nature portfolio

Peer Review File

Continuous operation of a coherent 3,000-qubit system

Corresponding Author: Professor Mikhail Lukin

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Referee #1

(Remarks to the Author)
The manuscript “Continuous operation of a coherent 3000 qubit system” reports a major breakthrough in the design of neutral-atom in tweezers experiment with wide and important implications for this community. Together with the first demonstration of a single atom prepared in an optical tweezers in 2001 (Nature) by the group of Grangier, and the “atom-by-atom assembly” technique showing the deterministic preparation of arrays of 50 atoms reported in 2016 (Science) by the groups of Browaeys and Lukin, this new performance will become one of the major landmarks of this field and promote ever more exciting developments for neutral atoms in quantum technologies than we have seen in the last 20 years.
This new achievement is that atomic arrays can now operate in a “continuous” mode, where fresh atoms are always made available for quantum experiments. This is to be contrasted to the “pulsed” mode in which these experiments have mostly always been conducted, where atoms are first prepared by loading them from a vapor to a MOT to tweezers, before being made available to the quantum phase of the experimental cycle. Operating in a “continuous” mode where both phases ran parallel is made challenging by the opposing requirements of the preparation and quantum phase. In the preparation phase, laser-induced dissipation is required to (1) cool the motion of atoms such that they can be trapped in a microscopic tweezers (2001 landmark), to (2) image them such that the entropy related to the random presence or absence of a single atom can be removed with a re-arrangement procedure (2016 landmark), and finally to (3) initialize them in a well-defined internal spin state (optical pumping). In contrast, in the quantum phase, the quantum information stored in the atom has to be absolutely protected from such laser-induced dissipation. This 2025 landmark now reports how both experimental phases can be run at the same time by a combination of simple and extremely effective design choices and tools (the dual-lattice conveyor, the dark loading, the shelving beam, the finite-field cooling/imaging). As well referred to in the text, some of these tools had already been demonstrated in other settings, but combining all of them and tuning them to the purpose of continuous coherent operation is a brilliant world-first.
Finally, going beyond the raw results, we would like to also state that the manuscript was a pleasure to read, with an introduction accessible and instructive to the wide readership of Nature, useful references and comparisons to recent related developments in continuous operation, and a wealth of details offered in the Methods for the specialists of the field. The figures (and the movie !) are impactful and will surely contribute to further inspire and attract the young researchers needed to keep pushing this field.
In conclusion, we strongly encourage publication of this work in Nature. We have only one minor comment about a section of the Methods that should be considered in the revision from the authors.

Comment 1: Methods “Tweezers Loading from the Lattice Reservoir”
The authors suggest 2 possible mechanisms to explain the loading of the atoms inside the tweezers from the reservoir. The question arises as tweezers are usually loaded in presence of cooling light (ensuring that an atom falling in the tweezers potential can get its kinetic energy efficiently removed and stuck at the bottom of the trap), while here the tweezers are loaded in “the dark”: no cooling light is applied.
For each mechanism, the authors write a formula for the loading probability. While the effort is appreciated and useful, I would like to ask the authors to go one step further and present their numerical evaluation with realistic parameters. This would allow the authors to conclude on which mechanisms dominate, and for the readership to extract a requirement for future experimental design on the reservoir density for efficient “dark loading”.
When we performed this evaluation with a tweezers size of Vtz ~ (1 um)^3 , we found out that only the “statistical loading” would give a sizable loading probability (of order 1), while the “collisional loading” is 2 orders of magnitude smaller (taking

the cross-section of Rb at 3x10^-16 m^2, a relative velocity of 0.2 m/s, a loading time of 1 ms).

Minor typo:
The main text states that a reservoir extraction “takes less than 2 ms”, while it is 2.5 ms in the ED Fig 5.

Referee #2

(Remarks to the Author)
I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Referee #3

(Remarks to the Author)
Chiu, et al demonstrate a new experimental architecture that adds rapid tweezer reloading to the well-established Rydberg atom array, enabling future neutral atom quantum processors to operate (in-principle) indefinitely. Neutral atom systems have been argued and demonstrated to be remarkably scalable compared to competing systems, however the trap lifetimes of neutral atoms have posed a potent challenge to implementing high-depth quantum circuits. Without an advance such as that described in this manuscript, circuit depth would come at the expense of circuit width, as some fraction of the atoms would be held in reserve to repopulate lost atomic qubits. And due to the exponential scaling of atom loss with time, such techniques are not attractive. Addressing this field-wide challenge, this manuscript demonstrates an effective way to laser cool and tweezer load new atoms while holding onto previously trapped atoms. Critically, the authors demonstrate that the existing atoms remain coherent (with dynamical decoupling) and so are reasonably unaffected by this potentially disruptive process.

The manuscript is well written and the data is of high quality. The primary contributions are the demonstration of an `enabling technology’ more than a demonstration/investigation of new science; however, in the details the authors describe several new techniques that enabled their success. Much work is still to be done to show the preservation of entanglement during reloading and the operation of a full gate-set. However, given that Nature has published other significant `enabling technology’ papers recently, I believe that this manuscript to be appropriate to the journal.

I do have one major concern with the presentation in the manuscript, regarding the frequent description of the reloading process as continuous. For instance, the abstract mentions “high-rate, continuous reloading”. Likewise in line 50, line 64, the section header line 69, titles of figures 1 and 4 etc. There are many more instances throughout the manuscript that would become tedious to list. The reloading process is a pulsed/cyclic process, as clearly shown in figure 4a, ED figure 1 and 9. The authors state in the methods section that their apparatus can, “deliver a fresh reservoir cloud to the science region every ≈ 150 ms,” and there is a 21 ms period every cycle during which no atoms are available in the reservoir. The authors can certainly claim that this enables continuous operation of the resulting array, since the pulsed reload is much faster than the qubit lifetime. However, in comparison to continuous-wave vs pulsed lasers or indeed truly continuous cold atom ensemble preparation (e.g. Cline…Thompson PRL 134, 013403, 2025), it appears just inaccurate to call this system continuous reloading.

