

The Best Qubits for Quantum Computing Might Just Be Atoms

In the search for the most scalable hardware to use for quantum computers, qubits made of individual atoms are having a breakout moment.



Mikhail Lukin (left) pioneered the idea of neutral-atom quantum computing and has recently made dramatic headway together with Markus Greiner.

Ken Richardson for Quanta Magazine



By

Philip Ball

Contributing Writer

March 25, 2024

Introduction

At the end of last year, the tech giant IBM announced what might sound like a milestone in quantum computing: the first-ever chip, called the Condor, with more than 1,000 quantum bits, or qubits. Given that this was barely two years after the company unveiled the Eagle, the first chip with more than 100 qubits, it looked as though the field was racing forward. Making quantum computers that can solve useful problems beyond the scope of even the mightiest of today's classical supercomputers demands scaling them up even more — to perhaps many tens or hundreds of thousands of qubits. But that's surely just a matter of engineering, right?

Not necessarily. The challenges of scaling up are so great that some researchers think it will require totally different hardware from the microelectronics used by the likes of IBM and Google. The

qubits in the Condor and in Google's Sycamore chip are made from loops of superconducting material. These superconducting qubits have so far been the hare in the race to full-scale quantum computing. But now there's a tortoise coming from behind: qubits made from individual atoms.


Recent advances have transformed these "neutral-atom qubits" from outsiders to leading contenders.

"The last two or three years have seen more rapid advances than any previous such period," said the physicist Mark Saffman of the University of Wisconsin, Madison, who counted at least five companies racing to commercialize neutral-atom quantum computing.

Like the bits in ordinary computers, qubits encode binary information — 1s and 0s. But whereas a bit is always in one state or the other, the information in a qubit can be left indeterminate, in a so-called "superposition" that gives weight to both possibilities. To carry out a computation, qubits are linked using the phenomenon called quantum entanglement, which makes their possible states interdependent. A particular quantum algorithm might demand a succession of entanglements between different sets of qubits, and the answer is read out at the end of the computation when a measurement is made, collapsing each superposition down to a definite 1 or 0.

The idea of using the quantum states of neutral atoms for encoding information this way was proposed in the early 2000s by the Harvard physicist Mikhail Lukin and colleagues, and also by a group led by Ivan Deutsch of the University of New Mexico. For a long time, the broader research community agreed that neutral-atom quantum computing was a great idea in principle, Lukin said, but that "it just doesn't work out" in practice.

“But 20 years later, the other approaches haven’t closed the deal,” Saffman said. “And the skill set and the techniques needed to make neutral atoms work have been gradually evolving to the point where they are looking very promising.”



In Lukin’s lab, an elaborate system of lasers is used to move around and change the states of atoms held inside a vacuum cell hidden in the scene.

Ken Richardson for Quanta Magazine

Introduction

Lukin’s lab at Harvard has been among those leading the way. In December, he and his colleagues reported that they created programmable quantum circuits with hundreds of neutral-atom qubits and had performed quantum computations and error correction with them. And this month, a team at the California Institute of Technology reported that they made an array of 6,100 atomic qubits. Such results are increasingly winning converts to this approach.

“Ten years ago I would not have included these [neutral-atom] methods if I were hedging bets on the future of quantum computing,” said Andrew Steane, a quantum information theorist at the University of Oxford. “That would have been a mistake.”

Battle of Qubits

A key issue in the contest between qubit types is how long each kind of qubit can maintain its superposition before it is altered by some random (for example, thermal) fluctuation. For superconducting qubits like IBM's and Google's, this "coherence time" is typically around a millisecond at best. All steps of a quantum computation must happen within that time frame.

One advantage of encoding information in the states of individual atoms is that their coherence times are typically far longer. Moreover, unlike superconducting circuits, atoms of a given type are all identical, so bespoke control systems aren't needed to input and manipulate subtly different quantum states.

And whereas the wiring used to link up superconducting qubits into quantum circuits can become horribly complicated — more so as the system scales up — no wiring is needed in the case of atoms. All the entangling is done using laser light.


This benefit initially presented a challenge. There's a well-developed technology for carving out complicated microelectronic circuitry and wires, and one probable reason IBM and Google invested initially in superconducting qubits is not because these were obviously the best but because they required the kind of circuitry such companies are used to, said [Stuart Adams](#), a physicist at Durham University in the United Kingdom who works on neutral-atom quantum computing. "Laser-based atom optics looked totally unfamiliar to them. All the engineering is completely different."

