Introduction: Where Statistical Physics Meets Computation

Allon G. Percus*, Gabriel Istrate† and Cristopher Moore‡

1 Background

Computer science and physics have been closely linked since the birth of modern computing. This book is about that link. John von Neumann's original design for digital computing in the 1940s was motivated by applications in ballistics and hydrodynamics, and his model still underlies today's hardware architectures. Within several years of the invention of the first digital computers, the Monte Carlo method was developed, putting these devices to work simulating natural processes using the principles of statistical physics. It is difficult to imagine how computing might have evolved without the physical insights that nurtured it. It is impossible to imagine how physics would have evolved without computation.

While digital computers quickly became indispensable, a true theoretical understanding of the efficiency of the computation process did not occur until twenty years later. In 1965, Hartmanis and Stearns [30] as well as Edmonds [20, 21] articulated the notion of computational complexity, categorizing algorithms according to how rapidly their time and space requirements grow with input size. The qualitative distinctions that computational complexity draws between algorithms form the foundation of theoretical computer science. Chief among these distinctions is that of polynomial versus exponential time.

A combinatorial problem belongs in the complexity class P (*polynomial time*) if there exists an algorithm guaranteeing a solution in a computation time, or number of elementary steps of the algorithm, that grows at most polynomially with input size. Loosely speaking, such problems are considered computationally *feasible*. An example might be sorting a list of n numbers: even a particularly naive and inefficient algorithm for this will run in a number of steps that grows as $O(n^2)$, and so sorting is in the class P. A problem belongs in the complexity class NP (*non-deterministic polynomial time*) if it is merely possible to test, in polynomial time, whether a specific presumed solution is correct. Of course, P $\subseteq$ NP: for any problem whose solution can be found in polynomial time, one can surely verify the validity of a presumed solution in polynomial time.

*percus@ipam.ucla.edu
†istrate@lanl.gov
‡moore@cs.unm.edu

However, finding a needle in a haystack involves a great deal more than just verifying that an object is a needle. For many problems in NP, even our best attempts at finding a solution have yielded algorithms that require exponential time. A famous example is the traveling salesman problem of finding a tour with n cities, such that the tour's total length is less than a fixed constant. One can readily confirm whether a proposed solution meets the desired criteria, but there is no known polynomial-time algorithm for locating a solution among the $n!$ possible orders in which we could visit the cities. Thus, some problems in NP appear to be computationally hard or *intractable*. To the extent possible, we would like to categorize which NP problems are also in P, and which are not.

In the early 1970s, Cook [14] and Karp [36] took a large step towards such a categorization with the notions of NP-hardness and NP-completeness. A problem $\mathcal{A}$ is NP-hard if any NP problem can be converted or *reduced* to it in polynomial time: if there were a polynomial-time algorithm for $\mathcal{A}$, all NP problems could be solved in polynomial time. A problem is NP-complete if it is both NP-hard and itself in NP. NP-complete problems are therefore the hardest among NP problems. If any of them could be solved by a polynomial-time algorithm, that would immediately imply P = NP. It is widely believed that this is not the case. Finding a solution in an exponentially large search space seems intrinsically harder than checking a proposed solution. But proving or disproving that P = NP remains an open question to this day, and is unquestionably the central unsolved problem in theoretical computer science [17].

Cook's main result was a proof that the problem of *satisfiability*, or deciding whether or not a propositional formula in Boolean logic can be satisfied, is NP-complete. Having established satisfiability as the founding NP-complete problem, one may then prove that other problems are NP-complete by showing that satisfiability can be reduced to them. This is indeed what Karp did, for a host of well-known combinatorial problems including graph coloring, vertex cover, number partitioning, and the traveling salesman problem. Since that time, many thousands of other problems have been proven NP-complete, ranging over an astonishing variety of problem domains, from industrial resource allocation to predicting how proteins fold. It appears highly unlikely that there are algorithms that can guarantee a solution to any of these problems in polynomial time.

The crucial term, however, is "guarantee." Computational complexity theory deals with the universal quantifier: can one solve the problem in polynomial time, for all possible instances of the problem? This is a worst-case notion. Computer scientists sometimes imagine a fictitious adversary who designs instances that are as hard as possible, deliberately trying to make our algorithms fail. But in many cases, we would be satisfied with a weaker guarantee. What if the problem is NP-complete, but for many types of instances it can still be solved in polynomial time? What if the hard instances are actually rather rare, and in practice the problem can almost always be solved in polynomial time? It did not case take researchers very long to discover that this is exactly what happens for certain forms of the satisfiability problem [25, 28].

Here is where physics comes back into the picture. Physicists are used to problems given to them by nature, not designed by a malicious adversary. To physicists, nature can be

astonishingly benevolent, often admitting beautiful and elegant solutions to its problems. While computer science has to work in some contexts where there really is an adversary—for instance, in cryptography—perhaps real-world instances of NP-complete problems are more like natural systems than maliciously designed ones.

There is another crucial cultural difference between computer science and physics. While computer scientists think of problem instances as given with complete specificity, statistical physicists try to study the *macroscopic* properties of a system and avoid explicit consideration of its microscopic details. The canonical nineteenth century example is Boltzmann's unification of thermodynamics with Newtonian dynamics, showing how bulk properties such as temperature and pressure emerge from the statistical behavior of atoms in a gas. This approach necessarily implies averaging over the local properties of individual atoms in order to obtain a broad statistical description, explaining how the system "typically" behaves at an appropriate level of resolution.

Kirkpatrick, Gelatt, and Vecchi [38] exploited the relation between computationally hard combinatorial problems and the principles of statistical physics in their 1983 paper introducing the simulated annealing method. The idea was as follows. Finding the solution to a combinatorial optimization problem, such as the shortest possible traveling salesman tour, is formally equivalent to finding the *ground state*, or lowest-energy state, of a physical system in thermal equilibrium. In both cases there is an objective function to be minimized, consisting of contributions from all the components of the system: the tour length is the sum of all link lengths in the tour, and the energy is the sum of the interaction energies among all atoms in the system. In statistical mechanics, given a certain physical temperature τ (measured in units of energy), the probability of the system finding itself in a specific state C with energy $E(C)$ is proportional to the Boltzmann factor:

$$\mathbf{Pr}[C] \propto e^{-E(C)/\tau} . \tag{1}$$

In the *zero-temperature* limit $\tau \to 0$, this probability is concentrated at the ground state. The third law of thermodynamics tells us that achieving $\tau = 0$ is a physical impossibility, but by *annealing* a system, or cooling it slowly enough, one may at least come close to reaching the ground state. The computational analogy for a combinatorial optimization problem is to pick as a starting point some valid solution—not necessarily the optimal one—along with some starting value of the temperature parameter τ, and iteratively update this solution by a Monte Carlo process that mimics thermal fluctuations. One then slowly reduces τ over the course of the simulation, in the hope of descending upon the solution that minimizes the objective function.

Simulated annealing offered a very practical recipe for finding solutions to NP-hard problems that are optimal or near-optimal. But the analogy it suggested opened the way to a far broader view of the connections between statistical physics and theoretical computer science. The key insight involves considering computational problems whose inputs are random combinatorial structures: for instance, satisfiability over suitably generated random formulas, or graph coloring over a particular random graph ensemble. As Kirkpatrick et al. noted in their paper [38], this view is entirely appropriate for problems with large input size n. The asymp-

totic limit $n \to \infty$ is analogous to the thermodynamic limit in statistical physics, where the law of large numbers should apply and relative fluctuations around the average case should go to zero. When this occurs, the system is said to be *self-averaging*. Of course, asymptotic analysis is no stranger to complexity theory either. When one speaks of an algorithm that runs in polynomial time, one refers to the way running time scales with n in precisely the same limit $n \to \infty$.

While statistical physicists initially studied models on regular lattices, such as the Ising model of magnetism, they subsequently considered *disordered* models in which the parameters or topology of the interactions vary randomly from site to site. One model of particular interest has been that of *spin glasses* [44], generalizations of the Ising model that describe how glassy materials behave. It was noticed early on that finding the ground state of a spin glass is an NP-hard combinatorial optimization problem. Inspired by this analogy, researchers became increasingly interested in applying techniques from spin glasses to other combinatorial problems. This resulted in some important successes. Using the replica method, which we outline later in this chapter, Mézard and Parisi [42] provided a closed-form analytical prediction for the optimal value of the objective function in the minimum-weight matching problem, over an ensemble of random weights in the asymptotic limit. They subsequently extended this analysis to the traveling salesman problem [41], giving numerical estimates for the asymptotic optimal tour length that later studies have largely confirmed [33, 39, 49].