If the authors can edit their manuscript to address the major concern listed above as well as the various additional comments below, I can recommend publication in Nature.

Additional comments

1. It is very impressive to mitigate MOT scattering on the qubit coherence time. But what are the coherence time numbers without dynamical decoupling? What is the impact of the dynamical decoupling-sequence on future qubit operations (i.e. is it compatible with all of the various gates that might be employed)? Or asked another way, does this impose assumptions/limits on future operations to be compatible with the dynamical decoupling sequence (especially in the preservation of entanglement)?
2. Is there a differential light shift on the qubit states due to the light-shifting beam? Presumably the DD-sequence can cancel this but I'm interested to know the magnitude of the shift that must be removed.
3. In general, I would like to know more about the light-shifting beam. How far do you light shift the excited state and at what intensity/Rabi frequency? Is the drive resonant? Or if detuned, how far detuned? How far must you detuned the excited state in order to suppress scattering? Some details are provided in the methods about overall power (but not beam area) and guesswork can be made about detuning due to the laser not being locked, but given the importance of this technique, more information is warranted.
4. Does it matter if you light-shift/shield the qubit atoms or the atoms in the preparation zone? It would seem to scale better if the storage atoms could be left alone while shifting only the atoms in the preparation zone and detuning the imaging to match.
5. Do atom lifetime and coherence time match an expectation due to Raman scattering at the operational trap depths? In general, the discussion of lifetimes and coherence times does not make reference to an ab-initio expectation but rather only makes comparisons between different experimental configurations, primarily to show that the reloading sequence does not make them dramatically worse. Is the source of finite lifetimes and coherence times in the absence of reloading fully understood? Are there future changes that could be made to improve them?
6. What is the temporal stability of this system? Various paths need careful relative alignment and may drift over time. For instance, I guess (though confirmation would be appreciated) that the “push-out pulse” mentioned in line 426 is formed by a

Hermite-Gaussian beam with a nodal plane aligned to the atom plane – small misalignments of this optical system could dramatically reduce atom occupation. Is any unusual or special effort taken to ensure long-term passive stability? Or perhaps this is recalibrated manually fairly frequently?
7. In figure 4b, there appears to be a residual upward-sawtooth fluctuation with a period around 20 or 25 seconds. Is this real or a statistical fluke? If it is real, what's going on?
8. In figure 4c, the exponential sawtooth guides-to-the-eye are obscuring some significant behavior. For instance, by late times there seems to be a systematic offset. For the red sawtooths starting around 2 seconds, the lightest red points over overlayed with a dark red curve. Presumably the period of the sub-array reloading is known precisely and hard-coded into these guides-for-the-eye. But then, how does this curve get offset by a full fifth of a period after only 4 cycles?
9. In figure 4d, what limits the coherence at short times? The lightest-blue curve appears lower contrast than expected from the lifetime dependence.
10. Line 492 typo: “regular tweezer spacings of 9 m” should be 9 um.

Referee #4

(Remarks to the Author)
The manuscript “Continuous operation of a coherent 3,000-qubit system” by Chiu et al. presents an experimental architecture for a large-scale atom array system with continuous operation, a high reloading rate, and coherent storage and manipulation of quantum information. In particular, persistent refilling of the atom array with atoms in either a spin-polarized or coherent superposition state is demonstrated while preserving the quantum state of the stored qubits. To accomplish this, the authors make use of two optical lattice conveyor belts to transport atom reservoirs between experimental zones, and the separation of these zones, along with shielding, allows for continuous preservation of the superposition state. The 87Rb atoms start in a MOT and are loaded into the first optical conveyor, in which they are transported through a differential pumping tube and into a separated science chamber. There, they are transferred into a second optical conveyor that takes them into the microscope field of view. Next, atoms are repeatedly loaded from the reservoir into an AOD-generated dynamic optical tweezer array in which qubits are prepared. Finally, they are assembled in a static, SLM-generated tweezer array using an iterative technique.

The authors demonstrate the ability to maintain over 3,000 atoms in the array for over two hours of continuous operation. Atoms are sequentially ejected and the sub-arrays are refilled while the lattice reservoir of atoms is replaced every other tweezer extraction cycle. This allows atom replenishment at a rate much faster than the tweezer-limited lifetime. Coherence of existing qubits is preserved by shifting the excited state using a shielding beam and by using a constant finite magnetic field in the preparation zone. In addition, they find that applying dynamic decoupling, combined with the MOT being far from the tweezer array area, allows for minor modification to the coherence time of the qubits.

The mid-circuit replacement of lost atoms is a key capability to enable complex, error-corrected quantum computations with neutral atom arrays. Beyond applications in quantum computing, this work also could be useful for overcoming limitations in quantum metrology and quantum networking.

This work is of immediate interest to the quantum computing community. It shows that scaled neutral atom systems don’t suffer from their need of reloading atoms, i.e. coherence can be preserved despite reloading. No other qubit technology offers the same scaling advantage.
The abstract and introduction are clear and accessible to a broader audience, and proper context for the work is given. There are no flaws in the methodology and the approach is valid. The data is presented in a clear manner, and statistics are used appropriately throughout the manuscript. Relevant literature has been cited, and there are no additional experiments or data necessary to strengthen the work. The conclusions and outlook of the paper and data interpretation are appropriate and valid.

The authors present impressive results written in a clear manner, and this work will definitely be useful to the community. The main results, that culminate in continuous existence of a superposition state, were obtained in parallel (as noted by the authors) by J. A. Muniz et al., arXiv:2506.09936 (2025) and Y. Li et al., arXiv:2506.15633 (2025). These results come as a combination of two established methods in the community: continuous large-scale reloading of atom arrays, demonstrated in M. A. Norcia et al., PRX Quantum 5, 030316 (2024) and F. Gyger et al., Phys. Rev. Research 6, 033104 (2024)., and replenishment while maintaining qubit coherence, demonstrated in M. A. Norcia et al., Phys. Rev. X 13, 041034 (2023).

