Let ΔE_j be the change in the system's internal energy caused by changing its jth degree of freedom. Then the system is defined to be **non-interacting** if and only if ΔE_j for any j is independent of all other degrees of freedom $k \neq j$.

For our system of N distinguishable spins, the only possible change we can make to a degree of freedom is to negate it, $s_j \to -s_j$, which corresponds to flipping its alignment relative to the external magnetic field. It is easy to check that the change in the internal energy resulting from such a spin flip satisfies our definition of a non-interacting system:

$$E = -H\sum_n s_n = -H\left(s_j + \sum_{k\neq j} s_k\right) \longrightarrow -H\left(-s_j + \sum_{k=j} s_k\right)$$

$$\Delta E_j = 2Hs_j \quad \text{independent of } s_k \text{ for } k\neq j$$

$\to$ non-interacting ✓

Now let's make things more interesting by considering a different spin system that also includes a simple two-spin contribution to the internal energy:

$$E_i = -\sum_{(jk)} s_j s_k - H\sum_{n=1}^{N} s_n. \tag{111}$$

The first sum runs over all pairs of nearest-neighbour spins in the lattice, denoted (jk). What is the change in energy ΔE_j from Eq. 111 upon negating $s_j \to -s_j$? Does this indicate an interacting or non-interacting system?

$$E = -s_j\sum_{k\in(jk)} s_k - \sum_{(mk)\not\ni j} s_m s_k - H\left(s_j + \sum_{k\neq j} s_k\right)$$

$$\to s_j\sum_{k\in(jk)} s_k - \sum_{(mk)\not\ni j} s_m s_k - H\left(-s_j + \sum_{k\neq j} s_k\right)$$

$$\Delta E_j = 2s_j\left(H + \sum_{k\in(jk)} s_k\right)$$

Depends on s_k with $k\neq j$

$\to$ interacting

25 Apr.

26 Apr.

The pictures on the next page illustrate nearest-neighbour pairs for simple cubic lattices in $d = 2$ and 3 dimensions, while also introducing some additional lattice terminology. Instead of drawing up- and down-pointing arrows, these pictures identify the spins with *sites* in the lattice represented as points (or larger filled circles). In simple cubic lattices, all sites are positioned in a regular grid, separated by a constant distance along each basis vector. We can also draw *links* as solid lines connecting these nearest-neighbour sites, with each link corresponding to a term in $\sum_{(jk)}$. The picture of a two-dimensional lattice on the left

highlights the four links (with red hatch marks) that correspond to the four nearest neighbours (circled in red) of a particular site (circled in blue). For $d \geq 2$, an elementary unit of surface area is called a *plaquette*, while for $d \geq 3$ the elementary unit of volume is called a *cube*.

