


Atom-based architectures may have a scalability advantage over other platforms in the quest to build more powerful quantum processors.


Susan Curtis

Quantum Computing Enters the Atomic Realm

An experimental scheme demonstrated by researchers at Princeton and Yale universities is able to convert physical noise into errors that can be corrected more easily.

F. Wojciechowski / Princeton University





A team led by Mikhail Lukin (right) at Harvard University, USA, pictured with lab member Dolev Bluvstein, created the first programmable logical quantum processor, capable of encoding up to 48 logical qubits.

J. Chase / Harvard Staff Photographer

California-based startup Atom Computing announced that it had populated a revamped version of its commercial system with almost 1,200 qubits—more than had yet been reported for any hardware platform. “It’s exciting to be able to showcase the solutions we have been developing for the past several years,” says Ben Bloom, who founded the company in 2018 and is now its chief technology officer. “We have demonstrated a few firsts along the way, but while we have been building, the field has been getting more and more amazing.”

Quantum computers built from arrays of ultracold atoms have recently emerged as a serious contender in the quest to create qubit-powered machines that can outperform their classical counterparts. While other hardware architectures have yielded the first fully functioning quantum processors to be available for programming through the cloud, recent developments suggest that atom-based platforms might have the edge when it comes to future scalability.

That scalability advantage stems from the exclusive use of photonic technologies to cool, trap and manipulate the atomic qubits. Side-stepping the need for complex cryogenic systems or the intricacies of chip fabrication, neutral-atom quantum computers can largely be built from existing optical components and systems that have already been optimized for precision and reliability.

“The traps are optical tweezers, the atoms are controlled with laser beams and the imaging is done with a camera,” says Jeff Thompson, a physicist at Princeton University, USA, whose team has been working to build a quantum computer based on arrays of ytterbium atoms. “The scalability of the platform is limited only by the engineering that can be done with the optical system, and there is a whole industry of optical components and megapixel devices where much of that work has already been done.”

Such ready availability of critical components and systems has enabled both academic groups and commercial companies to scale their quantum processors from tens of atomic qubits to several hundred in the space of just a few years. Then, in November 2023, the

Large-scale atomic arrays

Neutral atoms offer many appealing characteristics for encoding quantum information. For a start, they are all identical—completely free of any imperfections that may be introduced through fabrication—which means that they can be controlled and manipulated without the need to tune or calibrate individual qubits. Their quantum states and interactions are also well understood and characterized, while crucial quantum properties such as superposition and entanglement are maintained over long enough timescales to perform computational tasks.

However, early attempts to build quantum computers from neutral atoms met with two main difficulties. The first was the need to extend existing methods for trapping single atoms in optical tweezers to create large-scale atomic arrays. Although technologies such as spatial light modulators enable laser beams to be used to produce a regular pattern of microtraps, loading the atoms into the tweezers is a stochastic process—which means that the probability of each trap being occupied is 50%. As a result, the chances of creating a defect-free array containing large numbers of atoms becomes vanishingly small.

The solution came in 2016, when three separate groups—based at the Institut d’Optique, France, Harvard University, USA, and the Korea Advanced Institute of Science and Technology (KAIST), Republic of Korea—demonstrated a concept called rearrangement. In this scheme, an image is taken of the atoms when they are first loaded into the tweezers, which identifies which sites are occupied and which are empty. All the vacant traps are switched off, and then the loaded ones are moved

The scalability advantage of atom-based platforms stems from the exclusive use of photonic technologies to cool, trap and manipulate the atomic qubits.

to fill the gaps in the array. This shuffling procedure can be achieved, for example, by using acousto-optic deflectors to alter the positions of the trapping laser beams, creating dynamic optical tweezers that can be combined with real-time control to assemble large arrays of single atoms in less than a second.

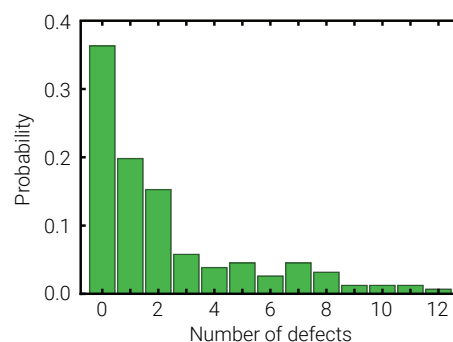

“Before that, there were lots of complicated ideas for generating single-atom states in optical tweezers,” remembers Thompson. “This rearrangement technique enabled the creation of large arrays containing one hundred or so single atoms without defects, and that has since been extended to much higher numbers.”

