

Peer Review File

Manuscript Title: Real-time quantum error correction beyond break-even.

Reviewer Comments & Author Rebuttals

Reviewer Reports on the Initial Version:

Referees' comments:

Referee #1 (Remarks to the Author):

The manuscript "Real-time quantum error correction beyond break-even" demonstrated a break-even-quantum error correction with the GKP encoding with 3D superconducting circuits. To achieve fault-tolerant quantum computers, precise quantum control is important whose precision is much higher than its threshold. For that purpose, the bosonic encoding plays an important role in the hardware efficient "pre"-error correction to overcome a physical limitation of the decoherence. The manuscript reports the break-even quantum error correction, where the coherence time of the corrected qubit is longer than its physical limitation. The method of the GKP encoding is so complicated that this procedure itself might generate large decoherence. The authors use a long coherent transmon ancilla and have investigated the optimization method to reduce the error of the error correction cycles.

The demonstration of break-even error correction is a really big milestone for fault-tolerant quantum computers. To achieve this, the authors have also investigated several techniques and they are well-studied based on many experimental data. The paper is also well-written and the results are clearly presented. I found these results are suitable for publication in nature. Still, I have some comments and question before publication:

1. In fig.2(a), what exactly are the subspace C1 - C3 and which kinds of errors are occurred? In the bosonic system, the dimension of the error-subspace should be infinite.

2. Page 4 line 7 in the first paragraph, "the agent corrects large errors in the initial parameter values, and in the subsequent few hundreds of epochs, it fine-tunes the circuit parameters to achieve the highest performance."

This

Does this qualitative explanation of behavior always hold? If so, is the optimization process dynamically changing?

3. Page 5 last paragraph, the authors evaluate the success probabilities (97%) of the error correction of their process. How is this high success probability consistent with the increased coherence time due to post-selection in the later paragraphs? If the error correction is perfect, the post-selection is supposed not to increase its coherence with correctable small errors. Is it right?

4. In method section and Supp. Infor S2F, the authors introduce the active oscillator cooling. The authors mention that this technique can also be realized in trapped ions. It might be true, but in the ion trap, the interaction of a $\sigma_+ + a^\dagger \sigma_-$ is easily achievable with irradiation of a laser beam, so called a red sideband transition. With the superconducting circuits, the cross-Kerr term is the original interaction of the dispersively interacted circuits. On the other hand, with the ion trap, the Jaynes-Cummings and anti-Jaynes-Cumming interaction is the origin between a bosonic mode and a qubit. This active cooling is widely used as "sideband cooling" in the ion trap.

5. What limits the error correction gain G for this experiment?

6. The coherence time of Pauli Y for the corrected qubit is shorter than other Pauli operators

because of the square shape of the grids. The maximum Pauli Y operator expectation value in Fig. 3c is smaller than others. Does it have same reason? This imperfection can be considered as a thermal excitation of the logical qubit. Then, can we prepare the better initial Pauli Y state from the initial Pauli X or Pauli Z state with oscillator manipulations?

7. Page 5 last paragraph, why the logical error probabilities per QEC cycle is the smallest for Y direction?

8. Are these processes applicable to the hexagonal grid state? Can the correction gain G improve with the hexagonal grid state?

9. How about the orthogonality of two QEC cycles? If the phase difference of two QEC steps is not exactly 90 degrees, then $R^{x_{\Delta}}$ might generate Z error.

10. Reference [36] should replace to the nature physics ver.

Referee #2 (Remarks to the Author):

The paper by Sivak et al reports on progresses on quantum error correction with Gottesman-Kitaev-Preskill codes, implemented in microwave cavities coupled to a transmon qubit.

The paper reports on an extended lifetime of the quantum information encoded in such stabilised GKP qubit by a factor of 2.27 as compared to the best passive qubit encoded in the same cavity. This improvement is substantial and yield a major breakthrough, as compared to compared to previous reported experimental results on GKP-quantum error-correction achieved in Ref.14, as well as on previous break even attempts with cat codes (Ofek et al). This improvement is made possible by the use of a new material (tantalum) for both transmon qubit and cavity, instead of aluminium, yielding higher coherence, together with the use of improved error-correction schemes (Refs 20,28), also combined with machine learning techniques for pulse design, reset of the auxiliary qubit, etc. I especially like that the long stabilisation time allows for testing the hypothesis of stationary errors, which seems to hold, with small deviations being accounted for transmon leakage to higher states.