In our view, the results presented in this manuscript are an impressive technical achievement. However, the main components have been shown by other groups before, so we have doubts as to whether Nature is the correct place to publish this work.

We have no major corrections, just a few minor comments, listed below:

In the section “Dual-Lattice Optical Transport”:
Are the atoms slowed down to standstill before the handover from Lattice 1 to Lattice 2 or is the handover happening “in flight”?

In Extended Data Figure 1b:
Why “anti-helmholtz” instead of simply “quadrupole”?

In the text it is mentioned that Lattice 2 is pulled 1cm back and then pushed forward again. This is not visible in the lattice speed graph. Is this because the speed is so slow compared to the usual transfer speed?
“while Lattice-2 is shipped to…where it serves” -> “while atoms in Lattice-2 are shipped to…where they serve”

Referee #5

(Remarks to the Author)
I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Referee responses and summary of revisions

Nature Manuscript number: 2025-06-15929 Chiu

July 30, 2025

We would like to thank all Referees for their thorough reading of our manuscript, helpful questions, and positive evaluations. In what follows we address comments point by point and outline changes made.

Referee: 1

The manuscript "Continuous operation of a coherent 3000 qubit system" reports a major breakthrough in the design of neutral-atom in tweezers experiment with wide and important implications for this community. Together with the first demonstration of a single atom prepared in an optical tweezers in 2001 (Nature) by the group of Grangier, and the "atom-by-atom assembly" technique showing the deterministic preparation of arrays of 50 atoms reported in 2016 (Science) by the groups of Browaeys and Lukin, this new performance will become one of the major landmarks of this field and promote ever more exciting developments for neutral atoms in quantum technologies than we have seen in the last 20 years.

This new achievement is that atomic arrays can now operate in a "continuous" mode, where fresh atoms are always made available for quantum experiments. This is to be contrasted to the "pulsed" mode in which these experiments have mostly always been conducted, where atoms are first prepared by loading them from a vapor to a MOT to tweezers, before being made available to the quantum phase of the experimental cycle. Operating in a "continuous" mode where both phases ran parallel is made challenging by the opposing requirements of the preparation and quantum phase. In the preparation phase, laser-induced dissipation is required to (1) cool the motion of atoms such that they can be trapped in a microscopic tweezers (2001 landmark), to (2) image them such that the entropy related to the random presence or absence of a single atom can be removed with a re-arrangement procedure (2016 landmark), and finally to (3) initialize them in a well-defined internal spin state (optical pumping). In contrast, in the quantum phase, the quantum information stored in the atom has to be absolutely protected from such laser-induced dissipation. This 2025 landmark now reports how both experimental phases can be run at the same time by a combination of simple and extremely effective design choices and tools (the dual-lattice conveyor, the dark loading, the shelving beam, the finite-field cooling/imaging). As well referred to in the text, some of these tools had already been demonstrated in other settings, but combining all of them and tuning them to the purpose of continuous coherent operation is a brilliant world-first.

Finally, going beyond the raw results, we would like to also state that the manuscript was a pleasure to read, with an introduction accessible and instructive to the wide readership of Nature, useful references and comparisons to recent related developments in continuous operation, and a wealth of details offered in the Methods for the specialists of the field. The figures (and the movie !) are impactful and will surely contribute to further inspire and attract the young researchers needed to keep pushing this field.

In conclusion, we strongly encourage publication of this work in Nature.

We thank the Referee for their detailed understanding and positive evaluation of our manuscript.

We have only one minor comment about a section of the Methods that should be considered in the revision from the authors.

Comment 1: Methods "Tweezers Loading from the Lattice Reservoir"
The authors suggest 2 possible mechanisms to explain the loading of the atoms inside the tweezers from the reservoir. The question arises as tweezers are usually loaded in presence of cooling light (ensuring that an atom falling in the tweezers potential can get its kinetic energy efficiently removed and stuck at the bottom of the trap), while here the tweezers are loaded in "the dark": no cooling light is applied. For each mechanism, the authors write a formula for the loading probability. While the effort is appreciated and useful, I would like to ask the authors to go one step further and present their numerical evaluation with realistic parameters. This would allow the authors to conclude on which mechanisms dominate, and for the readership to extract a requirement for future experimental design on the reservoir density for efficient "dark loading".

When we performed this evaluation with a tweezers size of Vtz $\sim$ (1 um)^3, we found out that only the "statistical loading" would give a sizable loading probability (of order 1), while the "collisional loading" is 2 orders of magnitude smaller (taking the cross-section of Rb at 3x10^{-16} m^2, a relative velocity of 0.2 m/s, a loading time of 1 ms).

We thank the Referee for this comment. Per our estimates, the average number of atoms statistically overlapping with the tweezer volume is ~ 5, and the lattice-averaged collision rate occurring within the tweezer volume is $\approx 300\,\mathrm{Hz}$. Here, similar to the Referee's response, we assume a scattering cross-section of

$\sigma = 3.5 \times 10^{-16}\,\mathrm{m}^2$ and relative velocity of $v_{\mathrm{rel}} \approx 0.25\,\mathrm{m/s}$, but a larger cubic tweezer volume $V_{\mathrm{twz}} = 8w_0^2 z_R \approx 10^{-17}\,\mathrm{m}^3$, with $w_0 = 800\,\mathrm{nm}$ the tweezer waist and z_R the Rayleigh length. For stochastic loading, we estimate a lattice density of $\langle n(r=0)\rangle_{\mathrm{lat}} \approx 5 \times 10^{11}\,\mathrm{cm}^{-3}$, where $\langle\,\rangle_{\mathrm{lat}}$ denotes averaging across the axial lattice sites. This is justified because the tweezer captures multiple lattice sites simultaneously ($2w_0 \gg \lambda_{\mathrm{lat}}/2$) which we account for by averaging axially. In the radial direction, the lattice extent is much larger than the tweezer waist, such that we can assume peak density at $r = 0$. Similarly, the density of atomic collisions is $\langle\gamma(r=0)\rangle_{\mathrm{lat}} = \frac{1}{2} v_{\mathrm{rel}}\sigma\langle n(r,z)\rangle_{\mathrm{lat}}^2 \approx 3 \times 10^{19}\,\mathrm{s}^{-1}\mathrm{m}^{-3}$. We have explicitly specified this calculation now in the Methods section:

Stochastic loading: **When turning on the tweezer, the number of reservoir atoms stochastically overlapped with the tweezer volume can be approximated as $N_{st} \propto V_{\mathrm{twz}}\langle n(r=0)\rangle_{\mathrm{lat}}$, where we assume peak density in the radial direction and $\langle\,\rangle_{\mathrm{lat}}$ denotes averaging across the axial lattice sites. These are valid approximations, as the tweezers are far smaller than the radial dimension of the reservoir but capture multiple lattice sites axially. For our reservoir parameters, we estimate an atom density of $\langle n(r=0)\rangle_{\mathrm{lat}} \approx 5 \times 10^{11}\,\mathrm{cm}^{-3}$ and tweezer volume $V_{\mathrm{twz}} \approx 10^{-17}\,\mathrm{m}^3$, which results in $\sim$5 atoms stochastically overlapped with the tweezer volume.**
Collisional loading**: An atom moving through a background gas collides with a rate $\Gamma(r,z) = n(r,z)v_{\mathrm{rel}}\sigma$, with v_{rel} the thermal relative velocity between two atoms, and σ the scattering cross section. The two-body collision density is given by $\gamma(r,z) = \frac{1}{2}n(r,z)^2 v_{\mathrm{rel}}\sigma$, such that the collision rate within the tweezer volume is proportional to $\Gamma_{\mathrm{twz}} \propto V_{\mathrm{twz}}\langle\gamma(r=0)\rangle_{\mathrm{lat}}$ using the same approximations as above. For our system parameters, we estimate $\langle\gamma(r=0)\rangle_{\mathrm{lat}} \approx 3 \times 10^{19}\,\mathrm{m}^{-3}\mathrm{s}^{-1}$ and therefore loading of $\sim$0.3 atoms per millisecond of tweezer-lattice overlap, assuming one atom trapped per atomic collision.**

These results confirm the Referee's conclusion that stochastic overlap of atoms with the tweezer volume is likely the dominant loading mechanism in our parameter regime. However, we would like to point out two simplistic assumptions this model is making: (a) While spatial overlap is necessary, it is not a sufficient condition for trapping (which would also require consideration of the atom's kinetic energy relative to the position in the potential of the tweezer trap); this leads to a decrease of stochastically captured atoms. (b) We expect the atomic collision rate within the tweezer volume to increase due to an acceleration of nearby atoms towards the tweezer trap center as the tweezer is turned on; this could result in a larger number of atoms loaded via collisions. As such, we expect a similar order of magnitude for both contributions. These dynamics are rich in physics, and we are currently working on a careful characterization with a follow-up publication planned. We have added the following clarifications to the corresponding Methods section:

While these order-of-magnitude estimates imply that stochastic loading is dominating, we would like to point out two simplistic assumptions this model is making: (a) While spatial overlap is necessary, it is not a sufficient condition for trapping; additionally considering kinetic energy constraints would lead to a decrease of stochastically captured atoms. (b) Due to the abruptly altered potential landscape on tweezer turn-on, nearby atoms accelerate toward the trap center and the atomic collision rate within the tweezer increases; this results in a larger number of atoms loaded via collisions[1]. For these reasons, we expect both mechanisms to contribute and further investigation is warranted.

Based on these revisions we have adapted the preceding text, and also shortened, reordered and added further clarifications in the text following which is included below:

In the regime of few tweezer extractions (i.e., saturated loading, see also Fig. 1b), we find that, on average, ≈ 5 atoms are lost from the reservoir per tweezer extraction, which is of the same order as the combined tweezer loading estimates above. This number is measured by comparing reservoir atom number with the number of atoms obtained in tweezers from repeated extraction, accounted for parity-projection. It should be noted that this serves as an upper bound and does not necessarily imply *loading* of 5 atoms per tweezer, since we expect the rather invasive extraction procedure to accelerate evaporative losses in the reservoir as well. After loading, captured atoms are transported out of the active reservoir perpendicular to the axial lattice potential. As shown in Extended Data Fig. 4b, we observe no difference in survival when moving atoms through the reservoir versus in free-space as a function of tweezer velocity[2].

Minor typo: The main text states that a reservoir extraction "takes less than 2 ms", while it is 2.5 ms in the ED Fig 5.

We thank the Referee for pointing this out and have corrected this sentence accordingly.

Referee: 3

Chiu, et al demonstrate a new experimental architecture that adds rapid tweezer reloading to the well-established Rydberg atom array, enabling future neutral atom quantum processors to operate (in-principle) indefinitely. Neutral atom systems have been argued and demonstrated to be remarkably scalable compared to competing systems, however the trap lifetimes of neutral atoms have posed a potent challenge to implementing high-depth quantum circuits. Without an advance such as that described in this manuscript, circuit depth would come at the expense of circuit width, as some fraction of the atoms would be held in reserve to repopulate lost atomic qubits. And due to the exponential scaling of atom loss with time, such techniques are not attractive. Addressing this field-wide challenge, this manuscript demonstrates an effective way to laser cool and tweezer load new atoms while holding onto previously trapped atoms. Critically, the authors demonstrate that the existing atoms remain coherent (with dynamical decoupling) and so are reasonably unaffected by this potentially disruptive process.