2 Phase Transitions

Let us return to the issue of worst-case versus average-case complexity, as that leads us to one of the primary subjects of this book. In 1991, a paper by Cheeseman, Kanefsky, and Taylor appeared in the artificial intelligence literature, entitled "Where the *really* hard problems are" [11]. The paper pointed out two empirical properties taking place in random instances of several NP-complete problems, including satisfiability and graph coloring.

- With an appropriate parametrization of the ensemble, there is a sharp phase boundary separating different problem instances. In the case of satisfiability, when the density α of logical constraints in the formula (specifically: the number of clauses per variable) lies below a certain critical value α_c, the formula is almost certainly satisfiable. This is shown by the solid curve in figure 1. When $\alpha > \alpha_c$, the formula is almost certainly unsatisfiable. Furthermore, the threshold becomes increasingly sharp as the problem size n increases. Such a phase boundary had in fact been predicted several years earlier, by Huberman and Hogg [31], in the context of artificial intelligence applications.

- These problems are relatively easy to solve as long as the input instance is clearly in one phase or the other. The hard instances are those near the boundary between the two regions: that is where search algorithms require the largest running time to find a solution, or determine that there is none. Increasing the control parameter α gives rise to an "easy-hard-easy" pattern, as one moves from the underconstrained region,

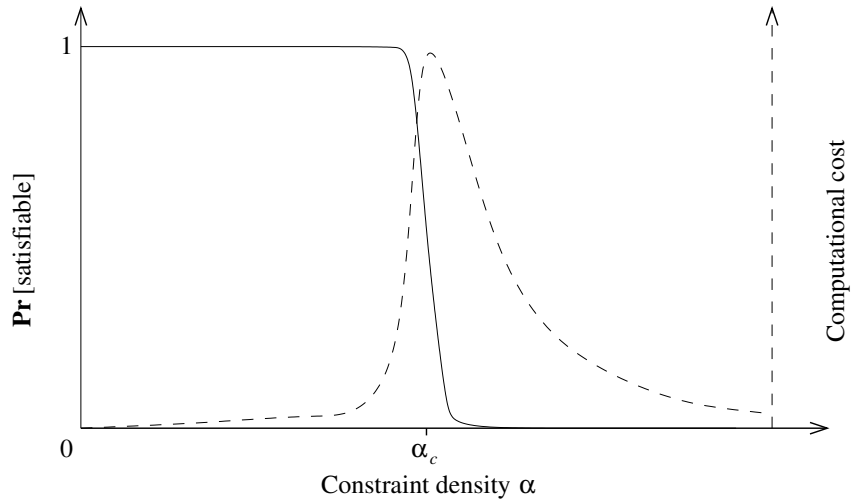

Figure 1: Schematic representation of phase transition in satisfiability. Solid curve denotes probability that a formula is satisfiable, falling rapidly from 1 to 0 at the threshold. Broken curve denotes computational cost for determining satisfiability, and displays easy-hard-easy pattern as constraint density is increased across the threshold.

across the critical threshold, and into the overconstrained region. This is shown by the broken curve in figure 1.

To a statistical physicist this looks conspicuously like a *phase transition*. As some macroscopic parameter crosses a critical threshold, the system undergoes a sudden change in its properties. The canonical example is when water freezes; when the temperature crosses the freezing point, the global behavior changes drastically even though the local interactions between water molecules stay the same. Furthermore, phase transitions are associated with the phenomenon of *critical slowing-down*, where relaxation times diverge and the system can take very long to reach equilibrium. Phase transitions have been a major area of study in mathematical physics since years before the introduction of computational complexity theory [10, 24, 34]. There is a well-developed theory of critical phenomena for modeling and analyzing the system's properties, particularly the types of nonanalyticities that appear at the critical point and how various quantities scale in its vicinity.

The study of phase transitions also has a distinguished history in graph theory. In their 1959 paper, Erdős and Rényi [23] proved that over an ensemble of random graphs with n vertices, for large n the global structure of the graph changes dramatically when the mean degree is increased from slightly below 1 to slightly above 1. Below the threshold, connected clusters of vertices are very small, and with high probability the largest one has size only $O(\log n)$. At the threshold, however, there is an abrupt change, and above the threshold a *giant component* of size $O(n)$ emerges. The probability that the graph possesses a giant component is governed by a *zero-one law*, jumping discontinuously from zero to one at the threshold.

Whether referred to as phase transitions, sharp thresholds, or zero-one laws, these phenomena are hardly anomalous. As discussed in chapter 2, one might argue that in view of the law of large numbers, when the size of the problem gets large enough, the probability distribution of an event over independent random inputs should become very sharply peaked. If one tosses a coin many times, and the coin is even very slightly biased, with high probability the results of the overall experiment will closely reflect the bias. If the coin comes up heads with probability p, and p is an adjustable parameter, then if one were to increase p from $1/2-\epsilon$ to $1/2+\epsilon$ the chances of seeing more heads than tails would jump abruptly from very low to very high. As the number n of coin tosses goes to infinity, this threshold becomes sharp. Admittedly, the example of coin tossing is trivial. But it motivates why one should not be altogether surprised if, say, the minimum fraction of constraints violated in a Boolean formula becomes sharply concentrated about its mean for large formulas. The critical value α_c of the constraint density would then be the value at which this fraction reaches zero and the formula becomes satisfiable.

To the extent that threshold phenomena in computationally hard problems mirror phase transitions as studied in statistical physics and graph theory, much existing work can be carried over to the realm of computer science. That is what Kirkpatrick and Selman did in 1994 [37], in their paper analyzing the satisfiability phase transition. Using the statistical physics techniques of finite-size scaling, they provided numerical estimates of several fundamental properties of the transition, including the control parameter's critical value and the *critical exponent* characterizing how quickly the critical window narrows with increasing problem size. The hope was that an improved understanding of phase transitions in this context would lead to an improved understanding of which instance classes of an NP-hard problem are truly hard, what *makes* them hard, and how one might design algorithms appropriately.

This hope is well on its way to being fulfilled. Over the past decade, the goal of understanding and exploiting the relationship between phase transitions and average-case complexity has mobilized computer scientists, mathematicians, and physicists alike. A flurry of activity has resulted, encompassing conjectures, theoretical insights, and numerical as well as rigorous results. These form the subject of the upcoming chapters, and fall into two broad and occasionally commingled categories: methods of analysis from statistical physics used to investigate algorithmic behavior near the threshold, and probabilistic techniques used to prove properties of the threshold. The first category contains recent work [47] inspired by the statistical mechanics of the satisfiability problem, exploring the connection between the precise nature of the phase transition and the problem's average-case complexity. It also includes the more detailed view of phase structure discussed in chapter 3, originating from spin glass theory, and resulting in algorithmic methods such as the *survey propagation* algorithm described in chapter 4. The second category contains a host of new results in probabilistic analysis. Examples include an exact mathematical characterization of the phase transition [9] for a specific variant of satisfiability (albeit one in P), as well as bounds on the location of the critical threshold for the NP-complete variant 3-SAT, discussed in chapter 7.

3 Basic Models

The satisfiability problem will occupy our attention for much of this book. It has become the combinatorial model of choice for investigating threshold behavior, due to its fundamental role in NP-completeness, its practical applicability in artificial intelligence, and its simple formulation. In this section we define satisfiability, as well as two more model problems, graph coloring and the spin glass. Other combinatorial problems that have attracted the attention of physicists, such as vertex cover and number partitioning, will be described in depth in subsequent chapters of the volume.

3.1 Satisfiability (SAT)

Consider n Boolean variables $x_1, \ldots, x_n$, where each x_i can be assigned the value TRUE or FALSE. Define a logical *formula* $\phi(x_1, \ldots, x_n)$ as a Boolean function of these variables, composed of the logical operators AND, OR, and NOT. The satisfiability problem is to determine whether there is a truth assignment for these variables such that $\phi(x_1, \ldots, x_n) = \text{TRUE}$.