Higher-fidelity entangled states

In these atomic arrays, the qubits are encoded in two long-lived energy states that are controlled with laser light. In rubidium, for example, which is often used because its well-understood atomic transitions can be manipulated relatively easily, the single outermost electron occupies one of two distinct energy levels in the ground state, caused by the coupling between the electron spin and the nuclear spin. The atoms are easily switched between these two energy states by flipping the spins relative to each other, which is achieved with microwave pulses tuned to 6.8 GHz.

While atoms in these stable low-energy levels offer excellent single-qubit properties, the gate operations that form the basis of digital computation require the qubits to interact and form entangled states. Since the atoms in a tweezer array are too far apart for them to interact while remaining in the ground state, a focused laser beam is used to excite the outermost electron into a much higher energy state. In these highly excited Rydberg states, the atom becomes physically much larger, generating strong interatomic interactions on sub-microsecond timescales.

One important effect of these interactions is that the presence of a Rydberg atom shifts the energy levels in its nearest neighbors, preventing them from being excited into the same high-energy state. This phenomenon, called the Rydberg blockade, means that only one of the atoms excited by the laser will form a Rydberg state, but it's impossible to know which one. Such shared excitations are the characteristic feature




Large defect-free arrays of single atoms can be created through the process of rearrangement. In this example, demonstrated by a team led by Antoine Browaeys of the Institut d’Optique, France, an ordered array of 324 atoms was created from 625 randomly filled traps.

Reprinted with permission from K.-N. Schymik, *Phys. Rev. A* **106**, 022611(2022); ©2022 by the American Physical Society

of entanglement, providing an effective mechanism for controlling two-qubit operations between adjacent atoms in the array.

Until recently, however, the logic gates created through two-atom entanglement were prone to errors. “For a long time, the fidelity of two-qubit operations hovered at around 80%—much lower than could be achieved with superconducting or trapped-ion platforms,” says Thompson. “That meant that neutral atoms were not really taken seriously for gate-based quantum computing.”

The sources of these errors were not fully understood until 2018, when breakthrough work by Antoine Browaeys and colleagues at the Institut d’Optique and Mikhail Lukin’s team at Harvard University analyzed the effects of laser noise on the gate fidelities. “People had been using very simple models of the laser noise,” says Thompson. “With this work, they figured out that



Jeff Thompson and his team at Princeton University, USA, have pioneered the use of ytterbium atoms to encode and manipulate quantum information.

S.A. Khan / Fotobuddy

phase fluctuations were the major contributor to the high error rates.”

At a stroke, these two groups showed that suppressing the laser phase noise could extend the lifetime of the Rydberg states and boost the fidelity of preparing two-qubit entangled states to 97%. Further enhancements since then have yielded two-qubit gate fidelities of more than 99%—the minimum threshold for fault-tolerant quantum computing.

The quest for fault tolerance

That fundamental advance established atomic qubits as a competitive platform for digital quantum computing, catalyzing academic groups and quantum startups to explore and optimize the performance of different atomic systems. While rubidium continues to be a popular choice, several groups believe that ytterbium could offer some crucial benefits for large-scale quantum computing. “Ytterbium has a nuclear spin of one half, which means that the qubit can be encoded purely in the nuclear spin,” explains Thompson. “While all qubits based on atoms or ions have good coherence by default, we have found that pure nuclear-spin qubits can maintain coherence times of many seconds without needing any special measures.”

Pioneering experiments in 2022 by Thompson’s Princeton group, as well as by a team led by Adam Kaufman at JILA in Boulder, CO, USA, first showed the potential of the ytterbium-171 isotope for producing long-lived atomic qubits. Others have followed their lead, with Atom Computing replacing the strontium atoms in its original prototype with ytterbium-171 in

the upgraded 1,200-qubit platform. “Strontium also supports nuclear qubits, but we found that we needed to do lots of quantum engineering to achieve long coherence times,” says Bloom. “With ytterbium, we can achieve coherence times of tens of seconds without the need for any of those extra tricks.”

The rich energy-level structure of ytterbium also provides access to a greater range of atomic transitions from the ground state, offering new ways to manipulate and measure the quantum states. Early experiments have shown, for example, that this additional flexibility can be exploited to measure some of the qubits while a quantum circuit is being run but without disturbing the qubits that are still being used for logical operations.

Indeed, the ability to perform these mid-circuit measurements is a critical requirement for emerging schemes to locate and correct physical errors in the system, which have so far compromised the ability of quantum computers to perform complex computations. These physical errors are caused by noise and environmental factors that perturb the delicate quantum states, with early estimates suggesting that millions of physical qubits might be needed to provide the redundancy needed to achieve fault-tolerant quantum processing.