The manuscript is well written and the data is of high quality. The primary contributions are the demonstration of an 'enabling technology' more than a demonstration/investigation of new science; however, in the details the authors describe several new techniques that enabled their success. Much work is still to be done to show the preservation of entanglement during reloading and the operation of a full gate-set. However, given that Nature has published other significant 'enabling technology' papers recently, I believe that this manuscript to be appropriate to the journal.

We thank the Referee for their careful reading and positive evaluation of our manuscript.

I do have one major concern with the presentation in the manuscript, regarding the frequent description of the reloading process as continuous. For instance, the abstract mentions "high-rate, continuous reloading". Likewise in line 50, line 64, the section header line 69, titles of figures 1 and 4 etc. There are many more instances throughout the manuscript that would become tedious to list. The reloading process is a pulsed/cyclic process, as clearly shown in figure 4a, ED figure 1 and 9. The authors state in the methods section that their apparatus can, "deliver a fresh reservoir cloud to the science region every $\approx$ 150 ms," and there is a 21 ms period every cycle during which no atoms are available in the reservoir. The authors can certainly claim that this enables continuous operation of the resulting array, since the pulsed reload is much faster than the qubit lifetime. However, in comparison to continuous-wave vs pulsed lasers or indeed truly continuous cold atom ensemble preparation (e.g. Cline. . . Thompson PRL 134, 013403, 2025), it appears just inaccurate to call this system continuous reloading.
If the authors can edit their manuscript to address the major concern listed above as well as the various additional comments below, I can recommend publication in Nature.

We thank the Referee for raising this point. We agree that the cyclic reloading protocol *enables* continuous operation of the array, and agree that this is an important distinction to make. We have carefully reviewed any reference of the text to "continuous" reloading and adapted the wording according to the Referee's suggestion.

Additional comments
1. It is very impressive to mitigate MOT scattering on the qubit coherence time. But what are the coherence time numbers without dynamical decoupling? What is the impact of the dynamical decoupling-sequence on future qubit operations (i.e. is it compatible with all of the various gates that might be employed)? Or asked another way, does this impose assumptions/limits on future operations to be compatible with the dynamical decoupling sequence (especially in the preservation of entanglement)?

We thank the Referee for this helpful question. Without dynamical decoupling, we measure a site-resolved coherence time $T_2^* \approx 5\,\text{ms}$, limited by finite atomic temperature and differential light shifts from the tweezer light [3]. Due to the short time scale, we observe no effect of MOT scattered light on the T_2^*.

In regards to future operation, we do not anticipate any fundamental incompatibilities between our operational protocols and the use of dynamical decoupling sequences. As demonstrated in [4], an entangled Bell pair can be preserved during tweezer transport when dynamical decoupling pulses are applied, maintaining both coherence and entanglement. The observed parity oscillations confirm that two-atom entanglement remains robust throughout the transport and dynamical decoupling process. Furthermore, in [5], dynamical decoupling pulses are applied consistently throughout the entire quantum circuit, including between entangling gates, transport, and storage periods.

2. Is there a differential light shift on the qubit states due to the light-shifting beam? Presumably the DD-sequence can cancel this but I'm interested to know the magnitude of the shift that must be removed.

We thank the Referee for this question. According to our calculations, the differential light shift on the qubit states due to the light-shifting beam is $\approx 5\,\text{Hz}$. This shift is much smaller than the differential light shift from tweezer trap inhomogeneity ($\approx 100\,\text{Hz}$), and both are negligible compared to the qubit Rabi drive ($\sim$ MHz). Therefore, in the context of single qubit gates, we observe no impact experimentally from differential light shift onto the qubit states due to the 1529 nm light-shifting beam.

3. In general, I would like to know more about the light-shifting beam. How far do you light shift the excited state and at what intensity/Rabi frequency? Is the drive resonant? Or if detuned, how far detuned? How far must you detuned the excited state in order to suppress scattering? Some details are provided in the methods about overall power (but not beam area) and guesswork can be made about detuning due to the laser not being locked, but given the importance of this technique, more information is warranted.

We thank the Referee for raising these questions. The 1529 nm light-shift beam is shaped with a passive holographic phase plate as described in Methods, which yields a $\pm 125\,\mu$m flattop region (vertical) with $\approx 50\,\mu$m Gaussian waist (horizontal). Given the $\sim 2.5\,$W of 1529 nm light on the atoms, this yields a peak intensity of $I_0 \sim 10^8\,\mathrm{W/m^2}$. The drive is near resonant (red-detuned from the $5P_{3/2} \rightarrow 4D_{5/2}$ by $\approx 1\,$GHz), which results in an Autler-Townes splitting $\pm\Delta_{\mathrm{AT}} \approx \pm 2\pi \times 10\,$GHz on the excited state. Although the laser is not frequency-locked, the measured frequency stability is $\pm 10\,$MHz (free-running) and therefore the variation of excited state splitting is minimal due to the relatively large detuning from the $4D_{5/2}$ resonance. We employed the ARC package for all calculations regarding excited state shifts [6].

The scattering suppression depends on the power and detuning of imaging/cooling light as well as the excited state light shift. For $\Delta_{\mathrm{cool}} = 2\pi \times 60\,$MHz cooling light detuning, we estimate a scattering rate suppression of $\approx (\Delta_{\mathrm{AT}}/\Delta_{\mathrm{cool}})^2 > 10{,}000$. We further note that, in practice, suppression could be larger due to dark resonances and EIT-type physics. Detailed understanding of these phenomena could be useful for further optimization of scattering rate suppression.

In the process of revisiting these calculations, we found to have underestimated the excited state light shift previously. Together with this revision, we have added the following details to the corresponding Methods section:

Throughout the experiment and particularly during qubit preparation, we protect the $5S_{1/2}$ ground-state qubits in the storage zone from near-resonant photon scattering with the $5S_{1/2} \rightarrow 5P_{3/2}$ transition by addressing the $5P_{3/2} \rightarrow 4D_{5/2}$ transition near 1529 nm[7]. This results in an Autler-Townes splitting $\pm\Delta_{\mathrm{AT}} \approx \pm 2\pi \times 10\,$GHz of the excited state, and therefore a suppression of scattering from nearby imaging/cooling light with detuning Δ_{cool} by a factor of $\sim (\Delta_{\mathrm{AT}}/\Delta_{\mathrm{cool}})^2 > 10{,}000$ ("shielding").

