

Automatic post-selection by ancillae thermalization

L. Wright^{1,*}, F. Barratt,² J. Dborin,³ G. H. Booth,¹ and A. G. Green³¹Department of Physics, King's College London, Strand, London WC2R 2LS, United Kingdom²Department of Mathematics, King's College London, Strand, London WC2R 2LS, United Kingdom³London Centre for Nanotechnology, University College London, Gordon St., London, WC1H 0AH, United Kingdom

(Received 12 October 2020; accepted 19 July 2021; published 13 August 2021)

Tasks such as classification of data and determining the ground state of a Hamiltonian cannot be carried out through purely unitary quantum evolution. Instead, the inherent nonunitarity of the measurement process must be harnessed. Post-selection and its extensions provide a way to do this. However, they make inefficient use of time resources—a typical computation might require $O(2^m)$ measurements over m qubits to reach a desired accuracy and cannot be done intermittently on current (superconducting-based) NISQ devices. We propose a method inspired by thermalization that harnesses insensitivity to the details of the bath. We find a greater robustness to gate noise by coupling to this bath, with a similar cost in time and more qubits compared to alternate methods for inducing nonlinearity such as fixed-point quantum search for oblivious amplitude amplification. Post-selection on m ancillae qubits is replaced with tracing out $O[\log \epsilon / \log(1 - p)]$ (where p is the probability of a successful measurement) to attain the same accuracy as the post-selection circuit. We demonstrate this scheme on the quantum perceptron, quantum gearbox, and phase estimation algorithm. This method is particularly advantageous on current quantum computers involving superconducting circuits.

DOI: [10.1103/PhysRevResearch.3.033151](https://doi.org/10.1103/PhysRevResearch.3.033151)

I. INTRODUCTION

Algorithms for classification of data, optimizing the energy to find the ground state properties of a Hamiltonian (and indeed optimizing classifiers for a given data set) require the use of nonlinear operations that cannot be achieved solely through unitary quantum evolution. When carrying out these tasks on a quantum computer we must use the nonunitarity of the measurement process. There are several ways in which to do this depending upon the relative abundance of resources, quantified by measures such as the number of qubits, coherence times, and gate fidelities. Post-selection and the related repeat-until-success algorithms are popular choices.

However post-selection makes inefficient use of time resources—a typical computation requires $O(2^m)$ measurements over m qubits to reach a desired accuracy. Nevertheless, it is a frequently used tool in atomic contexts where coherence times are long and manipulation timescales short, so that time is not the limiting resource. For superconducting circuits where coherence times are much shorter and where measurements of a subset of qubits is not possible while maintaining the coherent evolution of the remainder, this is more problematic.

Curiously, post-selection in fact uses classical nonlinearity—through a yes or no decision based upon a

measurement on ancillae qubits. This suggests how a more time-efficient scheme might be developed. Fundamentally, the nonlinearity of the classical world is induced by observation of only a portion of a larger quantum world. It is possible then to replace post-selection with a scheme where explicit measurement of ancillae qubits is not required, i.e., where they are traced out or simply ignored.

Eigenstate thermalization gives a clue as to how this can be achieved. Coupling a small system at high temperatures to a large, low temperature bath allows us to cool the small system. Eigenstate thermalization extends this notion to closed quantum systems. Coupling a large number of ancillae qubits in a low entropy state (e.g., $|00000\dots\rangle$) to a small system and evolving the total system under some unitary evolution allows entropy to flow from the small system of interest to the ancillae.

II. RESULTS

Inspired by this, we replace the classical yes or no nonlinearity of post-selection with a nonlinearity attained by tracing out ancillae. In the following we demonstrate the application of ancillae thermalization to the quantum perceptron, quantum gearbox and a ground state preparation algorithm. Robustness to noise is demonstrated for the quantum gearbox in simulation for two different noise models and on the “ibmq_virgo” quantum device.

Our scheme is a type of amplitude amplification—a generalization of Grover search [1] introduced by Brassard *et al.* [2]. Amplitude amplification comes in various flavors depending upon whether the state of the target qubits or probability of success at each iteration is known or not. For a direct comparison with our method, we focus upon an implemen-

*lewis.wright@kcl.ac.uk

TABLE I. Comparison of computational resources for unitary transformations. Computational resources in the of post-selection, ancillae thermalization, and $\frac{\pi}{3}$ fixed-point oblivious amplitude amplification (FP OAA) [3] for guaranteeing unitary transformations in the large initial success probability regime. We define $Q(O_N)$ as the number of operations required to implement the N -qubit gate, O , and p_0 as the initial probability of a successful measurement. In the asymptotic limit the number of operations required to implement controls on U_{n+m} is constant. Additionally, $Q(W_n) \sim O(1)$ and can be ignored in practice. Although ancillae thermalization has more operations, we observe lower susceptibility to gate errors. We suspect this is due to the exponentially fewer operations acting on each ancilla in ancillae thermalization, exposing them to less gate noise. Justification for these values can be found in Appendix A while a comparison of noise robustness can be found in Sec. III.

