The strong-correlations puzzle

A solution to one of the most famous problems in theoretical physics, formulated almost 50 years ago, may at last be within reach. But as **Jorge Quintanilla** and **Chris Hooley** explain, it relies not on theory, but on experiments with ultracold atoms trapped by beams of light

Jorge Quintanilla is at the ISIS facility of the Science and Technology Facilities Council's Rutherford Appleton Laboratory and at St Catherine's College, Oxford, UK, e-mail jorge.quintanilla@stfc.ac.uk. **Chris Hooley** is in the School of Physics and Astronomy, University of St Andrews, UK, a member of the Scottish Universities Physics Alliance

Quantum matter is everywhere, from the interiors of neutron stars to the electrons in everyday metals. Like ordinary, classical matter, it is made up of many interacting particles. In classical matter, however, it is possible to think of each particle as an individual entity, whereas in quantum matter Heisenberg's uncertainty principle prevents us from telling individual particles apart: their behaviour can only be described collectively. In spite of this, many types of quantum matter are well understood from a theoretical point of view. For example, the "electron liquid" that is responsible for the flow of electricity through ordinary metals, the magnetic properties of many insulating materials and the normal and superfluid phases of helium at very low temperatures have all succumbed to the probing of theorists.

But the behaviour of some forms of quantum matter has proved a much harder nut to crack. High-temperature superconductors, for example, are not really understood despite more than two decades of research since they were first discovered. Also mysterious are various exotic types of magnet: while the electrical resistance of most metals increases with the square of their temperature, T, for some metallic magnets like manganese silicide the resistance is proportional to $T^{1.5}$. And then there is the quark–gluon plasma, which occurs when neutrons are pressed together so tightly that their quarks lose their identity and form a single homogeneous liquid. Such a plasma is believed to have formed during the first few microseconds after the Big Bang, but has also recently been recreated in the Relativistic Heavy Ion Collider (RHIC) at the Brookhaven National Laboratory in the US, with further experiments planned at CERN's Large Hadron Collider.

All these forms of quantum matter have one thing in common: very strong – rather than weak – correlations between the particles from which they are composed. Materials with weak correlations are relatively easy to understand: as the component particles barely interact with each other, one can extrapolate the behaviour of one or two non-interacting particles (like those in an ideal gas) to get a good insight into how they behave *en masse*. Strong correlations, however, lead to qualitatively new behaviour. High-temperature superconductors, for example, display not only an unconventional superconducting phase but also mysterious "bad metal" and "pseudogap" behaviours.

In an ideal world, physicists would build theories to account for these data and then use these theories to suggest what experiments should be done to explore the phenomena further. Unfortunately, however, there is a scarcity of reliable, unbiased theoretical predictions to which the existing treasure-trove of data can be compared. The reason is that even the simplest models of strongly correlated systems have proved difficult to solve theoretically. In a sense, this should not be a surprise: one would naively expect any model of 10^{26} particles to be unsolvable!

We are sometimes saved by the fact that the model consists of lots of independent pieces, such as the (approximately) non-interacting atoms in a gas or the independent normal modes of vibration of a string. Decades of excellent physics have been built from the independent-particle starting point, but it is widely believed that the above phenomena lie beyond its explanatory range. To describe them, we must take interactions between particles into account in a hitherto untried way.

Perhaps the most famous approach to these problems is the "Hubbard model", formulated by the physicist John Hubbard in the 1960s (see box on page 35). In the almost five decades since then, it has become a "standard model" of condensed-matter physics and materials research. So what is Hubbard's model a model of, how did he come up with it, and why is it so important today? To answer these questions, we must trace the story back almost 50 years, to a laboratory in a rural corner of Oxfordshire just after the end of the Second World War.

A magnetic puzzle

After the wartime technological success of the American-led Manhattan Project, Britain was determined to build up its own civilian and military nuclear capabilities. This was crucial for the country's strategic efforts, since after the end of the Second World War the US was no longer sharing all nuclear information with its allies – including Britain and France as well as the Soviet Union. The main British nuclear-research project took place at Harwell in Oxfordshire. Between 1947 and 1968 some 14 nuclear reactors were built on the site.