The results reported are state-of-the-art and the technological improvements residing in the aspects listed above, yielding to such a substantial lifetime extension, are worth reporting in Nature.

The very extensive supplementary information section addresses in great detail the technical aspects of the paper and replies to most possible questions a reader might have. However I do have a couple of additional questions and a comment:

- In Fig 2 (c), why do the authors stop the training after slightly more than 500 epochs, while the reward has started to increase again and has not saturated?
- Do I understand correctly that the main limiting factor to push the lifetime extension to even higher gain factors resides in the leakage of the transmon? How does the squeezing of the GKP states stabilised come into play? Are there other limiting factors that prevent the authors from achieving $G = 10$ or 100 ?
- It would be interesting if the authors could comment on their perspectives regarding scalability, both in terms of the actual experimental platform (how to accommodate for several cavities? Or more generally how to encode several modes?), and as to whether the reinforcement learning protocol they use would be scalable to the scenario where multiple encoded qubits are considered.

- I would suggest to replace the nomenclature "ancilla"/"ancillary" with "auxiliary qubit" everywhere throughout manuscript and supplementary information

Author Rebuttals to Initial Comments:

We thank both Referees for their valuable time and for the detailed assessment of our work. We are excited that both Referees found our results suitable for publication in Nature. In this response letter, we answer their questions and describe the changes that we made to the paper. In the attached colored version of the manuscript, we highlighted the changes in red.

Response to the First Referee.

The manuscript "Real-time quantum error correction beyond break-even" demonstrated a break-even-quantum error correction with the GKP encoding with 3D superconducting circuits. To achieve fault-tolerant quantum computers, precise quantum control is important whose precision is much higher than its threshold. For that purpose, the bosonic encoding plays an important role in the hardware efficient "pre"-error correction to overcome a physical limitation of the decoherence. The manuscript reports the break-even quantum error correction, where the coherence time of the corrected qubit is longer than its physical limitation. The method of the GKP encoding is so complicated that this procedure itself might generate large decoherence. The authors use a long coherent transmon ancilla and have investigated the optimization method to reduce the error of the error correction cycles.

The demonstration of break-even error correction is a really big milestone for fault-tolerant quantum computers. To achieve this, the authors have also investigated several techniques and they are well-studied based on many experimental data. The paper is also well-written and the results are clearly presented. I found these results are suitable for publication in nature. Still, I have some comments and question before publication:

We are very grateful to the Referee for providing this comprehensive review. We are glad that the Referee found our results suitable for publication in Nature. We address the comments of the Referee below.

1. In fig.2(a), what exactly are the subspace C1 - C3 and which kinds of errors are occurred? In the bosonic system, the dimension of the error-subspace should be infinite.

We agree with the Referee that there are infinitely many subspaces in a bosonic system. In Fig.2(a), we depict a *cartoon* system with 4 subspaces only to illustrate the main idea of how the state is transferred between subspaces in a trickle-down fashion. This cartoon can be applied to any system of dimension 8 that satisfies Knill-Laflamme QEC conditions. In the Supp. Info. S4 B we describe the actual structure of the subspaces and their corresponding errors in our system: C1-C2 correspond to the photon loss and photon gain errors, C3-C5 correspond to higher-order errors that involve two oscillator ladder operators, etc.

We added an explanatory text in the description of Fig.2(a) that makes it clear that the figure depicts a cartoon system and that the error subspaces of our actual infinite-dimensional system are explained in Supp. Info. S4 B.

2. Page 4 line 7 in the first paragraph, "the agent corrects large errors in the initial parameter values, and in the subsequent few hundreds of epochs, it fine-tunes the circuit parameters to achieve the highest performance." Does this qualitative explanation of behavior always hold? If so, is the optimization process dynamically changing?

Such qualitative explanation is valid in most cases, but not always. We observed that the rapid initial progress in the reward is correlated with rapid initial change of certain parameters (those that are most mis-calibrated), hence we interpret this trend as correcting the initialization errors. When the parameter vector gets close to optimal, the progress slows down because it becomes more difficult to resolve how parameter changes affect the stochastic reward. However, if the parameter values are rather poorly initialized, the agent will not be able to start learning at all. Hence, we now added the qualifier "typically" in the text.

