

Q&A

CONDENSED-MATTER PHYSICS

Optical lattices

Markus Greiner and Simon Fölling

Optical lattices have rapidly become a favoured tool of atomic and condensed-matter physicists. These crystals made of light can be used to trap atoms at very low temperatures, creating a workshop in which to pore over and tinker with fundamental properties of matter.

Why make an optical lattice?

Imagine you are trying to understand a complex quantum-physical phenomenon such as high-temperature superconductivity, just one of many intriguing effects that occur in the 'quantum gas' of electrons that pervades a solid crystal. The crystal lattice of a solid has tiny dimensions — atoms are spaced about a tenth of a nanometre apart — and it would be great to have an enlarged version to enable us to look at its physics more closely. An optical lattice affords just that possibility: it is a crystal formed by interfering laser beams, with a typical dimension about 1,000 times larger than that of a conventional crystal. Ultracold atoms in the lattice play the part of electrons in the solid; they tunnel quantum-mechanically between lattice sites just as single or paired electrons (Cooper pairs) tunnel through the periodic potential wells created by positive ions in crystalline materials (Fig. 1).

Why not just use a real crystal?

Real solid materials are incredibly complex. They have an involved band structure of allowed energy states, and the Coulomb interactions of electrons are difficult to account

for. Disorder and the inevitable effects of vibrations of the crystal lattice add to the intricacy. It is not desirable, and usually impossible, to take account of all these effects simultaneously. The theorists' approach is to construct highly simplified models that focus on particular aspects of the system. But often even these cannot be solved reliably. That is particularly true of models involving fermions — the group of particles to which electrons, protons, neutrons and many atoms belong. And so a divide has opened up: on the one hand, real materials cannot be sufficiently described by theorists; on the other, experimentalists cannot test the (sometimes contradictory) theoretical results obtained from simple models.

How can optical lattices help?

They provide a way of realizing the simplified models of condensed-matter theory in experimental practice. Optical lattices thus implement Richard Feynman's pioneering idea of 'quantum simulation' — using one quantum system to investigate another. Feynman introduced the concept because a correct numerical description of a quantum system requires resources that scale exponentially

with the number of particles involved — one of the problems confronting theorists today. Using results from optical lattices, theorists want to test which of their models they can best rely on to, say, construct a phase diagram of a condensed-matter system, or assess the evolution of a physical parameter. In addition, optical lattices provide ways to control various factors, such as the strength of interatomic interactions, band structure, spin composition and levels of disorder, more easily than in real crystals — often even dynamically during the course of a single experiment.

Are there limits to what you can simulate with an optical lattice?

We can't realistically use them to simulate very complicated crystals, such as superconducting ceramics, in their entirety. But that's not the goal. To understand phenomena such as high-temperature superconductivity, we want to find the simplest possible system that shows the effect. Consequently, it is not a bad idea to start from simple models and add complexity step by step.

How exactly do you make an optical lattice?

Lattice potentials are created by making two coherent laser beams propagating in opposite directions interfere with each other. The result is a standing wave with a periodic pattern of dark and bright stripes. The light induces an electric dipole moment in the atoms of the ultracold gas, modifying their energy. Depending on the frequency of the light, atoms are pulled towards either the bright or the dark regions and are therefore confined to small areas in space. By using additional lasers from different directions, two- or even three-dimensional lattice structures can be constructed (Fig. 2).

Can you make a lattice for any kind of atom?