At a Glance: Strong correlations

- Many advanced materials, including organic conductors and high-temperature superconductors, are examples of "strongly correlated" matter
- Such materials are hard to describe theoretically because strong interactions between particles produce phenomena that cannot be predicted by studying the behaviour of individual particles alone
- The theoretical model widely used to describe strongly correlated matter, known as the Hubbard model, is almost 50 years old. Although the model is a useful tool, it has only been solved exactly for perfectly pure, 1D materials – unlike the messy, 3D systems found in nature
- Experiments on ultracold atoms may provide a new means of testing the 2D or 3D Hubbard model in a highly pure, controllable system – a "quantum simulator" for more complicated materials

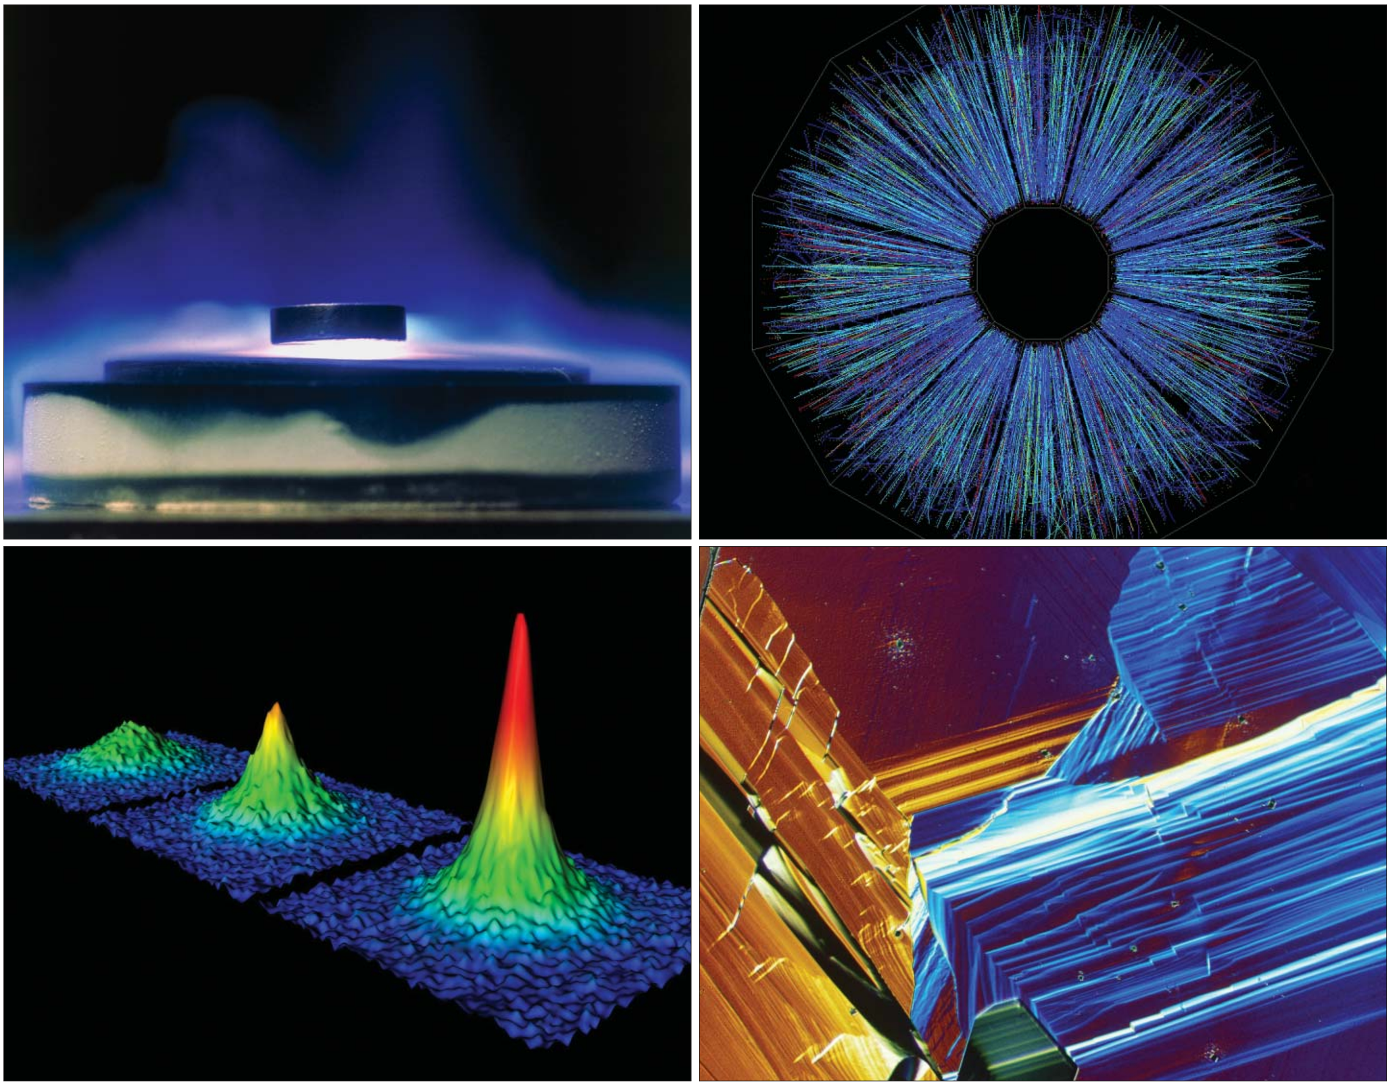
David Parker/IMI/University of Birmingham/Science Photo Library