More recently, however, it has become clear that fewer qubits may be needed if the physical system can be engineered to limit the impact of the errors. One promising approach is the concept of erasure conversion—demonstrated in late 2023 by a team led by Thompson and Shruti Puri at Yale University, USA—in which the physical noise is converted into errors with known locations, also called erasures.

In their scheme, the qubits are encoded in two metastable states of ytterbium, for which most errors will cause them to decay back to the ground state. Importantly, those transitions can easily be detected without disturbing the qubits that are still in the metastable state, allowing failures to be spotted while the quantum processor is still being operated. “We just flash the atomic array with light after a few gate operations, and any light that comes back illuminates the position of the error,” explains Thompson. “Just being able to see where they are could ultimately reduce the number of qubits needed for error correction by a factor of ten.”

While rubidium continues to be a popular choice, several groups believe that ytterbium could offer some crucial benefits for large-scale quantum computing.

Experiments by the Princeton researchers show that their method can currently locate 56% of the errors in single-qubit gates and 33% of those in two-qubit operations, which can then be discarded to reduce the effects of physical noise. The team is now working to increase the fidelity that can be achieved when using these metastable states for two-qubit operations, which currently stands at 98%.


Looking at logical qubits

Meanwhile, Lukin's Harvard group, working with several academic collaborators and Boston-based startup QuEra Computing, has arguably made the closest approach yet to error-corrected quantum computing. One crucial step forward is the use of so-called logical qubits, which mitigate the effects of errors by sharing the quantum information among multiple physical qubits.

Previous demonstrations with other hardware platforms have yielded one or two logical qubits, but Lukin and his colleagues showed at the end of 2023 that they could create 48 logical qubits from 280 atomic qubits. They used optical multiplexing to illuminate all the rubidium atoms within a logical qubit with identical light beams, allowing each logical block to be moved and manipulated as a single unit. Since each atom in the logical block is addressed independently, this hardware-efficient control mechanism prevents any errors in the physical qubits from escalating into a logical fault.


For more-scalable processing of these logical qubits, the researchers also divided their architecture into three functional zones. The first is used to store and manipulate the logical qubits—along with a reservoir of physical qubits that can be mobilized on demand—ensuring that these stable quantum states are isolated from processing errors in other parts of the hardware. Pairs of logical qubits can then be moved, or “shuttled,” into the second entangling zone, where a single excitation laser drives two-qubit gate operations with a fidelity of more than 99.5%. In the final readout zone, the outcome of each gate operation is measured without affecting the ongoing processing tasks.

The team also configured error-resistant quantum circuits to run on the logical processor, which in one example yielded a fidelity of 72% when operating on




Atom Computing's first-generation quantum computer exploited around 100 qubits of single strontium atoms, while its next-generation platform can accommodate around 1,200 ytterbium atoms.

Atom Computing



Schematic of the logical processor, split into three zones: storage, entangling and readout. Logical single-qubit and two-qubit operations are realized transversally with efficient, parallel operations.

D. Bluvstein et al. *Nature*, **626**, 58 (2024); CC-BY-NC 4.0



Researchers at the Boston-based startup QuEra, which collaborates on neutral-atom quantum computing with Mikhail Lukin's group at Harvard University, USA.

Courtesy of QuEra

10 logical qubits, increasing to 99% when the gate errors detected in the readout zone at the end of each operation were discarded. When running more complex quantum algorithms requiring hundreds of logical gates, the performance was up to 10 times better when logical qubits were used instead of their single-atom counterparts.

While this is not yet full error correction, which would require the faults to be detected and reset in real time, this demonstration shows how a logical processor can work in tandem with error-resistant software to improve the accuracy of quantum computations. The fidelities that can be achieved could be improved still further by sharing the quantum information among more physical qubits, with QuEra's technology roadmap suggesting that by 2026 it will be using as many as 10,000 single atoms to generate 100 logical qubits. "This is a truly exciting time in our field as the fundamental ideas of quantum error correction and fault tolerance start to bear fruit," Lukin commented. "Although there are still challenges ahead, we expect that this new advance will greatly accelerate the progress toward large-scale, useful quantum computers."

Scalability for the future

In another notable development, QuEra has also won a multimillion-dollar contract to build a version of this logical processor at the UK's National Quantum Computing Centre (NQCC). The QuEra system will be one of seven prototype quantum computers to be installed at the national lab by March 2025, with others including a cesium-based neutral-atom system from Infleqtion (formerly ColdQuanta) and platforms exploiting superconducting qubits and trapped ions.