Quantum Computing With Atoms

The basic architecture of neutral-atom quantum computers involves three zones. Lasers move the atoms between them.

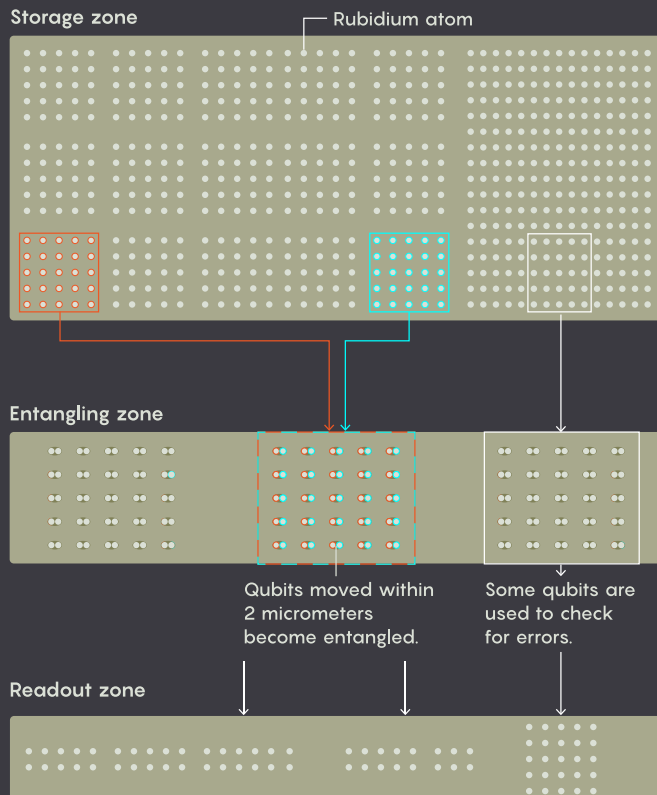
OPTICAL TWEEZERS

Focused laser beams are used to trap the atoms and move them around.



LOGICAL PROCESSOR ARCHITECTURE

In a quantum computation, optical tweezers move blocks of atomic qubits between a storage zone and an entangling zone, where they are made to interact. After a sequence of entanglements, qubits move into the readout zone, where their states are measured.



Qubits made of electrically charged atoms — known as ions — can also be controlled with light, and ions were long regarded as better qubit candidates than neutral atoms. Because of their charge, ions are relatively easy to trap in electric fields. Researchers have created ion traps by suspending the ions in a tiny vacuum cavity at ultralow temperatures (to avoid thermal jiggling) while laser beams switch them between different energy states to manipulate the information. Ion-trap quantum computers with dozens of qubits have now been demonstrated, and several startups are developing the technology for commercialization. “So far, the system with the highest performance in terms of fidelity, control and coherence has been trapped ions,” Saffman said.

Trapping neutral atoms is harder because there’s no charge to hold onto. Instead, the atoms are immobilized within fields of intense light created by laser beams, called optical tweezers. The atoms typically prefer to sit where the light field is most intense.

And there’s a problem with ions: They all have an electric charge of the same sign. That means the qubits repel one another. Jamming a lot of them into the same small space gets harder the more ions there are. With neutral atoms, there’s no such tension. This, researchers say, makes neutral-atom qubits much more scalable.

What’s more, trapped ions are arranged in a row (or, recently, a looping “racetrack”). This configuration makes it difficult to entangle one ion qubit with another that’s, say, 20 places along the row. “Ion traps are inherently one-dimensional,” Adams said. “You have to arrange them in a line, and it’s very hard to see how you get up to a thousand qubits that way.”

Neutral-atom arrays can be a two-dimensional grid, which is much easier to scale up. “You can put a lot in the same system, and they don’t interact when you don’t want them to,” Saffman said. His group and others have trapped over 1,000 neutral atoms this way.

“We believe we can pack tens or even hundreds of thousands in a centimeter-scale device,” he said.

Indeed, in their recent work, the team at Caltech created an optical-tweezer array of about 6,100 neutral cesium atoms, although they haven’t yet performed any quantum computations with them. These qubits also had coherence times of a whopping 12.6 seconds, a record so far for this qubit type.