Brookhaven National Laboratory

JILA

Michael W Davidson/Science Photo Library

Piecing it together Strong correlations between component particles are the common thread linking a number of fascinating physical systems, including (clockwise from top left) high-temperature superconductors, quark–gluon plasmas, organic superconductors and – under some conditions – clouds of ultracold atoms.

Such reactors produce neutrons as a fission by-product, and these subatomic particles carry part of the energy of the reaction to the outside world. One of the main lines of enquiry to be pursued using this new hardware (particularly the reactors named Dido and Pluto) was the study of how neutrons scatter when passing through different materials – information that scientists and engineers needed to make optimal choices of materials in reactor designs.

But in neutron scattering, the researchers got more than they had bargained for: the scattering of neutrons through materials turned out to be an excellent probe of the goings-on inside the materials themselves. Indeed, neutron scattering became, and remains to this day, one of the major methods of investigating quantum matter. True, the ways of producing neutrons are changing, with increased emphasis on spallation, which uses a particle accelerator instead of a fission reactor. (ISIS, the current neutron source at the Harwell site, is one of the world's most powerful pulsed spallation sources.) Nevertheless, the basic principles have remained the same: neutron-scattering experiments allow researchers to study events inside a material on the atomic scale. It was experiments of this type that suggested the Hubbard model in the first place.

The neutron carries no electric charge, but it has its own magnetic moment, or "spin", like a tiny bar magnet. Atoms also have spin, and bombarding magnetic materials with neutrons excites the spins of these materials' constituent atoms. One of the simplest types of excitation is called a "spin wave". It resembles a Mexican wave produced by a crowd of football fans, where each football fan represents an atom and the raising of hands corresponds to the atom's spin tilting off its original axis. Such spin waves can have a wide range of wavelengths, from 0.1 nm (10^{-10} m) up to centimetres – the size of the entire sample. Typically, however, longer-wavelength modes are more common at lower temperatures and lower incident-neutron energies.

