

Quantum Phase Transitions

Subir Sachdev

Department of Physics, Yale University, New Haven, CT 06520 USA

October 4, 2004

1 Introduction

We all observe phase transitions in our daily lives, with hardly a second thought. When we boil water for a cup of tea, we observe that the water is quiescent until it reaches a certain temperature (100°C), and then bubbles appear vigorously until all the water has turned to steam. Or after an overnight snowfall, we have watched the snow melt away when temperatures rise during the day. The more adventurous among us may have heated an iron magnet above about 760°C and noted the disappearance of its magnetism.

Familiar and ubiquitous as these and many related phenomena are, a little reflection shows that they are quite mysterious and not easy to understand: indeed, the outlines of a theory did not emerge until the middle of the 20th century, and although much has been understood since then, active research continues. Ice and water both consist of molecules of H_2O , and we can look up all the physical parameters of a single molecule, and of the interaction between a pair of molecules, in standard reference texts. However, no detailed study of this information prepares us for the dramatic change that occurs at 0°C . Below 0°C , the H_2O molecules of ice are arranged in a regular crystalline lattice, and each H_2O molecule hardly strays from its own lattice site. Above 0°C , we obtain liquid water, in which all the molecules are moving freely throughout the liquid container at high speeds. Why do 10^{23} H_2O molecules co-operatively “decide” to become mobile at a certain temperature, leading to the phase transition from ice to water?

We understand these phase transitions by the delicate balance between the diverging interests of the energy, E , and the entropy, S . The principles of

thermodynamics tell us that systems in thermal equilibrium seek to minimize their free energy $F = E - TS$, where T is the temperature measured in the absolute (Kelvin) scale. The energy is determined by the interactions between the H_2O molecules, and this is minimized in the crystalline structure of ice. The entropy, as explained by Boltzmann, is a measure of the degree of ‘randomness’ in a phase; more precisely, it is proportional to the logarithm of the number of microscopic arrangements of H_2O molecules available at a given total energy and volume—the entropy is clearly larger in the liquid water phase. It is now easy to see that at low T , $F = E - TS$ will be smaller in the ice phase, while at higher T the contributions of the entropy to F become more important, and the free energy of the liquid water phase is lower. The free energies of ice and liquid water cross each other at 0°C , accounting for the phase transition at this temperature.

So far we have described what are more completely referred to as *thermal* phase transitions, which are caused by the increasing importance of entropy in determining the phase of a system with rising temperatures. Let us now turn to the central topic of this article, *quantum* phase transitions. Such transitions occur only at the absolute zero of temperature, $T = 0^\circ\text{K}$, where thermodynamics tells us that the system should be in its lowest energy state (also called the ‘ground state’). In the simple classical model of H_2O discussed so far, we anticipate that there is some perfect crystalline arrangement which minimizes the inter-molecular interaction energy, and at $T = 0^\circ\text{K}$ all the H_2O molecules reside at rest on the sites of this lattice. This appears to be a unique quiescent state, and so where then is the possibility of a phase transition? The problem with this model is that it is incompatible with laws of quantum mechanics, which govern the behavior of all the microscopic constituents of matter. In particular, Heisenberg’s uncertainty principle tells us that it is not possible to simultaneously specify both the position and momentum of each molecule. If we do place the molecules precisely at the sites of a perfect crystalline lattice (thus determining their positions), then the momenta of the molecules are completely uncertain—they cannot be at rest, and their kinetic energy will add an unacceptable cost to the energy of this putative ground state. So determining the state of H_2O at $T = 0^\circ\text{K}$ becomes a delicate matter of optimizing the potential and kinetic energies, while maintaining consistency with Heisenberg’s uncertainty principle. As in our earlier discussion of thermal phase transitions, this delicate balance implies that, at least in principle, more than one phase is possible even at $T = 0^\circ\text{K}$. These phases have distinct macroscopic properties, while contain-

ing the same microscopic constituents; they are separated by quantum phase transitions. The key difference from thermal phase transitions is that the fluctuations required by the maximization of entropy at $T > 0^\circ\text{K}$ have now been replaced by quantum fluctuations demanded by the uncertainty principle. In addition to the familiar ice phase, H_2O exhibits numerous other phases at $T = 0^\circ\text{K}$ under strong applied pressure, each with a distinct crystal structure and separated from the other phases by quantum phase transitions. We will not describe these transitions here, because they are quite complicated, but focus on simpler examples of quantum-uncertainty induced transitions.

Our discussion so far would seem to indicate that a quantum phase transition is a theoretical abstraction, which couldn't possibly be of any relevance to an understanding of materials of technological importance. After all, it is impossible to cool any material down to 0°K in the laboratory, and it takes heroic effort even to get close to it. As we will discuss below, it has become clear in the last decade that these conclusions are quite wrong. Quantum transitions at $T = 0^\circ\text{K}$ leave characteristic fingerprints in the physical properties of materials at $T > 0^\circ\text{K}$, and these fingerprints are often clearly visible even at room temperature. A complete understanding of the physical properties of many materials emerges only upon considering the relationship of their $T > 0^\circ\text{K}$ phases to the distinct ground states and quantum phase transitions at $T = 0^\circ\text{K}$.

2 Interacting qubits in the laboratory

We begin our discussion of quantum phase transitions with a simple example. Rather than tackling the full complexity of atomic/molecular potential and kinetic energies, we consider the simplest possible quantum mechanical system, a set of interacting *qubits*. All computers store information as strings of (classical) 'bits', which are abstract representations of the two possible states of a classical electrical circuit, usually denoted 0 and 1. Similarly, a qubit is the simplest quantum degree of freedom, which can be in only one of two quantum states, which we denote $|\uparrow\rangle$ and $|\downarrow\rangle$. There has been much recent discussion, much of it quite abstract and theoretical, on using coupled qubits to perform computations; here we show that when many qubits are actually coupled with each other in the laboratory, they undergo a quantum phase transition.

A common physical realization of a qubit is provided by the spin of an

electron: each electron spins on its own axis, and its angular momentum about any fixed axis (say z) is only allowed to take the values $+\hbar/2$ (in the $|\uparrow\rangle$ state) or $-\hbar/2$ (in the $|\downarrow\rangle$ state); here \hbar is a fundamental constant of nature, Planck’s constant. Associated with this angular momentum is a magnetic moment, M_z , which produces a magnetic field that can be measured in the laboratory. Most materials have equal numbers of electrons with spin $+\hbar/2$ and $-\hbar/2$, and so the net magnetic moment is zero. However, certain materials (like iron) have a net excess of electrons with spin $+\hbar/2$ over $-\hbar/2$ (say), endowing them with a macroscopic magnetic moment—this is the explanation for its magnetic properties. Here, we are interested in separating the magnetic qubits on distinct lattice sites and controlling the interactions between them, to allow us to tune the qubits across a quantum phase transitions. Iron is not a suitable candidate for this because it is a metal, and its electrons move freely throughout the entire crystal lattice. We need to look at insulators, in which the individual ions have a net magnetic moment. Reasoning in this manner, in 1991 Thomas Rosenbaum at the University of Chicago, Gabriel Aeppli at NEC Research, and their collaborators undertook a series of experiments on the insulator LiHoF_4 , the simplest of which we shall describe below (see Fig 1). In this insulator, the Li^+ and F^- ions have equal numbers of up and down spin electrons and are non-magnetic. However, each Ho^{3+} ion has a net magnetic moment. This magnetic moment has contribution from both the electron spin on its own axis, and from the orbital motion of each electron around the Ho nucleus, and these motions are strongly coupled to each other via the ‘spin-orbit’ interaction. We do not need to enter these complexities here because, after the dust settles, the Ho^{3+} ion on site j has only two possible magnetic states, which we denote $|\uparrow\rangle_j$ and $|\downarrow\rangle_j$. Moreover, the crystalline structure also chooses a natural quantization axis, and these states have a magnetic moment, M_z , along the preferred crystalline axis. So each Ho^{3+} ion is a perfect physical realization of a qubit.

Before we look at the problem of many coupled qubits, let us dwell a bit more on the physics of an isolated qubit on a single Ho^{3+} ion. One of the fundamental principles of quantum mechanics is the *principle of superposition* of quantum states. For a qubit, this principle implies that its state can not only be $|\uparrow\rangle$ or $|\downarrow\rangle$, but also an arbitrary superposition

$$|\psi\rangle = \alpha |\uparrow\rangle + \beta |\downarrow\rangle \tag{1}$$

where α, β are arbitrary complex numbers; the normalization of the state $|\psi\rangle$ requires that $|\alpha|^2 + |\beta|^2 = 1$. The state $|\psi\rangle$ has no classical analog, and

so it is not possible to provide a picture consistent with our naive classical intuition. An observation of the qubit in state $|\psi\rangle$ will show that it is in state $|\uparrow\rangle$ with probability $|\alpha|^2$ and in state $|\downarrow\rangle$ with probability $|\beta|^2$. But this does not imply that the state $|\psi\rangle$ is a random statistical mixture of the $|\uparrow\rangle$ and the $|\downarrow\rangle$ states—for some measurements there is an additional quantum interference term which characterize $|\psi\rangle$ as a true superposition. Crudely speaking the qubit is both up and down “at the same time”. Of particular interest is the state of the qubit with $\alpha = \beta = 1/\sqrt{2}$, which we denote as

$$|\rightarrow\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle + |\downarrow\rangle). \quad (2)$$

