CONTENTS

INTRODUCTION: WHY SPIN GLASSES?

Spin glasses are disordered magnetic materials, and it's hard to find a less promising candidate to serve as a focal point of complexity studies, much less as the object of thousands of investigations. On first inspection, they don't seem particularly exciting. Although they're a type of magnet, they're not very good at being magnetic. Metallic spin glasses are unremarkable conductors, and insulating spin glasses are fairly useless as practical insulators. So why the interest?

Well, the answer to that depends on where you're coming from. In what follows we'll explore those features of spin glasses that have attracted, in turn, condensed matter and statistical physicists, complexity scientists, and mathematicians and applied mathematicians of various sorts. In this introduction, we'll briefly touch on some of these features in order to (we hope) spark your interest. But to dig deeper and get a real sense of what's going on—that can fill a book.

Spin glass research provides mathematical tools to analyze some interesting (and hard) real-world problems.

Suppose you're given the following easily stated problem. You're shown a collection of N points on the plane, which we'll call cities. You're asked to start at one of the cities (any one will

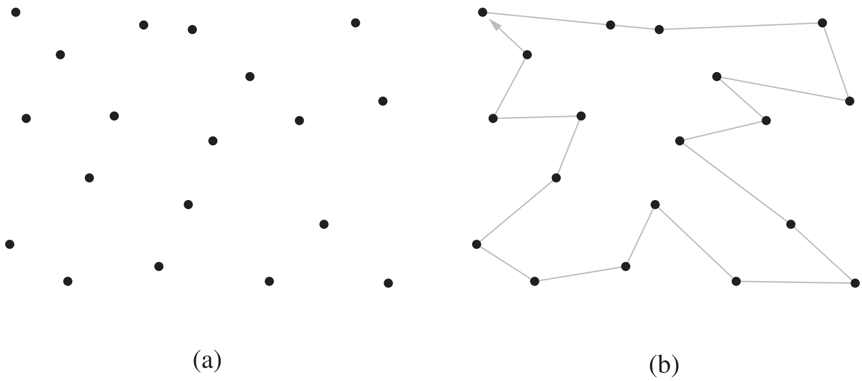

Figure I.1. (a) An instance of a TSP problem with 19 cities. (b) One possible tour.

do), draw an unbroken line that crosses each of the other cities exactly once, and returns to the starting point. Such a line is called a tour, and an example is shown in figure I.1. All you need to do is to find the shortest possible tour.

This is an example of the Traveling Salesman Problem, or TSP for short.[1] An "instance," or realization of the problem, is some specific placement of points on the plane (which a priori can be put anywhere). You should be able to convince yourself that the number of distinct tours when there are N cities is $(N-1)!/2$. The factor of two in the denominator arises because a single tour can run in either direction.

Notice how quickly the number of tours increases with N: for 5 cities, there are 12 distinct tours; for 10 cities, 181,440 tours; and for 50 cities (not unusual for a sales or book tour in real life), the number of tours is approximately 3×10^{62}. The seemingly easy (i.e., lazy) way to solve this is to look at every possible tour and compute its length, a method called exhaustive search. Of course, you're not about to do that yourself, but you have access to a modern high-speed computer. If your computer can check out—let's be generous—a billion tours every second, it would take it 10^{46} years to come up with the answer for a 50-city tour. (For comparison, the current age of the universe is estimated at roughly 1.3×10^{10} years.) Switching to the fastest supercomputer won't help you much. Clearly, you'll need to find a much more efficient algorithm.

Does this problem seem to be of only academic interest? Perhaps it is,[2] but the same issues—lots of possible trial solutions to be tested and a multitude of conflicting constraints making it hard to find the best one—arise in many important real-world problems. These include airline scheduling, pattern recognition, circuit wiring, packing objects of various sizes and shapes into a physical space or (mathematically similarly) encoded messages into a communications channel, and a vast multitude of others (including problems in logic and number theory that really are mainly interesting only to academic mathematicians).

These are all examples of what are called combinatorial optimization problems, which typically, though not always, arise from a branch of mathematics called graph theory. We'll discuss these kinds of problems in chapter 6, but what should be clear for now is that they have the property that the number of possible solutions (e.g., the number of possible tours in the TSP) grows explosively as the number N of input variables (the number of cities in the TSP) increases. Finding the best solution as N gets large may or may not be possible within a reasonable time, and one often has to be satisfied with finding one of many "near-optimal," or very good if not the best, solutions. Whichever kind of solution one seeks, it's clear that some clever programming is required. For both algorithmic and theoretical reasons, these kinds of problems have become of enormous interest to computer scientists.

What have spin glasses to do with all this? As it turns out, quite a lot. Investigations into spin glasses have turned up a number of surprising features, one of which is that the problem of finding low-energy states of spin glasses is just another one of these kinds of problems. This led directly from studies of spin glasses to the creation of new algorithms for solving the TSP and other combinatorial optimization problems. Moreover, theoretical work trying to unravel the nature of spin glasses led to the development of analytical tools that turned out to apply nicely to these sorts of problems. So, even in the early days of spin glass research, it became clear that they could appeal to a far greater class of researchers than a narrow group of physicists and mathematicians.

Spin glasses represent a gap in our understanding of the solid state.

Why is a crystalline solid (in which constituent atoms or molecules sit in an ordered, regular array) rigid? It may be surprising to learn that it wasn't until the twentieth century that we understood the answer to this question at a deep level.

Why is window glass (which does not have crystalline structure; the atoms sit in what look to be random locations) rigid? That's an even harder question, and you may be even more surprised to learn that we still can't answer that question at a deep level.

Of course, at what level you're satisfied with an explanation depends on your point of view: an answer that satisfies a chemist may not satisfy a physicist (and vice versa), and mathematicians are hard to convince of anything (so they're seldom satisfied). To be fair, at some level we've understood the nature of the solid state since the nineteenth century, when modern thermodynamics and statistical mechanics were developed by Gibbs, Boltzmann, and others. The basic idea is this. Atoms and molecules at close range attract each other, but they're never isolated from the rest of the world; consequently, the constituent particles of a system always have a random kinetic energy that we measure as temperature. At higher temperatures, entropy (roughly speaking, disorder induced by random thermal motions) wins out, and we have a liquid or gas. At lower temperatures the attractive forces win out, and the system assumes a low-energy ordered state—a crystalline solid. Liquids and crystals are two different phases of matter, and the transition from one to the other, not surprisingly, is called a phase transition.

If you've taken introductory-level physics or chemistry courses you know all this. But there are deeper issues, which enter because there are features accompanying the ordered state that aren't so easy to explain. One of these is what Philip Anderson calls "generalized rigidity" [13]: when you push on the atoms in a crystal at one end, the force propagates in a more or less uniform manner throughout the crystal so that the entire solid moves as a single entity.

This is something we all take for granted. Why is it mysterious? Well, interatomic forces are short range and typically extend only about 10^{-8} cm, whereas when you push on a solid at one end, the force you apply is transmitted in a perfectly uniform manner a billion or more times the range of the interatomic force. How does that happen? What changed from the liquid state, where the exact same forces are present? At the very least, why doesn't the solid crumple, or bend? (And for that matter, what new phenomena need to be invoked to explain crumpling and bending when they *do* happen?)

This phenomenon isn't unique to solids; the transmission of forces over long distances also occurs, for example, in liquid crystals. In fact, this property is widespread in a general sense: it occurs whenever there's a transition to an ordered state that possesses a symmetry (whose form may not always be obvious) that differs from the thermally disordered state. Without

INDEX