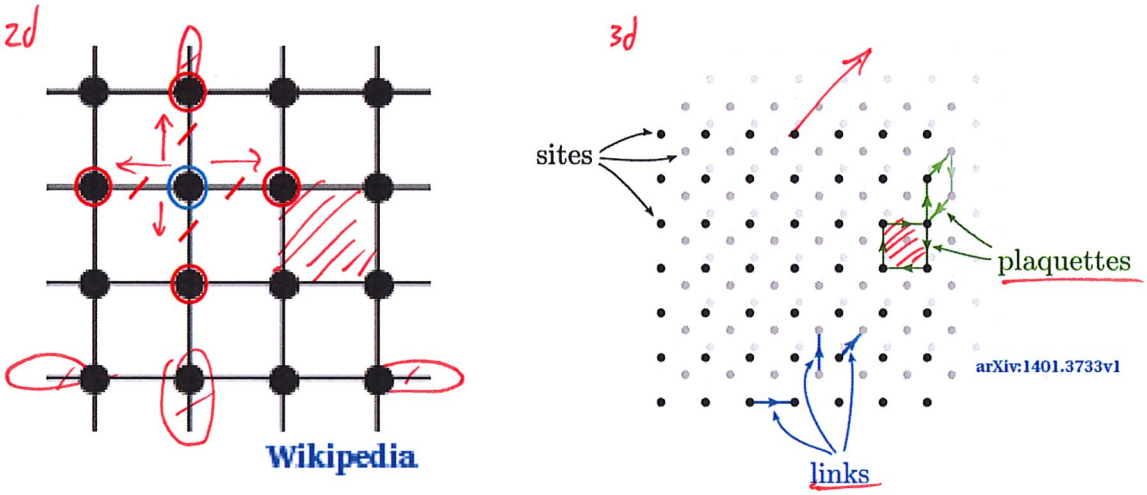

Computing the energy in Eq. 111 requires determining all of the nearest-neighbour pairs to be summed in the first term, which is equivalent to all of the links in the lattice, $\ell = (jk)$. When considering a finite lattice, this task is complicated by the need to consider the edges of the lattice. We can avoid this complication by imposing **periodic boundary conditions**, which remove these edges by adding links between each site on the left edge of the lattice and the corresponding site on the right edge, and similarly in all other dimensions. This is illustrated below for the simple case of the one-dimensional lattice, drawn as a circle to emphasize that all N sites remain separated by a constant distance. In higher dimensions, periodic boundary conditions produce flat (zero-curvature) d-dimensional tori that preserve the simple cubic lattice structure.

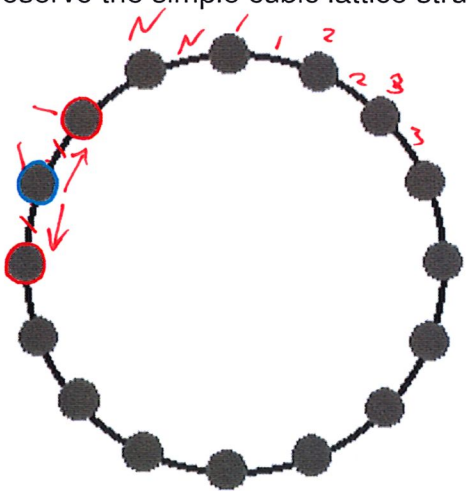

With periodic boundary conditions, we can easily see that the N-site one-dimensional lattice drawn above has N links. Each site has two links connecting it to its two nearest neighbours, and each of those links is shared between two sites, so that $\#\ell = 2N/2 = N$. Looking back to the two-dimensional lattice drawn

farther above, the four links per site produce $\#\ell = 4N/2 = 2N$. How many terms are there in the sum $\sum_{(jk)}$ in Eq. 111 for N-site lattices with periodic boundary conditions in d dimensions?

2d links per site
2 sites sharing each link
$\#\ell = \frac{2d \cdot N}{2} = d \cdot N$ links in total

The energy of interacting spins given by Eq. 111, with a lattice structure defined to specify which spins form the nearest-neighbour pairs (jk), defines a famous system known as the d-dimensional **Ising model**. Since the 1960s, the Ising model has been the basis of thousands of scientific studies analyzing everything from ferromagnetism to neural networks to urban segregation.[13] The model was proposed in 1920 by Wilhelm Lenz, whose PhD student Ernst Ising solved the one-dimensional system as a research project in 1924. Exactly solving the two-dimensional case (with $H = 0$) took another twenty years, culminating in renowned work by Lars Onsager in 1944. The three-dimensional Ising model remains an open mathematical question, with no known exact solution.

In this context, 'solving' the Ising model means deriving a closed-form expression for its canonical partition function,

$$Z(\beta, N, H) = \sum_{\{s_n\}} \exp\left[-\beta E(s_n)\right] = \sum_{\{s_n\}} \exp\left[\beta \sum_{(jk)} s_j s_k + \beta H \sum_{n} s_n\right].$$

($d \cdot N$ over the first sum, N over the second)

As in Section 3.4, the partition function sums over all possible spin configurations $\{s_n\}$, which amounts to a sum of 2^N exponential factors for N spins, with $\mathcal{O}(N)$ terms within each exponential. Now that the system is interacting, the partition function no longer factorizes into the N identical cosh factors of Eq. 41, making it extremely difficult to evaluate. This is why there is no known exact solution to the three-dimensional Ising model, and it also makes 'brute-force' numerical computations impractical. Even for a system of $N = 1023$ spins, twenty orders of magnitude smaller than our typical $N \sim 10^{23}$, there are roughly $2^{1023} \sim 10^{310}$ terms in the partition function, far beyond the capabilities of existing or foreseeable supercomputers.