This notation is suggestive: the $|\rightarrow\rangle$ state has the remarkable property that it has a magnetic moment, M_x pointing in the $+x$ direction. So a superposition of a state with magnetic moment up ($|\uparrow\rangle$) and down ($|\downarrow\rangle$) has a magnetic moment pointing right! This is a deep and special property of quantum mechanics, and is ultimately closely linked to Heisenberg’s uncertainty principle: states with definite M_z ($|\uparrow\rangle$ or $|\downarrow\rangle$) have uncertain M_x , and conversely a state with uncertain M_z ($|\rightarrow\rangle$) has a definite M_x . Similar to (2), we will also need the state

$$|\leftarrow\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle - |\downarrow\rangle) \quad (3)$$

which has a magnetic moment pointing in the $-x$ direction. One of the most important properties of a qubit, which has no analog in classical bits, is that the $|\rightarrow\rangle$ and $|\leftarrow\rangle$ states can also be used as a basis for expressing the state of a qubit, and this basis is as legitimate as the $|\uparrow\rangle$, $|\downarrow\rangle$ states we have used so far. The reader can easily see from (1,2,3) that the state $|\psi\rangle$ can also be written as a superposition of the $|\rightarrow\rangle$ and $|\leftarrow\rangle$ states:

$$|\psi\rangle = \alpha' |\rightarrow\rangle + \beta' |\leftarrow\rangle \quad (4)$$

with $\alpha' = (\alpha + \beta)/\sqrt{2}$ and $\beta' = (\alpha - \beta)/\sqrt{2}$. So we can view $|\psi\rangle$ as a state ‘fluctuating’ between up and down states (as in (1)), or as a state ‘fluctuating’ between right and left states (as in (4))—a remarkable fact completely at odds with our classical intuition, but a fundamental property of a qubit.

We are now ready to couple the qubits, residing on the Ho atoms in LiHoF_4 , to each other. This coupling is described by the Hamiltonian, H , which is a quantum-mechanical representation of the energy. This Hamiltonian will have two terms, which are the analogs of the kinetic and potential energies of the water molecules we discussed earlier. The quantum

phase transition occurs because of a delicate interplay between these energies, and this has a crude parallel to the interplay between E and S near a thermal phase transition. We reiterate, however, that our discussion here of a quantum phase transition is at $T = 0^\circ\text{K}$, and so we are always seeking the quantum-mechanical state with the lowest total energy. We schematically represent the Hamiltonian as

$$H = H_z + gH_x \tag{5}$$

where $H_{z,x}$ are the two announced components of the Hamiltonian, and g is a dimensionless parameter we will tune to move the qubits across the quantum phase transition—the role of g here will parallel that of T for a thermal phase transition. In the language of the $|\uparrow\rangle_j, |\downarrow\rangle_j$ representation of the qubits, the H_z term is a ‘potential’ energy *i.e.* it determines the optimum configuration of the M_z magnetic moments which will minimize the energy. The term H_x is a ‘kinetic’ energy, but unlike the case of the water molecules, it has a very simple form because there are only two possible states of each qubit. So the only ‘motion’ possible is a flipping of the qubit between the up and down states: it is precisely this up-down flipping, or ‘quantum tunnelling’ which is induced by H_x . An interesting and fundamental property of $H_{z,x}$, which will become clear from our discussion below, is that roles of kinetic and potential energies are reversed in the $|\rightarrow\rangle_j, |\leftarrow\rangle_j$ basis of the qubits. In this case, the H_x will be a ‘potential’ energy, while the H_z term will induce quantum tunnelling between the left and right states, and is thus a ‘kinetic’ energy. So quantum systems have this peculiar property of looking quite different depending upon the choice of observables, but there remains a unique underlying state which has these different physical interpretations.

Let us first discuss the meaning of H_z further—this will give us a good picture of the ground state for $g \ll 1$. In LiHoF_4 , H_z arises from the ‘magnetic dipole’ interaction. Each M_z magnetic moment produces a dipolar magnetic field (much like the familiar magnetic field patterns around a bar magnet), and this field will tend to align the other magnetic moment parallel to itself. For the LiHoF_4 , these dipolar interactions are optimized for a ‘ferromagnetic’ arrangement of the M_z moments. In other words, the ground state is

$$|\uparrow\rangle = \dots |\uparrow\rangle_{j_1} |\uparrow\rangle_{j_2} |\uparrow\rangle_{j_3} |\uparrow\rangle_{j_4} |\uparrow\rangle_{j_5} \dots \tag{6}$$

where the labels $j_{1\dots}$ extend over all the $\sim 10^{23}$ qubits in the lattice. There is nothing in the Hamiltonian or the crystal structure which distinguishes up

from down, so another ground state, with the same energy is

$$|\Downarrow\rangle = \dots |\downarrow\rangle_{j_1} |\downarrow\rangle_{j_2} |\downarrow\rangle_{j_3} |\downarrow\rangle_{j_4} |\downarrow\rangle_{j_5} \dots \quad (7)$$

The equivalence between (6) and (7) is therefore related to a *symmetry* between the up and down orientations of the magnetic moments. This last statement needs to be qualified somewhat in any realistic system. While the perfect and infinite crystal of LiHoF₄ does indeed have an unblemished up-down symmetry, any realistic crystal will always have a slight preference for up or down induced by imperfections and boundaries. Any slight preference is sufficient to break the symmetry, and so let us assume that the ground state has been chosen in this manner to be $|\Uparrow\rangle$. This is a simple example of a common phenomenon in physics: a *spontaneously broken symmetry*. The ferromagnetic arrangement of the magnetic in such a state will lead to a magnetic field much like that produced by everyday permanent magnets, and this is easy to detect in the laboratory.

The alert reader will have noticed by now that our description above of the states $|\Uparrow\rangle$ and $|\Downarrow\rangle$ surely cannot be the whole story: we have localized each qubit in the up direction, and this must lead to problems with Heisenberg's uncertainty principle. Won't each qubit tunnel to the down state occasionally, just to balance out the uncertainties in M_z and M_x ? The contribution H_x in (5) is one term which performs this tunnelling. For $g \ll 1$ it can be shown that our discussion of the $|\Uparrow\rangle$ and $|\Downarrow\rangle$ states is essentially correct, and the ferromagnetic phase is not disrupted by quantum tunnelling. However, the story is very different for $g \gg 1$. Now the state of the qubits is determined entirely by the optimization of this tunnelling, and this happens when the qubit is equally likely to be in the up or down states. We choose an arbitrary phase in our definition of the qubit states to have the tunnelling in H_x prefer the state $|\rightarrow\rangle_j$ for all j . Thus a large tunnelling makes all the qubits point to the 'right'. This result also suggests a simple way in which the value of g can be tuned in the laboratory: simply apply a magnetic field along the $+x$ or 'transverse' direction. This transverse field will enhance the tunnelling between the $|\uparrow\rangle_j$ and $|\downarrow\rangle_j$ states of the qubit, and is a powerful tool for 'tuning' the strength of quantum fluctuations. In a large transverse field, $g \rightarrow \infty$, and then the ground state of H is clearly

$$|\Rightarrow\rangle = \dots |\rightarrow\rangle_{j_1} |\rightarrow\rangle_{j_2} |\rightarrow\rangle_{j_3} |\rightarrow\rangle_{j_4} |\rightarrow\rangle_{j_5} \dots \quad (8)$$

Note that this state is very different from the $|\Uparrow\rangle$ and $|\Downarrow\rangle$ and while the relationship (2) holds for a single qubit, the analogous relationship for many

qubits does *not* hold *i.e.*

$$|\Rightarrow\rangle \neq \frac{1}{\sqrt{2}} (|\uparrow\rangle + |\downarrow\rangle); \quad (9)$$

The correct expression for $|\Rightarrow\rangle$ is derived by inserting (2) into (8) for each site j and then expanding out the products: one finds a very large number of terms with up and down oriented qubits, and only two of these terms appear on the right hand side of (9). The inequality (9), and the far more complicated behavior of the many qubit system, is what allows the appearance of a non-trivial quantum phase transition.

In passing, we note that the state on the right-hand-side of (9) is often referred to as ‘‘Schrödinger’s cat’’. It is a quantum superposition of two macroscopically distinct states (the $|\uparrow\rangle$ and $|\downarrow\rangle$ states), much like the quantum superposition between a dead cat and a live cat that Schrödinger speculated about. In practice, such ‘‘cat states’’ are exceedingly hard to create, because (as we have discussed above) even a tiny external perturbation will ‘collapse’ the system wavefunction into either $|\uparrow\rangle$ or $|\downarrow\rangle$.

We have now assembled all the ingredients necessary to understand the quantum phase transition of the many qubit system as a function of increasing g . For small g , as we discussed earlier, the qubits are in the state $|\uparrow\rangle$ —the most important property of this state is that it has a non-zero value of the average magnetic moment $\langle M_z \rangle$. We expect $\langle M_z \rangle$ to evolve smoothly as the value of g is increased. In contrast, at very large g , we have the state $|\Rightarrow\rangle$, in which $\langle M_z \rangle$ is strictly zero, and we expect this to be true for a range of large g . Now imagine the functional dependence of $\langle M_z \rangle$ on g : it is impossible to connect the two limiting regimes with a smooth function. There must be a singularity at at least one value $g = g_c$ as shown in Fig 2, where $\langle M_z \rangle$ first vanishes. This is the location of the quantum phase transition—with increasing g , this is the point at which the ferromagnetic moment vanishes, the up-down symmetry of the qubit system is restored, and we reach a ‘paramagnetic’ state. In a very real sense, the complicated many qubit ground state for $g < g_c$ is similar, and smoothly connected, to the state $|\uparrow\rangle$ for $g < g_c$, while for $g > g_c$ a corresponding satisfactory model is provided by $|\Rightarrow\rangle$. Only at $g = g_c$ do we obtain a truly different ‘critical’ state, in a very complex and ‘*entangled*’ qubit arrangement we will not describe in any detail here.