Any Boolean formula can be rewritten in *conjunctive normal form* (CNF), defined as follows. Let a *literal* be either a variable x_i or its complement $\overline{x_i} = \mathsf{NOT}\ x_i$. Let a *clause* be the *disjunction*, i.e., the OR, of a set of literals. A CNF formula is then the *conjunction*, i.e., the AND, of a set of clauses. Thus a possible CNF formula might be

$$\phi = (x_3\ \mathsf{OR}\ \overline{x_2}\ \mathsf{OR}\ x_5)\ \mathsf{AND}\ (\overline{x_1}\ \mathsf{OR}\ x_2)\ \mathsf{AND}\ (x_3\ \mathsf{OR}\ x_4\ \mathsf{OR}\ \overline{x_5}\ \mathsf{OR}\ \overline{x_6}). \tag{2}$$

In this particular formula, the clauses do not all have the same length (number of literals). When they do, and when the length is k, the formula is said to be k-CNF. Satisfiability on k-CNF formulas is a frequently discussed version of the problem, and is known as k-SAT.

It is possible to rewrite any CNF formula, and thus any Boolean formula, in 3-CNF form. For instance, by introducing two new variables z_1 and z_2, eq. (2) can be rewritten as

$$\begin{aligned}\phi \;=\; & (x_3\ \mathsf{OR}\ \overline{x_2}\ \mathsf{OR}\ x_5)\ \mathsf{AND}\ (\overline{x_1}\ \mathsf{OR}\ x_2\ \mathsf{OR}\ z_1)\ \mathsf{AND}\ (\overline{x_1}\ \mathsf{OR}\ x_2\ \mathsf{OR}\ \overline{z_1}) \\ & \mathsf{AND}\ (x_3\ \mathsf{OR}\ x_4\ \mathsf{OR}\ z_2)\ \mathsf{AND}\ (\overline{x_5}\ \mathsf{OR}\ \overline{x_6}\ \mathsf{OR}\ \overline{z_2})\,.\end{aligned}$$

Satisfiability is known to be NP-complete. It is relatively straightforward to transform any CNF formula into a 3-CNF formula in a number of steps that is polynomial in n, using the method above. It follows that 3-SAT, and more generally k-SAT for any $k \geq 3$, is NP-complete. By contrast, 2-SAT is in the complexity class P, and it is *not* generally possible to transform a CNF formula into a 2-CNF formula!

All of these forms of SAT are *decision problems*. The goal is to answer the yes/no question of whether the formula is satisfiable. But there is also an analogous optimization problem, MAX-SAT: find a truth assignment that *maximizes* the number of satisfied clauses. Since it is not clear how to verify that a proposed truth assignment is optimal, it is common to define an NP version of this problem in which we ask whether it is possible to satisfy more than

a certain number of clauses. Interestingly, this version of MAX-2-SAT, the optimization problem corresponding to 2-SAT, is NP-complete in spite of the fact that 2-SAT is solvable in polynomial time.

If we wish to construct *random* k-SAT instances with n variables and m clauses, a natural way to do so is as follows. Construct each clause by choosing k variables randomly and negating each one with probability 1/2. The *clause density* is then $\alpha = m/n$; note that we take the limits $m, n \to \infty$ simultaneously so that α is held constant.

The *satisfiability threshold conjecture* states that there is a critical α_c separating the underconstrained phase from the overconstrained phase. Specifically, it is believed that for each $k \geq 3$, there is a *sharp threshold* α_c with the following property: given any $\epsilon > 0$, in the limit $n \to \infty$ the probability that a random k-SAT formula is satisfiable tends to 1 if $\alpha < (1-\epsilon)\alpha_c$ and to 0 if $\alpha > (1+\epsilon)\alpha_c$. For $k = 2$, this is known to be true, and the threshold is exactly $\alpha_c = 1$. For $k = 3$, it has not yet been rigorously proven that a sharp threshold exists; but assuming it does, numerical evidence and arguments from statistical physics suggest that $\alpha_c \approx 4.27$ [43].

3.2 Graph Coloring (COL)

Consider a graph with n vertices and edges connecting certain pairs of vertices. The graph coloring problem is to assign colors to vertices so that no edge connects two vertices of the same color. When our palette is limited to q colors, the problem is known as q-COL.

Analogously to SAT, 3-COL is NP-complete whereas 2-COL can be solved in polynomial time. The optimization problem corresponding to the decision problem is to find the color assignment that minimizes the number of violated edges, connecting vertices of the same color. This is best described by a picture: figure 2 shows an example of a 3-COL assignment that results in two violated edges.

Now construct graphs randomly, so that an edge connects any given pair of the n vertices with probability p. This is the famous $\mathcal{G}_{n,p}$ ensemble of random graphs studied by Erdős and Rényi [23]. The mean degree of such a graph is $\alpha = p(n-1) \simeq pn$ in the limit $n \to \infty$. Just as there is a sharp threshold at $\alpha = 1$ where the giant component emerges, we believe that for any $q \geq 3$, q-COL has a sharp threshold where the probability of q-colorability drops abruptly from 1 to 0. Recent numerical estimates for the 3-COL threshold location give $\alpha_c \approx 4.69$ or 4.70 [8, 48].

3.3 Spin Glasses

Computer scientists typically describe problems using *hard constraints*, where for instance every clause must be satisfied or every edge must have endpoints of different colors. Physicists, on the other hand, often describe systems in terms of an energy function or *Hamiltonian*. To map one onto the other, we may assign a positive energy contribution to each violated constraint; then the system's ground-state energy will be zero if and only if it is possible to satisfy all constraints. This mapping has been especially fruitful in the case of the spin

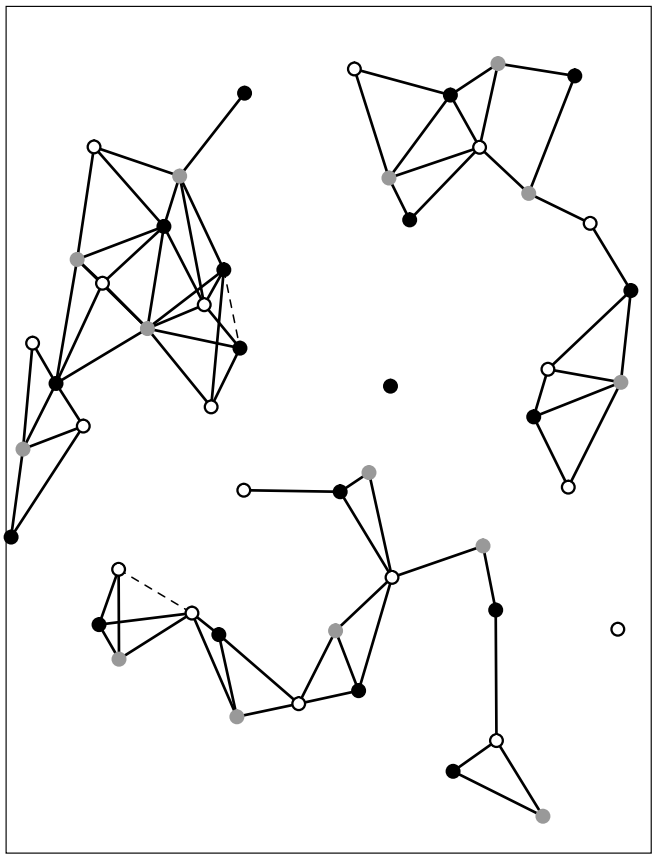

Figure 2: 3-COL assignment in a graph, with two edge violations. Colors assigned to vertices are represented by filled circles, shaded circles and empty circles. Dashed lines denote violated edges in the graph, connecting a pair of vertices of the same color.

glass Hamiltonian, a central topic of research among statistical physicists over the past two decades. Arguably, the spin glass model is as fundamental to the physics of disordered systems as satisfiability is to computational complexity.