Method	Measurements	Qubits	Gates
Post-selection	$O(p_0^{-1})$	$n + m$	$Q(U_{n+m})$
$\frac{\pi}{3}$ FP OAA	0	$O(n + m)$	$O\left(\frac{\log \epsilon}{\log(1 - p_0)}[m + Q(U_{n+m})]\right)$
Ancillae Thermalization	0	$O\left(n + \frac{\log \epsilon}{\log(1 - p_0)}m\right)$	$O\left(\frac{\log \epsilon}{\log(1 - p_0)}[m + Q(U_{n+m}) + Q(W_n)]\right)$

tation that does not require the knowledge of either and decreases the error monotonically, $\frac{\pi}{3}$ fixed-point search [3]. Used in the context of amplitude amplification, this algorithm which we call $\frac{\pi}{3}$ fixed-point oblivious amplitude amplification (FP OAA) is equivalent to the optimal “fixed-point quantum search” from Yoder *et al.* in the regime of large initial success probabilities [4]. In this regime we find ancillae thermalization more robust to the effects of gate errors, with a similar cost in time and more qubits. In addition, the different structure of our method allows implementation of a wider class of transformations while retaining no knowledge of the target state, including nonunitary transformations, the latter of which OAA algorithms alone do not allow [5]. We demonstrate this using a nonoptimal procedure for ground state preparation shown in Fig. 3. Table I and Fig. 4 show the trade-offs in resources between the various approaches for unitary transformations and robustness to gate noise, respectively, in the large success probability regime.

A. From post-selection to ancillae thermalization

Figure 1(a) demonstrates post-selection used in a repeat-until-success circuit. A unitary U applies a desired operation R to an input state $|\psi\rangle$, conditioned upon the outcome of a measurement of m ancillary qubits. For example the unitary

$$U|0\rangle^{\otimes m}|\psi\rangle = \sqrt{p_0}|0\rangle^{\otimes m}R|\psi\rangle + \sum_{k=1}^{2^m-1} \sqrt{p_k}|k\rangle E_k|\psi\rangle \quad (1)$$

achieves the rotation R on $|\psi\rangle$ and the state $|0\rangle^{\otimes m}$ on the ancillary qubits with probability p_0 . The states $E_k|\psi\rangle$ corresponds to incorrect transformations of the target qubits. If the procedure fails, all qubits are reset and the process is repeated. The probability of failure after N iterations of the algorithm is $\epsilon \sim (1 - p_0)^N$; the transformation R is implemented *exactly* with a finite probability.

Ancillae thermalization and more generally amplitude amplification take a different philosophy. The output from U on the target qubits is interpreted as a superposition of correctly transformed ($R|\psi\rangle$) and incorrectly transformed ($E|\psi\rangle$) states. Tracing out the ancillary qubits prepares the target qubits in a mixed state ρ that is *approximately correct*. The density matrix ρ has an overlap $\langle\psi|R^\dagger\rho R|\psi\rangle = 1 - \epsilon$

with $R|\psi\rangle$. For ancillae thermalization, this fidelity with the target state is obtained with an exponential reduction in measurements for an increased cost in ancillary qubits. This is achieved by iteratively entangling fresh ancillary qubits with the target qubits via unitary $V = UW$, where W is a reset gate that transforms $E_k|\psi\rangle \rightarrow |\psi\rangle$ for all k . The ancillae conditionally entangle with the incorrectly transformed parts of the system’s wave function through control gates on the previous




FIG. 1. Comparison between post-selection and ancillae thermalization. (a) Post-selection circuit acting on n -qubit target state $|\psi\rangle$ and m ancillae. Applying U and measuring $|0\rangle^{\otimes m}$ on the ancillae guarantees a successful transformation $R|\psi\rangle$ of the target qubits. Measuring $|1\rangle$ on any ancilla implies a failure and the procedure is repeated. (b) Ancillae thermalization circuit equivalent to $O(N)$ attempts at applying $R|\psi\rangle$ with $N = 2$. Note that each attempt of post-selection is exponential in measurements. Ancillae thermalization is a modification of post-selection. Measurements are replaced with gates and new ancillae entangle with the incorrectly transformed parts of the target qubits’ wave function. All ancillae are then traced out. The unitary V acts on $|\psi\rangle$ and m ancillae per iteration. It is factorized as $V = UW$ where W is a reset gate discussed in Appendix B. The circuit produces a finalized mixed target state ρ that has an overlap with desired state $R|\psi\rangle$ up to error ϵ , i.e., $\text{Tr}(\rho|\psi'\rangle\langle\psi'|) = 1 - \epsilon$ where $|\psi'\rangle = R|\psi\rangle$.

ancilla at each iteration. The circuit to achieve this is shown in Fig. 1(b). The overlap between the target qubits and $R|\psi\rangle$ increases exponentially with the number of iterations applied (see Appendix B for full details).