It is easy to imagine that a Mexican wave formed by football fans in an array of seats ought to look different from one performed by a crowd of figure skaters moving on ice. Similarly, it might seem sensible to assume that the magnetic properties of insulators should be rather different from those of metals: in insulators, the spins have fixed positions, whereas in metals they are free to move about the material. Yet the pioneering experiments that were carried out at

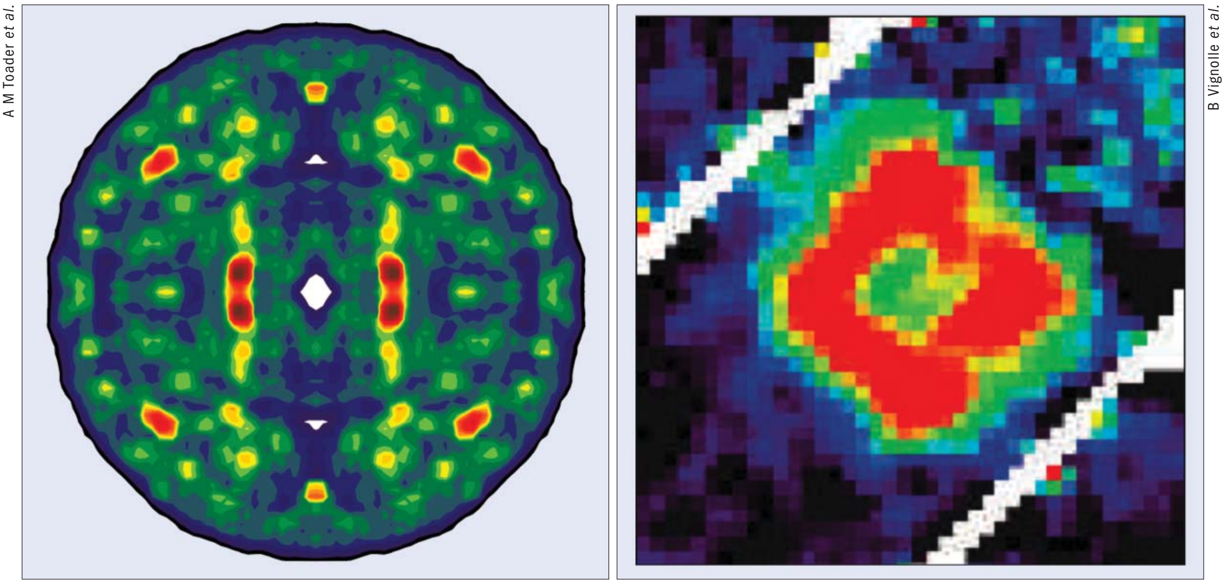

A M Toader *et al.*

B Vignolle *et al.*

Magnetic maps Many of today's most advanced electronic materials fall into the category of strongly correlated materials, including organic conductors, cuprate superconductors, the colossal-magnetoresistance materials and the "heavy electron" compounds – so-called because their quasiparticles behave very much like electrons, except with several hundred times the electron mass. The figure shows magnetic neutron-scattering patterns (under different conditions) from the insulator La_2CuO_4 (left) and the strongly correlated material obtained by doping it with strontium, which is a high-temperature superconductor (right). Our understanding of such strongly correlated materials could be greatly improved by a general solution of the Hubbard model.

Harwell in the late 1950s and early 1960s revealed that the spin-wave behaviour in some metals was the same as in ordinary insulating magnets. In particular, a series of papers by Ray Lowde and Roger Elliott established that the scattering of neutrons through iron was best described by a model designed with magnetic insulators in mind – even though iron is a metal. This was a puzzle that required an explanation.

The "simplest possible" model

Thanks to the work of the future Nobel-prize-winning physicist Nevill Mott and others, the inspiration behind what later became the Hubbard model was already in the air. The basic idea was that the models that theorists were using to think about metals, which featured a fluid of electrons, and magnetic insulators, which included magnetic moments at fixed positions, were actually two extremes of the same thing. It was Hubbard, then a shy and modest young theorist working at Harwell in the 1960s, who made this notion mathematically concrete. As he later wrote, in formulating the Hubbard Hamiltonian he "set up the simplest possible model containing the necessary ingredients" to describe the behaviour of correlated materials.