Before discussing the spin glass, let us define the simpler *Ising model*, which uses the language of spins to provide a basic explanation of the physical phenomenon of ferromagnetism. Consider n binary variables $s_1, \ldots, s_n$, where any s_i can be assigned the value $+1$ or -1. Physically, these variables represent the quantum mechanical *spin* state of a spin-1/2 particle. We call $s_i = +1$ "up" and $s_i = -1$ "down." The energy associated with interactions between particles is given by the Ising Hamiltonian

$$E_{\text{Ising}} = -J \sum_{\langle i,j \rangle} s_i \, s_j$$

where the notation $\langle i, j \rangle$ represents all pairs of particles i and j that are nearest neighbors on a square lattice. Recall that at a given temperature τ, the probability of the system adopting a specific configuration C of the spin variables is given by the Boltzmann factor, eq. (1). In the thermodynamic limit ($n \to \infty$), and for positive J, there is a phase transition at the *Curie temperature* τ_c. For $\tau > \tau_c$ there is no overall magnetization: there are as many up spins as down spins, and distant sites are independent of one another. But for $\tau < \tau_c$, the system becomes spontaneously magnetized, and the majority of spins align with one another. At $\tau = 0$ the system adopts a ground-state configuration, minimizing E_{Ising} by having all spins s_i taking on the same value: either up or down.

The Ising model is clearly an idealization or "toy model" of magnetic materials: in particular, it assumes that spins are arranged in a square lattice and only interact with their nearest neighbors. However, physicists have found that while the transition temperature τ_c depends on the details of a materials' topology and interactions, the *type* of transition does not. In particular, the critical exponents describing how correlations, relaxation times, and magnetizations scale near the phase transition depend only on the dimension of the lattice, and not, for instance, on whether it is triangular versus square, or on whether interactions extend to neighbors several steps away. This fortunate fact means, first of all, that the Ising model is a far more effective description of real ferromagnetic materials, such as iron, than we might have originally thought. Second of all, it inspires physicists to think of systems as being grouped into *universality classes* that abstract away their details and capture the qualitative aspects of their behavior. Such a qualitative classification is not altogether dissimilar to that of computational complexity theory, in which constraint satisfaction problems are generically NP-complete and only a few special cases are in P.

While the Ising model displays a nontrivial phase structure, as a combinatorial optimization problem it is uninteresting. But now, imagine that the material being modeled is not entirely homogeneous and that interactions are not only between lattice neighbors. Furthermore, imagine that not all interactions are ferromagnetic, encouraging spins to align in the same direction; some are also antiferromagnetic, encouraging them to align in opposite

directions. The Hamiltonian generalizes to

$$E = -\sum_{i<j} J_{ij}\, s_i\, s_j\,, \tag{3}$$

where the coupling constants J_{ij} can be positive, negative or zero. Equation (3) describes the Hamiltonian for the spin glass, so named because magnetic materials modeled in this way, such as copper-manganese alloys, can have a "glassy" phase displaying short-range order but long-range disorder.

Finding the ground state of a spin glass is, in general, an NP-hard problem [6]. Like the models we have considered up until now, it is appropriate to consider the spin glass over an ensemble of random inputs. Physically, this makes sense: glasses are distinguished by the randomness of the disorder spread through them, as distinct from the ordered structure of a crystalline solid. Thus, the theory of spin glasses has focused on studying these materials when the coupling constants J_{ij} are chosen randomly. One analytically tractable ensemble, introduced by Sherrington and Kirkpatrick in 1975 [51], considers J_{ij} for each unordered pair i, j to be chosen independently from a Gaussian distribution with mean zero. Another ensemble, the Edwards-Anderson $\pm J$ spin glass [22], considers J_{ij} to be nonzero only when i and j are lattice neighbors—as in the Ising model—in which case it is chosen from $+J_0$ or $-J_0$ with equal probability.

The basic problem we have posed for the spin glass is an optimization problem rather than a decision problem. Of course, the model is related in a formal sense to combinatorial problems that are well-studied outside of the physics community. For instance, a straightforward transformation of eq. (3) shows that finding the ground state of a spin glass is equivalent to the classic graph-theoretic problem of weighted MAX-CUT [6]. In this problem, the vertices of a weighted graph must be partitioned into two sets, while maximizing the sum of the weights along all edges that connect vertices in opposite sets.

Moreover, the NP-complete problems we have presented earlier can be captured by variants of the spin glass Hamiltonian. Consider first graph coloring. One may generalize the basic Ising model to the *Potts* model, where spins s_i are not binary but can take on any one of q values, or colors, in $\{1, \ldots, q\}$. The Potts Hamiltonian is then

$$E_{\text{Potts}} = -J \sum_{\langle i,j \rangle} \delta(s_i, s_j)\,,$$

where the Kronecker δ function gives 1 when $s_i = s_j$ and 0 otherwise. This model can be extended from a lattice to an arbitrary graph. If J is negative, the Potts model is antiferromagnetic and adjacent spins try to have different colors; since the energy is J times the number of violated edges, the ground-state energy is zero if and only if the graph is q-colorable.

Now consider satisfiability. For each Boolean variable x_i, define a corresponding spin $s_i = +1$ if $x_i = \textsc{true}$, and $s_i = -1$ if $x_i = \textsc{false}$. Define the clause matrix W such that $W_{ji} = +1$ if clause j includes literal x_i, $W_{ji} = -1$ if clause j includes literal $\overline{x_i}$, and $W_{ji} = 0$

otherwise. It follows that for k-CNF formulas, the indicator expression

$$V_j = \frac{1}{2^k} \prod_{i=1}^{n} (1 - W_{ji} s_i)$$

takes on the value 0 when clause j is satisfied and 1 when it is violated. The number of violated clauses is then given by

$$E_{k-\mathrm{SAT}} = \sum_{j=1}^{m} V_j = \frac{1}{2^k} \sum_{j=1}^{m} \prod_{i=1}^{n} (1 - W_{ji} s_i) \,. \tag{4}$$

Minimizing the number of violated clauses is the same as maximizing the number of satisfied clauses, so finding the ground state of this "k-SAT Hamiltonian" is equivalent to solving the MAX-k-SAT problem. The ground-state energy is zero if and only if the formula is satisfiable. The main difference between eq. (4) and the spin glass Hamiltonian of eq. (3) is that now there are interactions between groups of k spins, rather than just between pairs. Fortunately, models of this type are well known in statistical physics.

In satisfiability and graph coloring, we might be tempted to focus simply on finding a solution or confirming that none exists. Indeed, this is traditionally how computer scientists have framed such problems. But the analogy with spin glasses reveals that the system's behavior is far richer than that: the transition from satisfiability to unsatisfiability is only a part of the picture. As we will see in the next section, the *order parameter* relevant to the transition is not only the probability that a solution exists, but is in fact an entire probability distribution describing the solution structure. Certain techniques developed for spin glasses lend themselves particularly well to the study of k-SAT [46], leading to both analytical insights and new algorithms.

4 The View from Statistical Physics

The chapters of this volume discuss diverse methods of analysis that researchers have brought to bear on combinatorial problems, drawn from computer science, mathematics, and physics. In this section, we give a glimpse of how statistical physicists study the structure of these problems, focusing in particular on the replica method.

The replica method is a powerful analytical approach developed in spin glass theory and recently applied to a variety of NP-hard problems. We do not attempt to cover the technical details involved in the replica method; rather, we outline the method in the hope of demonstrating how analyzing the physics of the problem can provide valuable insight. Our discussion serves as background to much of the material in this book. In illustrating the replica method and related approaches, we use k-SAT as a model problem and follow the language of Martin et al. [40].

Consider a physical system in thermal equilibrium at temperature τ. As we have seen in eq. (1), the probability of being at a specific state C is proportional to the Boltzmann factor

$\exp(-E(C)/\tau)$. Normalizing this gives

$$\mathbf{Pr}[C] = \frac{e^{-E(C)/\tau}}{Z}\,, \qquad \text{where } Z = \sum_C e^{-E(C)/\tau}\,.$$

But the quantity Z, called the *partition function*, is far more than a mere normalizing constant: it embodies a considerable amount of information about the system. Z is a generating function from which one can directly obtain many thermodynamic quantities of interest. For instance, the thermally-averaged energy of the system,

$$\langle E\rangle = \sum_C E(C)\,\mathbf{Pr}[C] = \frac{1}{Z}\sum_C E(C)e^{-E(C)/\tau}\,,$$

is given by Z and its first derivative:

$$\langle E\rangle = \frac{\tau^2}{Z}\frac{dZ}{d\tau}\,.$$

In principle, the ground-state energy E_{GS} is then found by taking the limit $\tau \to 0$. There, one can show that $\langle E\rangle$ approaches $-\tau \log Z$.

