Quantum computing with neutral atoms

David S. Weiss and Mark Saffman

Citation: Physics Today **70**, 7, 44 (2017); doi: 10.1063/PT.3.3626
View online: http://dx.doi.org/10.1063/PT.3.3626
View Table of Contents: http://physicstoday.scitation.org/toc/pto/70/7
Published by the American Institute of Physics

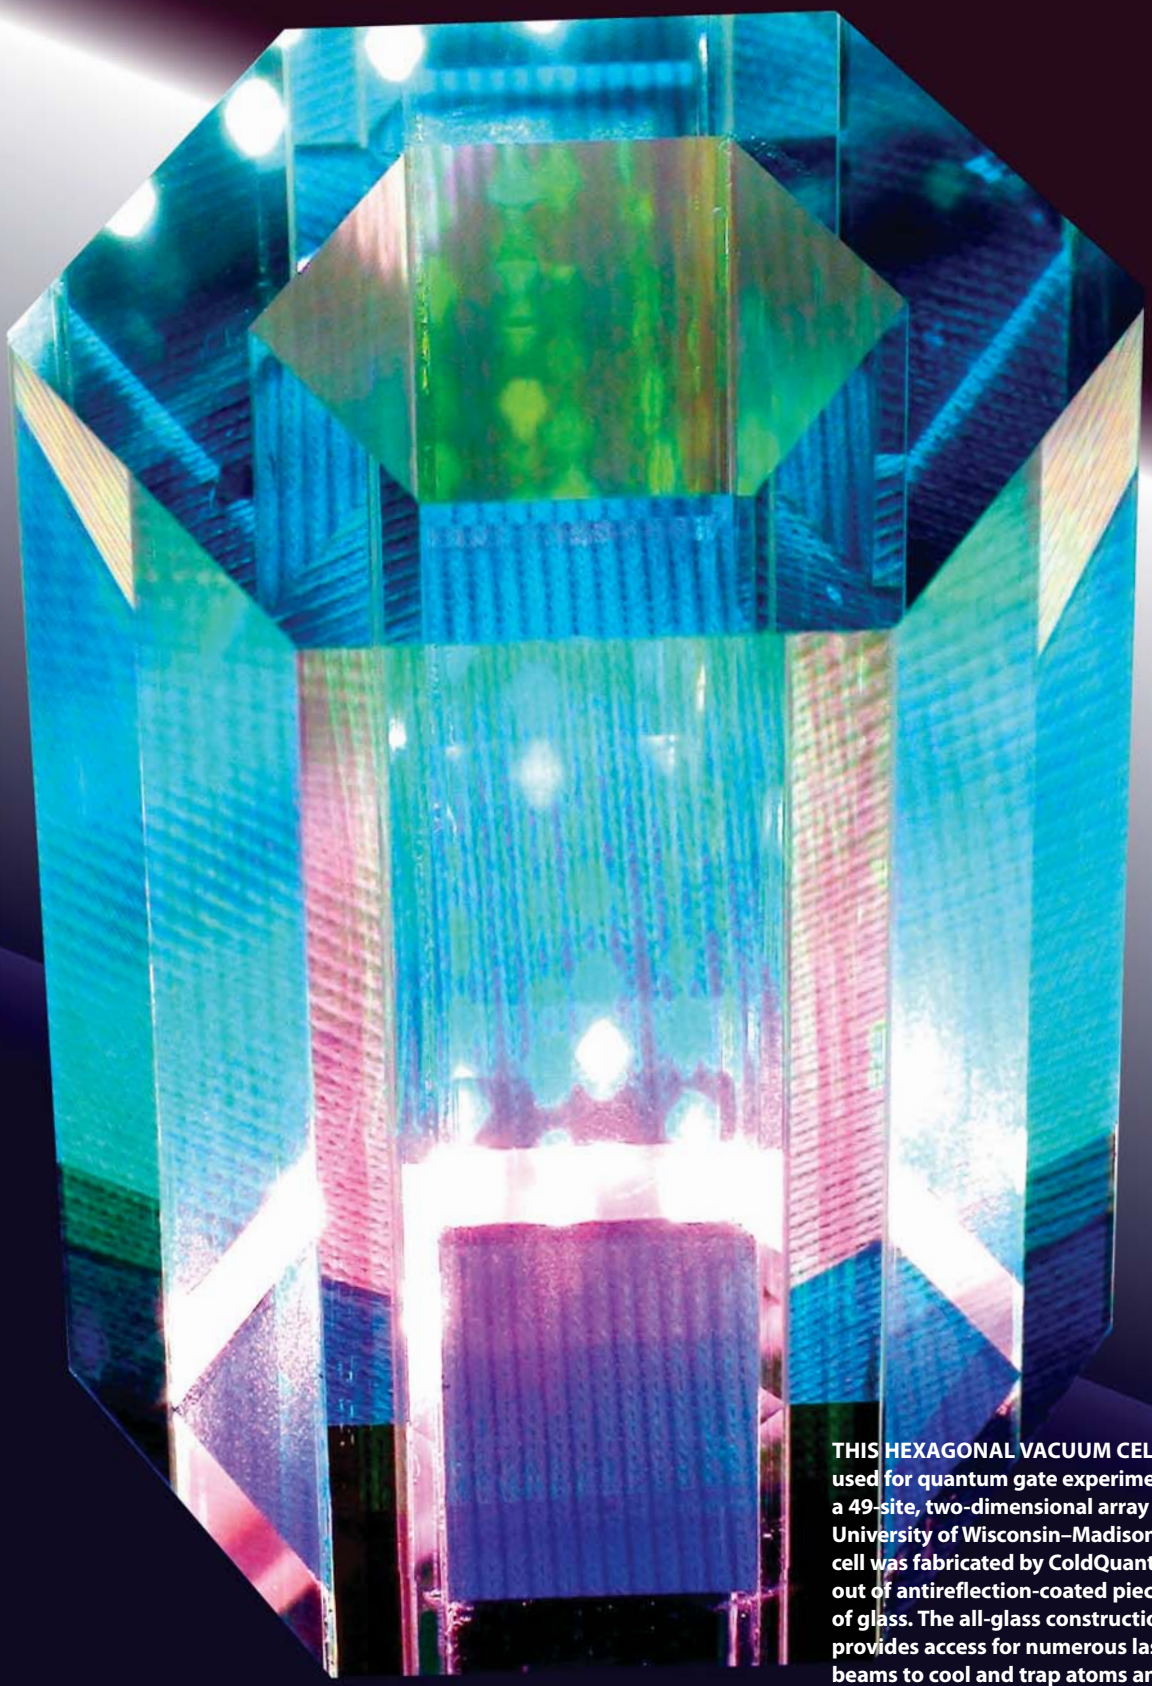

THIS HEXAGONAL VACUUM CELL was used for quantum gate experiments in a 49-site, two-dimensional array at the University of Wisconsin–Madison. The cell was fabricated by ColdQuanta Inc out of antireflection-coated pieces of glass. The all-glass construction provides access for numerous laser beams to cool and trap atoms and to control an array of qubits.

David Weiss is a professor of physics and associate head for research at the Pennsylvania State University in University Park, and **Mark Saffman** is a professor of physics at the University of Wisconsin–Madison.

QUANTUM COMPUTING with neutral atoms

David S. Weiss
and Mark Saffman

With their hyperfine states serving as two-level qubits, atoms can be packed into closely spaced, laser-cooled arrays and be individually addressed using laser pulses.

In 1982 Richard Feynman conceived of a "quantum mechanical computer."[1] His central idea was that a quantum device could take advantage of quantum entanglement and superposition to make calculations that are impossible on a classical computer. Two descendants have emerged from that conception: quantum simulation and quantum computing.

In a quantum simulator, one quantum system is used to model the behavior of a different quantum system. For the past 15 years, quantum simulators have taken advantage of the simplicity and controllability of ultracold atoms and ions to study a wide range of many-body physics, including aspects of Hubbard models, quantum magnetism, superconductivity, and the solutions to many other quantum models, both previously conceived and newly