arXiv.org/cond-mat/9605174>
- Krieger, Martin H. 1996. *Constitutions of Matter: Mathematical Modeling the Most Everyday of Physical Phenomena*. Chicago: University of Chicago Press.
- Laplace, Pierre Simon, Marquis de. 1820. *A Philosophical Essay on Probabilities*. New York: Dover Publications, 1952.
- Lenz, W. 1920. Beitrag zu Verständnis der magnetischen Erscheinungen in festen Körpern. *Physikalische Zeitschrift* 21:613–615.
- McCoy, Barry M., and Tai Tsun Wu. 1973. *The Two-dimensional Ising Model*. Cambridge, Mass.: Harvard University Press.
- Onsager, Lars. 1944. Crystal Statistics. I. A two-dimensional model with an order-disorder transition. *Physical Review* 65:117–149.
- Peierls, R. 1936. On Ising's model of ferromagnetism. *Proceedings of the Cambridge Philosophical Society, Mathematical and Physical Sciences* 32:477–481.
- Regge, Tullio, and Riccardo Zecchina. 2000. Combinatorial and topological approach to the 3D Ising model. *Journal of Physics A: Mathematical and General Physics* 33:741–761.
- Vichniac, Gérard Y. 1984. Simulating physics with cellular automata. *Physica D* 10:96.
- Wilson, Kenneth G. 1979. Problems in physics with many scales of length. *Scientific American* 241(2):158–179.