B. Quantum perceptron

The quantum perceptron [6] is the first explicit example of a quantum circuit fulfilling the requirements for a meaningful quantum neural network. It was introduced by Schuld *et al.* [7]. It is able to simulate a classical perceptron while taking advantage of quantum properties such as processing input data as a superposition. In general, quantum neural networks struggle to construct a nonlinear activation function due to their linear dynamics. The quantum perceptron uses a post-select circuit shown in Fig. 2(a) to achieve this nonlinearity. This circuit implements the transformation $|\psi\rangle \rightarrow \exp[-iq(\theta)Y]|\psi\rangle$ onto a target qubit with probability $p(\theta) \sim O(1/2)$. The angle of rotation $q(\theta) = \arctan[\tan^2(\theta)]$ is sigmoidal in shape and can be used to capture the nonlinear properties found in classical neural networks in a quantum setting.

Ancillae thermalization removes the need to post-select in order implement the quantum perceptron. The circuit shown in Fig. 2(b) achieves the same level of accuracy as $O(N)$, with $N = 2$ attempts of the post-selection circuit. To achieve a total overlap with the desired state $\exp[-iq(\theta)Y]|\psi\rangle$ within additive error ϵ , the process of applying V to fresh ancillae and the target qubit must be repeated $O[\log(1/\epsilon)]$ times. This achieves a fidelity between the finalized target qubit and the desired state given by

$$\begin{aligned} & \langle \psi | e^{iq(\theta)Y} \rho e^{-iq(\theta)Y} | \psi \rangle \\ &= 1 - \delta \frac{\log(1/\epsilon)}{\log(1/\delta)} + 1 - |\langle \psi | e^{-iq(\theta)Y} \psi \rangle|^2, \end{aligned} \quad (2)$$

where $\delta = 1 - p(\theta)$ and ϵ has been rescaled. The fidelity increases exponentially with the number of iterations.

Results of applying ancillae thermalization to the quantum perceptron obtained from IBMQ's "qasm-simulator" i.e., a simulator with no noise and "ibmq-oursense" quantum machine are shown in Fig. 2(c). As in other applications of NISQ devices, there is an optimum circuit depth that balances theoretical advantages of deeper circuits with the effects of noise. The quantum perceptron displays an increase in accuracy with increasing iterations up to a threshold where further operations increase exposure to finite gate fidelity leading to a decrease in accuracy. This point is emphasized in the subfigure which shows a lower fidelity for a higher number of iterations.

C. Phase estimation

Next, we apply our procedure to a ground state preparation algorithm. Although more efficient state preparation algorithms exist, see [8–10], this setting is still of interest since it reveals the role of the ancillae as an effective low-temperature bath in addition to demonstrating a FP OAA scheme for nonunitary transformations.

The quantum phase estimation algorithm shown in Fig. 3(a), computes the eigenvalue θ satisfying $A|A\rangle = \exp(2\pi i\theta)|A\rangle$. Post-selecting on the precision qubit register can prepare target qubits in the ground state of an n -qubit