The Hubbard model is indeed extremely simple. It features electrons that can hop between "sites", representing atoms, that are arranged in an ordered, crystalline pattern. Crucially, when two electrons are on the same site, they have to pay an energy penalty due to their mutual repulsion. This introduces additional correlations between the electrons beyond those due to the Pauli exclusion principle, which are always present. If the electron–electron repulsion is weak, then the electrons roam more or less freely around the material and it behaves like a metal. But if it is strong, they are forced to localize at fixed atomic positions and the material behaves like a magnetic insulator.

The models used to think about metals and magnetic insulators were two extremes of the same thing

Theorists were quick to appreciate the simplicity and power of the Hubbard model. Within a few years Elliott Lieb and Fa-Yueh Wu, working at Northeastern University in Boston in the US, had solved it exactly in the 1D case. What they obtained is the "phase diagram" of the model: a sort of chart, allowing us to predict the behaviour of the model (metal, insulator, etc) given the density of electrons and the strength of their mutual repulsion. The results were striking. At "half filling" (in other words, when there is one electron per site), the model predicts that a Mott insulator – a log-jam where each electron is hemmed in by its neighbours – will form. When the Mott insulator is doped, by either adding or removing a small number of electrons, it becomes metallic.

But this is no ordinary metal, where charge and spin are carried together by the constituent electrons. In a doped Mott insulator, the charge and spin are carried entirely separately, by two different forms of excitation called, respectively, "holons" and "spinons". Although these holons and spinons are in reality complicated collective modes of all the electrons in the system, they behave almost as if they are particles in their own right: they are quasi-particles. This strange metal has become a model case for studying strongly correlated quantum matter, precisely because the entities that best describe its behaviour (the quasi-particle holons and spinons) have little to do with the original particles that make up the system on the microscopic scale.

Unfortunately, few real materials are 1D, and none are perfectly pure. We do not know how to generalize the process by which electrons transform into spinons and holons – as seen in the 1D Hubbard model – to higher dimensions. Furthermore, since doping the system introduces disorder in the lattice, we do not know which of the effects observed are due to this disorder and which would occur anyway in a (mythical) perfectly

John Hubbard: the man behind the model

Those who knew John Hubbard describe him as a very shy man – to the point that others, who did not know him so well, may have perceived him as somewhat aloof. Born on 27 October 1931, Hubbard was educated first at Hampton Grammar school and then at Imperial College, London, where he obtained his PhD in 1958 under Stanley Raimes. Unusually for his time and social context, he lived with his parents in Teddington throughout his university education.

At the end of his PhD, Hubbard was recruited to the Atomic Energy Research Establishment in Harwell, Oxfordshire, by Brian Flowers, who was then heading the theory division. An anecdote from this period of Hubbard's career illustrates his retiring personality. While at Imperial, Hubbard had dealt with the project assigned to him for his PhD fairly quickly, and had then looked for a more challenging problem. At the time, quantum-field-theory methods, particularly Feynman diagrams, were being applied to problems in many-body theory. However, it was difficult to bring the same methods to bear on the many-electron problem – relevant to solid-state systems – because the Coulomb interaction between electrons made quantities like the total energy diverge.

Hubbard realized that these divergences could be controlled: the trick was to sum up an infinite series of a particular class of Feynman diagrams. When Hubbard arrived in Harwell, he mentioned this to Flowers, who wanted to see the paper. Alas, there was no paper, Hubbard explained, because when he was about to write it up he saw an article by other researchers who had introduced a different method to solve the same problem. Hubbard had found their method physically appealing, checked privately that their results coincided with his, and concluded there was no need for an additional publication on the topic. Flowers then issued an explicit order that Hubbard should publish his groundbreaking work.

