nature portfolio

Peer Review File

Experimental demonstration of logical magic state distillation

Corresponding Author: Dr Sergio Cantu

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)
Experimental demonstration of logical magic state distillation

Key results

The authors demonstrate magic state distillation with logical qubits, a key subroutine for fault-tolerant quantum computing. They use a neutral atom processor to prepare five logical qubits in the color code and then run a distillation circuit, obtaining distilled magic states with improved fidelity compared to the input magic states.

Originality and significance

The results presented are of immediate interest to many people in the quantum information community.
This is the first experimental demonstration of logical magic state distillation in any platform, and it underlines the advantages of neutral atom devices for fault-tolerant quantum computing.
The paper also contains theory results of independent interest such as the method for designing state injection circuits and the simulation approach for circuits with non-Clifford gates.

Data and methodology

The presentation is of excellent quality and the experimental results are well-presented.
My only real criticism of the work is that the details of the device are not included and are instead relegated to an "upcoming manuscript".
These are key details to omit and I think this detracts from the quality of the work.

Appropriate use of statistics

I find the treatment of statistics is appropriate.

Conclusions

The conclusions and the interpretation of the data are robust.

Suggested improvements

I have some very minor suggestions for improvements:
- Section "Extending to larger code distance": I find the sentence "We note that the stabilizer postselection rate required to achieve a comparably high fideltiy is lower compared with d=3, ...". I thought a higher postselection rate was required?
- Upon reading the main text (pre Methods), the difference between factory and stabilizer postselection was not entirely clear to me. I would recommend moving some of the explanation on page 14 (paragraph starting "Given either decoder, we first ...") earlier in the text.
- Page 13: in the second equation in the left column, what is $\mathcal V$? I cannot find a definition anywhere.
- Page 14: I find the notation in Eq. (6) difficult to parse and I am not sure the equation is well formed. The aim is to find some $\vec f$ that has syndrome $\vec \sigma$, but has a different logical action than $\vec e$? To be consistent with Eq. (4) then

you should have $P(\vec f)$ in Eq. (6)? But I am not convinced by having $\vec f$ on the LHS and RHS of the equation, as on the RHS it represents an arbitrary error but on the LHS it is a specific error (same issue with Eq. (4)).

References

https://arxiv.org/abs/quant-ph/0503243 for randomized benchmarking

Clarity and context

The abstract is clear and accessible and the introduction gives an appropriate overview.

Referee #2

(Remarks to the Author)
The authors present magic state distillation using neutral atoms encoded in color codes and the [[5,1,3]] code for distillation. The authors find a fidelity improvement in encoded magic states after the distillation protocol, but not in physical magical states after the distillation protocol.

Magic state distillation is a key element for fault-tolerant quantum computation. The authors experimentally show the advantage of magic state distillation using encoded qubits. The authors tradeoff quality of the final magic state fidelity with increased post-selection on the measured stabilizers of the logical qubits.

The data and methodology are clearly presented. The use of statistics is appropriate. The conclusions are reliable and the methods presented could be used by others.

Suggested improvements and questions:
In figure 3 d, is the stabilizer rejection rate constant for introduced angle error or does it change with angle? In an ideal setting it should be independent of angle. This pre-factory rejection rate should be mentioned in the caption.

The single qubit magic state fidelity is limited by the measurement error. Is it possible to understand the cross-over as a tradeoff between preparation fidelity and post-selected measurement fidelity?

References: The references are appropriate and the tweet reference is necessary.

Clarity: The paper is quite clear and the abstract lays out the key results of the work.

Referee #3

(Remarks to the Author)
The manuscript by Rodriguez et al. is an impressive demonstration of the use of dynamically reconfigurable optical tweezers to encode and perform quantum operations on many logical qubits. As also explained by the authors, a central challenge in realizing quantum advantage with quantum error-correcting codes is creating non-Clifford operations within the code framework. In this work, non-Clifford resources are realized at the logical level by logically encoding particular superposition states called magic states. The authors demonstrate the encoding of magic states (and also arbitrary superposition states) into a QEC code. The fidelity of the logical magic states is further increased by utilizing magic state distillation, following a prescription guided by Bravyi et al. (Physical Review A, 71(2):022316, 2005). The authors then demonstrate how their distillation protocol scales as a function of the data code distance.