The shielding light is sourced from a single-frequency fiber laser (Connet CoSF-D) outputting up to 10 Watts at 1529 nm, which is passed through an AOM (G&H 3165-1) for fast switching control and fiber-coupled onto the experimental table. Similar to the Raman beam, we shape the shielding beam with a holographic phase plate (HOLO/OR

ST-356) to create a flat-top beam profile of $\approx 250\,\mu\text{m}$ along the vertical extent of the storage zone array with Gaussian beam waist $\approx 50\,\mu\text{m}$ horizontally and approximately $2.5\,\text{W}$ projected onto the atoms. The beam tails are knife-edged in an intermediate imaging plane to ensure no shielding crosstalk onto the preparation zone. The knife-edged flat-top beam profile is shown in Extended Data Fig. 7b.

In Fig. 3c, we demonstrate the shielding effect in a spectroscopy scan by stepping the 1529 nm-wavelength during simultaneous imaging and low-power shielding of storage zone atoms, and resolve the $4D_{5/2}$ resonance by suppression of global imaging signal. To further optimize shielding performance, we maximize the T_2-time of storage qubits as a function of shielding wavelength while locally imaging qubits in the preparation zone. In practice, the shielding light is operated free-running $\approx 1\,\text{GHz}$ red-detuned from the $5P_{3/2} \rightarrow 4D_{5/2}$ transition with a measured frequency stability of $\pm 10\,\text{MHz}$.

4. Does it matter if you light-shift/shield the qubit atoms or the atoms in the preparation zone? It would seem to scale better if the storage atoms could be left alone while shifting only the atoms in the preparation zone and detuning the imaging to match

We thank the Referee for this important question. Shifting the atoms in the preparation zone may significantly impact the efficiency of our cooling schemes that rely on closed transitions and dark states (in our case, PGC and EIT during qubit preparation). This is because the 1529 nm light introduces an admixture of the $4D_{5/2}$ state into the D_2 manifold via near-resonant coupling, effectively dressing the $5P_{3/2}$ level, which breaks the coherence conditions required for dark-state cooling and optical pumping. Furthermore, the relatively large excited state splitting of $\approx 2\pi \times 10\,\text{GHz}$ applied simultaneously with our optical pumping schemes would require a light-shift beam homogeneity of $< 1\%$ to ensure resonant addressing of all preparation zone atoms. Therefore, we found that light-shifting the storage zone qubits (which, as mentioned previously, has negligible impact on the qubit states) is much more practical experimentally.

5. Do atom lifetime and coherence time match an expectation due to Raman scattering at the operational trap depths? In general, the discussion of lifetimes and coherence times does not make reference to an ab-initio expectation but rather only makes comparisons between different experimental configurations, primarily to show that the reloading sequence does not make them dramatically worse. Is the source of finite lifetimes and coherence times in the absence of reloading fully understood? Are there future changes that could be made to improve them?

We thank the Referee for this clarifying question. Our estimates of Raman scattering support the experimental observations. Using the Kramers-Heisenberg formula (Eq. 8 in Ref [8]), we estimate the Raman scattering rate from the tweezer light and find it to be consistent with the experimentally measured T_1 time (without reloading). We have added this to the caption of main Fig. 3:

b, A similar comparison probing depolarization of qubits initialized in $|1\rangle$ with the reference measurement $T_1 = 12.6(1)\,\mathrm{s}$ (gray), consistent with Raman scattering calculations due to the tweezer light[8].

The same equation is used to estimate the Raman scattering from our Raman light, which matches our experimentally observed scattering rate. A potential route to extending T_1 is to use more detuned tweezer light, as the Raman scattering rate scales as Δ^{-3} for a fixed trap depth in the far-detuned regime, where Δ is the detuning of the trap wavelength. As noted in [4], coherence can be further improved by increasing tweezer detuning (since differential light shifts scale as Δ^{-1} at constant trap depth), achieving lower atomic temperatures [9, 10], reducing trap depth [11], and shielding magnetic field fluctuations.

6. What is thhe temporal stability of this system? Various paths need careful relative alignment and may drift over time. For instance, I guess (though confirmation would be appreciated) that the "push-out pulse" mentioned in line 426 is formed by a Hermite-Gaussian beam with a nodal plane aligned to the atom plane – small misalignments of this optical system could dramatically reduce atom occupation. Is any unusual or special effort taken to ensure long-term passive stability? Or perhaps this is recalibrated manually fairly frequently?

We thank the Referee for this question. Indeed, all components of our system need careful and routine alignment to ensure optimal performance. In terms of addressing the preparation zone with imaging/cooling light as well as for the Talbot push-out pulse mentioned on line 426, we overlap light at multiple frequencies into an optical fiber, and then utilize a single beam path on the experiment table (one for the forward direction, one for the backward direction) as described in Methods. Therefore, all lights in the preparation zone (including the push-out pulse) are regular TEM_{00} Gaussians and the push-out pulse broadly addresses all atoms, both trapped in-plane and out-of-plane. Yet, the out-of-plane atoms are so weakly confined that they leave quickly ($< 30\,\mu\mathrm{s}$) while the in-plane atoms remain trapped when PGC is applied simultaneously. We have edited the corresponding Methods section to stress this point:

Additionally, we observe atoms trapped in weak out-of-plane potentials created by the SLM (Talbot effect), which we remove with a

brief resonant push-out pulse[11] applied to the entire array during PGC.