This is all for a given *realization* of the system, meaning a given setting of the couplings J_{ij} in a spin glass, or a given instance of graph coloring or k-SAT. We want to work over an ensemble of random instances. In the case of k-SAT, in order to study the problem's phase behavior for clause density α near α_c, we need to understand the distribution of the random variable E_{GS} over the ensemble. Notably, if we denote by $\overline{E_{GS}}$ the average over random instances at a given α—not to be confused with the thermal average above—then a typical instance will be satisfiable if $\overline{E_{GS}} = 0$ and unsatisfiable if $\overline{E_{GS}} > 0$.

Using the k-SAT Hamiltonian in eq. (4), $\overline{E_{GS}}$ for $m = \alpha n$ clauses can be written as

$$\overline{E_{GS}} = \overline{\min_{s_1,\ldots,s_n} \frac{1}{2^k}\sum_{j=1}^{\alpha n}\prod_{i=1}^{n}(1 - W_{ji}s_i)}\,.$$

Since this form does not clearly lend itself to analysis, let us make use of the partition function. We know that

$$\overline{E_{GS}} = -\lim_{\tau\to 0}\tau\,\overline{\log Z}$$

where

$$Z = \sum_{s_1,\ldots,s_n} \exp\left(-\frac{1}{\tau}\frac{1}{2^k}\sum_{j=1}^{\alpha n}\prod_{i=1}^{n}(1 - W_{ji}s_i)\right)\,.$$

Calculating the ensemble average of the logarithm of a function is difficult if not impossible. However, note that the logarithm satisfies the property

$$\log Z = \lim_{r\to 0}\frac{Z^r - 1}{r}\,,$$

so

$$\overline{E_{GS}} = -\lim_{\tau\to 0}\lim_{r\to 0}\tau\frac{\overline{Z^r}-1}{r}\,.$$

We now introduce the replica trick: proceed as if r were a positive integer, and calculate the rth moment of the partition function. Then, after doing so, perform an analytic continuation to real r and finally take the limit $r \to 0$. Proceeding in this way,

$$\begin{aligned}\overline{Z^r} &= \overline{\left[\sum_{s_1,\ldots,s_n}\exp\left(-\frac{1}{\tau}\frac{1}{2^k}\sum_{j=1}^{\alpha n}\prod_{i=1}^{n}(1-W_{ji}s_i)\right)\right]^r}\\ &= \sum_{s_1^1,\ldots,s_n^1}\cdots\sum_{s_1^r,\ldots,s_n^r}\overline{\prod_{j=1}^{\alpha n}\exp\left(-\sum_{a=1}^{r}\frac{1}{\tau}\frac{1}{2^k}\prod_{i=1}^{n}(1-W_{ji}s_i^a)\right)}\;.\end{aligned}$$

We can interpret the leading sums as being over the states of r copies, or *replicas*, of the system, all with the same random realization. Thus, for integer r, the rth moment $\overline{Z^r}$ is the average partition function for the r replicas.

Although it is much easier to calculate $\overline{Z^r}$ than $\overline{\log Z}$, it is by no means easy. Furthermore, even if we can calculate $\overline{Z^r}$ for all integer r, its analytic continuation may or may not make sense in the limit $r \to 0$. One can argue that for finite n, given all moments of Z, one can perfectly reconstruct its probability distribution and, therefore, justify the existence and uniqueness of the analytic continuation. Unfortunately, we are interested in the $n \to \infty$ limit where this argument does not hold. It is remarkable that in spite of the lack of rigorous justification, the replica method has been extremely successful in predicting effects later confirmed by other theoretical approaches and by numerical experiments.

After a certain amount of algebraic manipulation [40], writing $\overline{Z^r}$ in the limit $n \to \infty$ leads to a saddle-point integral. In order to write this integral, we introduce the notation $\vec{\sigma}$ to represent an r-dimensional binary vector, $\vec{\sigma} \in \{-1,1\}^r$, with components $\sigma^1,\ldots,\sigma^r$ over the r replicas. We furthermore introduce $\vec{u}$ to represent a real-valued 2^r-dimensional vector with components $\{u_{\vec{\sigma}} : \vec{\sigma} \in \{-1,1\}^r\}$, in turn, over the 2^r possible values of $\vec{\sigma}$. Then,

$$\begin{aligned}\overline{Z^r} &= \int\prod_{\vec{\sigma}} du_{\vec{\sigma}}\;\delta\left(\sum_{\vec{\sigma}}u_{\vec{\sigma}}-1\right)e^{nF(\vec{u})}\text{, where}\\ F(\vec{u}) &= \alpha\log\left[\sum_{\vec{\sigma}_1,\ldots,\vec{\sigma}_k}u_{\vec{\sigma}_1}\cdots u_{\vec{\sigma}_k}\exp\left(-\frac{1}{\tau}\sum_{a=1}^{r}\prod_{l=1}^{k}\frac{1+\sigma_l^a}{2}\right)\right]-\sum_{\vec{\sigma}}u_{\vec{\sigma}}\log u_{\vec{\sigma}}\,.\end{aligned}$$

In the replica formalism, $u_{\vec{\sigma}}$ has a physical interpretation as the fraction of spin indices i with the sequence of values $s_i^a = \sigma^a$ over the replicas $a \in \{1,\ldots,r\}$. Note that $u_{\vec{\sigma}}$ is constrained by the Dirac δ-function, representing a normalization condition. Counting the 2^r dimensions of $\vec{u}$ and the one constraint, $\overline{Z^r}$ is an integral over a space of 2^r-1 dimensions.

The saddle-point approximation for this integral gives $\overline{Z^r} \sim \exp(nF_{\max})$ in the large n limit, where F is maximized over all $\vec{u}$. F is symmetric under any permutation $\pi(a)$ of

the replicas: as long as a certain vector $\vec{u}^*$ maximizes F, so too does any vector $\vec{u}$ where $u_{\sigma^1,\ldots,\sigma^r} = u^*_{\sigma^{\pi(1)},\ldots,\sigma^{\pi(r)}}$ for all $\vec{\sigma}$, i.e., $\vec{u}^*$ with its components reshuffled. If, furthermore, F has a *unique* maximum, then $u^*_{\sigma^1,\ldots,\sigma^r} = u^*_{\sigma^{\pi(1)},\ldots,\sigma^{\pi(r)}}$ for all $\vec{\sigma}$, and we may restrict our space to vectors $\vec{u}$ that are themselves invariant under permutation. For a given $u_{\vec{\sigma}}$, we need only consider how many components of $\vec{\sigma}$ have the value 1 and how many have the value -1.

The assumption of a unique maximum for F is known as *replica symmetry*. Under such an assumption, we can write $u_{\vec{\sigma}}$ in the form

$$u_{\vec{\sigma}} = \int_{-1}^{1} d\mu\, P(\mu) \prod_{a=1}^{r} \frac{1+\mu\sigma^a}{2}\,, \tag{5}$$

which may be inserted into the equation for $F(\vec{u})$. We can now perform the analytic continuation of r to the real numbers and take the limit $r \to 0$. The result is a functional equation for $F[P(\mu)]$. Using Lagrange multipliers we optimize F over the function space of $P(\mu)$, leading to a self-consistent integral equation for $P(\mu)$. What is the meaning of $P(\mu)$ when $\tau \to 0$? Let μ_i be the magnetization, or average value of a given variable s_i over all possible ground-state configurations. It turns out that if the assumption of replica symmetry is correct, $P(\mu)$ gives the distribution of magnetizations μ_i. Thus, $P(\mu)$ is simply the probability density for a variable's value averaged over the ground states.

The self-consistent equation for $P(\mu)$ admits a family of solutions. All of these share an important qualitative feature in the $\tau \to 0$ limit. For α below a fixed α_c ($\alpha_c = 1$ for 2-SAT, $\alpha_c \approx 4.6$ for 3-SAT), $P(-1) = P(1) = 0$, meaning that the expected fraction of variables completely "frozen" to a value of TRUE or FALSE over all optimal assignments is zero. For $\alpha > \alpha_c$, however, the form of $P(\mu)$ changes and the distribution starts having nonzero weight at $\mu = \pm 1$: a "backbone" emerges, in which certain variables are constrained to be TRUE or FALSE. $P(\mu)$ is, therefore, not only a function, but also an order parameter that signals a phase boundary.