The Referee is correct that the optimization process is dynamically changing. We use the Adam optimizer (see [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)), which adapts the learning rate according to the changes in several lower-order moments of the stochastic reward.

3. Page 5 last paragraph, the authors evaluate the success probabilities (97%) of the error correction of their process. How is this high success probability consistent with the increased coherence time due to post-selection in the later paragraphs? If the error correction is perfect, the post-selection is supposed not to increase its coherence with correctable small errors. Is it right?

We agree with the Referee that if 1) only small correctable errors are allowed and 2) the QEC is perfect, the post-selection cannot improve the logical error rate. The fact that we demonstrated both high correction efficiency of physical errors and high improvement after post-selection implies that assumptions 1-2 are violated. In particular, the relaxation of an auxiliary qubit can induce large uncorrectable displacement errors, as described in Supp. Info. S4 I, hence the post-selection is expected to help in our setting.

The two numbers quoted in the main text: i) 97% correction efficiency, and ii) factor of 6.3 improvement after post-selection, provide independent information about the distribution of errors:

$$\frac{\text{\# of direct logical errors (they leave no syndrome trace)}}{\text{\# of physical errors that lead to logical errors}} = \frac{1}{6.3} \approx 0.16$$

$$\frac{\text{\# of physical errors that lead to logical errors}}{\text{\# of physical errors that leave syndrome traces}} = 0.03$$

From these equations we can infer that only a tiny fraction $0.16 \times 0.03 = 0.005$ of errors don't leave a trace in the measurement outcomes. This is consistent with our expectations, since the errors that leave no signature are only those that are exact finite-energy logical operations.

4. In method section and Supp. Infor S2F, the authors introduce the active oscillator cooling. The authors mention that this technique can also be realized in trapped ions. It might be true, but in the ion trap, the interaction of a $\sigma_+ + a^\dagger \sigma_-$ is easily achievable with irradiation of a laser beam, so called a red sideband transition. With the superconducting circuits, the cross-Kerr term is the original interaction of the dispersively interacted circuits. On the other hand, with the ion trap, the Jaynes-Cummings and anti-Jaynes-Cumming interaction is the origin between a bosonic mode and a qubit. This active cooling is widely used as "sideband cooling" in the ion trap.

Our cooling protocol is specified at the high level of abstraction, using two ingredients: conditional displacement gates, and auxiliary qubit resets. These ingredients can be realized in both platforms (superconducting circuits, and trapped ions), albeit by different means. As the Referee points out, the Hamiltonian of the trapped ion illuminated by the laser is different from our driven superconducting system. However, it still allows to implement the conditional displacement gate, known as the internal state-dependent force in that community, through sideband driving (instead of the large displacement approach, as in our case). Likewise, the auxiliary qubit reset could be realized by cycling the qubit excitation to another short-lived level which quickly decays to the ground state (instead of measurement-based feedback, as in our case). Both these ingredients were realized in a trapped ion GKP experiment Ref [20] [Nature Physics 18, 296–300](#) to implement a protocol which is similar to ours. However, this work has not applied the same tools to demonstrate ion cooling (they used conventional cooling methods). Therefore, we feel optimistic that this method could work, and hope that the trapped ions community will investigate it further.

In the Methods section, we added a brief clarification of how we expect the basic ingredients for this cooling to be realized in trapped ions.

5. What limits the error correction gain G for this experiment?

We refer to [our answer](#) to a similar question of the second Referee.

6. The coherence time of Pauli Y for the corrected qubit is shorter than other Pauli operators because of the square shape of the grids. The maximum Pauli Y operator expectation value in Fig. 3c is smaller than others. Does it have same reason? This imperfection can be considered as a thermal excitation of the logical qubit. Then, can we prepare the better initial Pauli Y state from the initial Pauli X or Pauli Z state with oscillator manipulations?

There are three main reasons why this expectation value is smaller: 1) In our experiment, the state preparation pulse for the Y eigenstate was longer than for X/Z eigenstates (it has one additional circuit layer), hence the state preparation fidelity is worse for this state due to extra decoherence; 2) The logical readout in the Y basis involves a longer conditional displacement, hence the logical readout quality in this basis is also worse. 3) Finally, even if both 1 and 2 were perfect, it is expected that logical readout of the finite-energy state in the Y basis using the phase estimation method will yield a lower value than in X/Z basis due to the roll-off of the envelope. We believe that these three effects account for the observation made by the Referee.