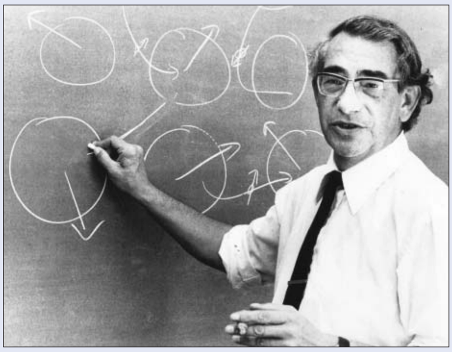
AIP Emilio Segrè Visual Archives

Hubbard's most famous papers are the series he wrote on his eponymous model, starting in 1963. He was not the only one working on the strong-correlations problem: some months earlier, Takeo Izuyama, working at Nagoya University, and Duk-Joo Kim and Ryogo Kubo, at the University of Tokyo, both in Japan, had argued that a proper description of correlations in metals with strong electron–electron interactions could explain the observed spin-wave spectrum. Martin Gutzwiller, who was then working at IBM's research laboratories in Zürich, had also produced essentially the same model. Yet it was Hubbard's calculations that showed that the model that now bears his name could in fact describe both the metallic and insulating behaviour as two extremes of the same thing. His application of a Green's function technique to the model was a template for many other works in condensed-matter theory, and his papers from that time contain many crucial insights, such as the existence of so-called Hubbard bands that are a main feature of our current understanding of Mott insulators.

Eventually, Hubbard became the leader of the solid-state theory group at Harwell, and Walter Marshall succeeded Flowers as head of the theory division. Unlike the shy Hubbard, Marshall, who was also an excellent theorist, was very proactive in hunting for personnel and for funding. This was a blessing in disguise for Hubbard, as Marshall ignored Hubbard's reticence completely and kept "parachuting" postdocs into his group.

Hubbard left the UK for the US in 1976, following Marshall's promotion to director of the Atomic Energy Research Establishment and a subsequent major reform of its facilities in Harwell. He joined Brown University and the IBM Laboratories in San José, California, where his research focused on the study of critical phenomena: phase transitions near which universal behaviour, independent of material-specific properties, is observed. He died, aged just 49, in San José on 27 November 1980. (Main source: Stephen Lovesey, private communications)

clean system. The problem is particularly acute in the case of high-temperature superconductors, where some researchers claim that disorder effects are masking phase transitions that are crucial to understanding the origin of the superconducting behaviour. This double difficulty – not being able to make the samples perfectly clean, and not knowing theoretically what they would do if they were – is at the core of much of modern condensed-matter physics.

Enter cold atoms

The task of predicting the behaviour of the Hubbard model in two or three dimensions is daunting. The 1D case is special because in order for two electrons to pass each other, they must actually pass through each other. This simplifies the problem in the same way that queuing simplifies the post office: it allows theorists (or post-office staff) to deal with one interaction event at a time. This simplification allows theorists to formulate a very large number of conservation laws, and the solution of the 1D problem is built on these.

The 2D problem is qualitatively different. It has so far resisted exact solution, and the approximations that theorists are forced to make to "solve" it are quite crude. For example, a common approach is to assume that "fast" electrons in the material are moving through a "slow" magnetic background. But in the Hubbard model there is really only one intrinsic timescale, so it is difficult to justify these techniques. It is also far from clear that these methods are sufficient to capture the essential physics. Conventional computer simulations also face formidable obstacles, as the complexity of the problem grows very quickly with the size of the system. In practice, only a few lattice sites containing a handful of particles can be simulated directly; even with the fastest supercomputers, the full Hubbard model (without approximations) can only be simulated in simple systems like 16 atoms arranged in a 4 × 4 lattice.

But help may be at hand from an unlikely quarter: atomic physics. Ultracold atoms trapped in crossed laser beams (an "optical lattice") can, under certain circumstances, also be described by the Hubbard model. In such cold-atom systems, atoms play the role of electrons, and the optical lattice supplies the periodic potential in which they move – an "artificial crystal of light", as atomic physicist Immanuel Bloch of the University of Mainz in Germany described it in *Physics World* (April 2004 pp25–29). The same quantum-mechanical rules that govern electrons in a metal also apply to the atoms in the "crystal". This means that these atomic systems could in principle be used as a kind of analogue computer to examine the behaviour of