9.2 Ising model phases and phase transition

Despite the insolubility of the Ising model in an arbitrary number of d dimensions, we can still make robust predictions for its large-scale behaviour by considering the simplified limits of high and low temperature, much as we did for non-interacting spin systems in Section 3.4. We can also simplify the system by

[13] For a brief discussion, see Charlie Wood, "The Cartoon Picture of Magnets That Has Transformed Science", *Quanta Magazine*, 2020.

setting $H = 0$ in this section, and considering just

$$E_i = -\sum_{(jk)} s_j s_k \qquad\qquad Z(\beta, N) = \sum_{\{s_n\}} \exp\left[\beta \sum_{(jk)} s_j s_k\right]. \tag{112}$$

We will see that the behaviour of this zero-field Ising model is qualitatively different at high temperatures compared to low temperatures. In other words, the system exhibits at least two distinct phases for different temperatures. This is a necessary but not sufficient condition for there to be a true phase transition — it leaves open the possibility of a gradual *crossover* between these two phases, as opposed to a rapid transition. In this section we will use the Ising model to more rigorously define what exactly constitutes a phase transition, and how this can be distinguished from a crossover.

First, though, let's consider the **high-temperature** limit $\beta \to 0$, where the Ising model partition function becomes extremely simple:

$$\lim_{\beta\to 0} Z(\beta, N) = \sum_{\{s_n\}} \exp(0) = 2^N$$

In this limit, all 2^N spin configurations are adopted with the same probability $p_i = 1/2^N$, regardless of their internal energy from Eq. 112. In effect, that energy has become negligible compared to the temperature.

Rather than computing the expectation value of that internal energy, there is a simpler observable that we can consider to characterize this high-temperature phase. This is the **magnetization** $M = n_+ - n_-$, retaining our definition of $n_\pm$ as the number of spins with value ± 1, even without an external field for them to align with or align against. It is convenient to normalize the magnetization by the number of spins,

$$m \equiv \frac{M}{N} = \frac{n_+ - n_-}{n_+ + n_-}, \tag{113}$$

so that $-1 \le m \le 1$ for any value of N. In addition, without an external field to distinguish between ± 1 spins, it is also convenient to consider the absolute magnitude $0 \le |m| \le 1$.

Our task is now to determine the expectation value of the magnetization at high temperatures. Above we found that all spin configurations are equally probable in this regime, so $\langle |m| \rangle$ will be determined by how many of these equally-probably micro-states have a particular magnetization. For example, there are only two micro-states with $|m| = 1$, corresponding to $(n_+, n_-) = (N, 0)$ and $(0, N)$. In general, just as we saw in Eq. 23, there are

$$\binom{N}{n_+} = \binom{N}{n_-} = \frac{N!}{n_+!\, n_-!}$$

$$\langle m \rangle = \sum_i m_i p_i = \frac{1}{2^N} \sum_i m_i$$

equally probable micro-states with a given $n_+ = N - n_-$. For large $N \gg 1$ this binomial coefficient has a factorially narrow peak around

$$n_+ = n_- = \frac{1}{2}N \qquad \longrightarrow \qquad |m| = 0.$$

This characterizes a **disordered phase** with similar numbers of up- and down-pointing spins producing a small magnetization. In the *thermodynamic limit* $N \to \infty$, the expectation value of the magnetization in the disordered phase vanishes exactly, $\langle |m| \rangle \to 0$.