In regards to system stability, the beams addressing the preparation zone are aligned via optimization of imaging signal as part of our daily calibration protocol, which also includes alignment of the Raman beam to the storage zone array via row-by-row inhomogeneity in fitted Rabi frequency (see Extended Data Fig. 7a). While, in practice, we have found that beam alignment is relatively stable, we observe that overlap of tweezer arrays generated by AODs and SLMs can drift significantly dependent on the experiment duty cycle, likely due to thermalization of the tweezer objective. In this work, this was particularly apparent when operating continuously for multiple hours (see Fig. 2c). In the future, continuous operation for durations beyond what is demonstrated will require automated calibration within an experimental cycle, which, for example, can be done with a small part of the array reserved for "calibration atoms" used for live feedback. In the conclusion of the main text, we note that:

In addition, while the present experiments demonstrate continuous operation for over two hours, achieving much longer operation would benefit from active stabilization of the SLM-AOD overlap or automated beam alignment procedures.

7. In figure 4b, there appears to be a residual upward-sawtooth fluctuation with a period around 20 or 25 seconds. Is this real or a statistical fluke? If it is real, what's going on?

We thank the Referee for this observation. We believe that this is indeed simply a statistical effect, and it does not seem to appear in any other relevant measurements.

8. In figure 4c, the exponential sawtooth guides-to-the-eye are obscuring some significant behavior. For instance, by late times there seems to be a systematic offset. For the red sawtooths starting around 2 seconds, the lightest red points over overlayed with a dark red curve. Presumably the period of the sub-array reloading is known precisely and hard-coded into these guides-for-the-eye. But then, how does this curve get offset by a full fifth of a period after only 4 cycles?

We thank the Referee for this observation. The guide-to-eye curve is an exponential sawtooth function of the form $1 + f(x)$ (for starting in qubit state $|0\rangle$)

or $f(x)$ (for starting in $|1\rangle$), with $f(x)$ defined as

$$f(x) = a \left[\exp\left(-\frac{(x-b) \bmod c}{T_2} \right) - 1 \right], \tag{1}$$

with a, b, c, T_2 determined by our experiment. In this version, we have updated the guide-to-eye as suggested by the Referee such that, now, it is easier to observe slight differences in coherence times between subarrays. In particular, the last subarray (lightest red/blue curve) in a full replenishment cycle appears to have slightly longer T_2 coherence time than the other subarrays. These variations in relative coherence between subarrays are a point of further exploration with several possible causes, including slightly different AOD-SLM tweezer overlap precision or homogeneity across each individual subarray.

9. In figure 4d, what limits the coherence at short times? The lightest-blue curve appears lower contrast than expected from the lifetime dependence.

We thank the Referee for this clarifying question. After halting the reloading protocol under dynamical decoupling, all curves in Fig. 4d undergo an *additional* dynamical decoupling (DD) sequence of the form $\mathrm{X}_{\pi/2} - (\mathrm{XY16\text{-}64}) - \mathrm{X}_{-\pi/2,\phi}$, for which the final $\pi/2$ pulse phase ϕ is varied. This is described in the figure caption:

d, After multiple rounds of reloading under dynamical decoupling, we apply a final DD sequence and vary the phase of the last $\pi/2$-pulse to read out in different qubit bases. Complementary to (c), the observed coherence contrast varies for each subarray (color shading) due to the time-offset in subarray replenishment.

Consequently, the lightest-blue curve in Fig. 4d corresponds to qubits that have undergone a full 80-ms DD cycle already and experience expected dephasing according to the T_2-time. Basically, the contrast comparison between Fig. 4d and Fig. 4c should be drawn not from $t = 0.4\,\mathrm{s}$ but $t = 0.5\,\mathrm{s}$ in Fig. 4c, where the dark blue data point should instead be imagined at the bottom of its guide-to-the-eye curve.

10. Line 492 typo: “regular tweezer spacings of 9 m” should be 9 um.

We thank the Referee for bringing this point to our attention and have corrected the manuscript accordingly.

Referee: 4

The manuscript "Continuous operation of a coherent 3,000-qubit system" by Chiu et al. presents an experimental architecture for a large-scale atom array system with continuous operation, a high reloading rate, and coherent storage and manipulation of quantum information. In particular, persistent refilling of the atom array with atoms in either a spin-polarized or coherent superposition state is demonstrated while preserving the quantum state of the stored qubits. To accomplish this, the authors make use of two optical lattice conveyor belts to transport atom reservoirs between experimental zones, and the separation of these zones, along with shielding, allows for continuous preservation of the superposition state. The 87Rb atoms start in a MOT and are loaded into the first optical conveyor, in which they are transported through a differential pumping tube and into a separated science chamber. There, they are transferred into a second optical conveyor that takes them into the microscope field of view. Next, atoms are repeatedly loaded from the reservoir into an AOD-generated dynamic optical tweezer array in which qubits are prepared. Finally, they are assembled in a static, SLM-generated tweezer array using an iterative technique.

The authors demonstrate the ability to maintain over 3,000 atoms in the array for over two hours of continuous operation. Atoms are sequentially ejected and the sub-arrays are refilled while the lattice reservoir of atoms is replaced every other tweezer extraction cycle. This allows atom replenishment at a rate much faster than the tweezer-limited lifetime. Coherence of existing qubits is preserved by shifting the excited state using a shielding beam and by using a constant finite magnetic field in the preparation zone. In addition, they find that applying dynamic decoupling, combined with the MOT being far from the tweezer array area, allows for minor modification to the coherence time of the qubits.

The mid-circuit replacement of lost atoms is a key capability to enable complex, error-corrected quantum computations with neutral atom arrays. Beyond applications in quantum computing, this work also could be useful for overcoming limitations in quantum metrology and quantum networking.

This work is of immediate interest to the quantum computing community. It shows that scaled neutral atom systems don't suffer from their need of reloading atoms, i.e. coherence can be preserved despite reloading. No other qubit technology offers the same scaling advantage. The abstract and introduction are clear and accessible to a broader audience, and proper context for the work is given. There are no flaws in the methodology and the approach is valid. The data is presented in a clear manner, and statistics are used appropriately throughout the manuscript. Relevant literature has been cited, and there are no additional experiments or data necessary to strengthen the work. The conclusions and outlook of the paper and data interpretation are appropriate and valid.