invented. A quantum computer is conceptually akin to a digital classical computer, with classical bits replaced by quantum bits, called qubits, that can exist in superpositions of states and can be mutually entangled. The requirements for a quantum computer are more demanding than those for a quantum simulator, so in contrast to quantum simulation, experimental quantum computing is still in its infancy.

Quantum computing seeks to solve numerical problems using a sequence of logic gates, which can operate on either one or two qubits at a time and change their states. Although any calculation that can be performed on a classical digital computer could also be performed on a quantum computer, doing so would be foolish for most problems. It is much harder to manipulate and measure qubits than it is bits. But hard computational problems exist for which no efficient classical algorithms are known. Quantum computing took off as a field of study in 1994 when applied mathematician Peter Shor (then at Bell Labs) showed that a quantum computer could be used to efficiently factor large numbers, a problem whose classical intractability is behind most modern cryptography schemes. Dozens of quantum algorithms have since been developed, with applications to database searching, pattern classification, multivariate optimization, and solving large systems of coupled equations.

Within a year of the promulgation of Shor's factoring algorithm, scientists demonstrated the first quantum gate at NIST using two trapped ions as qubits.[2] The necessary attributes of qubits are well understood. Qubits must be well isolated from the environment to prevent decoherence of their fragile quantum state. They must be prepared with high fidelity—that is, the actual state must closely resemble the intended state. And they must be accurately measured. Finally, to reliably implement any quantum code—that is, to make a universal quantum