We next need to determine $\langle |m| \rangle$ in the **low-temperature** limit $\beta \to \infty$. In this regime, as we saw in Section 3.4.1, the Boltzmann factor $\exp\left[\beta \sum_{(jk)} s_j s_k\right]$ makes it exponentially more likely for the system to adopt micro-states with lower energies. In particular, we can expect the ground state to dominate the expectation value of the magnetization, $\langle |m| \rangle$, up to exponentially suppressed corrections from higher-energy excited states. With $H = 0$, the Ising model has two degenerate ground states corresponding to the two ways all the spins can be aligned with each other: $(n_+, n_-) = (N, 0)$ and $(0, N)$. What is the ground-state energy of the N-site Ising model in d dimensions?

$$E_0 = -\sum_{(jk)} s_j s_k = -\sum_{(jk)} (\pm 1)^2 = -d \cdot N$$

-1 from every link

As mentioned above, both of these degenerate ground states have the maximal magnetization $|m| = 1$. Let's check what effect the first excited state would have on the overall magnetization of the system. The first excited energy level is obtained by flipping a single spin — negating its value. Starting from the two degenerate ground states, this produces all possible micro-states with $(n_+, n_-) = (N-1, 1)$ and $(1, N-1)$. Because any one of the N spins in the lattice could be flipped, the degeneracy of this first excited energy level grows with N:

$$\binom{N}{1} + \binom{N}{N-1} = 2N.$$

At the same time, as N increases the magnetization of each of these micro-states gets closer to that of the ground state,

$$|m| = \frac{N-2}{N} = 1 - \frac{2}{N}.$$

The key factor is the probability for the system to be in one of these micro-states, which depends on the value of the energy E_1 for the first excited energy level. What is this E_1 for the N-site Ising model in d dimensions?

$E_1 = 2d - (dN - 2d) = -(dN - 4d)$

+1 From $2d$ links connected to flipped spin

-1 From all other $dN - 2d$ links

Let's bring everything together by computing the relative probability for the d-dimensional Ising model to be in its ground state with $|m| = 1$ compared to its first excited state with $|m| = 1 - \frac{2}{N}$. This probability is the product of the degeneracy of each energy level times the Boltzmann factor that governs the probability of the system adopting any of these degenerate micro-states:

$$\frac{P(E_0)}{P(E_1)} = \frac{2 \cdot \exp\left[\beta d \cdot N\right]}{2N \cdot \exp\left[\beta\left(d \cdot N - 4d\right)\right]} = \frac{\exp\left[4\beta d\right]}{N}.$$

For any fixed N, a sufficiently low temperature will cause the ground state to dominate, with exponentially suppressed contributions from higher energy levels, just as we previously found for simpler non-interacting systems. This characterizes an **ordered phase** with essentially all spins aligned in the same direction, producing a large expectation value for the magnetization, $\langle |m| \rangle \to 1$.

26 Apr.

We have now seen how the behaviour of the magnetization $\langle |m| \rangle$ distinguishes the high- and low-temperature phases of the zero-field Ising model in d dimensions. In the high-temperature disordered phase, the magnetization is small and $\langle |m| \rangle \to 0$ in the thermodynamic limit $N \to \infty$. In the low-temperature ordered phase, the magnetization is large and $\langle |m| \rangle \to 1$ as $T \to 0$.

This contrast between ordered and disordered phases is typical behaviour for interacting statistical systems. These two phases are distinguished by an **order parameter** — an observable (related to a derivative of the free energy) that is zero in the disordered phase but non-zero in the ordered phase.[14] The magnetization is the order parameter for the Ising model, which we will connect to the free energy in the next section. Note that the order parameter need not reach its maximum value in the ordered phase — in the case of the Ising model, we don't need complete domination by the fully ordered ground state. So long as there is a tendency towards order, mathematically defined by a non-zero order

[14] There are atypical (but interesting and important) *topological phase transitions* that are not characterized by such an order parameter. The most famous example is the BKT phase transition named after Vadim Berezinskii, J. Michael Kosterlitz and David Thouless, which was awarded the 2016 Nobel Prize in Physics. It is also possible for a single system to have multiple distinct phase transitions, each characterized by a different order parameter.