The authors present impressive results written in a clear manner, and this work will definitely be useful to the community. The main results, that culminate in continuous existence of a superposition state, were obtained in parallel (as noted by the authors) by J. A. Muniz et al., arXiv:2506.09936 (2025) and Y. Li et al., arXiv:2506.15633 (2025). These results come as a combination of two

established methods in the community: continuous large-scale reloading of atom arrays, demonstrated in M. A. Norcia et al., PRX Quantum 5, 030316 (2024) and F. Gyger et al., Phys. Rev. Research 6, 033104 (2024)., and replenishment while maintaining qubit coherence, demonstrated in M.A. Norcia et al., Phys. Rev. X 13, 041034 (2023).
In our view, the results presented in this manuscript are an impressive technical achievement. However, the main components have been shown by other groups before, so we have doubts as to whether Nature is the correct place to publish this work.

We thank the Referee for their positive evaluation of our manuscript.

We have no major corrections, just a few minor comments, listed below: In the section "Dual-Lattice Optical Transport": Are the atoms slowed down to standstill before the handover from Lattice 1 to Lattice 2 or is the handover happening "in flight"?

We thank the Referee for this clarifying question and the interesting idea of "in-flight" atom handover. In our experiment, the Lattice 1 is brought to a complete standstill before atoms are handed over to Lattice 2.

In Extended Data Figure 1b:
Why "anti-helmholtz" instead of simply "quadrupole"?

We thank the Referee for this suggestion, and agree that "quadrupole" improves readability. We have updated Extended Data Fig. 1b accordingly.

In the text it is mentioned that Lattice 2 is pulled 1cm back and then pushed forward again. This is not visible in the lattice speed graph. Is this because the speed is so slow compared to the usual transfer speed?

We thank the Referee for this question. Originally, the "pull-back and push-forward" lattice movement technique in order to avoid atom spilling from the reservoir during preparation zone operations was included in Extended Data Fig. 1b but later omitted for clarity, indeed mostly because this "pull-and-push" movement applied to Lattice 2 is so short and at low speeds that it is barely visible if depicted accurately. Furthermore, this movement is more closely related to the local operations in the preparation zone; in contrast, Extended Data Fig. 1b is intended to depict the sequence related to the loading and transport stages of the conveyor belt lattices. We have edited the caption of

Extended Data Fig. 1b to reflect this stylistic choice (see next question).

"while Lattice-2 is shipped to... where it serves" → "while atoms in Lattice-2 are shipped to... where they serve"

We thank the Referee for this great suggestion. Together with the previous comment, we have implemented the following changes in the caption for Extended Data Fig. 1b accordingly:

Vacuum chamber and dual-lattice transport sequence. a, [...] b, Summary of the dual-lattice transport sequence, including the MOT stage, loading and cooling of Lattice-1 as well as lattice transport and handover. After the lattice handover, we restart the lattice loading procedure in the MOT chamber, while atoms in Lattice-2 are shipped to the science region where they serve as an atomic reservoir for tweezer extraction. In the Lattice-2 velocity graph, the brief back-and-forth movement used to avoid atom spilling during qubit preparation is omitted for clarity. The gray-shaded regions indicate the previous/next lattice loading cycle.

Addendum

We additionally thank Referees 2 and 5 for their comments and positive evaluation of our manuscript.

Additional changes

Here we outline other minor changes:

1. We updated the citation for [Bluvstein et al. 2025] and added journal information to other references where applicable.

2. We improved wording at various places throughout the manuscript, and added further cross-references to Figures to improve readability.

3. Where applicable, we added a disclaimer about error bars in the figure captions ("Error bars represent the standard error of the mean").

5. We revised our data availability statement and updated the COI to state the submission of a provisional patent.

References

[1] Comparat, D. et al. Optimized production of large Bose-Einstein condensates. Phys. Rev. A, 73:043410, Apr 2006.

[2] Gyger, F. et al. Continuous operation of large-scale atom arrays in optical lattices. Phys. Rev. Res., 6:033104, 2024.

[3] Levine, H. et al. Dispersive optical systems for scalable Raman driving of hyperfine qubits. Phys. Rev. A, 105:032618, Mar 2022.

[4] Bluvstein, D. et al. A quantum processor based on coherent transport of entangled atom arrays. Nature, 604(7906):451–456, Apr 2022.

[5] Bluvstein, D. et al. Logical quantum processor based on reconfigurable atom arrays. Nature, 626(7997):58–65, 2024.

[6] N. Šibalić, J. D. Pritchard, K. J. Weatherill, and C. S. Adam. ARC: An open-source library for calculating properties of alkali Rydberg atoms. Computer Physics Communications, 220:319–331, 2017.

[7] Hu, B. et al. Site-selective cavity readout and classical error correction of a 5-bit atomic register. arXiv, 2024.

[8] I. D. Moore, W. C. Campbell, E. R. Hudson, M. J. Boguslawski, D. J. Wineland, and D. T. C. Allcock. Photon scattering errors during stimulated Raman transitions in trapped-ion qubits. Phys. Rev. A, 107:032413, Mar 2023.

[9] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. Analysis of dephasing mechanisms in a standing-wave dipole trap. Phys. Rev. A, 72:023406, Aug 2005.

[10] Tamara Dordević, Polnop Samutpraphoot, Paloma L. Ocola, Hannes Bernien, Brandon Grinkemeyer, Ivana Dimitrova, Vladan Vuletić, and Mikhail D. Lukin. Entanglement transport and a nanophotonic interface for atoms in optical tweezers. Science, 373(6562):1511–1514, 2021.

[11] Manetsch, H. J. et al. A tweezer array with 6100 highly coherent atomic qubits. arXiv, 2024.