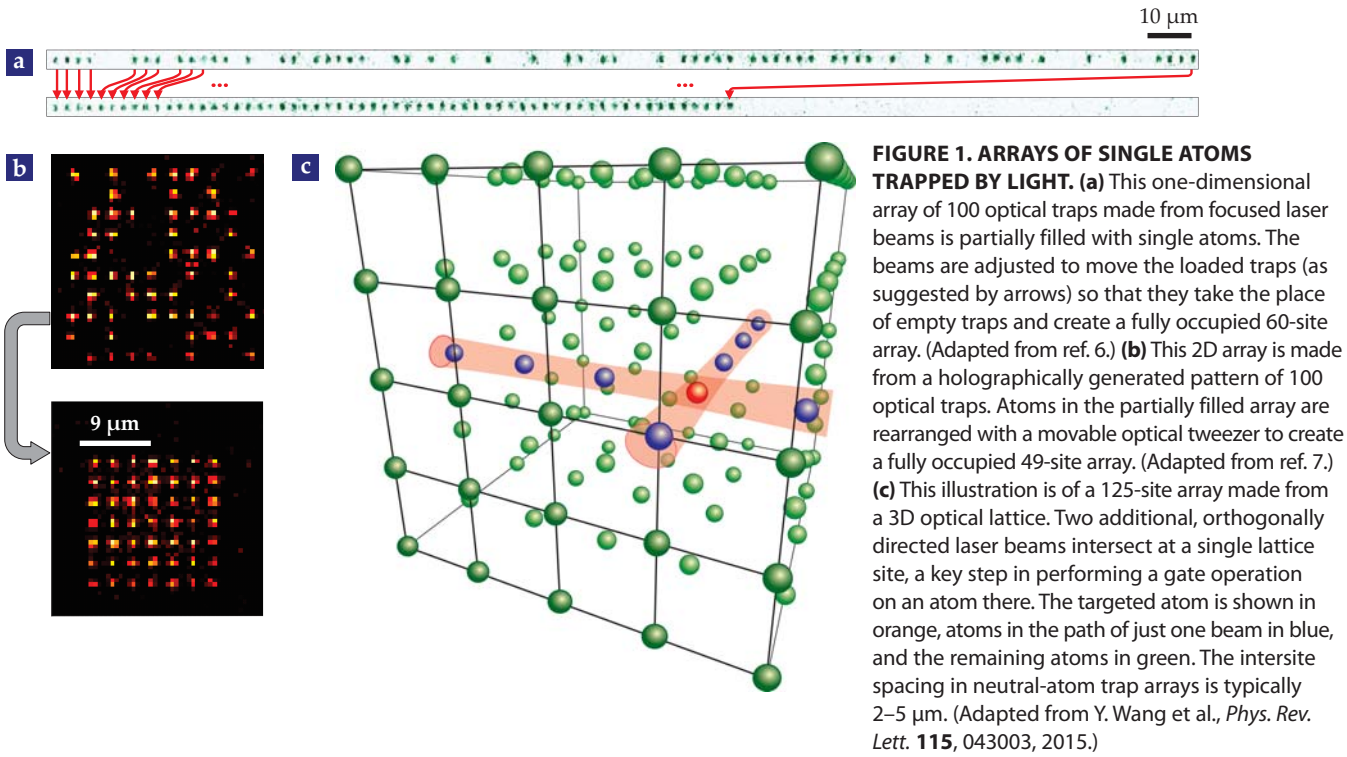

FIGURE 1. ARRAYS OF SINGLE ATOMS TRAPPED BY LIGHT. (a) This one-dimensional array of 100 optical traps made from focused laser beams is partially filled with single atoms. The beams are adjusted to move the loaded traps (as suggested by arrows) so that they take the place of empty traps and create a fully occupied 60-site array. (Adapted from ref. 6.) **(b)** This 2D array is made from a holographically generated pattern of 100 optical traps. Atoms in the partially filled array are rearranged with a movable optical tweezer to create a fully occupied 49-site array. (Adapted from ref. 7.) **(c)** This illustration is of a 125-site array made from a 3D optical lattice. Two additional, orthogonally directed laser beams intersect at a single lattice site, a key step in performing a gate operation on an atom there. The targeted atom is shown in orange, atoms in the path of just one beam in blue, and the remaining atoms in green. The intersite spacing in neutral-atom trap arrays is typically 2–5 μm. (Adapted from Y. Wang et al., *Phys. Rev. Lett.* **115**, 043003, 2015.)

computer—quantum error correction must be incorporated into qubit memory and gate operations (see PHYSICS TODAY, February 2005, page 19).

For the past 20 years, researchers have been racing to build systems of entangled qubits to meet those challenges. Despite remarkable experimental developments, no one has yet built a quantum computer that can perform a calculation that cannot be simulated on a classical computer. The achievement would take about 50 qubits and about 10^4 gates. And even that would fall far short of what is needed to factor a classically intractable number. Nonetheless, a handful of physical platforms have emerged as prime contenders for quantum-computing hardware. They include trapped ions, superconductors, neutral atoms, photons, quantum dots, and spins in solid-state hosts.[3] (See, for example, the article by Ignacio Cirac and Peter Zoller, PHYSICS TODAY, March 2004, page 38.)

Each of those platforms has proven capable of basic quantum-logic operations, albeit with varying degrees of fidelity. Recent experiments with about 10 qubits in ion and superconducting systems, for example, have demonstrated multiparticle entanglement, small-scale quantum algorithms, and partial quantum error correction. Progress across the board has been such that not only are governments worldwide investing heavily in quantum science and technology, but private corporations have taken notice. R&D is under way at major information technology companies, including Microsoft, Google, Intel, and IBM, and at several smaller companies. (One company, D-Wave, sells devices specifically marketed as quantum computers. The degree to which the devices use quantum entanglement to solve problems is debatable, but they are clearly not universal quantum computers because they do not allow for complete control of the quantum state of individual qubits.)

Although all aspects of qubit manipulation need improvement, scaling up to larger numbers of qubits is a central challenge. Unlike in classical computers, whose bit overhead for error correction is a few tens of percent, error correction schemes in quantum computers require that each logical qubit be encoded in an ensemble of as many as 100 physical qubits. It is beyond the scope of this article—and our threshold for controversy—to compare the many ideas for scaling in each qubit platform. Instead, we describe the state of the art of neutral-atom quantum computing and try to offer a sense of why we think that platform is especially promising from the perspective of scalability.