In calculating $\overline{Z^r} \sim \exp(nF_{\max})$ under the assumption of replica symmetry, one finds that the structural change in $P(\mu)$ is mirrored by a structural change in $\overline{E_{GS}}$. For $\alpha < \alpha_c$, $\overline{E_{GS}} = 0$, whereas for $\alpha > \alpha_c$, $\overline{E_{GS}} > 0$. The latter is certainly not surprising: if a finite fraction of the variables could be frozen to specific values over all satisfying assignments, the addition of a single new clause would with finite probability lead to a violation. Interestingly, though, the nature of the transition in $P(\mu)$ is different for 2-SAT and for 3-SAT [47]. In the case of 2-SAT, there is a continuous transition. At $\alpha = \alpha_c$, $P(\pm 1) = 0$ and then increases continuously with increasing α. In the case of 3-SAT, there is a discontinuous transition. Already at $\alpha = \alpha_c$, a nonzero backbone spontaneously emerges with $P(\pm 1) > 0$.

The replica symmetric (RS) solution is a very convenient one, and helps provide a valuable physical understanding of the phase structure. But it only tells part of the story. The manner in which the various limits are taken lacks mathematical rigor, and there is no guarantee that replica symmetry holds: F may not have a unique maximum. Indeed, while the RS solution is exact for 2-SAT, it does not give the correct threshold location for 3-SAT. Empirically, we know that α_c is closer to 4.27 than to 4.6. We do not yet know how to provide a firm mathematical foundation for the replica method, but we can, on the other hand, improve

upon the RS solution by explicitly introducing a form of *replica symmetry breaking* (RSB) into $F(\vec{u})$ [7].

The intuition for RSB is as follows. Consider two different optimal assignments for a 3-SAT instance, and look at the relative Hamming distance d between them: the fraction of variables that are set to TRUE in one configuration but FALSE in the other. A consequence of the RS assumption and eq. (5) is that, for large n, over the space of optimal configuration pairs the Hamming distance distribution $p(d)$ is sharply peaked at a single value

$$d_{RS} = \frac{1}{2} - \frac{1}{2} \int_{-1}^{1} \mu^2 P(\mu)\, d\mu \,.$$

Under RSB, however, one allows $p(d)$ to take on a non-trivial structure. The simplest form of RSB, called "one-step replica symmetry breaking," assumes that the distribution has two sharp peaks. One might imagine optimal assignments being contained within clusters, so that any pair within a given cluster is separated by Hamming distance d_0, but pairs from different clusters are separated by Hamming distance d_1. More complex forms of RSB impose multiple peaks in $p(d)$, corresponding to a hierarchical construction of clusters within clusters. The limiting case of this is the *full RSB* scheme, introduced by Parisi for the Sherrington-Kirkpatrick spin glass [44], where $p(d)$ is a continuum.

For 3-SAT, one-step RSB appears to be sufficient. The resulting analysis [7, 43] suggests a more subtle phase diagram than the one seen in figure 1. For α below $\alpha_{RSB} \approx 3.92$, the RS solution appears to be correct. The distribution $p(d)$ only displays a single peak, at d_{RS}, and with high probability instances are satisfiable. In this "easy satisfiable" phase, one is very likely to find a satisfying assignment with minimal computational effort. At α_{RSB}, the system undergoes a replica symmetry breaking transition, and for $\alpha_{RSB} < \alpha < \alpha_c$, $p(d)$ displays both peaks d_0 and d_1. This is the "hard satisfiable" phase, where with high probability instances are satisfiable, but the satisfying assignments are separated into clusters, creating many local optima and making it difficult to find a solution. Finally, at α_c, the system becomes overconstrained, and enters the unsatisfiable phase where satisfying assignments are unlikely to exist.

This clustering picture is one of the topics of chapter 3. The most successful approach for analyzing it is based on the *cavity method*, which uses techniques closely related to the replica method but not identical. In addition to providing a refined understanding of the k-SAT phase structure, the cavity method results in a prediction for the threshold location that is believed to be exact: $\alpha_c \approx 4.27$ for $k = 3$. Even more significantly, the method gives rise to an efficient algorithmic procedure for finding satisfying assignments in the hard satisfiable phase. The procedure, known as *survey propagation*, is discussed and analyzed in chapter 4. It is a perfect example of how improved insights into critical phenomena in combinatorial problems can lead to direct improvements in algorithmic performance.

5 The View from Computer Science

Theoretical tools from the physics community, such as the replica method, have been instrumental in improving our understanding of the fundamentals of computing. At the same time, advances in theoretical computer science—both from a probabilistic and from an algorithmic perspective—have expanded the frontiers of mathematical physics. One clear difference between these two approaches is in the methods of analysis. But another very fundamental cultural difference involves the questions that are asked. Statistical physicists have been deeply interested in understanding the structure of computational problems and algorithms, characterizing critical behavior through appropriate order parameters and critical exponents. Theoretical computer scientists have focused on proving threshold properties with the aim of obtaining increasingly tight bounds on algorithmic performance. One of the great successes of the past decade has been the interaction between these two groups. Let us now highlight some of the main results that have come from the computer science community, and how these have motivated further results among physicists.

Two key challenges related to the phase transition in satisfiability are proving that a unique sharp threshold exists for random k-SAT, and identifying its location. For $k \geq 3$, both of these remain open problems. Friedgut [26] has shown that there exists a *function* $\alpha_c(n)$ such that the threshold becomes arbitrarily sharp about $\alpha_c(n)$ in the large n limit. However, it has not been proven that $\alpha_c(n)$ converges to a constant α_c as $n \to \infty$, in spite of overwhelming belief and a large body of experimental and nonrigorous evidence such as the survey propagation method mentioned above. Much effort has, therefore, been devoted to establishing upper and lower bounds on the threshold location, discussed at length in chapter 7. For $k = 3$, the best lower bound [29, 35] and upper bound [19] to date give $3.52 < \alpha_c < 4.506$, still leaving a considerable gap. For large k, the best lower bounds [1, 2] and upper bound [25] give $\alpha_c = (1 - o(1))2^k \log k$, confirming nonrigorous results from the replica method [46] and survey propagation [43].

The case of $k = 2$ is quite different. As mentioned above, random 2-SAT is known to have a sharp threshold at $\alpha_c = 1$ [12, 27]. Moreover, a lot is now understood about this transition [9], including exact values for the critical exponents that characterize the sharpness of the threshold. This has put on a firm mathematical footing the original 1994 numerical study of Kirkpatrick and Selman [37], at least for $k = 2$. In addition, Bollobás et al. [9] have proven an essential property concerning the nature of the 2-SAT transition: an order parameter called the *spine* goes to zero continuously as α approaches α_c from above. Since the spine is an upper bound on the backbone mentioned above, a corollary of this result is that the replica prediction of a continuous backbone transition at the 2-SAT threshold is correct.

Rigorous results have been obtained for a number of other variants of satisfiability. One example is the problem of *1-in-k* satisfiability, a form of k-SAT where a clause is satisfied only if *exactly* one literal in the clause is TRUE. For $k \geq 3$, 1-in-k-SAT has been shown [4] to have a sharp threshold at $\alpha_c = 1/\binom{k}{2}$. Interestingly, even though 1-in-k-SAT is NP-complete, what makes the precise threshold value possible to obtain here is the structural similarity of the

problem to 2-SAT. Just as in 2-SAT, if one maps the formula to a hypergraph, unsatisfiability emerges as a direct consequence of contradictory cycles of implications. Another example of satisfiability—albeit not an NP-complete one—with a rigorously determined threshold is the XORSAT problem. Here, the logical XOR operator, or *exclusive or*, replaces the OR operator in a CNF formula. This is equivalent to solving a linear system of equations modulo 2, and is solvable in polynomial time by Gaussian elimination. The problem has attracted the interest of statistical physicists as being a form of SAT that is particularly amenable to analysis via the replica method. A one-step RSB calculation [50] predicts a threshold location of $\alpha_c \approx 0.918$ for 3-XORSAT. Recent work in theoretical computer science [18] has shown rigorously that the replica prediction is indeed exact, justifying not only the numerical results of the replica method but also the physical picture that it provides. In particular, this work has proven the existence of a hard satisfiable phase like the one proposed for 3-SAT, where solutions exist but are grouped into clusters with large Hamming distances between them.