The Rydberg Blockade

For two or more qubits to become entangled, they need to interact with one another. Neutral atoms “feel” one another’s presence via so-called van der Waals forces, which arise from the way one atom responds to fluctuations in the cloud of electrons in another atom nearby. But these feeble forces are felt only when atoms are extremely close together. Manipulating normal atoms to the required precision using light fields just can’t be done.

As Lukin and his colleagues pointed out in their original proposal back in 2000, the interaction distance can be dramatically increased if we boost the size of the atoms themselves. The more energy an electron has, the further it tends to roam from the atomic nucleus. If a laser is used to pump up an electron into an energy state far greater than those usually found in atoms — called a Rydberg state after the Swedish physicist Johannes Rydberg, who in the 1880s studied the way atoms emit light at discrete wavelengths — the electron can roam thousands of times farther out from the nucleus than usual.

This boost in size enables two atoms held several micrometers apart — perfectly feasible in optical traps — to interact.

This movie created by the Harvard team shows the execution of a 48-logical-qubit circuit, which they say is the most advanced circuit ever executed on a quantum computer. Groups of eight atomic qubits are first brought together and entangled into error-corrected logical qubit blocks, indicated by red ovals. These blocks are then entangled with each other to create a circuit with hundreds of logical gate operations.

Lukin Lab/Harvard

Introduction

To implement a quantum algorithm, the researchers first encode quantum information in a pair of atomic energy levels, using lasers to switch electrons between the levels. They then entangle atoms' states by switching on Rydberg interactions between them. A given atom may be excited to a Rydberg state or not, depending on which of the two energy levels its electron is in — only one of these sits at the right energy to resonate with the frequency of the excitation laser. And if the atom is currently interacting with another, this excitation frequency shifts slightly so that the electron won't resonate with the light and won't be able to make the jump. This means only one or the other of a pair of interacting atoms can sustain a Rydberg state at any moment; their quantum states are correlated — or in other words, entangled. This so-called Rydberg blockade, first proposed by Lukin and colleagues in 2001 as a way of entangling Rydberg-atom qubits, is an all-or-nothing effect: Either there's a Rydberg blockade or there isn't. "The Rydberg blockade makes interactions between atoms digital," Lukin said.

At the end of the computation, lasers read out atoms' states: If an atom is in the state that is resonant with the illumination, the light is scattered, but if it's in the other state, there's no scattering.

In 2004, a team at the University of Connecticut demonstrated a Rydberg blockade between rubidium atoms, trapped and cooled to just 100 microkelvins above absolute zero. They cooled the atoms by using lasers to "suck out" the atoms' thermal energy. The approach means that, unlike superconducting qubits, neutral atoms require no cryogenic cooling and no cumbersome refrigerants. These systems can therefore be made very compact. "The apparatus as a whole is at room temperature," Saffman said. "One

centimeter away from these super cold atoms, you have a room-temperature window.”

In 2010 Saffman and his co-workers reported the first logic gate — a fundamental element of computers, in which one or more binary input signals generate a particular binary output — made from two atoms using the Rydberg blockade. Then, crucially, in 2016, Lukin’s team and research groups in France and in South Korea all independently figured out how to load many neutral atoms into arrays of optical traps and move them around at will. “This innovation brought new life to the field,” said Stephan Dürr of the Max Planck Institute of Quantum Optics in Garching, Germany, who uses Rydberg atoms for experiments in light-based quantum information processing.

Much of the work so far uses rubidium and cesium atoms, but the physicist Jeff Thompson at Princeton University prefers encoding the information in the nuclear spin states of metal atoms such as strontium and ytterbium, which have even longer coherence times. Last October, Thompson and colleagues reported two-qubit logic gates made from these systems.


And Rydberg blockades don’t have to be between lone atoms. Last summer, Adams and his co-workers showed that they could create a Rydberg blockade between an atom and a trapped molecule, which they made artificially by using optical tweezers to pull a cesium atom next to a rubidium atom. The advantage of hybrid atom-molecule systems is that atoms and molecules have very different energies, which could make it easier to manipulate one without affecting others. What’s more, molecular qubits can have very long coherence times. Adams stresses that such hybrid systems are at least 10 years behind all-atom systems, and entanglement of two such qubits has yet to be achieved. “Hybrid systems are really hard,” Thompson said, “but we will probably be forced to do them at some point.”

High-Fidelity Qubits

No qubit is perfect: All can incur errors. And if these go undetected and uncorrected, they scramble the result of the computation.