As qubits, neutral atoms boast several attractive features. They are all identical and can readily be prepared by optical pumping in well-defined initial states. Their qubit states can be precisely measured using fluorescence. And in some cases they can be well isolated from the environment, which allows for long decoherence times; last year, more than seven seconds was demonstrated for an array of single atoms.[4] The qubit states can also be rapidly and accurately controlled with electromagnetic fields. Most notably, many atoms can be trapped in close proximity without affecting each other's quantum states unless they are called on to do so.

Trapping single atoms

Several research groups trap neutral atoms using either magnetic fields or light, but light traps have received the most attention for quantum computing. Atoms are polarizable, and the oscillating electric field of a light beam induces an oscillating electric dipole moment in the atom. The associated energy shift in an atom from the induced dipole, averaged over a light-oscillation period, is called the AC Stark shift.

When the light frequency is detuned from an atomic resonance, little spontaneous emission occurs and the light creates a conservative potential for the atoms. Atoms are attracted to

BOX 1. ADDRESSING ATOMS WITH TWO SINGLE-QUBIT GATES

Just as there are several ways to trap single atoms, there are several ways to independently target, or address, them. One technique, recently demonstrated at the Pennsylvania State University, faces head-on the problem of cross talk among neighboring atoms by performing single-qubit gates on individual cesium atoms in a crowded, three-dimensional array. The technique was inspired by nuclear magnetic resonance (NMR), which manipulates spins in ensembles of molecules, and by magnetic resonance imaging, which yields spatial resolution for NMR. But important differences distinguish the new technique: Its spatial resolution comes from light-intensity gradients instead of magnetic-field gradients, and the targets are individual atoms instead of many atoms.

Two laser beams with small waists address an atom by intersecting each other at its lattice site, as shown in figure 1c of the main text. The target atom (orange) experiences a shift in its resonance frequency by the same (AC Stark) effect that creates the potential that traps the atom in the site. The atom's resonance frequency is shifted away from that of neighboring spectator atoms (green) by twice as much as that of so-called line atoms (blue) in the path of only one beam. Here, panel a shows the fraction of atoms (color coded as in figure 1c) that make a transition out of the qubit basis—consisting of two Cs hyperfine sublevels—as a function of microwave frequency.[4]

Logic gates are performed using a combination of optical beams and microwave pulses. Unlike the laser beams, the microwave pulses are spatially broad; all atoms see the same microwave intensity. A microwave pulse with a frequency between the target and line atoms' resonances is applied to the array. At that frequency, none of the atoms leave the qubit basis, but the off-resonant pulse shifts the phase of the qubit states. The sign of the phase shift depends on whether the microwaves are above or below resonance. The target atom thus gets a phase shift with an opposite sign to that imparted to all other atoms.

To avoid cross talk, two atoms in two different planes are addressed in a four-stage sequence, shown in panel b. In the top line, the black pulses represent the global microwave pulses resonant with the qubit transition, and the purple pulses produce the phase shifts. The addressing beams are turned on and off in the pattern shown on the second line. The third and fourth lines show the path of the addressing beams in two planes. Two of the stages target an atom and two are dummy stages, with pairs of addressing beams that do not cross. Between each stage, a black pulse flips all the qubit states.[4]

Thus all untargeted atoms experience each of two different beam configurations twice, with an odd number of spin flips in between. The sum of all the phase shifts experienced by the untargeted atoms is zero, so their stored quantum information is unchanged by the sequence of stages. The two targeted atoms, in contrast, experience an additional sign change associated with their change from target atom to line atom. As a result, they get a net phase shift, the magnitude of which can be controlled by the intensity of the microwaves. Using additional global microwave pulses, one can turn those site-dependent phase shifts into any single-qubit gate of interest.