The manuscript is clear, very well-written, and coupled with excellent figures and data presentation. The reported fidelities are exceptional for neutral atom quantum devices. The results from this manuscript are highly relevant in the march towards universal quantum computation and quantum advantage with quantum error-corrected systems. These results will likely influence future developments within the quantum computing and quantum information science community. The quality of presentation is excellent, and I have no additional comments related to typos, errors in calculations, or the methodologies used in the paper.

The main experimental platform is modeled after the experimental architecture used in Bluvstein et al. (Nature 626, 58-65, 2024), which shares many of the same authors of this manuscript. This was a landmark paper that first demonstrated the use of reconfigurable optical tweezers to encode and perform quantum operations on many logical qubits. This work also featured logical encoding of a superposition state (in a surface code), realization of non-Clifford operations (using small 3D codes), and measurement of the magic associated with non-Clifford operations. These previous achievements (from many of the same authors) tone down the significance of the results of the present manuscript. The magic state distillation demonstrated in the present paper is undoubtedly impressive. Still, one can argue that the new tools and insight needed for this breakthrough were demonstrated in Fig. 3. of Bluvstein et al. (where the fidelity of a logical GHZ state was improved using post-selection on logical ancilla qubit measurements).

While I think the results of this manuscript are certainly outstanding and one avenue for realizing a transversal universal gate set, I contextualize the present results as a smaller-scale, specific application of the new insights, tools, and results already demonstrated in Bluvstein et al. (Nature, 2024). This being said, Nature has published at least two recent manuscripts on logical magic state encoding with flag qubits: Lukas Postler, et al. Nature, 605(7911):675– 680, 2022 and Riddhi S. Gupta, et al. Nature, 625(7994):259–263, 2024, indicating that magic state preparation is of interest to its readers. Compared with flag qubits, I agree with the authors that performing magic state distillation at the logical level, as demonstrated in this work, provides substantial protection and is better suited for the low error rates needed in future quantum information processors. Given past precedent, I think it would be a disservice to the readership of Nature not to highlight magic state distillation performed at the logical level. Taking everything into consideration, I recommend this article for publication.

Referees' comments:

Referee #1 (Remarks to the Author):

Experimental demonstration of logical magic state distillation

Key results

The authors demonstrate magic state distillation with logical qubits, a key subroutine for fault-tolerant quantum computing.
They use a neutral atom processor to prepare five logical qubits in the color code and then run a distillation circuit, obtaining distilled magic states with improved fidelity compared to the input magic states.

Originality and significance

The results presented are of immediate interest to many people in the quantum information community.
This is the first experimental demonstration of logical magic state distillation in any platform, and it underlines the advantages of neutral atom devices for fault-tolerant quantum computing.
The paper also contains theory results of independent interest such as the method for designing state injection circuits and the simulation approach for circuits with non-Clifford gates.

Data and methodology

The presentation is of excellent quality and the experimental results are well-presented.
My only real criticism of the work is that the details of the device are not included and are instead relegated to an "upcoming manuscript".
These are key details to omit and I think this detracts from the quality of the work.

Appropriate use of statistics

I find the treatment of statistics is appropriate.

Conclusions

The conclusions and the interpretation of the data are robust.

We thank the referee for their positive and careful evaluation of our work. We are in the final stages of wrapping up our detailed technical characterization paper of our experimental system, and hope to post it on arXiv before the present manuscript is officially published. However, since the system design and control approach shares many similarities to prior work from our collaborators at Harvard (Bluvstein et al., Nature 2024), we believe that there is sufficient information for the community to reproduce our results. We have added an explicit reference to papers describing the dynamically reconfigurable architecture there, and included the most important experimental parameters in the present manuscript.