Pretty much, although alkali atoms from group I of the periodic table are generally easiest to work with — their single valence electron gives them

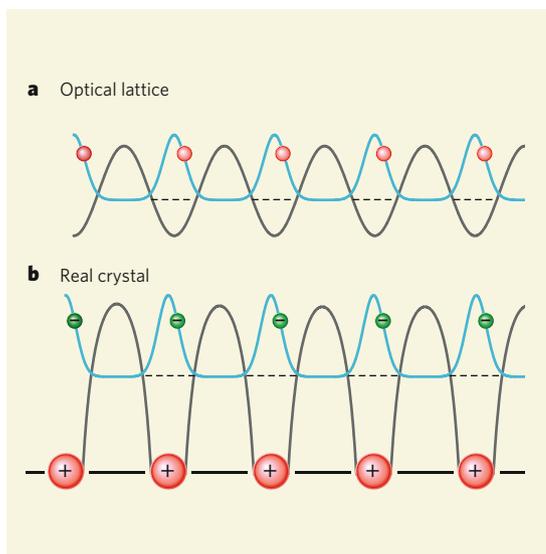


Figure 1 | Crystal simulation.

Ultracold atoms in an optical lattice can simulate condensed-matter phenomena that usually occur only in the 'electron gas' of a solid-state crystal. In an optical lattice (a), atoms are trapped in a sinusoidal potential well (grey) created by a standing-wave laser beam. The atoms' wavefunctions (blue) correspond to those of valence electrons in a real crystal (b). Here, the periodic potential is caused by the attractive electrostatic force between the electrons (–) and the ions (+) forming the crystal. The motion and interaction of the particles, whether ultracold atoms or electrons, determine the physics of the material. Thus, for example, superfluidity in a gas of ultracold atoms corresponds to superconductivity in an electron gas.

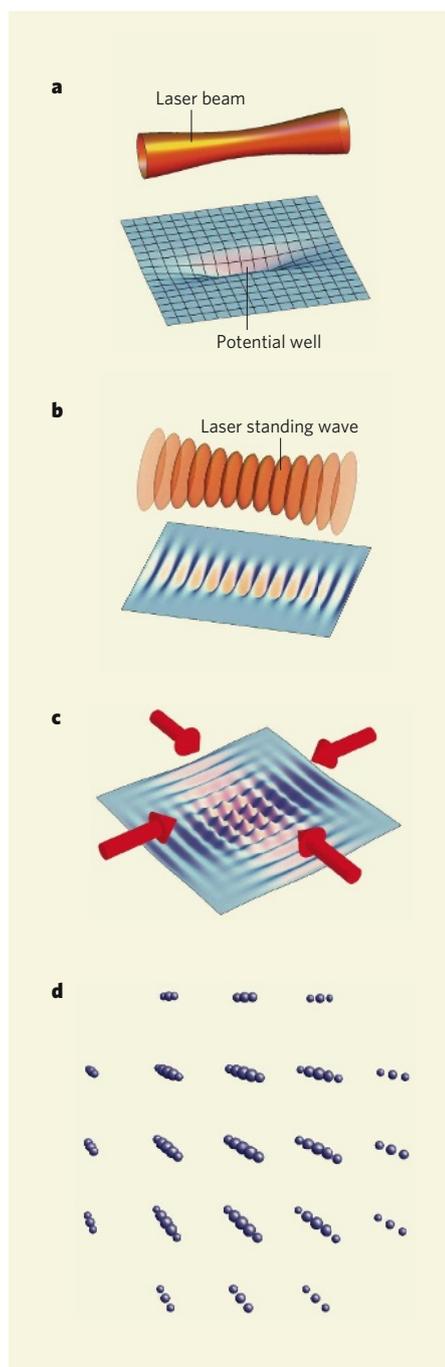


Figure 2 | Potential lattice. Laser light creates the peculiar potential landscapes of optical lattices. **a**, Laser light creates a repulsive or an attractive potential that is proportional to the laser's intensity along its axis of propagation. **b**, By allowing two counter-propagating laser beams to interfere, a sinusoidal standing wave can be formed. Ultracold atoms can be trapped in the potential minima that occur every half-wavelength, separated from the neighbouring minimum by a potential 'wall'; this is the basis of a one-dimensional optical lattice. **c**, **d**, Adding more laser beams at right angles to the first one creates a two-dimensional (**c**) and finally a three-dimensional cubic 'crystal' in which atoms (**d**) are trapped in the minima of the lattice potential. Real optical lattices can have millions of sites, and by changing the angles and wavelength of the laser beams, different lattice geometries can be created.