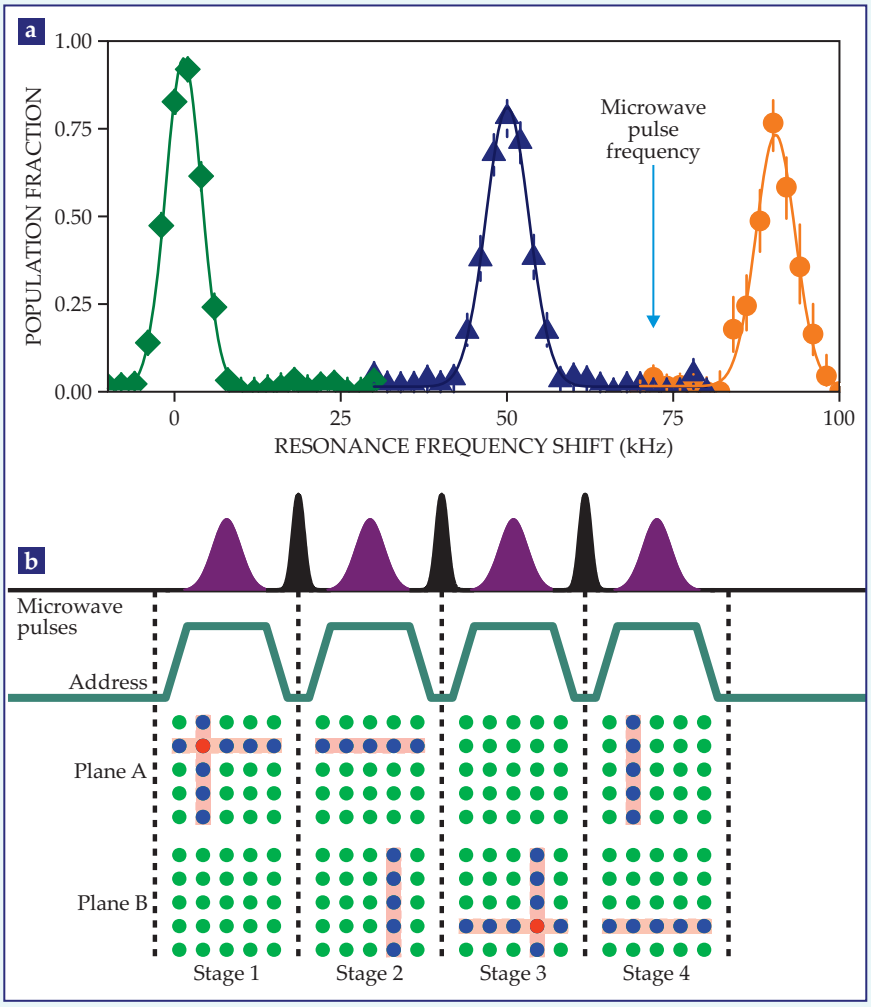

light below the resonance frequency (red detuned) and repelled by light above it (blue detuned). The AC Stark shift is proportional to the light's intensity. Thus the shape of the intensity field is the shape of an associated atom trap.

The simplest light trap, called an optical tweezer, is a red-detuned laser beam that holds atoms at its focus. Arrays of optical tweezers have been created using both conventional and holographic optics. They can also be made in the form of an optical lattice, whose sites are the nodes or antinodes of a standing wave produced by coherent light interfering with itself. To date, one-, two-, and three-dimensional arrays, such as the ones shown in figure 1, have each trapped up to 50 individual atoms in apparatuses designed for quantum computing. That's a large number for present-day qubit systems, but there is

clearly room to grow. Given a typical 5 μm spacing between sites, ten thousand atoms can be trapped in a 0.5 mm 2D array, and a million atoms can be trapped in a 0.5 mm 3D array. The possibility of such dramatic scaling is at the root of our optimism about using neutral atoms as qubits.

The first step in trapping atoms is almost always laser cooling them to microkelvin temperatures. Conveniently, spontaneous emission during the cooling process scatters light, which can be imaged to detect individual atoms.

But as an array of traps is being loaded with atoms, the cooling light tends to cause atoms at the same site to collide in pairs with each other and be kicked out of the array, which leaves a random half of the traps empty and half of the traps filled with exactly one atom. As long as the vacancy locations are known, it may be possible to reconfigure a quantum computation to account for them. But such a reconfiguration might itself require a quantum computer, so the preference is to start with exactly one atom in every trap.

One elegant way to help achieve a filled array is to start with a Bose–Einstein condensate and then slowly turn on an optical lattice. As the potential wells of the lattice deepen, the atoms undergo a transition from a superfluid to a Mott-insulator state, with one atom per site. The drawbacks of that approach are that nonzero initial temperatures can lead to vacancies, and the intersite atomic tunneling required for the transition works best in a lattice whose sites are more closely spaced than is typically desired for quantum-computing schemes. Another way is to use additional laser beams to suppress the loss of colliding atoms, a technique that can fill up to 90% of the sites in an array.[5] But perhaps the most promising approach is to determine where the vacancies are and then rearrange the atoms to completely fill them and thus form a smaller but filled array.

Several labs are starting to rearrange atoms. At Harvard University, Mikhail Lukin and colleagues filled a 60-site 1D optical-tweezer array by turning off the empty traps and moving the filled ones next to each other[6] (see figure 1a). At the Institut d'Optique near Paris, Antoine Browaeys and colleagues used a mobile optical tweezer to shuttle 50 atoms among sites in a 2D tweezer array[7] (see figure 1b). At the University of Bonn, Andrea Alberti, Dieter Meschede, and their colleagues combined site-selective atomic-state changes and state-selective translations of the lattice potentials to adjust the spacings between atoms in a 1D optical lattice.[8] At the Pennsylvania State University (PSU), one of us (Weiss) is pursuing a similar site- and state-selection scheme in three dimensions, in which the number of required steps scales efficiently, as the cube root of the number of atoms.

One unique and inconvenient difficulty presented by neutral atomic qubits is that they are vulnerable to collisions with residual background-gas atoms that knock them out of their traps. Those collisions occur about once every 100 seconds per atom in a standard room-temperature vacuum system. Lifetimes that exceed tens of minutes are possible in cryogenic vacuum systems. Infrequent atom loss, like other qubit errors, can be repaired with quantum error correction, by using atom-rearrangement hardware to replace a lost atom.