Suggested improvements

I have some very minor suggestions for improvements:
- Section "Extending to larger code distance": I find the sentence "We note that the stabilizer postselection rate required to achieve a comparably high fideltiy is lower compared with d=3, ...". I thought a higher postselection rate was required?

We meant that the rate of samples passing post-selection is lower, which is the same as what the referee means. We appreciate that our prior terminology can lead to different interpretations, so we have modified this sentence to now read “We note that more stabilizer postselection is required to achieve a comparably high fidelity as d=3, since our gate fidelities are not yet past the circuit threshold for the data code.”

- Upon reading the main text (pre Methods), the difference between factory and stabilizer postselection was not entirely clear to me. I would recommend moving some of the explanation on page 14 (paragraph starting "Given either decoder, we first ...") earlier in the text.

We thank the referee for this suggestion. We have now added a definition of both the factory postselection and stabilizer postselection in P4 of the main text.

- Page 13: in the second equation in the left column, what is $\mathcal V$? I cannot find a definition anywhere.

We thank the referee for catching this typo. The equation is calculating the probability that *m* and *n* are observed, conditional on the true Bloch vector being *v*. The $\mathcal V$ part is a typo and should be removed.

- Page 14: I find the notation in Eq. (6) difficult to parse and I am not sure the equation is well formed. The aim is to find some $\vec f$ that has syndrome $\vec \sigma$, but has a different logical action than $\vec e$? To be consistent with Eq. (4) then you should have $P(\vec f)$ in Eq. (6)? But I am not convinced by having $\vec f$ on the LHS and RHS of the equation, as on the RHS it represents an arbitrary error but on the LHS it is a specific error (same issue with Eq. (4)).

We thank the referee for this suggestion. We have updated the notation to address the issue that *e* meant different things on the two sides of the equation. In the updated version, we explicitly denote the most-likely error and second-most-likely error with a subscript for clarify.

References

https://arxiv.org/abs/quant-ph/0503243 for randomized benchmarking

We have added a reference to this paper.

Clarity and context

The abstract is clear and accessible and the introduction gives an appropriate overview.

Referee #2 (Remarks to the Author):

The authors present magic state distillation using neutral atoms encoded in color codes and the [[5,1,3]] code for distillation. The authors find a fidelity improvement in encoded magic states after the distillation protocol, but not in physical magical states after the distillation protocol.

Magic state distillation is a key element for fault-tolerant quantum computation. The authors experimentally show the advantage of magic state distillation using encoded qubits. The authors tradeoff quality of the final magic state fidelity with increased post-selection on the measured stabilizers of the logical qubits.

The data and methodology are clearly presented. The use of statistics is appropriate. The conclusions are reliable and the methods presented could be used by others.

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

Suggested improvements and questions:
In figure 3 d, is the stabilizer rejection rate constant for introduced angle error or does it change with angle? In an ideal setting it should be independent of angle. This pre-factory rejection rate should be mentioned in the caption.

The stabilizer acceptance rate is roughly constant, varying between 13-16%. The stabilizer acceptance rate is 13.2% for all data sets with added rotations, which were taken in an interleaved fashion, and 16% for the data set with no added rotation, which was taken on a different day with slightly better calibration. We have added a clarification of this acceptance rate in the caption now.

The single qubit magic state fidelity is limited by the measurement error. Is it possible to understand the cross-over as a tradeoff between preparation fidelity and post-selected measurement fidelity?

We thank the referee for this interesting question.

Both the input logical magic state and the distilled logical magic state undergo the same error correction procedure in our analysis, and we expect them to have the same sensitivity to measurement errors - the output logical magic state is generally intended to be used in a larger circuit, so we do not post-select on its stabilizers when evaluating the fidelity. The crossover between input logical fidelity and output logical fidelity occurs because by post-selecting on stabilizers of the distillation syndrome logical qubits, we can catch errors and improve the distilled output fidelity past the input fidelity without directly measuring the output logical qubit.

Compared to the direct preparation of a physical magic state, the error-corrected implementation does indeed have a higher measurement fidelity due to its ability to correct single measurement errors. When the physical fidelity is high enough, the preparation fidelity after distillation can also be better than a physical magic state, due to the quadratic suppression. In our case, performing post-selection on the distillation syndrome logical qubits improves the output fidelity of distillation enough to outperform the single physical qubit magic state fidelity.