The resolution that the Referee proposes might work in practice, but the required oscillator manipulations are complex enough that they would likely introduce more cost than benefit. An additional rotation gate from X/Z to Y basis can be realized using a circuit with 4 layers, see [Nature Physics 18, 1464–1469](#), while in the current implementation the Y eigenstate can be prepared with only 1 additional layer compared to X/Z eigenstates.

7. Page 5 last paragraph, why the logical error probabilities per QEC cycle is the smallest for Y direction?

As explained in Supp. Info. S4 I, the decay of an auxiliary qubit during the big conditional displacement is among the dominant sources of logical errors. In our protocol, which cycles between two orthogonal quadratures, such a decay can only generate X or Z errors (depending on the cycle), but not Y errors. We believe this is the reason for having an overall lower probability of Y errors. Equivalently, the lifetime of the Y eigenstates is shorter, because only the dominant X/Z errors affect these states. Note that similar behavior was observed in the previous GKP experiments.

8. Are these processes applicable to the hexagonal grid state? Can the correction gain G improve with the hexagonal grid state?

Our protocol is fully applicable to the hexagonal grid. The only change would be to do the virtual rotation gate by an angle $\pi/3$ instead of $\pi/2$. We added a clarification of this point in Supp. Info. S4 C.

Given that the auxiliary qubit relaxation during the conditional displacement gate is among the dominant sources limiting the logical lifetime, we expect that the hexagonal code would perform slightly worse in terms of the QEC gain, because the average length of the big conditional displacement is greater in the hexagonal grid. Note that this explanation is consistent with the results in [Nature 584, 368–372](#), where hexagonal code was shown experimentally.

9. How about the orthogonality of two QEC cycles? If the phase difference of two QEC steps is not exactly 90 degrees, then R^x_{Δ} might generate Z error.

We thank the Referee for bringing up this other possible way of introducing an error by the protocol, and we fully agree with their statement. The virtual rotation gate in our protocol is responsible for ensuring the orthogonality, its calibration is described in Supp. Info. S4 C. In the ideal case, the rotation angle would be simply $\pi/2$, but in practice it needs to also compensate for the spurious rotation due to the dispersive coupling (hence, it depends on the measured auxiliary qubit state). This parameter is learned by the agent (see Fig. S14 d), and we identified it as one of the most critical parameters of the protocol.

10. Reference [36] should replace to the nature physics ver.

We thank the Referee for catching this outdated citation, it is updated now.

Response to the Second Referee.

The paper by Sivak et al reports on progresses on quantum error correction with Gottesman-Kitaev-Preskill codes, implemented in microwave cavities coupled to a transmon qubit.

The paper reports on an extended lifetime of the quantum information encoded in such stabilised GKP qubit by a factor of 2.27 as compared to the best passive qubit encoded in the same cavity. This improvement is substantial and yield a major breakthrough, as compared to previous reported experimental results on GKP-quantum error-correction achieved in Ref.14, as well as on previous break even attempts with cat codes (Ofek et al). This improvement is made possible by the use of a new material (tantalum) for both transmon qubit and cavity, instead of aluminium, yielding higher coherence, together with the use of improved error-correction schemes (Refs 20,28), also combined with machine learning techniques for pulse design, reset of the auxiliary qubit, etc. I especially like that the long stabilisation time allows for testing the hypothesis of stationary errors, which seems to hold, with small deviations being accounted for transmon leakage to higher states.

The results reported are state-of-the-art and the technological improvements residing in the aspects listed above, yielding to such a substantial lifetime extension, are worth reporting in Nature.

The very extensive supplementary information section addresses in great detail the technical aspects of the paper and replies to most possible questions a reader might have. However I do have a couple of additional questions and a comment:

We thank the Referee for their time and effort, and for their support and enthusiasm for our results. We reply to the Referee's questions below.

- In Fig 2 (c), why do the authors stop the training after slightly more than 500 epochs, while the reward has started to increase again and has not saturated?