The final step in preparing atom qubits for a computation will likely be some variant of sideband laser cooling. That technique was first invented for trapped ions (see PHYSICS TODAY, October 2005, page 24) and can cool atoms to their vibrational ground state up to 90% of the time,[9] cold enough for most quantum-computing schemes. It is no doubt possible to start even colder; however, to handle the inevitable heating that will arise in a long quantum computation, it may be necessary to transfer quantum states to other cold atoms while the original atoms are being recooled.

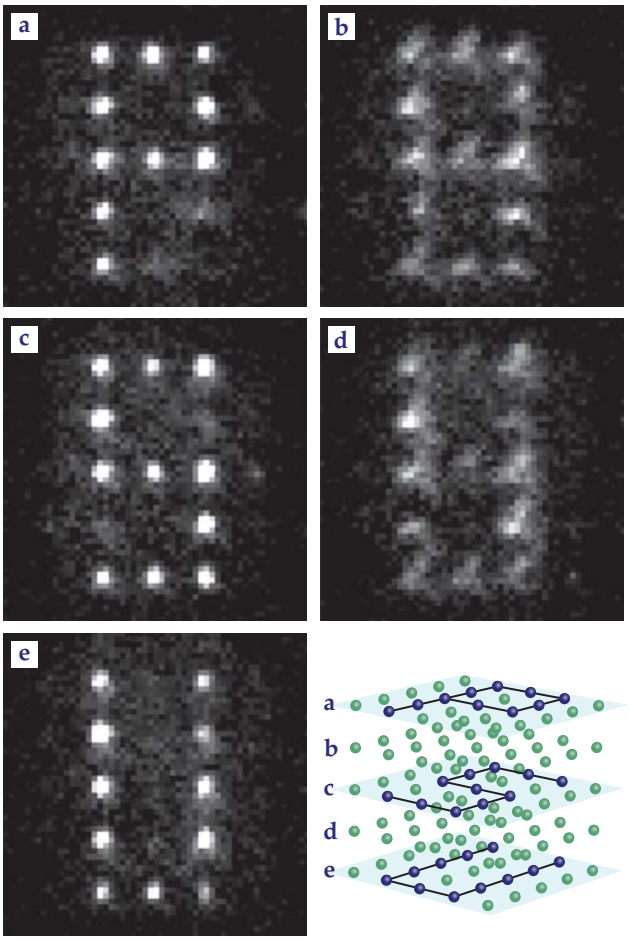

FIGURE 2. TARGETED SINGLE-QUBIT GATES produce the letters P, S, and U in the planes **(a-e)** of a three-dimensional $5 \times 5 \times 5$ optical lattice. In this experiment, all cesium atoms in the array are put into equal quantum superpositions of their hyperfine qubit states. The atoms that make up the target pattern are selectively phase shifted using site-resolved single-qubit gates. A final microwave pulse brings the targeted atoms into the lower qubit state and untargeted atoms into the upper qubit state. A subsequent optical pulse removes the untargeted atoms from the lattice, and each plane is imaged in turn via fluorescence from the remaining (targeted) atoms. No targeted atoms reside in planes b and d. The hazy pattern collected in those planes is the out-of-focus fluorescence of atoms in adjacent planes. Note that because Cs atoms randomly occupy only 40% of the lattice sites in this experiment, the images are the average of 20 implementations. (Adapted from ref. 4.)

Quantum gates

Measurement accuracy, controllability, and the fact that all atoms of one species are identical constitute major strengths of atom- and ion-based quantum computing. Those strengths are well

BOX 2. A TWO-QUBIT GATE

Neutral atoms separated by more than a few angstroms interact very weakly unless they are both excited into a Rydberg state—a highly energetic state with large, delocalized electronic wavefunctions. The resulting dipole–dipole interaction can act over tens of microns and lead to large resonant frequency shifts.

The Rydberg blockade mechanism can be used to entangle two qubits. Imagine two atoms held in separate optical traps a few microns apart. Control $|c\rangle$ and target $|t\rangle$ qubits are initialized in a superposition of stable, hyperfine ground states $|0\rangle$ and $|1\rangle$ of the two atoms, as shown in the energy level diagram of panel a. State $|1\rangle$ is resonantly coupled to a Rydberg state $|r\rangle$ via an optical transition that is far from resonance with state $|0\rangle$.

A three-pulse sequence on the two atoms produces a controlled-phase (C_z) gate:[11] First, a π pulse applied to the control atom transfers it from state $|1\rangle$ to $|r\rangle$. Second, a 2π pulse is applied to the target atom that drives it from $|1\rangle$ to $|r\rangle$ and back again, provided no nearby atom already occupies a Rydberg state. And third, a π pulse is applied again to the control atom to return its state to $|1\rangle$. Nothing happens to the part of the input wavefunction that is in $|00\rangle$. For $|01\rangle$ or $|10\rangle$, one atom is Rydberg excited but not the other. In those cases, driving an atom to and from a Rydberg state is analogous to rotating a spin-½ object by 2π, so the two-qubit state acquires a phase shift of π radians.

The crucial part of the entangling gate occurs to the $|11\rangle$ part of the input wavefunction. The presence of the control atom in the Rydberg state shifts the target atom out of resonance by an amount B. The target atom thus remains in the ground state and fails to acquire a π phase shift. And because the target atom's phase shift depends on the control atom's state, the final two-body state is entangled. Together with arbitrary single-qubit gates, the entangling operation forms a universal quantum-computing gate set. What's more, because there are never two simultaneous Rydberg excitations, the atoms never exert significant forces on each other even as they become entangled.