So far we have described the physics of a quantum phase transition at $T = 0$, but this is not of direct relevance to any practical experiment. It is essential to describe the physics at $T > 0$, when the qubits will no longer

reside in their lowest energy state. Describing the structure of the very large number of higher energy (or ‘excited’) states would seem to be a hopelessly complicated task—there are an exponentially larger number of ways in which the qubits can rearrange themselves. Fortunately, over most of parameter space, a powerful conceptual tool of quantum theory provides a simple and intuitive description: the *quasiparticle*. A quasiparticle appears in any experiment just like an particle: it is a point-like object which moves while obeying Newton’s Laws (more precisely, their quantum generalizations). However, rather than being a fundamental degree of freedom in its own right, a quasiparticle emerges from the collective behavior of many strongly coupled quantum degrees of freedom (hence the ‘quasi-’): it is a ‘lump’ of excited qubits which, when viewed from afar, moves around just like a particle. Moreover, the spectrum of excited states of the qubits can be usefully decomposed into states describing multiple quasiparticles moving around and colliding with each other.

A simple description of the quasiparticle states is also possible. Consider first the limit $g \gg 1$, where the ground state is $|\Rightarrow\rangle$ as in (8). To create an excited state, and we must flip qubits from right to left, and because it costs a large energy ($\sim g$) to flip a qubit, let us flip just one at the site j ; this yields the state

$$|j\rangle = \dots |\rightarrow\rangle_{j_1} |\rightarrow\rangle_{j_2} |\leftarrow\rangle_j |\rightarrow\rangle_{j_3} |\rightarrow\rangle_{j_4} |\rightarrow\rangle_{j_5} \dots \quad (10)$$

The left qubit is just a stationary object at the site j above, but corrections from H_z endow it with all the characteristics of a particle: it becomes mobile and acquires an energy which depends its velocity (its ‘dispersion relation’). We can now also create left-oriented qubits on other sites, and these behave much like a gas of particles: the quasiparticles are relatively independent of each other when they are far apart, and they collide and scatter of each other should their paths intersect. Also, for large enough g , it should be evident that this quasiparticle picture provides a description of all the lower energy excited states above the $|\Rightarrow\rangle$ ground state. A somewhat more subtle fact is that this quasiparticle description works for *all* points, $g > g_c$, on the paramagnetic side of the quantum phase transition. As we lower g , each quasiparticle is no longer a single flipped qubit as in (10), but becomes a more diffuse object localized around the site j ; furthermore, there is a lowering of the upper bound on the energy below which the quasiparticle picture applies. Ultimately, as we reach $g = g_c$ from above, the size of the quasiparticle diverges, and the upper bound on the quasiparticle picture reaches the ground

state energy. Only strictly at $g = g_c$ does the quasiparticle picture fail completely,¹ but quasiparticles can be usefully defined and identified at all $g > g_c$.

Let us now turn to a description of the excited states for $g < g_c$. Here again, we find a quasiparticle description, but in terms of an entirely new type of quasiparticle. The description of this quasiparticle is simplest for the case in which the qubits are aligned along a one-dimensional chain, rather than a three dimensional crystal, and so we restrict our discussion here to this case. In the ground states $|\uparrow\rangle$ or $|\downarrow\rangle$, at $g = 0$, the interactions between them in the magnetic dipoles aligns the qubits all parallel to each other. The simplest excited state is then the state with exactly one pair of qubits which are anti-aligned, and this leads to a defect (or a ‘domain wall’) between sites j and $j + 1$ separating perfectly aligned qubits (we are numbering the sites consecutively along the chain):

$$|j, j + 1\rangle = \dots |\uparrow\rangle_{j-2} |\uparrow\rangle_{j-1} |\uparrow\rangle_j |\downarrow\rangle_{j+1} |\downarrow\rangle_{j+2} |\downarrow\rangle_{j+3} \dots \quad (11)$$

With a small non-zero value of g , the qubits are allowed to tunnel between the up and down states, and it is then not difficult to see that this defect becomes mobile *i.e.* it is a quasiparticle. However there is no simple relationship between this quasiparticle and that in (10), and the qubits are organized in very different superposition states. The behavior of the quasiparticle in (11) with increasing g parallels that below (10) with decreasing g : the domain wall becomes increasingly diffuse as g increases, and the quasiparticle description holds only below an energy which goes to zero precisely at $g = g_c$.

Collecting all the information above, we can draw a ‘phase diagram’ of the coupled qubit model as a function of the transverse field, which is tuned by g , and the temperature T . This is shown in Fig 3. There is a quantum phase transition at the critical point $g = g_c$, $T = 0$, flanked by ferromagnetic and paramagnetic states on its sides. On both sides of the critical point, there is a range of temperatures below which the quasiparticle description is appropriate, and this upper bound in temperature corresponds is related to the upper bound in energy discussed above by Boltzmann’s constant k_B . Furthermore, there is a ‘fan’ emanating from the quantum critical point where the quasiparticle description is not appropriate and a fundamentally new ap-

¹A technical aside: for the simple model under consideration here, this statement is strictly true only in spatial dimensions $d < 3$. In $d = 3$, a modified quasiparticle picture can be applied even at $g = g_c$

proach based upon the entangled state at $g = g_c$, and its non-quasiparticle states has to be developed.

The description of this new quantum-critical regime has been the focus of much research in recent years. Some models display a remarkable universality in their properties in this regime *i.e.* some observable properties of the qubits are independent of the precise couplings between them. For example, we can ask for the value of the qubit relaxation rate, Γ_R which is the inverse time it takes for a perturbation from thermal equilibrium to die away; this rate is given by

$$\Gamma_R = \mathcal{C}_Q \frac{k_B T}{\hbar} \quad (12)$$

where \hbar is Planck's constant and \mathcal{C}_Q is a dimensionless universal number which depends only on some known, gross features of the model (such as its spatial dimensionality), but is independent of the details of the couplings in H . So remarkably, a macroscopic collective dynamic property of 10^{23} qubits is determined only by the absolute temperature and by fundamental constants of nature. It is worth noting here that it is only in this quantum-critical region that the qubits are strongly entangled with each other, in a sense similar that required for quantum computation. The dynamics of the qubits here is actually quite 'incoherent' here, and overcoming such decoherence effects is one of the major challenges on the road to building a quantum computer in the laboratory.

Extensions of the simple phase diagram in Fig 3, and of its physical properties, have been explored in the experiments of Rosenbaum, Aeppli and collaborators. A good understand now exists of very clean crystals of LiHoF_4 . However, crystals which have been intentionally doped with impurities have a far more complicated interactions between the qubits, and their study remains an active topic of research.

3 Squeezing the Bose-Einstein condensate

Prompted by communications with S. N. Bose in 1924, Albert Einstein considered the problem of cooling a gas in a container to very low temperatures. At room temperature, a gas such as helium consists of rapidly moving atoms, and can be visualized as classical billiard balls which collide with the walls of the container and occasionally with each other. As the temperature is lowered, the atoms slow down, and their quantum-mechanical characteristics

become important: de Broglie taught us that each atom is represented by a wave, and the de Broglie wavelength of the helium atoms becomes larger as the temperature is lowered. This has dramatic macroscopic consequences when the wavelength becomes comparable to the typical distance between the atoms. Now we have to think of the atoms as occupying specific quantum states which extend across the entire volume of the container. We are faced with the problem of cataloging all such many-atom states. The atoms are indistinguishable from each other, and quantum mechanics requires that we interpret two states which differ only by the exchange of the position of a pair of atoms as not being two distinct states at all, but as components of a single state. Furthermore, if the atoms are ‘bosons’ (any atom with an even total number of electrons, protons, and neutrons is a boson, as is helium) an arbitrary number of them can occupy any single quantum state *i.e.* there is no exclusion principle as there is for ‘fermions’ such as electrons. If the temperature is low enough then the many atom system will search for its lowest energy state, and for bosons this means that *every* atom will occupy the *same* lowest energy wavelike quantum state extending across the entire container. A macroscopic number of atoms occupying a single microscopic state is a Bose-Einstein condensate. Einstein showed that the Bose-Einstein condensate appeared below a critical temperature which was roughly determined by the condition that the de Broglie wavelength of an atom equal the mean atomic spacing.

Unbeknownst to Bose and Einstein, the Bose-Einstein condensate had actually already been discovered in the laboratory over a decade before their theoretical work, but several more decades would pass before this connection between theory and experiment was clearly understood—a striking example of how convoluted the progress of science can often be, and how things that seem obvious in retrospect can go unnoticed for a long time. It had been noted by Kammerlingh Onnes that liquid helium at very low temperatures displayed the remarkable property of superfluidity: the ability to flow without any appreciable viscosity. However, it was not until 1938 that Fritz London first proposed that superfluidity was linked to the formation of a Bose-Einstein condensate of helium atoms. This proposal was met with skepticism: the Bose-Einstein theory was for an ideal gas of non-interacting bosons, while there was no doubt that the helium interacted strongly with each other. Subsequent theoretical developments have since shown that London was essentially correct: interactions between the bosons do deplete the condensate (not all atoms are in the lowest energy single boson state), but

the condensate remains intact in the sense that it retains a finite fraction, *i.e.* a macroscopic number of atoms.