particularly suitable optical properties. At the very low energy scales (temperatures) at which optical-lattice models operate, most atoms are actually very similar. The most important property is whether they are bosons (with integer spin) or fermions (with half-integer spin). Many bosons can exist in the same quantum state, but quantum mechanics absolutely forbids this for fermions, according to the rule known as the Pauli exclusion principle. Whether a particular atomic isotope is bosonic or fermionic depends on the number of its (fermionic) constituents: protons, neutrons and electrons. If this number is even, the total spin is an integer, and the atom is a boson; if it is odd, the total spin is a half-integer, and the atom is a fermion. It's easier to work with bosons: one consequence of the Pauli exclusion principle is that fermions are tougher to cool. Today, optical-lattice experiments are being pursued with many different atoms — bosonic rubidium-87, sodium-23, potassium-39 and caesium-133; and fermionic potassium-40, lithium-6 and strontium-87.

How do you put the atoms into the lattice?

Usually, by approaching the problem from the other direction: you put the lattice into the atoms. First, you create a Bose–Einstein condensate — an ensemble of bosons all in the same quantum state — or a cold gas of fermionic atoms. Then you slowly ramp up the lasers to create the periodic lattice potential, and the atoms reorder to adapt to their new environment. Similarly, the whole 'crystal' can be removed from the atoms simply by ramping down the lasers, thus liberating the atoms into free space once more.

How did the work with optical lattices start?

The first generation of optical lattices came on the scene in the 1990s. They were mainly used for the laser-cooling of atoms, as the lattice potential increases the efficiency of some methods of optical cooling. Still, the temperatures at that point were too high and the occupation numbers too low (many lattice sites were left unfilled) for many-body quantum physics to feature. Nevertheless, fundamental phenomena such as Bloch oscillations of atoms within the lattice and the Bragg scattering of light on the atoms could already be observed.

When did people start using lattices to investigate many-body physics?

A major boost came with the creation of the first Bose–Einstein condensates in the mid-1990s. This meant the availability of a dense and extremely cold sample of atoms on which a lattice could simply be superimposed. It became possible to study superfluid behaviour in the lattice, emulating for example the physics of arrays of Josephson junctions — tunnel junctions in superconductors that have important applications in quantum-mechanical circuits. But experiments with complex many-body

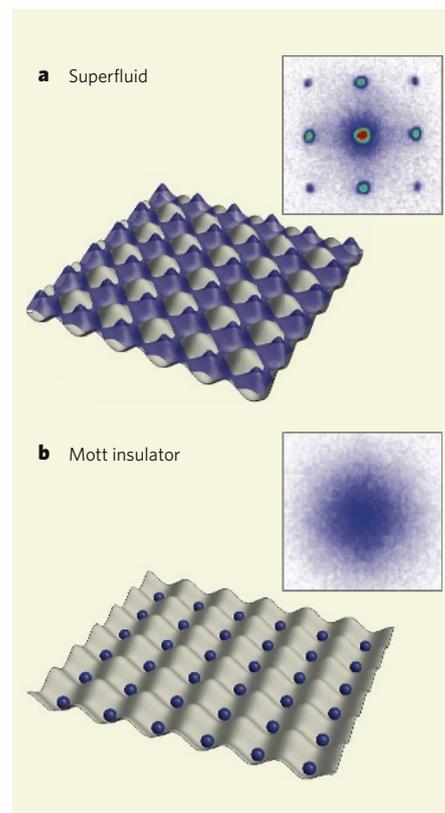


Figure 3 | Superfluids and insulators. In a quantum simulation of a bosonic Hubbard model in an optical lattice, two distinct ground states exist. **a**, For weak interaction strengths the ground state is a superfluid. The atoms have a common macroscopic wavefunction (purple wave), with particles fully delocalized throughout the available space. After the atoms are released from the lattice, the coherent wave nature becomes evident in the fact that a distinct diffraction pattern arises (inset). **b**, If the interactions become very strong, the ground state is a 'Mott insulator' in which each particle becomes localized to one specific lattice site. In that state, the number of atoms in each site is fixed, and phase coherence, as well as superfluidity, disappears.

states, relevant for simulating condensed-matter systems, started only around 2000, first with Bose–Einstein condensates in three-dimensional lattices, and later with ultracold Fermi gases.

