Physics 114 Statistical Mechanics Spring 2021

Week 9 Conceptual Overview

Concept checklist from Readings:

- Spins in a paramagnetic lattice do not interact with each other but only with an external field, B, so each spin has energy $\epsilon = \pm\mu B$. The sign depends on if it is spinning antiparallel or parallel to an applied field. The macrostate is given by how many spins, n are parallel to the field. We thus have a microcanical ensemble where energy is $E = -n\mu B + (N-n)\mu B$.

- For the microcanonical paramagnet, there are $\Omega = 2^N$ total microstates, each equally likely. The multiplicity of macrostates is $\Omega(n) = \frac{N!}{n!(N-n)!}$. We use familiar arguments to find $\frac{1}{T} = \frac{\partial S(E)}{\partial E}$.

- As usual, canonical statistics may seem easier :-) The partition function yields all other quantities of interest. For a paramagnet, $Z_1 = 2cosh\beta\mu B$ and $Z_N = Z_1^N$. Canonical and microcanonical treatments agree that $E = -N\mu B\ tanh(\beta\mu B)$.

- As with other systems, $F = -kT\ lnZ_N$, $\bar{E} = -\frac{\partial\ lnZ_N}{\partial\beta}$, and $C_B = (\partial E/\partial T)_B$ can be found. The specific heat is, for example, $C_B = kN\ (\beta\mu B)^2\ sech^2(\beta\mu B)$.

- Of special interest for magnetic systems are the expected *magnetization*

  $$M = \mu\ \Sigma_i \bar{s}_i = -\frac{\partial F}{\partial B}$$

  In the paramagnetic case, $M = N\mu\ tanh(\beta\mu B)$.

- Also important is the *magnetic susceptibility* ... which describes how willing the system is to change its magnetization M in response to an external field. The susceptibility is

  $$\chi = \frac{\partial M}{\partial B}$$

  and for a paramagnet, $\chi = N\beta\mu^2\ sech^2(\beta\mu B)$. At high temperatures, we see the *Curie law* where susceptibility drops as $1/T$.

- $\chi = \beta(< M^2 > - < M >^2)$, so *fluctuations* in magnetization determine the susceptibility, just as fluctuations in energy determine the specific heat.

- The intensive quantity of magnetization per spin is $m = M/N$. From here on, we tend to drop the symbol μ ... and treat magnets as if they are just spins of size ± 1.

- The thermodynamics of magnetic systems is not very intuitive for most of us. G&T focus on the H field because it is what we control. The analogy to a fluid is M is like pressure, P and field H is like volume V. (Definitely not intuitive.) If we accept this, we can write $G(T, P, N)$ and $F(T, H, N) = G(T, M, N) - HM$, the usual Legendre transformation. Then
$$M = -(\frac{\partial F}{\partial H})_T \ ; \ \chi = (\frac{\partial M}{\partial H})_T$$

- Paramagnets can be artificially set up at a *negative temperature*. $T < 0$ is actually hotter than $T = \infty$. See the review reading in Schroeder Ch. 3.3 for this interesting situation.

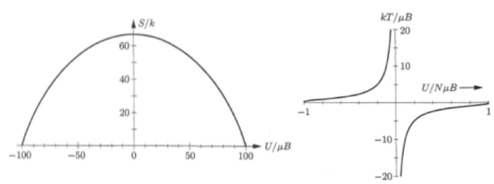

When U > 0, dS/dE<0 so T < 0

- The Ising model for a ferromagnet comprised of N distinguishable, quantized spins. The Hamiltonian for a microstate of the N spins looks like
$$E(\{s_1, ..., s_N\}) = -\Sigma_{i,j \ neighbors} \ \ Js_is_j - \Sigma_i \ s_iH$$
where $s_i = \pm 1$ are the values allowed for any spin. The canonical partition function is thus $Z_N = \Sigma_{microstates} e^{-\beta E(microstate)}$

- Writing Z_N down does not mean solving it in closed form. Now that we have a system with interactions, the partition function does not decompose into a product : $Z_N \neq Z_1^N$ as it did if we had $J = 0$ and were back to solving a paramagnetic system. Before we get into math details, let's get comfortable with the conceptual landscape.

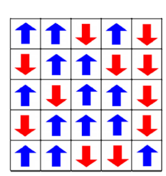

Figure 1: A microstate of ising spins. An external H field would point up or down.

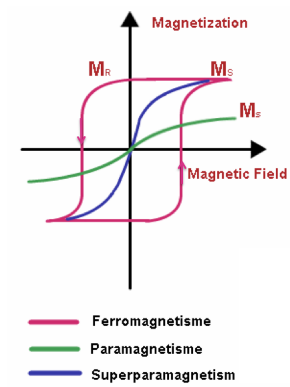

Figure 2: Green is a paramagnet. Magnetization is zero at $H = 0$ for a magnet that doesn't remember its history like a (green) paramagnet or (blue) superparamagnet (not part of our course, but cool.) The pink curve is a ferromagnet, which has had H increased and decreased repeatedly ... it remembers its past and has leftover magnetization, showing "hysteresis". For example, $M(H = 0) \neq 0$. However, if we started an experiment at $H = 0$ with a ferromagnet having $M = 0$, and then we increased (or decreased) H from zero, it would do what the blue curve does. Take home message: Because neighboring spins want to align in a ferromagnet, $M(H)$ has the same general shape, but rises more strongly than in a paramagnet.