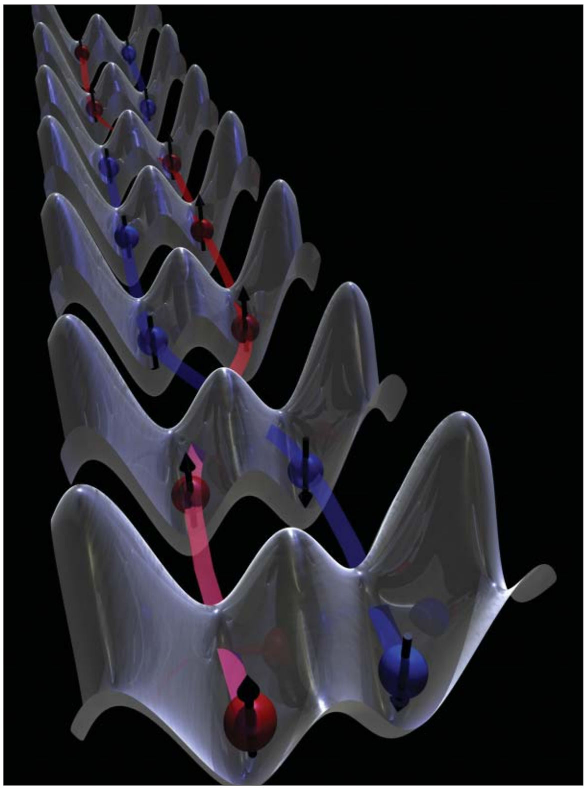

Controlled interactions Experiments with ultracold atoms trapped in optical lattices hold great promise for simulating the way particles behave in more complicated strongly correlated systems, including high-temperature superconductors.

Immanuel Bloch, University of Mainz

the Hubbard model in two or more dimensions. Moreover, by tweaking the experimental set-up, other more complex models can also be realized, for example the Falicov–Kimball model, which describes the behaviour of two types of particle – one heavy, one light – moving around on a lattice.

Ultimately, this process of quantum simulation could help theorists solve many outstanding problems in materials and condensed-matter physics. In order to exploit this possibility fully, however, some key differences between the cold-atom systems and the models of materials need to be addressed.

First, the electrons in real materials are fermions (the type of quantum particles that have half-integer spin and that obey the Pauli exclusion principle), while most cold-atom experiments to date have been performed with bosons (the other type of quantum particle, with integer spin and the "anti-Pauli" tendency to clump together). Researchers prefer bosons because the final stage of cooling in a cold-atom experiment requires the atoms to rethermalize by colliding with each other, and such collisions are more likely to take place between bosons than between fermions.

The second difference is that experiments on materials are normally carried out at constant temperature. This is possible because the experiments are strongly coupled to their environment – for example, the materials may be in contact with liquid helium and thus kept at a constant temperature of 4 K. Cold-atom systems, in contrast, must be isolated from their environment because their quantum nature only becomes apparent at temperatures as low as a few billionths of a degree above absolute zero. This isolation means that adding energy to the system – for example by varying the depth of the optical lattice to explore the atoms' behaviour – tends to heat the atoms, since the additional energy cannot escape. The consequences of the atoms' isolation are therefore twofold: first, the experiments are carried out at constant entropy rather than constant temperature; and second, the low temperatures required for interesting quantum behaviour to occur can be difficult to reach.

Finally, the electromagnetic fields used to trap the atoms inside the optical lattice create forces on the atoms that vary strongly with position, whereas the forces on particles within a material are homogeneous above the atomic scale.

Yet in spite of these difficulties, progress is being made at a rapid pace. The problem of cooling fermions to the necessary temperature was cracked a decade ago by Brian de Marco and Deborah Jin at JILA in Boulder, Colorado. More recently, independent experiments led by Tilman Esslinger at ETH Zürich and Bloch in Mainz have managed to load an ultracold Fermi gas into an optical lattice, and to combine this with a magnetic field that induces repulsive interactions between the atoms. The result is a very good realization of the 3D Hubbard model. These first experiments have revealed tantalizing evidence of the metal–insulator transition, a sudden change of the system's state whereby the atoms (standing in for electrons) localize on particular lattice sites as a consequence of their mutual repulsion. This is the 3D version of the Mott insulating behaviour that Lieb and Wu found in their solution of the 1D Hubbard model (see above). Finally, several theoretical proposals have been made that may help overcome the problems associated with temperature as well as the inhomogeneous electromagnetic field.