Finally, several connections have been suggested between the nature of the phase transition and the complexity of various classes of algorithms. Recall that in 2-SAT, the backbone order parameter marking the transition is continuous (though nonanalytic) across the threshold, whereas in 3-SAT, the replica method indicates that it is discontinuous. Now consider a broad class of resolution algorithms that work by successively assigning values to variables and backtracking to an earlier assignment when a constraint is violated. This class of algorithms, known as Davis-Putnam or DPLL [15, 16], is described in chapter 3. To analyze the relation between the continuity of the order parameter and the complexity of DPLL algorithms on random formulas at the transition, Monasson et al. [47] have considered the problem $(2+p)$-SAT, for real $p \in [0, 1]$, that interpolates between 2-SAT and 3-SAT. Random instances of $(2+p)$-SAT with constraint density α are generated by drawing $\alpha(1-p)n$ random clauses of length 2 and αpn random clauses of length 3. Based on the replica method, the transition appears continuous for p close to zero but becomes discontinuous when p exceeds a critical value $p_0 = 2/5$. At the same p_0, the complexity of DPLL algorithms seems to change from polynomial to exponential.

Recent results in the computer science literature have both supported this picture and helped clarify its limitations. It has been confirmed [5] that for $p \in [0, 2/5]$, the threshold behavior of $(2 + p)$-SAT resembles that of random 2-SAT, and the location of the threshold is simply the constraint density at which the 2-CNF component of the random formula becomes unsatisfiable. Furthermore, the spine order parameter that was proven continuous for the 2-SAT threshold [9] has indeed been shown to be discontinuous for k-SAT when $k \geq 3$ [32]. This is a weaker result than showing that the backbone is discontinuous, and hence does not in itself confirm the replica prediction, but motivates considering the spine as the order parameter of interest. The connection with the complexity of DPLL algorithms closely involves the study of *proof complexity*, of key interest in automated theorem proving: one attempts to bound the number of steps needed to confirm or refute a proposition. A combination of old and new work in proof complexity [3, 13, 32] shows that a discontinuity in the k-SAT spine directly correlates with exponential resolution complexity, meaning

that verifying the unsatisfiability of a formula takes DPLL an exponential number of steps. Conversely, a continuous spine implies that the resolution complexity must be smaller than any exponential—as for instance in 2-SAT and 1-in-k-SAT, where contradictions are verified in $O(n)$ steps. These results can be extended from k-SAT to more general constraint satisfaction problems [45], and lead to the hope that a better understanding of a problem's resolution complexity might help identify the threshold location.

6 Outline of the Volume

In this introductory chapter, we have sketched some of the models, methods of analysis, and results at the intersection of complexity theory and statistical mechanics. These have given rise to the rich set of scientific interactions that form the subject of this book as a whole. Our broad aim is to provide the reader with an appreciation of how physical approaches have contributed to the study of computationally hard problems, and how advances in probabilistic and algorithmic techniques have made the connection such a fruitful one. To this end, the volume draws on contributions from authors in the computer science, mathematics, and physics communities—roughly in equal measure, and often in close collaboration. We hope to make apparent how the constituent disciplines have worked together to create a new and flourishing field of research. Among the products of these research efforts have been new algorithmic methods, new combinatorics, and new physics.

Our main focus here is on phase transitions and threshold phenomena in the context of random combinatorial structures. The book is composed of four parts:

- **Part 1: Fundamentals.** This part includes two chapters, introducing the reader to the basic combinatorial and physical concepts of the volume. The present chapter has given the essential background. Chapter 2 provides a deeper view of the mathematical foundations of threshold phenomena. The authors explain the origins of sharp thresholds, and guide the reader through an explicit description of how they relate to complexity theory and mathematical physics, as well as applications in other fields of science such as economics. They highlight a number of important open problems that are echoed in subsequent chapters.

- **Part 2: Statistical Physics and Algorithms.** The four chapters of this part are devoted to the role of physical analysis in computational problems as well as the use of algorithmic methods in understanding the structure of physical problems. Chapter 3 studies the dynamics of DPLL and local search algorithms, showing how both the critical threshold and a different "dynamical" transition affect algorithmic complexity. Chapter 4 presents the algorithmic framework of survey propagation, describes the theory, and applies it to satisfiability and graph coloring. Chapter 5 discusses number partitioning, analyzing the phase transition by exploiting its similarities to a tractable physical problem. Chapter 6 considers the spin glass model, and shows how a well-studied algorithmic method can lead to a physical understanding of the dynamics of glassy systems.

- **Part 3: Identifying the Threshold.** This part contains three chapters that demonstrate the probabilistic and numerical techniques used to argue for the existence of thresholds, as well as to identify their location. Chapter 7 reviews recent improvements on upper bounds for α_c in the case of k-SAT, illustrating the methods by which these bounds have been derived. Chapter 8 presents a powerful methodology for proving conditional randomness that enables the analysis of search heuristics and leads to improved lower bounds on α_c. Chapter 9 discusses the phase transition for a form of satisfiability called HornSAT, using numerical finite-size scaling techniques together with related analytical models to investigate the nature of the threshold.

- **Part 4: Extensions and Applications.** The final part of the volume contains four chapters that connect the foregoing discussion to a range of applications extending far beyond model problems. Chapter 10 considers phase transitions in the context of quantum computing, and shows how the behavior of quantum search algorithms relates to problem structure. Chapter 11 relates computational complexity to physical models of surface growth, analyzing the scalability of parallel simulation processes by means of techniques from statistical mechanics. Chapter 12 introduces a biological motivation, investigating a model of RNA folding by way of random graph analysis and the threshold where a giant component emerges. Chapter 13 extends the concept of typical-case complexity beyond ensembles of random instances: the chapter proposes a framework in which realistic instances of a problem can be considered, and provokes thoughts on future directions for computational complexity.

The entire collection of chapters is intended to form a cohesive volume, rather than simply a set of technical articles. The chapters have been arranged to form a logical and pedagogical progression, although they may also be appreciated individually and do not necessarily have to be read in order.

A major challenge in editing a volume of this kind is anticipating the audience and adjusting the level of discussion to this audience. We expect our subject to interest a wide range of researchers in computer science, mathematics, and physics. Given the different scientific cultures and backgrounds, it is an ambitious exercise to make statistical physics understandable to computer scientists and computer science understandable to physicists. Nevertheless, we have worked hard to strike the right balance between phenomenology and theory. Wherever possible, we have standardized notation and terminology across the entire volume. We assume only basic literacy in the research tools of discrete mathematics and physics.

Of course, the intersection of computational complexity and statistical physics is vast. It is impossible to do it justice in a single volume. By focusing primarily on threshold phenomena, we necessarily omit other important and exciting research topics, such as the probabilistic analysis of Markov chain algorithms and the study of network dynamics. Nevertheless, we hope that this volume will serve both as a reference on an emerging cross-disciplinary field, and as a snapshot of the state of the field at this point in time.

References

[1] Achlioptas, D., and C. Moore. "The Asymptotic Order of the Random k-SAT Threshold." In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS '02)*, 779–788. New York: ACM Press, 2002.

[2] Achlioptas, D., and Y. Peres. "The Threshold for Random k-SAT is $2^k \ln 2 - O(k)$." In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC '03)*, 223–231. New York: ACM Press, 2003.

[3] Achlioptas, D., P. Beame, and M. Molloy. "A Sharp Threshold in Proof Complexity." In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC '01)* 337–346. New York: ACM Press, 2001.

[4] Achlioptas, D., A. Chtcherba, G. Istrate, and C. Moore. "The Phase Transition in NAESAT and 1-in-k SAT." In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '01)*, 721–722. New York: ACM Press, 2001.

[5] Achlioptas, D., L. M. Kirousis, E. Kranakis, and D. Krizanc. "Rigorous Results for Random $(2 + p)$-SAT." *Theor. Comp. Sci.* 265 (2001): 109–129.