In the last decade, a beautiful new realization of the Bose-Einstein condensate has been created in the laboratory. In 1995, Eric Cornell, Carl Wiemann and collaborators succeeded in trapping and cooling a gas of rubidium (Rb) atoms to a high enough density and low enough temperature to induce them to form a Bose-Einstein condensate. Unlike the truly macroscopic condensate of 10^{23} atoms in liquid helium, this condensate is a much more fragile and delicate object: it exists only in a trap containing 10^{3-6} atoms carefully isolated from the world around it and cooled to temperatures extremely close to absolute zero. Indeed, it is not a simple matter to know that atoms have formed a Bose-Einstein condensate. In their experiments, Cornell and Wiemann released the atoms from their trap, and then measured the velocities of the escaping atoms: at low enough temperatures, they found a large enhancement in the number of atoms with velocities close to zero, as would be the case for atoms which emerged from the Bose-Einstein condensate. See Fig 4. A more quantitative description of the velocity distribution function is in good agreement with the theory of the interacting Bose gas, and so Fig 4 is convincing evidence for the formation of a Bose-Einstein condensate of trapped Rb atoms.

Nothing we have said so far in this section relates to a quantum phase transition: any collection of bosons, when cooled to sufficiently low temperatures, will always reach the same quantum state—a Bose-Einstein condensate. Strong repulsive interactions between the bosons can deplete the condensate, but they never annihilate it, and the system as a whole remains a perfect superfluid. Is there any way to reach a different ground state of bosons and possibly also a quantum phase transition?

The key to the formation of the Bose-Einstein condensate is that some of the atoms move freely throughout the entire system in the same quantum state. This suggests that if we were able to disrupt the motion of the atoms by a set of microscopic barriers, we may be able to destroy the condensate and reach a new ground state. The precise implementation of this idea was discussed by Sebastian Doniach in 1981 in a slightly different context, and extended to the boson systems of interest here by Matthew Fisher and collaborators in 1989. For the trapped gas of Rb atoms we have discussed here, this idea was demonstrated in a remarkable recent experiment by Immanuel Bloch and collaborators in Munich, and so we will discuss the idea using their implementation.

Bloch and collaborators impeded the motion of the Rb bosons by applying a periodic array of barriers within the trap; more precisely they applied a periodic potential to the Rb atoms. In the two-dimensional view shown in Fig 5, we can visualize the Rb atoms as moving across an egg-carton: there is an array of sites at which the Rb atoms prefer to reside to lower their potential energy, and they have to go over a barrier to move from any site to its neighbor. In Bloch's laboratory, this periodic potential was created by standing waves of light: along each direction there were two counter-propagating laser beams (and so a total of six lasers) whose interference produces an effective periodic potential for the Rb atoms.

In the presence of a weak periodic potential, Bloch observed only a minor, but significant, change in the state of the Rb atoms. The atoms were released from the trap, as in the experiment of Cornell and Wiemann, and then their velocity distribution function was observed. Bloch observed the distribution shown in Fig 6. Note that it has a large peak near zero velocity, as Fig 4, and this is a tell-tale signature of the Bose-Einstein condensate. However, there is also a lattice of satellite peaks, and these record the influence of the periodic potential. Crudely speaking, the quantum state into which the atoms condense has 'diffracted' off the periodic potential, and the satellite peaks in Fig 6 represent the diffraction pattern of the periodic potential. So the shape of the condensate has adjusted to the periodic potential, but otherwise the Bose-Einstein condensate retains its integrity: this will change at stronger periodic potentials, as we shall discuss shortly.

In preparation for our description of the quantum phase transition, it is useful to explicitly write down the quantum state of the Bose-Einstein condensate. Let us assume that the atoms are only able to occupy a single quantum state in each minimum of the periodic potential (this is actually an excellent approximation). If atom number n (we number the atoms in some order) occupies the state in the j 'th minimum of the periodic potential (we also number the minima of the periodic potential in some order), we denote this state by $|j\rangle_n$. In the Bose-Einstein condensate, each atom will actually occupy the same state: the state which is a linear superposition of the states in every well. So the Bose-Einstein condensate (BEC) of bosons numbered $n_1, n_2, n_3 \dots$ in potential wells $j_1, j_2, j_3 \dots$ is

$$|BEC\rangle = \dots \left(\dots |j_1\rangle_{n_1} + |j_2\rangle_{n_1} + |j_3\rangle_{n_1} \dots \right) \\ \times \left(\dots |j_1\rangle_{n_2} + |j_2\rangle_{n_2} + |j_3\rangle_{n_2} \dots \right)$$

$$\times \left(\dots |j_1\rangle_{n_3} + |j_2\rangle_{n_3} + |j_3\rangle_{n_3} \dots \right) \dots \quad (13)$$

Upon expanding out the products above, note that the number of bosons in any given well fluctuates considerably *e.g.* one of the terms in the expansion is $\dots |j_1\rangle_{n_1} |j_1\rangle_{n_2} |j_1\rangle_{n_3} \dots$ in which all three of the bosons numbered $n_{1,2,3}$ are in the well j_1 and none in the wells j_2 or j_3 . The $|BEC\rangle$ state is a quantum mechanical superposition of all these states in which the number of bosons fluctuate strongly in each well. This is another way of characterizing the Bose-Einstein condensate, and is indeed the essential property which leads to an understanding of its superfluidity: the strong number fluctuations mean that particles are able to flow across the system with unprecedented ease, and without resistance.

Now turn up the strength of the periodic potential. Theory tells us that there is a critical strength at which there is a quantum phase transition, and beyond this point there is no superfluidity—the Bose-Einstein condensate has disappeared. What has taken its place? As in our discussion in Section 2 we can understand the structure of this state by looking at a limiting case: the state should be qualitatively similar to that in the presence of a very strong periodic potential. In such a situation, we expect the tunnelling between neighboring minima of the periodic potential to be strongly suppressed. This, combined with the repulsive interactions between the bosons in the same potential well should lead to a strong suppression of the number fluctuations in the quantum state. An idealized state, with no number fluctuations has the form

$$\begin{aligned} |I\rangle &= \dots |j_1\rangle_{n_1} |j_2\rangle_{n_2} |j_3\rangle_{n_3} \dots + \dots |j_1\rangle_{n_1} |j_2\rangle_{n_3} |j_3\rangle_{n_2} \dots \\ &+ \dots |j_1\rangle_{n_2} |j_2\rangle_{n_1} |j_3\rangle_{n_3} \dots + \dots |j_1\rangle_{n_2} |j_2\rangle_{n_3} |j_3\rangle_{n_1} \dots \\ &+ \dots |j_1\rangle_{n_3} |j_2\rangle_{n_2} |j_3\rangle_{n_1} \dots + \dots |j_1\rangle_{n_3} |j_2\rangle_{n_1} |j_3\rangle_{n_2} \dots + \dots \end{aligned} \quad (14)$$

Now each site $j_{1,2,3}$ has exactly one of the particles $n_{1,2,3}$ in all the terms. Indeed, any one of the terms above suffices to describe the configuration of the particles in the ground state, because the particles are indistinguishable and we can never tell which particular particle is residing in any given well—we need all the permutations in (14) simply to ensure that physical results are independent of our arbitrary numbering of the particles. Note also that all terms in $|I\rangle$ are also present in $|BEC\rangle$ —the difference between the states is that $|BEC\rangle$ has many more terms present in its quantum superposition, representing the number fluctuations. In this sense, the inequality between

$|BEC\rangle$ and $|I\rangle$ is similar to the inequality (9) between the states on the either side of the transition in the coupled qubit model.

What is the physical interpretation of $|I\rangle$? The complete suppression of number fluctuations now means that the bosons are completely unable to flow across the system. If the bosons were electrically charged, there would be no current flow in response to an applied electric field *i.e.* it is an insulator. Alternatively, in the atom trapping experiments of Bloch, if we were to ‘tilt’ the optical lattice potential (in two dimensions, we tilt the egg-carton of Fig 5), there would be no flow of bosons in the state $|I\rangle$. In contrast, bosons would flow without any resistance in the state $|BEC\rangle$. The quantum phase transition we have described in this section is therefore a *superfluid-insulator* transition.

It has not yet been possible to observe flows of atoms directly in the experiments by Bloch and collaborators. However, as in the state $|BEC\rangle$, the velocity distribution function of the atoms in the state $|I\rangle$ can be easily measured, and this is shown in Fig 7. Now there is no sharp peak near zero velocity (or at its diffractive images): the absence of a Bose-Einstein condensate means that all the particles have large velocities. This can be understood as a consequence of Heisenberg’s uncertainty principle: the particle positions are strongly localized within single potential wells, and so their velocities become very uncertain.

For future experimental studies, and analogous to our discussion in Fig 3 for the coupled qubits, we can now sketch a phase diagram as a function of the strength of the periodic potential and the temperature: this is shown in Fig 8. For weak periodic potentials, we have a superfluid ground state: above this, at finite temperature, there are excitations involving superflow over long length scales. At large periodic potentials, we have an insulating ground state: the excitations now involve small number fluctuations between neighboring pairs of sites, which cost a finite energy per number fluctuation. In the intermediate quantum-critical regime, we have behavior analogous to that in Fig 3. The characteristic relaxation rate for boson number fluctuations obeys an expression analogous to (12): the dynamics of the bosons is strongly dissipative (with velocities proportional to applied forces) and so we have christened this yet-to-be-observed regime as *Bose molasses*.