Having performed multiple trainings during the development of this project, we noticed that the progress typically stalls after about 500 epochs, hence in our final working routine the training is automatically interrupted to save time for performing other experiments with the trained protocol. However, sometimes we observed resurgences of progress much later, even after 1000 epochs. Hence, we don't exclude the possibility that a slightly higher QEC gain could be achieved with improved or prolonged training.

- Do I understand correctly that the main limiting factor to push the lifetime extension to even higher gain factors resides in the leakage of the transmon? How does the squeezing of the GKP states stabilised come into play? Are there other limiting factors that prevent the authors from achieving $G = 10$ or 100 ?

The Referee is correct that the main limiting factor resides in the auxiliary qubit (transmon), since the QEC protocol is not designed to correct its errors. However, leakage to non-computational states contributes only about 10% to the error budget (as inferred by post-selecting trajectories with no leakage, see Supp. Info. S4 G). The largest identified limiting factor is the bit-flip errors of the auxiliary qubit, which propagate to the GKP state either as rotation errors (if bit flip happens during the readout) or as large displacement errors (if bit flip happens during the conditional displacement). Note that the squeezing of the peaks does not affect this error channel. In Supp. Info. S4 I, we estimated that this source accounts for about 40% of the logical errors. The remaining part of the error budget remains unidentified and needs to be understood in the future research. We strongly suspect, based on the data in Supp. Info. S4 J, that the lifetime of the auxiliary qubit degrades during the execution of these operations, and hence the 40% estimate (based on the lifetime of un-driven transmon) might be too optimistic (i.e., the actual contribution is larger).

Note that the photon loss in the cavity, often claimed to be the “dominant error source” in such systems, does not significantly contribute to logical errors. Based on the simulation in [Phys. Rev. Lett. 125, 260509](#), we would expect a gain of $G \sim 70$ if the photon loss was the only source of errors in our system. The squeezing of the peaks comes into play when this error source is considered. If the peaks are too squeezed, the state becomes more sensitive to the photon loss and other types of oscillator errors (e.g., Kerr nonlinearity). The RL agent chose the squeezing (9.4 dB) which, we believe, is nearly optimal for the real error sources of our system.

Hence, as briefly mentioned in our Conclusion, the primary direction for improving the gain further towards $G=10$ or 100 , lies in the development of a novel auxiliary qubit (something other than transmon), better suited for the type of operations required by our QEC protocol.

- It would be interesting if the authors could comment on their perspectives regarding scalability, both in terms of the actual experimental platform (how to accommodate for several cavities? Or more generally how to encode several modes?), and as to whether the reinforcement learning protocol they use would be scalable to the scenario where multiple encoded qubits are considered.

The co-authors have different opinions regarding the scalability of the experimental platform based on the 3D superconducting cavities. However, a possible path towards scalable computing with GKP qubits might rely on high-quality 2D resonators or other compact 2D structures, instead of 3D cavities. There is a continuous progress in the quality of 2D resonators, see [PRX Quantum 3, 020312](#). Although it is doubtful that they will match the quality of 3D cavities, Purcell effect due to the coupling to an auxiliary qubit would likely set the limit for the oscillator lifetime. Hence, 2D resonators might be an appealing scalable alternative to 3D cavities. The physics and most experimental techniques would remain the same in that case.

Regarding the scalability of learning, we expect that a processor with multiple GKP qubits would require only local tune-up for each such qubit and its neighborhood. A similar problem is already successfully addressed in the automated tune-up of larger devices, such as Google's Sycamore processor. In such a “factorizable” case, our current RL algorithm could be sufficient for larger applications. However, if we consider the scaling that incorporates more components into a logical qubit, for example, as in the multi-mode GKP code [PRX Quantum 3, 010335](#), accompanied by an increase in the number of parameters, a more sophisticated RL algorithm might be required. We feel optimistic that problems with even a thousand parameters could be solved using modern RL. To draw a suitable parallel with other domains, an RL system recently learned to play StarCraft (computer strategy game with stochasticity and extremely high-dimensional action space), see [Nature 575, 350–354](#). Our future research plans include the adoption of such more advanced RL techniques.

- I would suggest to replace the nomenclature "ancilla"/"ancillary" with "auxiliary qubit" everywhere throughout manuscript and supplementary information

We thank the Referee for pointing out a negative connotation of this word, it is now replaced everywhere throughout the manuscript.