What are three-dimensional lattices good for?

This geometry turns out to be a virtually perfect rendering of the Hubbard model, a basic model that describes particles in a crystalline lattice. It takes into account the atoms' mutual repulsion or attraction and the rate at which they hop from site to site. Experiments with bosonic atoms in optical lattices confirmed results arrived at using models, in particular that ultracold bosons undergo a transition between two very distinct quantum states. The first is the Bose–Einstein condensate, a superfluid state in which the gas can move without friction and particles are delocalized over

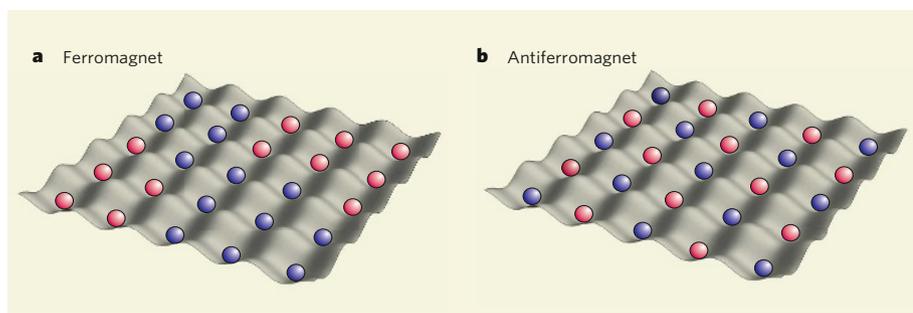


Figure 4 | Magnetic effects. When atoms in two different spin states (purple or red) are simultaneously present in the lattice, they can either favour or disfavour identical spins on neighbouring sites. **a**, In the first case, they distribute in such a way that compact areas of identical spins form — magnetic ‘domains’, which are the hallmark of ferromagnetism in solids. **b**, Conversely, in the antiferromagnetic state, a chessboard pattern of alternating spins develops. The underlying force driving the system into one of these states might be direct, long-range interactions, for example through the magnetic dipole of the atoms. Surprisingly, it can also arise from the tunnelling itself in combination with the peculiar quantum statistical properties of bosons and fermions. This ‘quantum magnetism’ does not require any direct magnetic interaction between the particles, and should be observable in an optical lattice implementing a simple fermionic Hubbard model.

the whole lattice. When the lattice ‘depth’ (the height of the potential barriers) is increased, the role of interactions between particles increases compared with that of the tunnelling. Eventually, the interactions bring the system into what is known as a Mott insulator state, in which free movement becomes impossible and each particle is localized to one lattice site only (Fig. 3).

Can optical lattices go beyond merely confirming theory?

So far, experimenters have mostly studied the Bose–Hubbard model: thanks to numerical simulations, we understand this model well, so it’s easy to compare experimental and numerical results. But there’s a lot of physics in boson interactions that isn’t covered by this ‘plain vanilla’ model, and that is not so well understood theoretically. For example, introducing more than one atomic spin state leads to magnetic effects. As different types of interactions compete, each favouring a different type of magnetic ordering, exciting new states such as ‘spin liquids’ are expected to arise. Then there are fermionic Hubbard models. As mentioned before, fermions behave fundamentally differently from bosons because of the Pauli exclusion principle. To form a superfluid, which requires all the particles to be in the same, coherent quantum state, fermions must ‘pretend’ to be bosons by, for example, teaming up to form Cooper pairs. Several laboratories are working on reproducing the behaviour of fermions in an optical lattice. Promising results, such as the creation of insulating and superfluid behaviours, are already being reported.