The circuit shown in panel b maps the phase entanglement onto probability amplitudes by placing the C_z gate (labeled Z in the circuit) between single-qubit rotation operations. The mapping of input to output probability amplitudes corresponds to a controlled-NOT gate operation, which flips the target qubit only if the control qubit is in the $|1\rangle$ hyperfine state. The upper and lower detectors register the output states of the control and target qubits.

Panel c shows the experimental truth table for the controlled-NOT gate. The probability of each of the four possible final, or output, states of the two qubits is plotted against their initial, or input, states.[14]

Since the Rydberg interaction is strong and long range, it is possible to entangle not just neighboring atoms but also those separated by several lattice sites. The blockade can also inhibit the excitation of multiple atoms at a time, a feature that allows researchers to efficiently build multiqubit gates and multiparticle entangled states.[18]

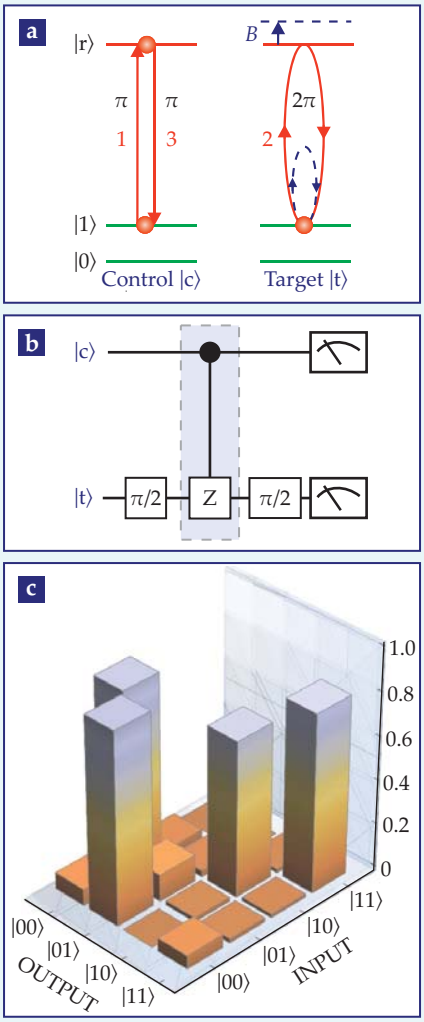

known in the context of atomic clocks, and they account for the clocks' superlative precision. In fact, viewed from the perspective of quantum computing, atomic clocks are built from single-qubit gates implemented in parallel on a large ensemble of atoms.

However, making precise single-qubit gates for a quantum computer requires the ability to execute operations on individual atoms. It also requires eliminating any cross talk, which can change the quantum states of nearby, untargeted atoms. One way to achieve that level of control is to drive atomic transitions between hyperfine levels using laser beams that are tightly focused on individual atoms. Such a gate can be fast—on the scale of a microsecond or less—and can have high fidelity if the light intensity is sufficiently stable. Another way is to drive atomic transitions with more easily controlled microwaves. Because microwave wavelengths are far too large to distinguish individual atoms, the atoms must be addressed by selectively shifting their resonance frequencies, either with magnetic fields or focused laser beams. Using crossed laser beams and a sequence of microwave pulses, one can reach into a 3D ensemble to perform a gate on a target atom without cross talk, as exemplified in box 1. Because single-qubit gate times range from fractions of a microsecond to a few hundred microseconds, as many as 10^5 gates can be performed within the longest demonstrated decoherence times.

Figure 2 illustrates a series of targeted single-qubit gates acting on a 3D array of cesium atoms, each of which is put in a quantum superposition of two hyperfine states. The experiment, which produces the letters P, S, and U in a fluorescent pattern, demonstrates independent control over atoms at all 125 lattice sites, which (given 40% site occupancy) amounts to 50 qubits at a time.

A two-qubit gate, which creates entangled states on demand, requires strong interactions among qubits. But neutral atoms in their electronic ground states interact weakly—the very reason they can be packed close together in a quantum computer. One solution is to temporarily move atoms so closely together that they experience a "controlled, cold collision."[10] When using that approach, however, it is difficult to control the

moving traps and the motional states of the atoms accurately enough to produce a high-fidelity gate. A more widely pursued solution, proposed nearly two decades ago by Dieter Jaksch (now at Oxford University) and colleagues in a seminal paper,[11] is to temporarily transfer atoms to Rydberg states, in which an electron is excited far from the atomic nucleus.

Rydberg atoms have strong mutual dipolar interactions. Indeed, the interaction between two Rydberg atoms separated by 5 μm and with principal quantum number $n = 100$ is about 12 orders of magnitude larger than that between ground-state atoms. The strong interaction produces a so-called Rydberg blockade—a phenomenon analogous to the well-known Coulomb blockade, by which only one electron at a time can move through a narrow channel. The Rydberg blockade prevents more than one atom in a small volume from being simultaneously excited to a Rydberg state. For instance, if an atom is excited to a Rydberg state, an attempt to excite a second, nearby atom with the same laser frequency will fail, because the long-range interaction between the atoms shifts the resonance condition for excitation of the second atom. The upshot: The blockade interaction can entangle two or more nearby qubits using a simple three-pulse sequence, as explained in box 2.