[6] Barahona, F. "On the Computational Complexity of Ising Spin Glass Models." *J. Phys. A* 15 (1982): 3241–3253.

[7] Biroli, G., R. Monasson, and M. Weigt. "A Variational Description of the Ground State Structure in Random Satisfiability Problems." *Eur. Phys. J. B* 14 (2000): 551–568.

[8] Boettcher, S., and A. G. Percus. "Extremal Optimization at the Phase Transition of the Three-Coloring Problem." *Phys. Rev. E* 69 (2004): 066703.

[9] Bollobás, B., C. Borgs, J.T. Chayes, J. H. Kim, and D. B. Wilson. "The Scaling Window of the 2-SAT Transition." *Random Structures and Algorithms* 18 (2001): 201–256.

[10] Broadbent, S. R., and J. M. Hammersley. "Percolation Processes: I. Crystals and Mazes." *Proc. Cambridge Phil. Soc.* 53 (1957): 629–641.

[11] Cheeseman, P., B. Kanefsky, and W. M. Taylor. "Where the *Really* Hard Problems Are." In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI '91)*, 331–337. Sydney, Australia, 1991.

[12] Chvátal, V., and B. Reed. "Mick Gets Some (the Odds are on his Side)." In *Proceedings of the 32nd Symposium on Foundations of Computer Science (FOCS '02)* 620–626. Pittsburgh, PA, 1992.

[13] Chvátal, V., and E. Szemerédi. "Many Hard Examples for Resolution." *J. ACM* 35 (1988): 759–768.

[14] Cook, S. "The Complexity of Theorem-Proving Procedures." In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing (STOC '71)*, 151–158. New York: ACM Press, 1971.

[15] Davis, M., and H. Putnam. "A Computing Procedure for Quantification Theory." *J. ACM* 7 (1960): 201–215.

[16] Davis, M., G. Logemann, and D. Loveland. "A Machine Program for Theorem Proving." *Communications of the ACM* 5 (1962): 394–397.

[17] Devlin, K. *The Millenium Problems: The Seven Greatest Unsolved Mathematical Puzzles of our Time.* New York: Basic Books, 2002.

[18] Dubois, O., and J. Mandler. "The 3-XORSAT Threshold." In *Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS '02)* 769–778. Vancouver, Canada, 2002.

[19] Dubois, O., Y. Boufkhad, and J. Mandler. "Typical Random 3-SAT Formulae and the Satisfiability Threshold." In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '00)*, 126–127. New York: ACM Press, 2000.

[20] Edmonds, J. "Maximum Matchings and a Polyhedron with 0,1-Vertices." *J. Res. Nat'l. Bureau of Standards (Section B)* 69B (1965): 125–130.

[21] Edmonds, J. "Paths, Trees and Flowers." *Canad. J. Math.* 17 (1965): 449–467.

[22] Edwards, S. F., and P. W. Anderson. "Theory of Spin Glasses." *J. Phys. F: Metal Phys.* 5 (1975): 965–974.

[23] Erdős, P., and A. Rényi. "On Random Graphs I." *Publ. Math. Debrecen* 6 (1959): 290–297.

[24] Fisher, M. E. "The Theory of Equilibrium Critical Phenomena." *Rep. Prog. Phys.* 30 (1967): 615–730.

[25] Franco, J., and M. Paull. "Probabilistic Analysis of the Davis-Putnam Procedure for Solving the Satisfiability Problem." *Discrete Appl. Math.* 5 (1983): 77–87.

[26] Friedgut, E. "Sharp Thresholds of Graph Properties, and the k-SAT Problem." *J. Amer. Math. Soc.* 12 (1999): 1017–1054.

[27] Goerdt, A. "A Threshold for Satisfiability." *J. Comp. & Sys. Sci.* 53 (1996): 469–486.

[28] Goldberg, A. "On the Complexity of the Satisfiability Problem." Courant Computer Science Report No. 16, New York University, 1979.

[29] Hajiaghayi, M., and G. B. Sorkin. "The Satisfiability Threshold of Random 3-SAT is at least 3.52." Technical Paper RC22942, IBM, 2003. Available at http://arxiv.org/abs/math.CO/0310193 (accessed August 5, 2005).

[30] Hartmanis, J., and R. Stearns. "On the Computational Complexity of Algorithms." *Trans. Am. Math. Soc.* 117 (1965): 285–306.

[31] Huberman, B., and T. Hogg. "Phase Transitions in Artificial Intelligence Systems." *Art. Intel.* 33 (1987): 155–171.

[32] Istrate, G., S. Boettcher, and A. G. Percus. "Spines of Random Constraint Satisfaction Problems: Definition and Connection with Computational Complexity." *Ann. Math. and Art. Intel.* 44 (2005): 353–372.

[33] Johnson, D. S., L. A. McGeoch, and E. E. Rothberg. "Asymptotic Experimental Analysis for the Held-Karp Traveling Salesman Bound." In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '96)*, 341–350. New York: ACM Press, 1996.

[34] Kadanoff, L. P., W. Götze, D. Hamblen, R. Hecht, E. A. S. Lewis, V. V. Palciauskas, M. Rayl, J. Swift, D. Aspnes, and J. Kane. "Static Phenomena near Critical Points: Theory and Experiment." *Rev. Mod. Phys.* 39 (1967): 395–431.

[35] Kaporis, A. C., L. M. Kirousis, and E. G. Lalas. "Selecting Complementary Pairs of Literals." *Electronic Notes in Discrete Mathematics* 16 (2003).

[36] Karp, R. "Reducibility among Combinatorial Problems." In *Complexity of Computer Computations*, ed. R. E. Miller and J. W. Thatcher, 85–103. New York: Plenum Press, 1972.

[37] Kirkpatrick, S., and B. Selman. "Critical Behavior in the Satisfiability of Random Boolean Expressions." *Science* 264 (1994): 1297–1301.

[38] Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. "Optimization by Simulated Annealing." *Science* 220 (1983): 671–680.

[39] Krauth, W., and M. Mézard. "The Cavity Method and the Travelling-Salesman Problem." *Europhys. Lett.* 8 (1989): 213–218.

[40] Martin, O. C., R. Monasson, and R. Zecchina. "Statistical Mechanics Methods and Phase Transitions in Optimization Problems." *Theor. Comp. Sci.* 265 (2001): 3–67.

[41] Mézard, M., and G. Parisi. "A Replica Analysis of the Travelling Salesman Problem." *J. Physique* 47 (1986): 1285–1296.

[42] Mézard, M., and G. Parisi. "Replicas and Optimization." *J. Physique Lett.* 46 (1985): L771–L778.

[43] Mézard, M., and R. Zecchina. "Random K-Satisfiability Problem: From an Analytic Solution to an Efficient Algorithm." *Phys. Rev. E* 66 (2002): 056126.

[44] Mézard, M., G. Parisi, and M. A. Virasoro. *Spin Glass Theory and Beyond.* Singapore: World Scientific, 1987.

[45] Molloy, M. "Models for Random Constraint Satisfaction Problems." *SIAM J. Comp.* 32 (2003): 935–949.

[46] Monasson, R., and R. Zecchina. "Statistical Mechanics of the Random K-Satisfiability Model." *Phys. Rev. E* 56 (1997): 1357–1370.

[47] Monasson, R., R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. "Determining Computational Complexity from Characteristic 'Phase Transitions'." *Nature* 400 (1999):133–137.

[48] Mulet, R., A. Pagnani, M. Weigt, and R. Zecchina. "Coloring Random Graphs." *Phys. Rev. Lett.* 89 (2002): 268701.

[49] Percus, A.G., and O.C. Martin. "The Stochastic Traveling Salesman Problem: Finite Size Scaling and the Cavity Prediction." *J. Stat. Phys.* 94 (1999): 739–758.

[50] Ricci-Tersenghi, F., M. Weigt, and R. Zecchina. "Simplest Random K-Satisfiability Problem." *Phys. Rev. E* 63 (2001): 026702.

[51] Sherrington, D., and S. Kirkpatrick. "Solvable Model of a Spin-Glass." *Phys. Rev. Lett.* 35 (1975): 1792–1796.