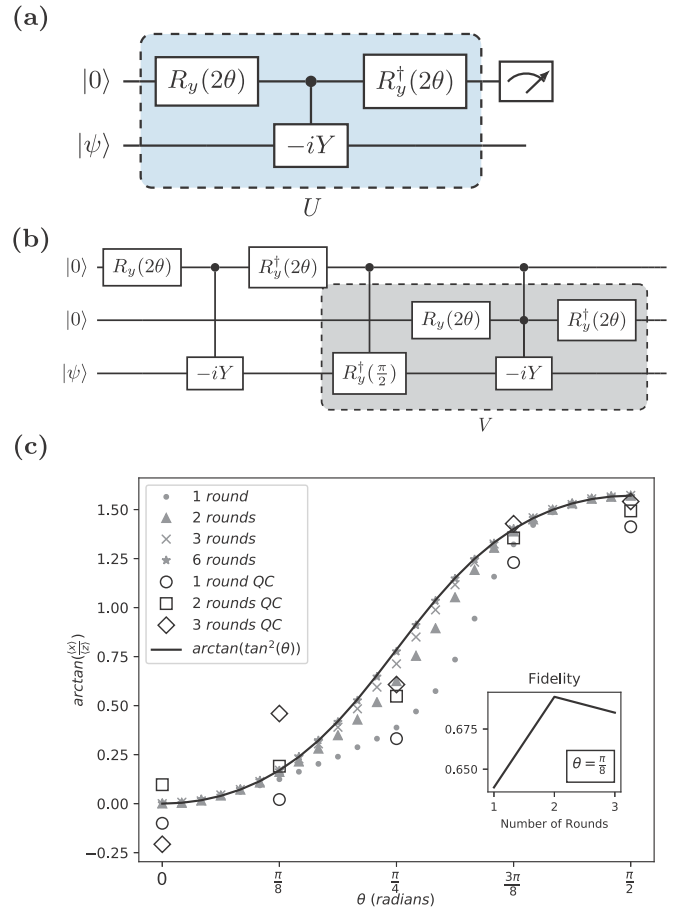


FIG. 2. Quantum perceptron/nonlinear activation function $q(\theta)$. (a) Post-select circuit for implementing angle $q(\theta)$ in the quantum perceptron, acting on an ancilla and the target qubit. A successful transformation of $\exp[-iq(\theta)Y]|\psi\rangle$ corresponds to measuring $|0\rangle$ on the ancilla with probability $p(\theta) = \cos^4(\theta) + \sin^4(\theta)$. Upon failure, when $|1\rangle$ is measured on the ancilla, the target qubit is guaranteed to transform as $\exp(-i\pi/4)|\psi\rangle$. As a result, the target qubit can be reset by rotation $R_y(\pi/2)$ and the circuit is repeated. (b) Ancillae thermalization circuit for an equivalent $O(N)$ attempts of post-selection, with $N = 2$ applications of the circuit in (a). The reset and second instance, which acts on a new ancilla, are both conditioned by the state of the first ancilla. (c) Angle $q(\theta)$ obtained by ancillae thermalization for different number of iterations and θ . The values were obtained from IBMQ's "qasm-simulator" (symbols) (no noise) and "ibmq_oursense" quantum computer (rings). The subfigure shows the fidelity from ibmq_oursense between the finalized target qubits and ground state for different numbers of iterations.

Hamiltonian. Ancillae thermalization achieves the same effect by tracing out ancillae qubits—the ancillae effectively provide a low entropy reservoir into which the excess energy of the target state can be transferred.

The circuit shown in Fig. 3(b) achieves the same level of accuracy as $O(N)$ for $N = 2$ attempts of phase estimation with post-selection. In a similar manner to the quantum perceptron, a total overlap with the ground state within additive error ϵ is achieved by applying V to fresh ancillae and the target qubits $O[2^m \log(1/\epsilon)]$ times. After tracing out the ancillae qubits, the fidelity between the finalized mixed target state ρ and the