The crucial feature of these atom-based experiments is that they are not only extremely tuneable (experimentalists can change the lattice in which the atoms move, as well as the strength of their mutual interactions) but they are also extremely clean, in that they are devoid of lattice imperfections and impurities. Thus, cold-atom simulations promise to show us unambiguously how the pure Hubbard model behaves.

This may not, of course, constitute a solution to the strong-correlation problem in real materials such as high-temperature superconductors, which are neither metals nor insulators. It is not even clear whether the Hubbard model is appropriate for the case of high-temperature superconductors – and even if it is, the Hubbard model will certainly not be applicable for many other materials of interest. Nonetheless, knowing with certainty the long-sought-after phase diagrams of the 2D and 3D Hubbard models will be a vital step in disentangling the multiplicity of effects at play in real materials.

Simple enough?

The Hubbard model remains the simplest non-trivial model of strong correlations, which are now being observed in cold-atom systems as well as in conventional materials. Indeed, given its simplicity, readers might feel concerned at the length of time it is taking researchers to capture its behaviour. But in truth this should be no surprise. Its essential feature is *feedback*: the inevitability that fluctuations in the system will interact with each other. Such feedback is also the distinguishing feature of systems that are traditionally considered much more complicated, including the

The Hubbard model remains the simplest non-trivial model of strong correlations, which are now being observed in cold-atom systems as well as in conventional materials

human brain. As the artificial-intelligence pioneer Ken Hill famously said, "If the brain were simple enough for us to understand it, we'd be so simple we couldn't"; the same principle holds true for strong correlations. However, with the advent of cold-atom experiments, we have moved a big step closer to having a proper grasp of this intriguing branch of modern physics.

After that step, of course, will come another: whatever is discovered in the first generation of optical-lattice simulations will be fed back to the theory of real materials. For example, if cold-atom systems behaving in accordance with the 2D Hubbard model become superconducting at low enough temperatures, as the Nobel laureate Philip Anderson of Princeton University has argued, this will have direct implications for theories of high-temperature superconductors. This in turn will provoke new experiments on those materials, including new neutron-scattering studies, and so on. In the end, one can imagine the very distinction between atomic and condensed-matter physics beginning to blur – a symbiosis that might be only the first step in the emergence of a new field of study, stretching all the way from astrophysics down to subatomic particles: quantum matter.

More about: Strong correlations

V L Campo Jr *et al.* 2007 Quantitative determination of the Hubbard model phase diagram from optical lattice experiments by two-parameter scaling *Phys. Rev. Lett.* **99** 240403
M Cramer *et al.* 2008 Do mixtures of bosonic and fermionic atoms adiabatically heat up in optical lattices? *Phys. Rev. Lett.* **100** 140409
B De Marco and D S Jin 1999 Onset of Fermi degeneracy in a trapped atomic gas *Science* **285** 1703
T-L Ho and Q Zhou 2007 Obtaining phase diagram and thermodynamic quantities of bulk systems from the densities of trapped gases *Phys. Rev. Lett.* **99** 120404
J Hubbard 1963 Electron correlations in narrow energy bands *Proc. R. Soc. A* **276** 238–257
R Jördens *et al.* 2008 A Mott insulator of fermionic atoms in an optical lattice *Nature* **455** 204
U Schneider *et al.* 2008 Metallic and insulating phases of repulsively interacting fermions in a 3D optical lattice *Science* **322** 1520–1525
A M Toader *et al.* 2005 Spin correlations in the paramagnetic phase and ring exchange in La_2CuO_4 *Phys. Rev. Lett.* **94** 197202
B Vignolle *et al.* 2007 Two energy scales in the spin excitations of the high-temperature superconductor $La_{2-x}Sr_xCuO_4$ *Nature Physics* **3** 163–167