References: The references are appropriate and the tweet reference is necessary.

Clarity: The paper is quite clear and the abstract lays out the key results of the work.

Referee #3 (Remarks to the Author):

The manuscript by Rodriguez et al. is an impressive demonstration of the use of dynamically reconfigurable optical tweezers to encode and perform quantum operations on many logical qubits. As also explained by the authors, a central challenge in realizing quantum advantage with quantum error-correcting codes is creating non-Clifford operations within the code framework. In this work, non-Clifford resources are realized at the logical level by logically encoding particular superposition states called magic states. The authors demonstrate the encoding of magic states (and also arbitrary superposition states) into a QEC code. The fidelity of the logical magic states is further increased by utilizing magic state distillation, following a prescription guided by Bravyi et al. (Physical Review A, 71(2):022316, 2005). The authors then demonstrate how their distillation protocol scales as a function of the data code distance.

The manuscript is clear, very well-written, and coupled with excellent figures and data presentation. The reported fidelities are exceptional for neutral atom quantum devices. The results from this manuscript are highly relevant in the march towards universal quantum computation and quantum advantage with quantum error-corrected systems. These results will likely influence future developments within the quantum computing and quantum information science community. The quality of presentation is excellent, and I have no additional comments related to typos, errors in calculations, or the methodologies used in the paper.

The main experimental platform is modeled after the experimental architecture used in Bluvstein et al. (Nature 626, 58-65, 2024), which shares many of the same authors of this manuscript. This was a landmark paper that first demonstrated the use of reconfigurable optical tweezers to encode and perform quantum operations on many logical qubits. This work also featured logical encoding of a superposition state (in a surface code), realization of non-Clifford operations (using small 3D codes), and measurement of the magic associated with non-Clifford operations. These previous achievements (from many of the same authors) tone down the significance of the results of the present manuscript. The magic state distillation demonstrated in the present paper is undoubtedly impressive. Still, one can argue that the new tools and insight needed for this breakthrough were demonstrated in Fig. 3. of Bluvstein et al. (where the fidelity of a logical GHZ state was improved using post-selection on logical ancilla qubit measurements).

While I think the results of this manuscript are certainly outstanding and one avenue for realizing a transversal universal gate set, I contextualize the present results as a smaller-scale, specific application of the new insights, tools, and results already demonstrated in Bluvstein et al. (Nature, 2024). This being said, Nature has published at least two recent manuscripts on logical magic state encoding with flag qubits: Lukas Postler, et al. Nature, 605(7911):675– 680, 2022 and Riddhi S. Gupta, et al. Nature, 625(7994):259–263, 2024, indicating that magic state preparation is of interest to its readers. Compared with flag qubits, I agree with the authors that performing magic state distillation at the logical level, as demonstrated in this work, provides substantial protection and is better suited for the low error rates needed in future quantum information processors. Given past precedent, I think it would be a disservice to the readership of Nature not to highlight magic state distillation performed at the logical level. Taking everything into consideration, I recommend this article for publication.

We thank the referee for their positive recommendation and detailed summary and comparisons. Indeed, while our demonstration builds on top of many of the same technical ingredients demonstrated in Bluvstein et al., it applies them to an important and novel setting (logical magic state distillation) of high relevance to future large-scale quantum computers, so we believe that it meets the high standards of publication for Nature.

Other changes:

In addition to the aforementioned changes to address the referee’s concerns, we have also improved our data analysis of the fidelities by incorporating the more widely accepted Bures measure for fidelity estimation (see Schmied et al., https://arxiv.org/pdf/1407.4759), which is known to address the bias of the uniform prior (Hilbert-Schmidt measure) towards low values of fidelity when the number of samples is low. This has led to a slight increase in the estimated fidelities, most notably for d=5 distillation, which has the smallest number of samples. We have updated the figures and description, and added an additional extended data figure illustrating the difference between various fidelity estimation methods.