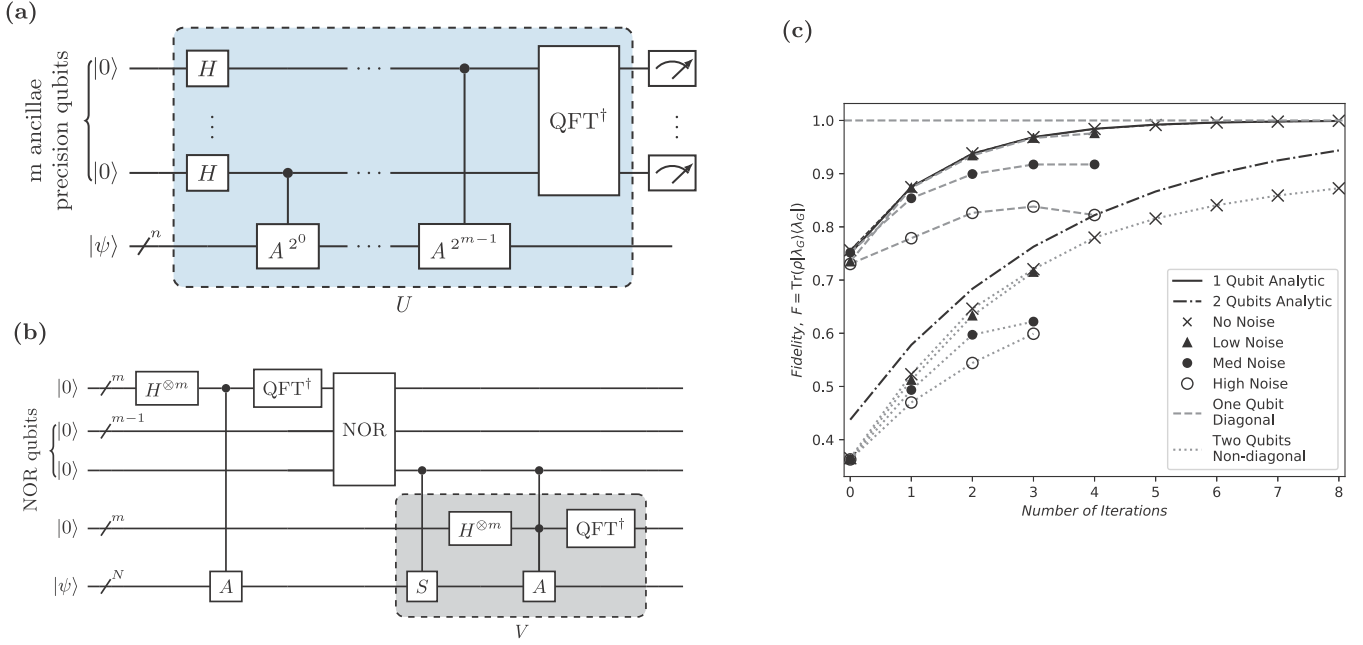


FIG. 3. Ground state preparation of a Hamiltonian. (a) Post-select circuit for ground state preparation in the form of the quantum phase estimation algorithm. The algorithm prepares the ground state of a Hamiltonian onto target qubits inputted as an equal superposition of all possible bit strings. The circuit consists of a synthesized unitary $A = \exp(-iH\tau)$ acting on the target qubits while being controlled by m precision ancillae. Measuring $|0\rangle^{\otimes m}$ on the ancillae indicates a successful ground state preparation onto the target qubits. Any other measurement indicates that an excited state has been prepared and failed. Upon failure of ground state preparation, all qubits are reset and the circuit is repeated. (b) Ancillae thermalization circuit for an equivalent $O(N)$ attempts of post-selection, with $N = 2$ applications of the circuit given in (a). The NOR gate compiles conditions from the ancillae while the reset gate S redistributes the weights of the incorrectly prepared states of the target qubits onto all bit strings. Each iteration acts on the scrambled states of the target qubits and is controlled by the output of the last NOR ancilla. A complete description for the NOR and scrambling gate can be found in Appendix B. (c) Fidelity between the finalized target qubits and ground state of the Hamiltonian using ancillae thermalization. We show results for different numbers of iterations. The results were obtained from IBMQ's "qasm-simulator" for $H_1 = 0|0\rangle\langle 0| - \frac{3\pi}{2}|1\rangle\langle 1|$ and $H_2 = \sum_{i,j=1}^4 a_{ij}|i\rangle\langle j|$. We simulate different levels of noise based upon a model of thermal relaxation between the qubits and their environment. The range of data points for the noise based simulations was restricted due to computational limitations. Additionally, an approximation was made on the initialization and scrambling operations in the nondiagonal two qubit case. Further details of these experiments can be found in Appendix B.

ground state is given by

$$(\lambda_G|\rho|\lambda_G) = (1 - J^{\log(\epsilon)/\log(J)+1}), \quad (3)$$

where $J = (N - 1)/N$, $N = 2^n$ is the number of eigenstates, and m is the minimum number of precision qubits required to distinguish between all energy values without imperfections. Therefore the precision, i.e., the number of ancillae used in the phase estimation circuit, dictates an upper bound on the fidelity. Ancillae thermalization shows an exponential increase in fidelity as the number of iterations increase compared to the fidelity attained with the same number of attempts of the post-select circuit. We assume that the value of the ground state energy is known up to precision 2^{-m} . Additionally, a preprocessing procedure has occurred which shifts all energy values by this amount such that correctly preparing the ground state is indicated by measuring $|0\rangle^{\otimes m}$.

Results of applying the ancillae thermalization to ground state preparation obtained from IBMQ's qasm-simulator with the addition of simulated noise are shown in Fig. 3(c). The fidelity was computed between the finalized target qubits and the ground states of the one qubit Hamiltonian $H_1 = 0|0\rangle\langle 0| - \frac{3\pi}{2}|1\rangle\langle 1|$ and the two qubit Hamiltonian $H_2 = \sum_{i,j=1}^4 a_{ij}|i\rangle\langle j|$ (a random set of parameters a_{ij} were chosen

in the latter case as described in Appendix C) for different numbers of iterations of the ancillae thermalization circuit. As in the case of the quantum perceptron, an increase in fidelity with the number of iterations reaches an upper bound when the circuit depth leads to too great an exposure to gate noise.