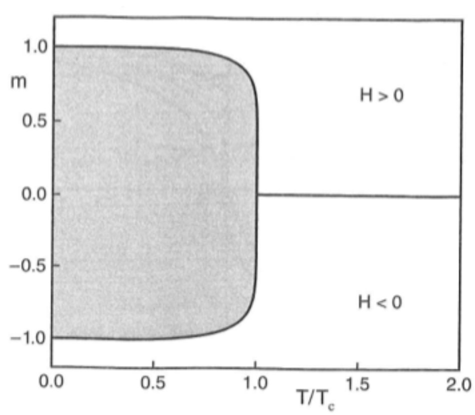

Figure 3: A phase diagram, M(T) for an Ising ferromagnet. First follow the black line ... it has $H = 0^{\pm}$. Above a critical temperature $T > T_c$ there is zero magnetization (entropy wins). For $T < T_c$ there is nonzero magnetization, that increases as T decreases, and spins become successively more aligned (energy wins). The white regions are best explained by looking back at the blue line in Figure 2. There is a nonzero M when $H \neq 0$ at all temperatures. Spins always want to align with H. The shaded regions are forbidden.

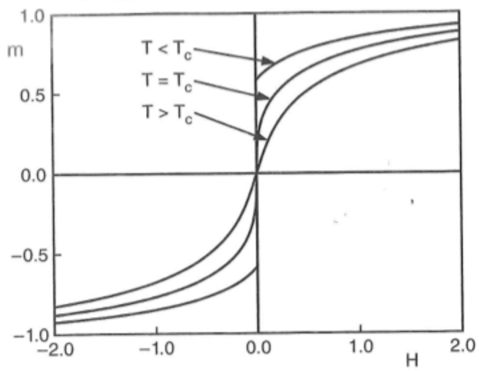

Figure 4: Plot of M(H) showing the difference between passing through $H = 0$ if the Ising ferromagnet is above, at, or below the critical temperature T_c. For $T \geq T_c$, M is continuous. But for $T < T_c$, as soon as H rises from zero, M jumps up to a finite value. (This is one case that isn't shown in Figure 2.) Another way to look at this jump is to look at Figure 3. The jump involves crossing the "forbidden" shaded region when going between positive and negative H values.

- One can make small models of interacting spins. Example 28.9 in B&B uses a 4x4 lattice and shows how to count states, leading to Z and $< E >$. There are some good lessons here: the ground state is *degenerate*; and there's a "crossover" from ordered to disordered macrostate as a function of kT/J.

- The simplest analytically-solvable case in the thermodynamic limit (i.e. $N \to \infty$) is a 1d Ising model with $H = 0$. It is tractable in a couple of ways. G&T section 5.5 talk about solving this *Ising chain* by directly counting states. They find find $Z_N = 2(2cosh(\beta J))^{N-1}$. If we close the chain so that the N^{th} spin interacts with the first, there is a small difference which is immaterial in the $N \to \infty$ limit: $Z_N = (2cosh(\beta J))^N$. From Z_N of course, we find energy, free energy, magnetization, specific heat, and susceptibility.

- Schroeder Ch. 8.2 points out that this Z_N for the ferromagnet has exactly the same mathematical form as Z_N for the paramagnet, if one replaces J with μB. This is only true in 1d.

- χ diverges as $T \to 0$. A phase transition!? It is a matter of definition ... since it doesn't occur at a nonzero temperature. Both B&B and G&T make an argument about "domain walls", which in 1d is just the place where a row of spins changes alignment. It costs a bit of energy and a lot of entropy in 1d, which supports the idea that in the $N \to \infty$ limit, the 1d Ising model is paramagnetic for all nonzero T.

- There is one more trick that works for the 1d Ising model for both zero and nonzero H. This is the *transfer matrix* method. We can write $Z_N = tr(\tilde{T}^N) = \lambda_+^N + \lambda_-^N$ where

$$T = \begin{pmatrix} e^{\beta(J+H)} & e^{-\beta J} \\ e^{-\beta J} & e^{\beta(J-H)} \end{pmatrix} \quad \lambda_\pm \textit{ are eigenvalues}$$

This exact solution not only lets us calculate all thermodynamic quantities, but supports the result that the $H = 0$ ferromagnetic transition does not exist save at $T_c = 0$.

- Let's now talk about higher dimensions. For $H = 0$ in 2d, 3d, ... there is definitely an exciting kind of transition... an "order-disorder" transition. T_c is called a *critical point*. Below T_c, the system exhibits *spontaneous magnetization*. Quantities like C and χ are singular, or even divergent with so-called *critical exponents*. E.g. $\chi \propto |T - T_c|^{-\gamma}$.