4 The cuprate superconductors

Electricity is the flow of electrons in wires. The wires are made of metals, like copper, in which each metallic atom ‘donates’ a few of its electrons, and these electrons move freely, in quantum-mechanical wavelike states, through the entire wire. However, single electrons are not bosons, but fermions which obey the exclusion principle, and so more than one electron cannot occupy the same quantum state to form a Bose-Einstein condensate. The motion of electrons in a metal has a small but finite resistance; this arises from scattering of the electrons off the ever-present defects in the crystalline arrangement of atoms in a metal. Consequently metals are conductors but not *superconductors*.

However, in 1911, Kammerlingh Onnes, cooled the metal mercury below 4.2°K and found that its electrical resistance dropped precipitously to an immeasurably small value—he had discovered a superconductor. An explanation of this phenomenon eluded physicists until 1957 when John Bardeen, Leon Cooper and Robert Schrieffer proposed a theory which again invoked the Bose-Einstein condensate. However the condensate was not of single electrons, but the analog condensate of *pairs* of electrons (now called Cooper pairs): each electron in the metal finds a partner, and the pairs then occupy the same quantum state which extends across the system. While individual electrons are fermions, pairs of electrons obey Bose statistics at long length scales, and this allows them to form an analog of the Bose-Einstein condensate. Strong fluctuations in the numbers of the Cooper pairs in different regions of the wire are then responsible for superconductivity, just as we found in Section 3 that superfluidity was linked to fluctuations in the number of Rb atoms between different minima of the periodic potential.

The phenomenon of superconductivity clearly raises the possibility of many exciting technological applications: resistanceless flow of electricity would have tremendous impact on electrical power transmission, high frequency electrical circuits behind the wireless communication revolution, and in medical applications like magnetic resonance imaging (MRI) which require large magnetic fields, to name but a few. Many such applications are already in the marketplace, and the main obstacle to more widespread usage is the low temperature required to obtain superconductor. Raising the maximum critical temperature below which superconductivity appears has been a central research goal for physicists since Kammerlingh Onnes.

A dramatic improvement in the temperature required for superconductiv-

ity came in 1987 in a breakthrough by Georg Bednorz and Alex Müller (see Fig 9). These IBM researchers discovered a new series of compounds (the ‘cuprates’) which became superconductors at temperatures as high as 120°K; the highest temperature on record prior to their work had been 15°K. This revolutionary discovery sparked a great deal of theoretical and experimental work, and many important questions remain subject to debate today. Here we will discuss the current understanding of the quantum ground states of these compounds, the role played by quantum phase transitions, and point out some of the open questions.

The best place to begin our discussion is the insulating compound La_2CuO_4 , whose crystal structure is shown in Fig 10. Despite its complexity, most of its electronic properties are controlled by a simple substructure—a single layer of Cu ions which reside on the vertices of a square lattice. The crystal consists of an infinite number of such layers stacked on top of each other, but the coupling between neighboring layers is small enough that we can safely ignore it, and concentrate on a single layer. A further simplification is that we need focus on only a single $3d$ orbital on each Cu ion. In La_2CuO_4 each such orbital has precisely one electron. The Coulomb repulsion between the electrons prevents them from hopping between Cu sites, and this is the reason La_2CuO_4 is a very good insulator. The reader should notice a simple parallel to our discussion of Rb atoms in Section 3: there, when the periodic potential was strong enough, each Rb atom was localized in a single well, and the ground state was an insulator. Here, the periodic potential is provided by the existing lattice of positively charged Cu ions, and one electron is trapped on each such ion.

So how do we get a superconductor? Generalizing the discussion of Section 3, the naive answer is: force La_2CuO_4 to undergo a insulator-to-superfluid quantum transition. However, unlike Section 3, we do not have any available tools to tune the strength of the periodic potential experienced by the electrons. Bednorz and Müller accomplished the equivalent by chemical substitution: the compound $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$ has a concentration δ of trivalent La ions replaced by divalent Sr ions, and this has the effect of removing some electrons from each square lattice of Cu ions, and creating ‘holes’ in the perfect insulating configuration of La_2CuO_4 . With the holes present, it is now possible to easily transfer electrons from site to site, and across the crystal, without having to pay the expensive price in energy associated with putting two electrons on the same site. This should be clear from the ‘egg-carton’ model of the insulator in Fig 5: the reader can verify

that if a carton is not full of eggs, we can move any egg across the carton only by ‘hops’ of eggs between neighboring sites, and without ever putting two eggs on the same site. With holes present, it becomes possible for electrons to move across the crystal, and when the concentration of holes, δ , is large enough, this material is a ‘high temperature’ superconductor. We now know that $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$ undergoes an insulator-to-superfluid quantum phase transition with increasing δ at the hole concentration $\delta = 0.055$. In some phenomenological aspects, this transition is similar to the superfluid-insulator transition discussed in Section 3, but it should be clear from our discussion here that the microscopic interpretation is quite different.

However there is a complication here which makes the physics much more involved than that of a single insulator-superfluid quantum phase transition. As we noted in Section 2, each electron has two possible spin states, with angular momenta $\pm\hbar/2$, and these act like a qubit. It is essential to also pay attention to the magnetic state of these spin-qubits in $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$. We discussed in Section 2 how such coupled qubits could display a quantum phase transition between magnetic and paramagnetic quantum ground states in LiHoF_4 . The spin-qubits in $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$ also display a related magnetic quantum phase transition with increasing δ ; the details of the magnetic and paramagnetic states here are, however, quite different from those on LiHoF_4 , and we will describe them further below. So with increasing δ , a description of the quantum ground state of $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$ requires not one, but at least two quantum phase transitions. One is the insulator-superfluid transition associated with the motion of the electron charge, while the second is a magnetic transition associated with the electron spin. (Numerous other competing orders have also been proposed for $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$ and related compounds, and it is likely that the actual situation is even more complicated). Furthermore, these transitions are not independent of each other, and much of the theoretical difficulty resides in understanding the interplay between the fluctuations of the different transitions.

For now, let us just focus on just the state of electron spin qubits in the insulator La_2CuO_4 , and discuss the analog of the qubit quantum phase transition in Section 2. In this insulator, there is precisely one such qubit on each Cu site. However, unlike the qubit on the Ho site in LiHoF_4 , there is no preferred orientation of the electron spin, and the qubit of a single Cu ion is free to rotate in all directions. We know from experiments that in the ground state of La_2CuO_4 , the ‘exchange’ couplings between the qubits are such that they arrange themselves in an *antiferromagnet*, as shown in Fig 11.

Unlike the parallel ferromagnetic arrangement of qubits in LiHoF_4 , now the couplings between the Cu ions prefer that each qubit align itself anti-parallel to its neighbor. The checkerboard arrangement in La_2CuO_4 optimizes this requirement, and each qubit has all its neighbors anti-parallel to it. While this arrangement satisfactorily minimizes the potential energy of the spins, the reader will not be surprised to learn that Heisenberg's uncertainty relation implies that all is not perfect: there is an energy cost in localizing each qubit along a definite spin direction, and there are ever-present quantum fluctuations about the quiescent state pictured in Fig 11. In La_2CuO_4 these fluctuations are not strong enough to destroy the antiferromagnetic order in Fig 11, and the resulting checkerboard arrangement of magnetic moments can and has been observed by scattering magnetic neutrons off a La_2CuO_4 crystal.

These neutron scattering experiments can also follow the magnetic state of the spin qubits with increasing δ . For large enough δ , when the ground state is a superconductor, it is also known that there is no average magnetic moment on any site. So there must be at least one quantum phase transition at some intermediate δ at which the antiferromagnetic arrangement of magnetic moments in La_2CuO_4 is destroyed: this is the antiferromagnet-paramagnet quantum transition. This transition is clearly coupled to the onset of superconductivity in some manner, and a simple description is therefore not straightforward.

With the insulator-superfluid and antiferromagnet-paramagnet quantum phase transitions in hand, and with each critical point being flanked by two distinct types of quantum order, we can now envisage a phase diagram with at least 4 phases. These are the antiferromagnetic insulator, the paramagnetic insulator, the antiferromagnetic superfluid, and the paramagnetic superfluid (see Fig 12). Of these, La_2CuO_4 is the antiferromagnetic insulator, and the high temperature superconductor at large δ is the paramagnetic superfluid; note that these states differ in *both* types of quantum order. The actual phase diagram turns out to be even more complex, as there are strong theoretical arguments implying that at least one additional order must exist, and we will briefly review one of these below.

One strategy for navigating this complexity is to imagine that, like the simpler system studied in Section 2, we did indeed have the freedom to modify the couplings between the spin qubits in undoped La_2CuO_4 , so that we can tune the system across a quantum phase transition to a paramagnet, while it remains an insulator. This gives us the luxury of studying the

antiferromagnet-paramagnet quantum phase transition in a simpler context, and decouples it from the insulator-superfluid transition. It is also reasonable to expect that the insulating paramagnet is a better starting point for understanding the superconducting paramagnet, because these two states differ only by a single type of order, and so can be separated by a single quantum phase transition.