The fidelity is lower in the two-qubit case than predicted analytically. This is due to an approximation made on the initialization and scrambling operations on the target qubits. Furthermore, due to the inclusion of Toffoli gates in the NOR gate, the ancillae thermalization modification of post-selection is highly sensitive to noise. This is exacerbated for larger Hamiltonians due to the increase in number of gates required to act on the target qubits and slower convergence of fidelity with the number of iterations. A more detailed discussion of these effects and of the parameter values used in the simulations can be found in Appendix C.

OAA schemes alone cannot deterministically implement nonunitary transformations. However, recent developments in block encoding [11] and quantum signal processing [12] allow us to embed an approximate n -qubit projector in the upper-left corner of an $n + k$ qubit unitary, where k is the number of ancillae needed for the encoding. Amplitude amplification is then used to deterministically implement this approximate

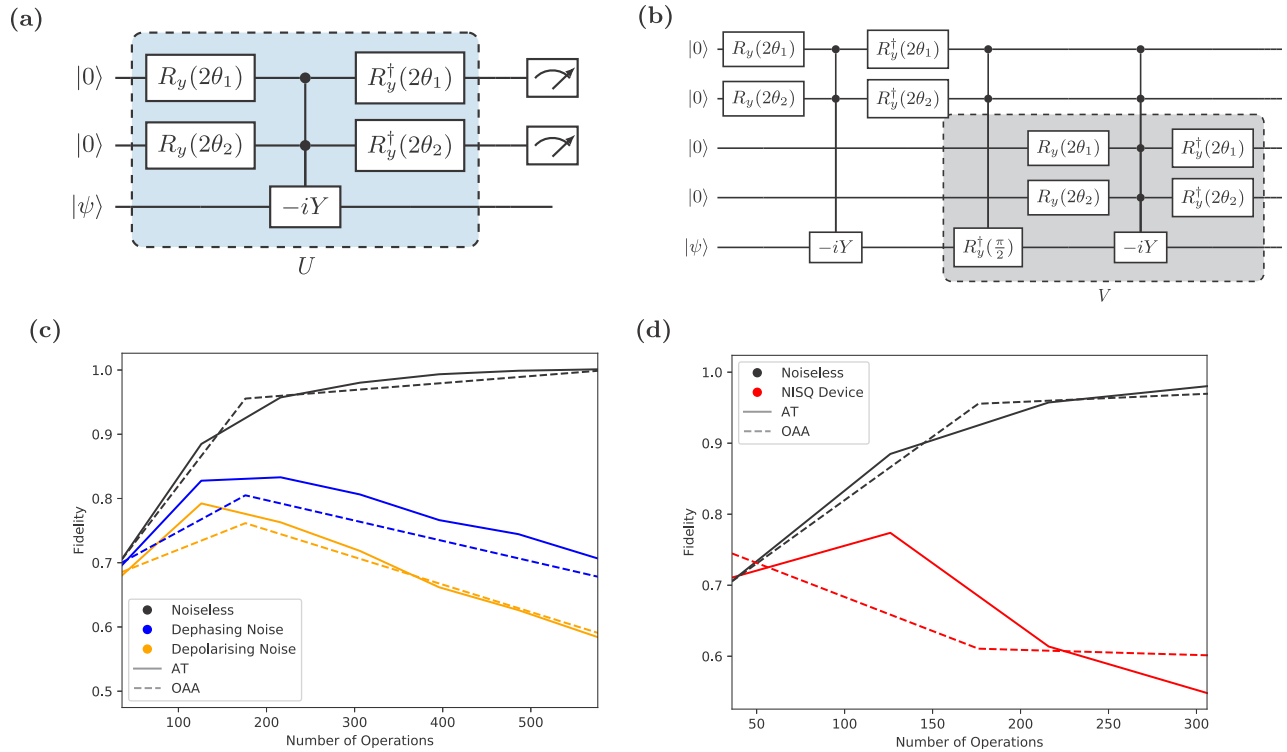


FIG. 4. Noise robustness for the quantum gearbox. (a) Quantum gearbox circuit for $m = 2$ ancillary qubits [13]. This circuit is a generalization of the quantum perceptron found in Fig. 2 with $m = 1$ ancilla. Measuring $|0^{\otimes m}\rangle$ on the ancillae with probability $p(\theta) = \cos^4(\theta) + \sin^4(\theta)$ corresponds to the successful transformation of $\exp[-iq(\theta)Y]$ on the target qubit where $q(\theta) = \arctan[\tan^2(\theta)]$ and $\sin(\theta) = \sin(\theta_1) \cdots \sin(\theta_m)$. Measuring $|1\rangle$ on any ancillae corresponds to applying $R_y(-\pi/2)$ onto the target qubits, and thus can always be reset by applying $R_y(\pi/2)$. For the simulation $\theta_1 = \theta_2 = \pi/4$. (b) Ancillae thermalization circuit for an equivalent $O(N)$ attempts of post-selection with $N = 2$ applications of the circuit in (a). Unitary V , which acts on the fresh ancillae and incorrect states of the target qubit, is conditioned by the state of all ancillae from the previous iteration. (c) and (d) Fidelity between the finalized target qubit and desired state $\exp[-iq(\theta)Y]|\psi\rangle$ using ancillae thermalization and $\frac{\pi}{3}$ FP OAA for a different number of operations, i.e., exposure to gate noise. Both dephasing and depolarizing noise models were simulated in (c) while (d) is run on the “ibmq_vigo” quantum device, in addition to a noiseless simulation.

projector onto the target qubits [10]. A further comparison of resource costs between these methods can be found in Table II in Appendix A. State-of-the-art algorithms for ground state preparation assume the initial target qubit state has a nontrivial overlap with the ground state. For ease of demonstration we initialize the target state in an equal superposition of all eigenstates and construct a reset gate, i.e., the scrambling gate, which scrambles every eigenstate such that the output has an equal overlap with all other eigenstates. In order for ancillae thermalization to have a competitive ground state preparation algorithm, a more sophisticated initialization and reset procedure must be implemented.

III. ROBUSTNESS TO NOISE