But a big obstacle to all quantum computing is that errors can't be identified and corrected in the way they are for classical computers, where an algorithm simply keeps track of what states the bits are in by making copies. The key to quantum computing is that the qubits' states are left undetermined until the final result is read out. If you try to measure those states before that point, you terminate the computation. How, then, can qubits be protected from errors that we can't even monitor?

One answer is to spread information over many physical qubits — constituting a single “logical qubit” — so that an error in one of them doesn't corrupt the information they collectively encode. This only becomes practical if the number of physical qubits needed for each logical qubit isn't too great. That overhead depends in part on what error-correcting algorithm is used.



From left: Simon Evered, Sophie Li, Alexandra Geim, Mikhail Lukin, Dolev Bluvstein and Markus Greiner peer at the experimental setup in Lukin's lab.

Ken Richardson for Quanta Magazine

Introduction

Error-corrected logical qubits have been demonstrated with superconducting and trapped-ion qubits, but until recently it hasn't been clear if they can be made from neutral atoms. That changed in December, when the Harvard team unveiled arrays of several hundred trapped rubidium atoms and ran algorithms on 48 logical qubits, each made from seven or eight physical atoms. The researchers used the system to conduct a simple logical operation called a controlled NOT gate, in which a qubit's 1 and 0 states are flipped or left unchanged depending on the state of a second "control" qubit. To conduct the computations, the researchers moved the atoms between three distinct regions in the trapping chamber: an array of atoms, an interaction region (or "gate zone") where specific atoms were dragged and entangled using the Rydberg blockade, and a readout zone. It's all made possible, Adams said, because "the Rydberg system offers you all this ability to shuffle qubits around and decide who's interacting with who, which gives you a flexibility that superconducting qubits don't have."

The Harvard team demonstrated error-correction techniques for some simple logical-qubit algorithms, although for the largest ones, with 48 logical qubits, they merely achieved error detection. According to Thompson, those latter experiments showed that "they can preferentially reject measurement outcomes with errors, and therefore identify a subset of outcomes with lower errors." This approach is called post-selection, and while it can play a role in quantum error correction, it doesn't by itself solve the problem.

Rydberg atoms might lend themselves to novel error-correcting codes. The one used in the Harvard work, called the surface code, "is very popular but also very inefficient," Saffman said; it tends to require many physical qubits to make one logical qubit. Other, more efficient proposed error-correcting codes require longer-range interactions between qubits, not just nearest-neighbor pairings.

Practitioners of neutral-atom quantum computing think long-range Rydberg interactions should be up to the task. “I am extremely optimistic that experiments over the next two to three years will show us that the overhead need not be as bad as people thought,” Lukin said.

Though there’s still more to be done, Steane considers the Harvard work “a step change in the degree to which error-correction protocols have been realized in the laboratory.”

Spinning Off

Advances like these have Rydberg-atom qubits drawing even with their competitors. “The combination of high-fidelity gates, the large numbers of qubits, high-accuracy measurements and flexible connectivity allows us to deem the Rydberg-atom array as a real competitor to the superconducting and trapped-ion qubits,” Steane said.

Compared to superconducting qubits, the technology comes at a fraction of the investment cost. The Harvard group has a spinoff company called QuEra, which has already made a 256-qubit Rydberg quantum processor called Aquila — an analog “quantum simulator,” which can run simulations of systems of many quantum particles — available on the cloud in partnership with Amazon’s Braket quantum computing platform. QuEra is also working to advance quantum error correction.

Saffman joined a company called Infleqtion, which is developing the neutral-atom optical platform for quantum sensors and communications as well as quantum computing. “I wouldn’t be surprised if one of the big IT companies goes into some sort of partnership with one of these spinoffs soon,” Adams said.

“Doing scalable error correction with neutral-atom qubits is definitely possible,” Thompson said. “I think 10,000 neutral-atom qubits is clearly possible within a few years.” Beyond that, he thinks that practical limitations on laser power and resolution will necessitate modular designs in which several distinct atom arrays are linked together.

If that happens, who knows what will come of it? “We don’t even know yet what we can do with quantum computing,” Lukin said. “I really hope these new advances will help us answer these questions.”

Correction: March 29, 2024.

Sophie Li, a member of the Harvard research team, was originally misidentified in a photo caption as Sophie Lee.