To study the insulating antiferromagnet-paramagnet quantum phase transition we need to modify the couplings between the qubits on the Cu sites, while retaining the full symmetry of the square lattice—in this manner we perturb the antiferromagnetic ground state in Fig 11. As this perturbation is increased, the fluctuations of the qubits about this simple checkerboard arrangement will increase, and the average magnetic moment on each site will decrease. Eventually, we will reach a quantum critical point, beyond which the magnetic moment on every site is precisely zero *i.e.* a paramagnetic ground state. What is the nature of this paramagnet? A sketch of the quantum state of this paramagnet, predicted by theory, is shown in Fig 13. The fundamental object in Fig 13 is a *singlet valence bond* between qubits—this valence bond has many similarities to the covalent bond which is the basis of organic chemistry. As in Section 2, we represent the states of the electron spin on site j , by $|\uparrow\rangle_j$ and $|\downarrow\rangle_j$, and the valence bond between spins on sites i and j is a state in which the spin qubits are *entangled* as below:

$$|V\rangle_{ij} = \frac{1}{\sqrt{2}} \left(|\uparrow\rangle_i |\downarrow\rangle_j - |\downarrow\rangle_i |\uparrow\rangle_j \right) \quad (15)$$

Note that in both states in the superposition in (15), the spins are anti-parallel to each other—this minimizes the potential energy of their interaction. However, each qubit fluctuates equally between the up and down states—this serves to appease the requirements of Heisenberg uncertainty principle, and lowers the total energy of the pair of qubits. Another important property of (15) is that it is invariant under rotations in spin space; it is evident that each qubit has no definite magnetic moment along the z direction, but we saw in Section 2 that such qubits could have a definite moment along the x direction. That this is not the case here can be seen by the identity

$$|V\rangle_{ij} = \frac{1}{\sqrt{2}} \left(|\rightarrow\rangle_i |\leftarrow\rangle_j - |\leftarrow\rangle_i |\rightarrow\rangle_j \right); \quad (16)$$

the reader can verify that substituting (2) and (3) into (16) yields precisely (15). So the singlet valence band has precisely the same structure in terms

of the $\pm x$ oriented qubits as between the $\pm z$ oriented qubits—it is the unique combination of two qubits which does not have *any* preferred orientation in spin space. This is one of the reasons it is an optimum arrangement of a pair of qubits. However, the disadvantage of a valence bond is that a particular qubit can form a valence bond with only a single partner, and so has to make choice between its neighbors. In the antiferromagnetic state in Fig 11, the arrangement of any pair of spins is not optimal, but the state has the advantage that each spin is anti-parallel to all four of its neighbors. The balance of energies shifts in the state in Fig 13, and now the spins all form singlet bonds in the rotationally invariant paramagnet. For the state illustrated in the figure, the spins have chosen their partner in a regular manner, leading to a crystalline order of the valence bonds. Detailed arguments for these, and related orderings of the valence bonds, appear in the theoretical works, and some of the physics is briefly discussed in the Side Panel. A significant property of the resulting bond-ordered paramagnet is that while spin rotation invariance has been restored, the symmetry of the square lattice has been broken. This *bond order* is an example the promised third quantum order. The interplay between the spin-rotation and square lattice symmetries in the antiferromagnet-paramagnet quantum transition makes its theory far more complicated than that for the ferromagnet-paramagnet transition described in Section 2.

Tests of these theoretical ideas for the quantum transition between the states in Fig 11 and 13 have been quite difficult to perform. Most computer studies of models of such antiferromagnets suffer from the ‘sign’ problem *i.e.* the interference between different quantum histories of the qubits makes it very difficult to estimate the nature of the quantum state across the transition. In 2002, Anders Sandvik and collaborators succeeded in finding a model for which the sign problem could be conquered. This model displayed a quantum transition between a magnetically ordered state, as in Fig 11, and a paramagnet with bond order, as in Fig 13.

To summarize, the physics of the spin-qubits residing on the Cu sites is a competition between two very different types of insulating quantum states. The first is the antiferromagnet in Fig 11 which is found in La_2CuO_4 : there is an average magnetic moment on each Cu site, and the ‘up’ and ‘down’ moments are arranged like the black and white squares on a chess board. The second state is a paramagnet in which the spins are paired in valence bonds described by the quantum state (15).

With this understanding in hand, we can now address the physics of the

electron charge, and the insulator-superconductor quantum transition with increasing δ in $\text{La}_{2-\delta}\text{Sr}_\delta\text{CuO}_4$. A picture of the insulating state was already shown in Fig 11. We show an analogous picture of the superconductor in Fig 14. Note first that there are Cu sites with no electrons: these are the holes, and as we have discussed earlier, they are able hop rapidly across the entire crystal without paying the price associated occupancy of two electrons on the same Cu site. The motion of these holes clearly perturbs the regular arrangement of magnetic moments at $\delta = 0$ in Fig 11, and this enhances their quantum fluctuations and reduces the average magnetic moment on every site. Eventually, the magnetic moment vanishes, and we obtain a paramagnet in which the electrons are all paired in singlet bonds, as has been shown in Fig 14. Unlike the valence bonds in Fig 13, the valence bonds in Fig 14 are mobile and can transfer charge across the system—this is illustrated in Fig 15 where the motion of the holes has rearranged the valence bonds and resulted in the net transfer of electronic charge in direction opposite to the motion of the holes. This suggests an alternative interpretation of the valence bonds: they are the *Cooper pairs* of electrons, which can move freely, as a pair, across the system with holes present. Their facile motion suggests that these composite bosons will eventually undergo Bose-Einstein condensation, and the resulting state will be a superconductor. We have therefore presented here a simple picture of the crossover from the insulator to a superconductor with increasing δ .

The essence of our description of the insulator-superconductor transition was the ‘dual life’ of a valence bond pair. In the insulator, these valence bonds usually prefer to crystallize in regular bond-ordered states, as shown in Fig 13, and in the results of Sandvik and collaborators. The valence bonds do resonate among themselves, and thus change their orientation, but this resonance does not lead to motion of any charge. In the superconductor, the valence bonds become mobile, and rather longer-ranged, and transmute into the Cooper pairs of the theory of Bardeen, Cooper, and Schrieffer. The Bose-Einstein condensation of these pairs then leads to high temperature superconductivity.

While this picture is appealing, does it have any practical experimental consequences? Can the dual interpretation of a valence bond be detected experimentally? One theoretical proposal was made in 2001 by Eugene Demler, Kwon Park, Ying Zhang and the author, reasoning as follows. Imagine an experiment in which we are locally able to impede the motion of the Cooper pairs. In the vicinity of such a region, the valence bonds will become less

mobile, and so should behave more like the valence bonds of the paramagnetic insulator. As we have seen in Fig 13, such ‘stationary’ valence bonds are likely to crystallize in a regular bond ordered arrangement. An experimental probe which can investigate the electronic states at the scale of the Cu lattice spacing (which is a few angstroms; an angstrom is 10^{-10} meters) should be able to observe a regular modulation associated with the bond order. We also proposed a specific experimental mechanism for impeding the motion of the Cooper pairs: apply a strong magnetic field transverse to the sample—this induces vortices around which there is a supercurrent of Cooper pairs, but the cores of the vortices are regions in which the superconductivity is suppressed. These cores should nucleate a halo of bond order. Matthias Vojta had studied the relative energies of different types of bond order in the superconductor, and found that in addition to the modulation with a period of 2 lattice spacings shown in Fig 13, states with a period of 4 lattice spacings were stable over an especially large range of doping.

On the experimental side, Jennifer Hoffman, Seamus Davis and collaborators carried out experiments at the University of California in Berkeley examining the electronic structure in and around the vortex cores on the surface of the high temperature superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ (closely related experiments have also been carried out by Craig Howald, Aharon Kapitulnik and collaborators at Stanford University). In a remarkable experimental tour-de-force, they were able to obtain sub-angstrom scale resolution of electronic states at each energy by moving a ‘scanning tunnelling microscope’ across a clean surface of $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. A picture of the observed structure of the electronic states around a single vortex is shown in Fig 16. Notice the clear periodic modulation, with a period which is approximately four lattice spacings. Is this an experimental signal of the bond order we have discussed above? While this is certainly a reasonable possibility, the issue has not yet been conclusively settled. There continues to be much debate in the community on the proper interpretation of these experiments, and on the relative role of single quasiparticle and valence bond excitations. Whatever the final answer turns out to be, it is exciting to be involved in a debate which involves fundamental theoretical issues of the physics of quantum phase transitions coupled with remarkable experiments at the cutting edge of technology.

5 Conclusions

This chapter has given a simple overview of our efforts to understand ‘*quantum matter*’. Quantum mechanics was originally invented as a theory of small numbers of electrons interacting with each other, with many features that jarred our classical intuition. Here, we have hopefully convinced the reader that many of these counter-intuitive features of quantum mechanics are also apparent in the bulk properties of macroscopic matter. We have discussed macroscopic phases whose characterizations are deep consequences of the quantum mechanical principle of superposition, and which can undergo phase transitions driven entirely by Heisenberg’s uncertainty principle. Theoretical work on the classification of distinct phases of quantum matter continues today: many examples which make intricate use of quantum superpositions to produce new states of matter have been found. Some of this work is coupled with experimental efforts on new materials, and this bodes well for much exciting new physics in the century ahead.