Resource costs such as the number of qubits and gate operations are a good indication of an algorithm’s efficiency. On near-term quantum devices, however, an algorithm’s robustness to noise is a much more practical measure. Ancillae thermalization is more robust than alternative schemes.

Demonstrating robustness to noise: Figs. 4(c) and 4(d) demonstrates ancillae thermalization’s robustness to noise compared with $\frac{\pi}{3}$ FP OAA for the quantum gearbox—an extension of the quantum perceptron with two ancillary

qubits [13]. The circuits for the quantum gearbox, ancillae thermalization, and $\frac{\pi}{3}$ FP OAA are given in Figs. 4(a) and 4(b) and Fig. 6 in Appendix C. The multicontrol gates in both circuits were implemented without additional ancillae. Specifically, the three-qubit control gate was implemented using six C-NOT gates and seven single-qubit control gates, while the four-qubit control gate was implemented using two single-qubit control gates and three three-qubit gates [14]. Simplified multiqubit Toffoli gates were used to further reduce operational cost [15].

Full state tomography was performed on the target qubit with 24 576 shots for both circuits. The fidelity was then computed between the target qubit and desired transformation $R = e^{-iY\theta}$, where $\sin(\theta) = \sin(\theta_1)\sin(\theta_2)$. Note that in practice full state tomography is not required as the target qubit will be assumed to have a sufficiently large fidelity with the desired state. Additionally, the number of single qubit and C-NOT operations were measured per iteration of each algorithm.

In addition to demonstration on IBMQ’s quantum machine, simulations of both circuits were performed with depolarizing and thermal relaxation noise. The depolarizing noise error parameter $\lambda = 0.001, 0.01$, for all the single-qubit and two-qubit gates, respectively. Details of the latter noise model are

given by the “high” noise level in Table III in Appendix C. All circuits run on IBMQ’s quantum machine were compiled using the OpenQASM backend [16] without additional error mitigation techniques.

Origin of noise robustness: We believe that the robustness to noise of ancillae thermalization arises because of the intrinsic robustness of the thermalization to changes in the coupling of the system to the bath. In the special case of the quantum perceptron, some of the controls—which are the proxy for the system-bath interaction—can be removed entirely without any detriment to the performance. This can be seen in Fig. 2(b) where no controls are placed upon the $R_y(2\theta)$ rotations of the fresh ancillae. We have preliminary evidence of robustness to reducing controls in other circumstances. This strongly suggests that ancillae thermalization does indeed inherit robustness to noise from independence upon the bath-system interaction. A thorough analysis will be the subject of a future work.

An additional consideration in the comparison between ancillae thermalization and $\frac{\pi}{3}$ FP OAA is the number of operations acting on each ancilla qubit. This number is fixed in ancillae thermalization by the depth of the post-select unitary, regardless of the finalized accuracy, exposing ancillae to less gate noise.