Once built, these development platforms will be used to understand and benchmark the capabilities of different hardware architectures, explore the types of applications that suit each one, and address the key scaling challenges that stand in the way of fault-tolerant quantum computing. "We know that much more practical R&D will be needed to bridge the gap between currently available platforms and a fully error-corrected neutral-atom quantum computer with hundreds of logical qubits," says Nicholas Spong, who leads the NQCC's activities in tweezer-array quantum computing. "For neutral-atom architectures, the ability to scale really depends on engineering the optics, lasers and control systems."

One key goal for hardware developers will be to achieve the precision needed to control the spin rotations of individual atoms as they become more closely packed into the array. While global light fields and qubit shuttling provide efficient and precise control mechanisms for bulk operations, single-atom processes must typically be driven by focused laser beams operating on the scale of tens of nanometers.

To relax the strict performance criteria for these local laser beams, Thompson's group has demonstrated an alternative solution that works for divalent atoms such as ytterbium. "We still have a global gate beam, but then we choose which atoms experience that gate by using a focused laser beam to shift specific atoms out of resonance with the global light field," he explains. "It doesn't really matter how big the light shift is, which means that this approach is more robust to variations in the laser. Being able to control small groups of atoms in this way is a lot faster than moving them around."

Another key issue is the number of single atoms that can be held securely in the tweezer array. Current roadmaps suggest that arrays containing 10,000 atoms could be realized by increasing the laser power, but scaling to higher numbers could prove tricky. "It's a challenge to get hundreds of watts of laser power into the traps while maintaining coherence across the array," explains Spong. "The entire array of traps should be identical, but imperfect optics makes it hard to make the traps around the edge work as well as those in the center."

Researchers working in the field believe the pace of progress is already propelling the technology toward the day when a neutral-atom quantum computer will outperform a classical machine.

Replenishing the atom supply

With that in mind, the team at Atom Computing has deployed additional optical technologies in its updated platform to provide a pathway to larger-scale machines. “If we wanted to go from 100 to 1,000 qubits, we could have just bought some really big lasers,” says Bloom. “But we wanted to get on a track where we can continue to expand the array to hundreds of thousands of atoms, or even a million, without running into issues with the laser power.”

The solution for Atom Computing has been to combine the atomic control provided by optical tweezers with the trapping ability of optical lattices, which are most commonly found in the world’s most precise atomic clocks. These optical lattices exploit the interference of laser beams to create a grid of potential wells on the subwavelength scale, and their performance can be further enhanced by adding an optical buildup cavity to generate constructive interference between many reflected laser beams. “With these in-vacuum optics, we can create a huge array of deep traps with only a moderate amount of laser power,” says Bloom. “We chose to demonstrate an array that can trap 1,225 ytterbium atoms, but there’s no reason why we couldn’t go much higher.”

Importantly, in a modification of the usual rearrangement approach, this design also allows the atomic array to be continuously reloaded while the processor is being operated. Atoms held in a magneto-optical trap are first loaded into a small reservoir array, from which they are transferred into the target array that will be used for computation. The atoms in both arrays are then moved into the deep trapping potential of the optical lattice, where rapid and low-loss fluorescence imaging determines which of the sites are occupied. Returning the atoms to the optical tweezers then allows empty sites within the target array to be filled from the reservoir, with multiple loading cycles yielding an occupancy of 99%.

Repeatedly replenishing the reservoir with fresh atoms ensures that the target array is always full of qubits, which is essential to prevent atom loss during



A quantum engineer measures the optical power of a laser beam at Atom Computing’s research and development facility in Boulder, CO, USA.

Atom Computing

the execution of complex quantum algorithms. “Large-scale error-corrected computations will require quantum information to survive long past the lifetime of a single qubit,” Bloom says. “It’s all about keeping that calculation going when you have hundreds of thousands of qubits.”

While many challenges remain, researchers working in the field believe the pace of progress in recent years is already propelling the technology toward the day when a neutral-atom quantum computer will be able to outperform a classical machine. “Neutral atoms allow us to reach large numbers of qubits, achieve incredibly long coherence times and access novel error-correction codes,” says Bloom. “As an engineering firm, we are focused on improving the performance still further, since all that’s really going to matter is whether you have enough logical qubits and sufficiently high gate fidelities to address problems that are interesting for real-world use cases.” [OPN](#)

Susan Curtis (sxcurtis@gmail.com) is a freelance science and technology writer based in Bristol, UK.