Why should we care about the fermionic Hubbard model?

One reason is what it might tell us about high-temperature superconductivity. This puzzling phenomenon — electrical conduction without resistance at temperatures of up to 130 kelvin — occurs in complex cuprate (copper oxide) materials, and has recently been reported in

a second, unrelated class of material. Despite 20 years of effort, there is still no consensus as to what causes either the formation of the ‘d-wave’ Cooper pairs that seem to lie behind the superconductivity, or other anomalous properties of these materials. The Hubbard model could be a good place to start: high-temperature superconducting phases seem always to be found near phases with antiferromagnetic order, in which particles of opposing spin states arrange themselves in a chessboard pattern (Fig. 4). This is a typical ground state of a fermionic Hubbard model.

So the Hubbard model holds the key to high-temperature superconductivity?

There are many reasons to believe that the Hubbard model contains most — but not all — of the ingredients necessary for understanding high-temperature superconductivity in the cuprates. A full explanation is likely to require additional phenomena, such as the interactions between electrons and lattice vibrations (phonons), and inhomogeneities at the mesoscopic level, although there is no consensus about which, if any, of these ingredients are important. But that is just where the remarkable control over ultracold atoms in an optical lattice comes into its own — we expect that we can approach the problem by realizing the bare Hubbard model first and then adding the other ingredients in a controlled way. This really would be a quantum simulator as Feynman first envisaged it.

How might we add to this quantum-simulation tool-box in the future?

A current focus is on extending the technique to further atomic species, each with different interaction properties, as well as on decreasing the temperature of the trapped atoms still further through new cooling techniques. An intense line of research is also to introduce long-range interactions into optical lattices by trapping, not atoms, but molecules with a dipolar

ground state that interact through electric dipole–dipole forces. And then there are efforts to improve how we can analyse the quantum state, for example by directly reading out every single lattice site or by probing characteristic quantum fluctuations of the whole ensemble.

Are there applications for optical lattices beyond condensed-matter simulations?

Because as few as one or two atoms can be isolated in each of the small traps, the traps can also act as miniature ‘test tubes’ in which atomic interactions can be studied with high precision. Even ‘micro-chemistry’ can be performed, with the ability, for example, to precisely control the formation of a molecule from two atoms. Equally, with just one atom in each site, interactions can be completely suppressed. That is seen as a way of improving the precision of optical atomic clocks still further. These create highly accurate time references through the extremely fast oscillations that occur between the energy states of the atoms. Those energy states are disturbed by atom–atom interactions that can be suppressed in an optical lattice. And last but not least, there are intriguing proposals to use atoms trapped at single lattice sites as ‘qubits’ for a so-called universal quantum computer.

How would a universal quantum computer work?

The first step would be to create large-scale entanglement between all atoms in a lattice, by letting each atom interact only with its next neighbour in a way that depends on its state. This technique has already been demonstrated experimentally. In principle, a universal quantum computer can be realized just by reading out and manipulating these single entangled qubits following such an entanglement step. The challenge is to do all the steps — entanglement, manipulation, readout and error correction — with high enough fidelity. We might have to wait years to see a useful universal quantum computer. But in the meantime, quantum simulators in the spirit of Feynman are already a specialized type of quantum computer: simulated condensed-matter models in optical lattices look set to become the first practically useful application of a quantum-computational system. ■

Markus Greiner and Simon Fölling are in the Department of Physics, Harvard University, 17 Oxford Street, Cambridge, Massachusetts 02138, USA.

e-mails: greiner@physics.harvard.edu; foelling@physics.harvard.edu

FURTHER READING

Feynman, R. P. *Int. J. Theor. Phys.* **21**, 467–488 (1982).

Bloch, I. *Nature Phys.* **1**, 23–30 (2005).

Auerbach, A. *Interacting Electrons and Quantum Magnetism* (Springer, Berlin, 2006).

Bloch, I., Dalibard, J. & Zwerger, W. Preprint at <http://arxiv.org/pdf/0704.3011> (2007).