Further reading

Technical information and original references to the literature can be found in the article by S. L. Sondhi, S. M. Girvin, J. P. Carini, and D. Shahar, *Reviews of Modern Physics* **69**, 315-333 (1997) and in the book *Quantum Phase Transitions*, S. Sachdev, Cambridge University Press, Cambridge (1999). A more detailed discussion of the application of the theory of quantum phase transition to the cuprate superconductors appears in *Order and quantum phase transitions in the cuprate superconductors*, S. Sachdev, *Reviews of Modern Physics*, July 2003, <http://arxiv.org/abs/cond-mat/0211005>. A related perspective on ‘fluctuating order’ near quantum critical points is given in *How to detect fluctuating order in the high-temperature superconductors*, by S. A. Kivelson, E. Fradkin, V. Oganesyan, I. P. Bindloss, J. M. Tranquada, A. Kapitulnik, and C. Howald, <http://arxiv.org/abs/cond-mat/0210683>. For a review of quantum phase transitions in metals see *Quantum phase transitions* by Thomas Vojta, <http://arxiv.org/abs/cond-mat/0010285>.

Side panel: Resonance between valence bonds

The first example of ‘resonance’ in chemistry appeared in the planar hexagonal structures proposed in 1872 by Kekulé for the benzene molecule C_6H_6 —note that this predates the discovery of quantum mechanics by many years. The structures are represented by the bonding diagrams in Fig 17. Of particular interest is (what we now call) the π -orbital bond between the neighboring C atoms. This bond involves two electrons and its quantum mechanical state is similar to the valence bond in (15) or (16) (the main difference is that the electronic charges are not so well localized on the two C atoms, and there are also contributions from states with both electrons on the same C atom). There are two possible arrangements of this bond, as shown in Fig 17. With the knowledge of quantum mechanics, and the work of Linus Pauling on the physics of the chemical bond, we know that the bonds tunnel back-and-forth between these two configuration, and this enhances the stability of the benzene structure. Moreover, the final structure of benzene has a perfect hexagonal symmetry, which is not present in either of the two constituent structures.

In 1974, Patrick Fazekas and Philip Anderson applied the theory of resonance between valence bonds to quantum antiferromagnets; specific applications to the high temperature superconductors were proposed in 1987 by G. Baskaran and Anderson and also by Steven Kivelson, Daniel Rokhsar, and James Sethna. The idea was that singlet valence bonds form between the spin qubits on neighboring Cu ions, and these resonate between different possible pairings of nearest-neighbor qubits, as shown in Fig 18. It was argued by Nicholas Read and the author that in many cases the consequences of these resonances on the square lattice are very different from those in benzene. Whereas in benzene, the final ground state has full hexagonal symmetry, here the resonances actually lead to *bond order* with structures like those in Fig 13, which break the symmetry of the square lattice. The reason for this is illustrated in Fig 18: regular bond-ordered configurations have more possibilities of plaquettes for resonance, and this lowers the energy of states with higher symmetry. Of course, each bond-ordered state has partners related to it by the symmetry of the square lattice (*e.g.* the state in Fig 13 has 3 partners obtained by rotating it by multiples of 90° about any lattice point), but the tunnelling amplitude between these macroscopically distinct states is completely negligible.

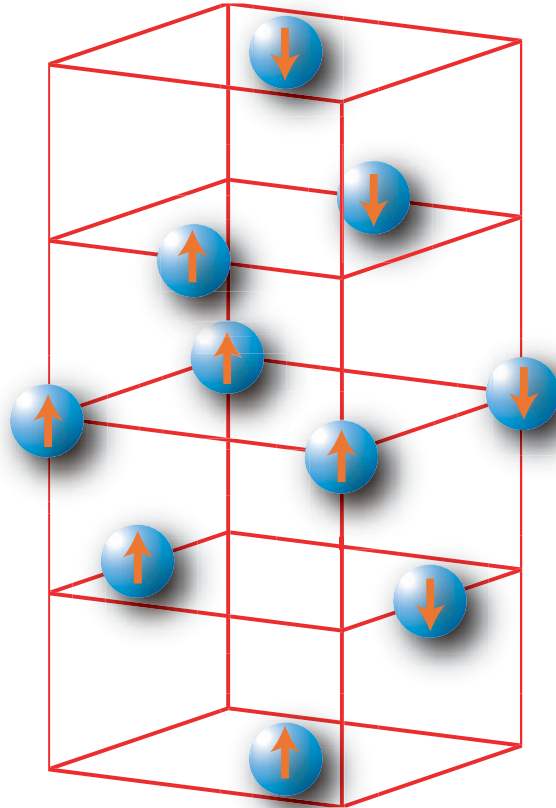


Figure 1: Locations of the Ho ions in the insulator LiHoF_4 : the Li and F ions are not shown, and are located in the spaces in between. Each Ho ion has a magnetic moment which can be oriented 'up' (in the $+z$ direction) or 'down'. This acts a qubit. The magnetic fields degenerated by the qubits couple them together to create an interacting qubit system. A quantum phase transition is induced in this system by applying a transverse magnetic field (oriented in the $+x$ direction).




Figure 2: Dependence of the average magnetic moment along the z axis on the strength of the applied transverse magnetic field at $T = 0$. There is a quantum phase transition at $g = g_c$. For $g < g_c$, the ground state is a ferromagnet, with a wavefunction qualitatively similar to $|\uparrow\rangle$ in Eq. (6) or to $|\downarrow\rangle$ in Eq. (7): small imperfections choose between the nearly equivalent possibilities, and thus break the up-down symmetry of the crystal. For $g > g_c$, the ground state is a paramagnet, with a state qualitatively similar to $|\Rightarrow\rangle$ in Eq. (8).

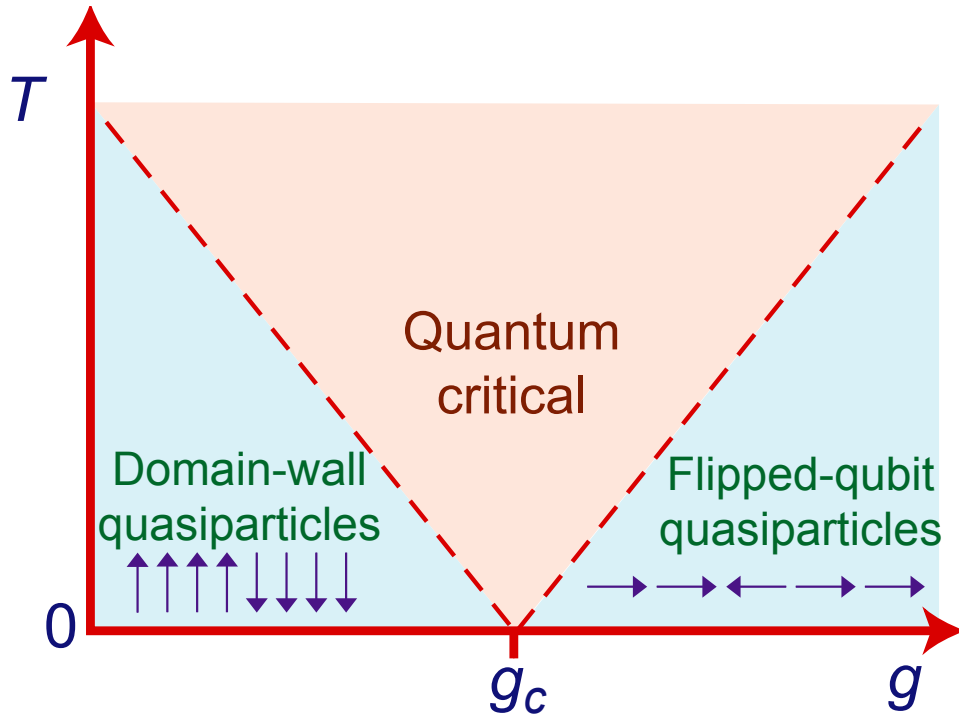


Figure 3: Phase diagram for a one-dimensional chain of coupled qubits as a function of the parameter g and the absolute temperature, T . The ground state at $T = 0$ is as described in Fig 2. The flipped qubit quasiparticles have a state similar to (10), while the domain wall quasiparticles have a state similar to (11). No quasiparticle picture works in the quantum-critical region, but we have relaxational dynamics characterized by (12).




Figure 4: Velocity distribution function of Rb atoms measured after the atoms were released from a trap at very low temperatures. The left plot is at 400×10^{-9} °K, the middle at 200×10^{-9} °K and the right plot at 50×10^{-9} °K. The large peak near zero velocity at the lower temperatures is a tell-tale signature of a Bose-Einstein condensate. (Figure from M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman and E. A. Cornell, *Science* **269**, 198 (1995) - need permission.




Figure 5: Two-dimensional section of Rb atoms moving in a periodic potential created by standing waves of laser light. The Rb atoms are like eggs which prefer to reside only at specific sites of an ‘egg-carton’. The superfluid is a superposition of states with strong fluctuations in the number of atoms in any given minimum of the periodic potential, as is expressed in (13). The insulator has a fixed number of atoms in each minima, as expressed in (14). (From H.T.C. Stoof, *Nature* **415**, 25 (2002), permission needed).




Figure 6: Velocity distribution function of Rb atoms released from a trap in the presence of a weak periodic potential. The large central peak is the signal of the Bose-Einstein condensate. The satellite peaks result from the diffraction of this condensate off the periodic potential. This signal represents observation of the state $|BEC\rangle$ in (13). Figure courtesy I. Bloch from M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, *Nature* **415**, 39 (2002).




Figure 7: As in Fig. 6 but with a stronger periodic potential. Now there are no sharp peaks, indicating the absence of a Bose-Einstein condensate, and the formation of an insulator in the state in $|I\rangle$ in (14). Figure courtesy I. Bloch from M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, *Nature* **415**, 39 (2002).