- The spin-spin correlation function $G(r)$ measures the correlations between spin directions when the spins are separated by distance r. $G(r = 0) = \overline{m^2} - \overline{m}^2 \propto \chi$. For arbitrary r, $G(r)$ usually dies off exponentially, as $e^{-r/\xi(T)}$ where ξ is the *correlation length*. But as one approaches the critical point T_c, $\xi(T) \rightarrow \infty$. At distances less than ξ will die off as $G(r) \propto \frac{1}{r^{d-2+\eta}}$. This slower, algebraic die-off near the phase transition suggests that spins "communicate" over large distances.

- The 2D Ising model has an exact solution thanks to Onsager (and later by Yang). The solution begins with the transfer matrix, but then treats spins using the Pauli spin matrices of quantum mechanics, creation/annihilation operators for spins ... beyond our scope.

- A very accessible technique that yields T_c and critical exponents (though not the correct ones) is mean field theory ... also known as "Weiss molecular field theory". In this theory, we replace $\Sigma_{neighbors\ j}\ s_j$ with $q\ m$. Here, q is the number of neighbors of any spin. Self-consistency in the definition of m leads to

  $$m = tanh(\beta\mu H + \beta q J m)$$

  This in turn leads to (for $H = 0$) $T_c = qJ/k$.

- G&T problem 5.18 shows that one can also write a *mean-field theory expression for the free energy*. This is an even more fundamental thing than the Weiss theory described above. This expression for free energy is called a *Landau theory*. For the ferromagnet:

  $$f(m) = a - Hm + b(1 - \beta qJ)m^2 + cm^4$$

  This free energy can be used to find out which solution of the self-consistent equation for m is the stable one!

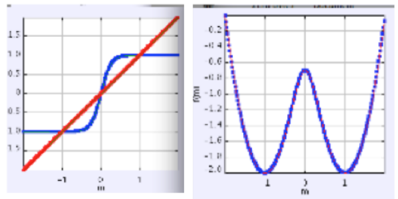

Mean field treatment of m (left) and F (right) for H=0

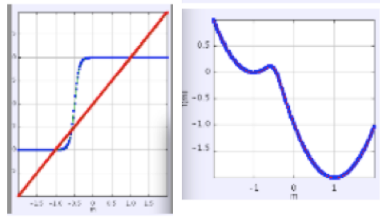

Mean field treatment of m (left) and F (right) for H> 0

- Monte Carlo again?! Yes, because another very accessible technique to study magnets is Monte Carlo (MC) simulation. Please know that MC creates a *trajectory*, a sequence of states of the system. MC sampling involves finding the average of a quantity of interest ... call it G:
$< G >_T = \ (1/T)\ \ \Sigma_j G_j$ where on step j of the trajectory of length T, the value of G is G_j.

- Being even more careful, we mights find partial averages of the quantity e.g. $< G >^{(Li)}$ over the i^{th} set of L steps along the trajectory. Taking many successive sets of L steps allows us to find the average:
$< G >_T = (\ L/T)\ \Sigma_{Li} < G >^{(Li)}$. This gives us a good estimate of the true $< G >$ as the trajectory length T grows. This also lets us estimate uncertainties in our estimate by finding the sum of squares of deviations of the set of $< G >^{(Li)}$. When you click "zero averages" for the i^{th} time in a G&T simulation and then compile data for L more steps, you are finding a partial average in this way.

- The nuts and bolts of MC simulation in its very simplest form involves just trying configurations at random. A slightly less simple form which lets us avoid wasting time in configurations of low probability is *importance sampling* via an *acceptance/rejection* algorithm we've met before: *Metropolis algorithm*. Now, the probability distribution we want to sample is the Maxwell-Boltzmann probability which governs the Ising system: $prob\ \propto\ e^{-\beta E(\{s_1,\ldots,s_N\})}$.

- G&T and Schroeder Ch. 8.2 both give us the rules for sampling with the Metropolis algorithm in order to achieve the canonical distribution. We make a trial move from spin microstate a to state b say, and then accept or reject it so that

$$prob_{a\to b}/prob_{b\to a} = e^{\beta(E_a - E_b)}$$

If E_a and E_b are the Hamiltonians associated with two different spin configurations, a and b respectively, this allows us to simulate the Ising model.

- *Critical slowing down* is an enemy of Metropolis Monte Carlo calculations near a critical point. B&B mentions the Wolff algorithm, and G&T describe it in the description page that shows up when you launch their simulation of the 2d, square lattice, Ising model: If we are interested only in the static properties of the Ising model, the algorithm used to sample the states is irrelevant as long as the transition probability satisfies what is known as detailed balance. The Wolff algorithm flips a cluster of spins rather than a single spin, and is an example of a global algorithm. The utility of the Wolff algorithm is that it allows us to sample states efficiently near the critical temperature; that is, it does not suffer as much from critical slowing down.