From the view of an experimental physicist, the task of exerting precise control over a large number of individual quantum particles is a grand challenge.

Research groups[12] at the University of Wisconsin–Madison and at the Institut d'Optique first demonstrated entanglement between individual rubidium atoms in hyperfine ground states using a Rydberg blockade in 2010. (See PHYSICS TODAY, February 2010, page 13, and February 2009, page 15.) The initial experiments just barely passed the threshold between separable and entangled states. Since then, experimental improvements and variations on the basic entanglement protocol[13,14] have led to entanglement fidelities of about 75%. That's well below what has been achieved with trapped-ion and superconducting qubits. But the fidelities are likely to improve when ground-state cooling and other advancements are incorporated into Rydberg gate experiments.

Another issue that can affect fidelity is the extreme sensitivity of atomic Rydberg states to background electric fields created by charge fluctuations in an experimental apparatus. Future advances may require better ways of mitigating those fields or the Rydberg state's sensitivity to them. A more intrinsic difficulty is a Rydberg state's finite radiative lifetime, which leads to decoherence during a gate operation. Obtaining optimal fidelity will be a trade-off between radiative decay when the gate is too slow and a compromised blockade when the gate is too fast. A recent numerical optimization of pulse shapes to minimize off-resonant excitation, even for fast 50 ns gates,[15] has shown the feasibility of achieving 99.99% entanglement fidelity. Such performance is a prerequisite for implementing error correction.

Error correction requires measuring the states of selected atoms without disturbing the quantum states of their neighbors. Several possible approaches take advantage of the flexibility of cold-atom experiments. For instance, two Rydberg-interacting atomic species could be used, one for storage and one for measurement. Or optical shielding techniques could be used to protect qubits that are not being measured. Those and other ideas are just starting to be developed,[16] so truly scalable quantum computation remains a problem to be tackled in the next generation of experiments.

Perspective

From the view of an experimental physicist, the task of exerting precise control over a large number of individual quantum particles is a grand challenge. A quantum computation requires preparing atoms in well-defined quantum states, controlling the atoms' interactions to carry out logical operations, and measuring the resulting states to extract the computational result—all while maintaining near-perfect isolation from the environment. The technical achievements needed to meet that challenge are likely to reap a broad range of benefits beyond the central goal of quantum computing.

For instance, advances in neutral atom quantum computing may broaden the capabilities of metrology and time keeping. Atomic clocks already operate with exceedingly low uncertainties, near one part in 10^{18}. At the frontiers of clock research are efforts to introduce entanglement and quantum-logic protocols to circumvent the standard quantum limits for uncorrelated particles.[17] The ability to create, control, and probe many-particle quantum states with high precision will enrich the experimental study of quantum mechanics at the border between classical and quantum dynamics.

As for the central goal, we anticipate that over the next five years several qubit platforms, including neutral atoms, are likely to reach a sufficient size and fidelity that they can perform quantum calculations that cannot be modeled on a classical computer. It is much harder to predict how long it will take before a quantum computer is able to factor intractably large numbers.

REFERENCES

1. R. P. Feynman, *Found. Phys.* **16**, 507 (1986).
2. C. Monroe et al., *Phys. Rev. Lett.* **75**, 4714 (1995).
3. T. D. Ladd et al., *Nature* **464**, 45 (2010).
4. Y. Wang et al., *Science* **352**, 1562 (2016).
5. A. V. Carpentier et al., *Laser Phys. Lett.* **10**, 125501 (2013); B. J. Lester et al., *Phys. Rev. Lett.* **115**, 073003 (2015).
6. M. Endres et al., *Science* **354**, 1024 (2016).
7. D. Barredo et al., *Science* **354**, 1021 (2016).
8. C. Robens et al., *Phys. Rev. Lett.* **118**, 065302 (2017).
9. X. Li et al., *Phys. Rev. Lett.* **108**, 103001 (2012); A. M. Kaufman, B. J. Lester, C. A. Regal, *Phys. Rev. X* **2**, 041014 (2012).
10. M. Anderlini et al., *Nature* **448**, 452 (2007); A. M. Kaufman et al., *Nature* **527**, 208 (2015).
11. D. Jaksch et al., *Phys. Rev. Lett.* **85**, 2208 (2000).
12. X. L. Zhang et al., *Phys. Rev. A* **82**, 030306 (2010); T. Wilk et al., *Phys. Rev. Lett.* **104**, 010502 (2010).
13. Y.-Y. Jau et al., *Nat. Phys.* **12**, 71 (2016).
14. K. M. Maller et al., *Phys. Rev. A* **92**, 022336 (2015).
15. L. S. Theis et al., *Phys. Rev. A* **94**, 032306 (2016).
16. M. Saffman, *J. Phys. B* **49**, 202001 (2016).
17. L. I. R. Gil et al., *Phys. Rev. Lett.* **112**, 103601 (2014); M. Khazali et al., *Phys. Rev. A* **94**, 023408 (2016).
18. M. Ebert et al., *Phys. Rev. Lett.* **115**, 093601 (2015); J. Zeiher et al., *Phys. Rev. X* **5**, 031015 (2015).

PT