Figure 8: Phase diagram of a gas of trapped bosonic atoms in the presence of a periodic potential—this diagram is analogous to that for the qubit chain in Fig 3. The Bose-Einstein condensate is present below the critical temperature T_c , which is the only true phase transition at $T > 0$.




Figure 9: History of materials with high critical temperatures (T_c) below which they are superconducting, as of early 2003. The cuprate superconductors were discovered on Jan 27, 1986, and led to a dramatic increase in T_c . Since then, apparently unrelated compounds (MgB_2) with a moderately high T_c have also been discovered. Also indicated are temperatures at which N_2 and H_2 liquefy—these gases are most commonly used to cool materials. The search for new materials with higher critical temperatures continues, and significant future increases may well appear. Figure courtesy H. R. Ott.




Figure 10: Crystal structure of the insulator La_2CuO_4 . Red sphere are Cu, orange are O, and blue are La. The Cu ions are on the vertices of square lattices that extend in the horizontal plane. (downloaded from the web—should be redrawn)

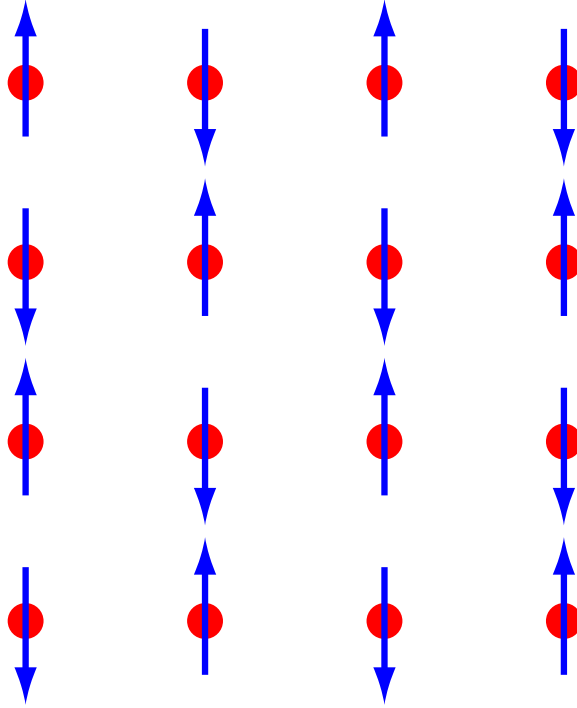


Figure 11: La_2CuO_4 , an insulating antiferromagnet. The spin qubits reside on the Cu sites (red circles), and their average magnetic moments arrange themselves in the checkerboard arrangement shown above. Unlike the qubits in LiHoF_4 in Fig 1, the spins above are free to rotate in any direction, provided the spins on neighboring Cu sites retain an antiparallel orientation.

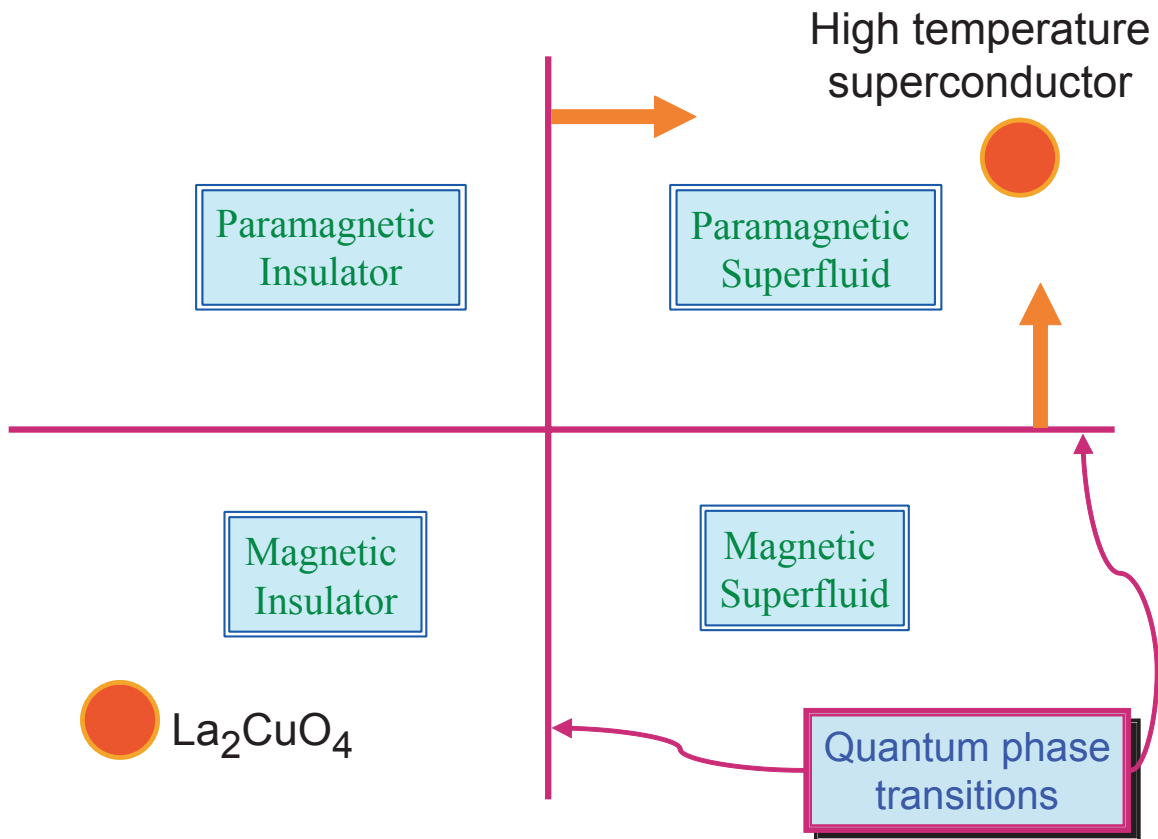


Figure 12: A minimal schematic theoretical phase diagram for the cuprate superconductors. The axes represent two suitable parameters which can be varied so that the lowest energy state traverses across the quantum phase transitions shown. The insulating magnet, La_2CuO_4 , is well understood. A theory of the high temperature superconductor can be developed by understanding the quantum phase transitions, and then moving along the orange arrows.




Figure 13: A paramagnetic insulator obtained by perturbing the antiferromagnet in Fig 11. The couplings between the qubits are modified so that quantum fluctuations of the average magnetic moment on each site are enhanced, and the insulator undergoes a quantum phase transition to a state with zero magnetic moment on every site. Each ellipse represents a singlet valence bond as in (15) or (16). The state above has bond order because the valence bonds are more likely to appear on the links with ellipses shown above, than on the other links.

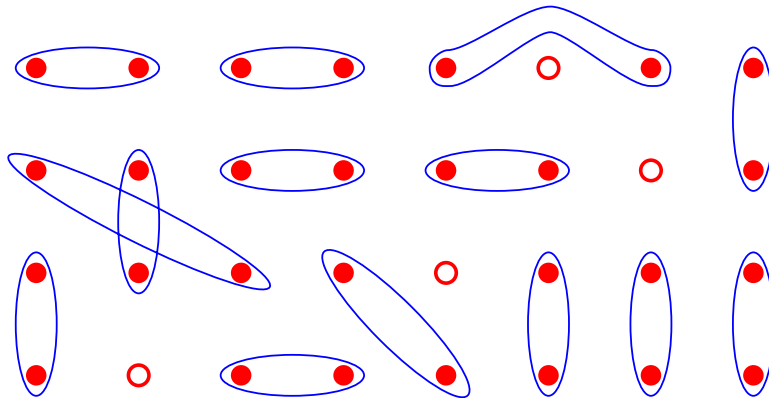


Figure 14: A snapshot of the superconducting state in $\text{La}_{2-\delta}\text{Sr}_{\delta}\text{CuO}_4$ for large δ . The open circles represent the holes, or Cu sites with no electrons. The spins on the remaining sites are paired with each other in the valence bonds of Fig 13. We argue in the text that the valence bonds in the superconductor possess vestiges of the bond order of Fig 13, and this bond order is enhanced and experimentally detectable near vortices in the superflow of Cooper pairs.




Figure 15: As in Fig 14, but the two holes on the lower left of Fig 14 have moved to the upper right. Alternatively, we can view the motion of the holes as the counterflow of the charge of the singlet valence bonds: this interpretation allows us to view a valence bond as a bosonic Cooper pair, which can undergo Bose-Einstein condensation.




Figure 16: Measurement of Hoffman, Davis and collaborators of the local density of states on the surface of the high temperature superconductor $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$ by a scanning tunneling microscope. A magnetic field has been applied perpendicular to the plane of the paper, and the dark regions are centered on the vortices in the superflow of the Cooper pairs; the diameter of each dark region is about 100 angstroms. There is a checkerboard modulation in each dark region with a period of approximately 4 Cu lattice spacings, and this is possibly a vestige of bond order, as discussed in the text. (Figure from J.E. Hoffman *et al.* Science **295**, 466 (2002) - need permission).




Figure 17: The two structures of benzene proposed by Kekulé. The second of the double lines represent valence bonds between π orbitals which are similar to the valence bonds in (13). The stability of benzene is enhanced by a quantum mechanical resonance between these two states.




Figure 18: Resonance of valence bonds around the plaquettes of a square lattice. Shown are resonances for two different valence bond arrangements are the plaquettes marked with a star. Note that there are many more possibilities of resonance in (a) than in (b), and this effect selects ground states with bond order.