We expect our intuition to be applicable to a variety of circuit and quantum machine archetypes. The fixed-point quantum search proposed by Yoder *et al.* has been shown to have an exponential decrease in query complexity over $\frac{\pi}{3}$ fixed-point quantum search in the regime of small initial success probabilities. Currently, it is unknown whether fixed-point quantum search for OAA has an increased robustness to gate noise compared to ancillae thermalization in this regime. In the large initial success probability regime, however, it is known that the operational costs of fixed-point quantum search coincide with $\frac{\pi}{3}$ FP OAA for an equivalent finalized success probability. Therefore, although a direct comparison has not been made, we expect ancillae thermalization to have the highest overall robustness to gate noise within the large initial success probability regime.

Mitigating qubit costs and control complexity: One drawback of ancillae thermalization is the use of resource intensive control gates. However, the same considerations that suggest robustness to gate noise also motivate ways to mitigate control costs and complexity. Inspired by the fact that a system thermalizes when only a subset of its modes are coupled to a heat bath, we have preliminary evidence that the number of control qubits and complexity of controls can be reduced by conditioning on a subset of factors of the unitary. In the case of thermalization, the coupling to the bath can be simple providing that the Hamiltonian is sufficiently scrambled. The scrambling transfers energy to the bath-coupled modes where it is dissipated. We find that it is sufficient to control simple factors of the unitary with the more complex factors playing the role of scrambling. In the context of repeat-until-success, such a mitigation scheme would correspond to controlling on an imperfect error flag. The result would no longer be guaranteed a success. This is less useful than its application in ancillae thermalization where even an imperfect correction of error increases the amplitude for success exponentially in time. The speedup from conditioning on a subset of operations

does not change the linear-in-time scaling of the number of qubits, only the prefactor to this scaling. Moreover, the time to reach the desired accuracy increases, reflecting the increased time to thermalize if the coupling to the bath is weakened. Nevertheless, this ability to tension these costs against one another will likely prove useful in near-term applications.

IV. DISCUSSION

The nonlinearity of the classical world can be understood by the observation of a minor part of a quantum system—the unobserved part of the system acting as an environment. The environment can be interpreted as a heat bath extracting entropy from our system, or equivalently an entanglement bath which gradually and selectively entangles with a subset of our system. A simple and effective model of a heat bath is to assume no back reaction so that each mode of the heat bath interacts exactly once with the system of interest. It forms such a small fraction of the overall size of the bath that the bath distribution is unaltered. At the same time, the fact that the system never interacts with this mode again means that the back reaction effects are not felt. We have used these two ideas to allow a set of ancillae qubits initialized in some low entropy state to extract entropy from our system. The free evolution of our ancillae is with a zero energy Hamiltonian—ensuring that entropy only flows from the system of interest to the ancillae and each ancilla interacts only once with the system corresponding to a no back reaction condition. The resulting algorithm is a type of FP OAA for unitary and nonunitary transformations, achieving nonlinearity by tracing out auxiliary degrees of freedom. Its structure is rather different from its counterparts, which require fewer qubits. The $\frac{\pi}{3}$ FP OAA scheme achieves an optimal amplitude amplification through a cunning cancellation of phases. It is a fundamental observation of statistical mechanics however that the nature of a heat bath does not determine the thermal equilibrium state (provided suitably weak coupling). Our scheme effectively harnesses this universality to obtain a degree of robustness to gate infidelity in addition to deterministically implementing a wider class of transformations without knowledge of the target state. Moreover, it seems possible to reduce the qubit cost and control gate complexity at the expense of longer times. This gives additional freedom to operate within the NISQ constraints of qubit count and gate fidelity. Which scheme is optimal is contingent upon the particular system to which the algorithm is applied. Still, it is gratifying that there exists a regime where a simple physically motivated scheme such as the one we present can outperform other methods.

ACKNOWLEDGMENTS

We acknowledge support from the EPSRC: L.W. and F.B. through EP/L015854/1, F.B. and A.G.G. through EP/S005021/1, J.D. through EP/S021582/1. G.H.B. acknowledges support from the Royal Society via a University Research Fellowship, as well as funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 759063).

APPENDIX A: RESOURCE SCALING

The entries in Table I of the main paper show a comparison of resources required to implement the techniques discussed in this paper compared to ancillae thermalization. Here we give more detail on how these entries are obtained.

1. Unitary transformation

Measurements: For a post-select unitary U acting on m ancillae and target qubits $|\psi\rangle$, up to $O(2^m)$ measurements on the ancillae are required to implement R onto $|\psi\rangle$. This assumes all measurements are independent of each other. On the other hand, no measurements are required in $\frac{\pi}{3}$ OAA and ancillae thermalization.

Number of qubits: All qubits can be reused upon failure in post-select circuits since the ancillae state collapses upon measurement. The state $E|\psi\rangle$ is independent of R and can always be reset. Therefore U acting on m ancillae and n -target qubits requires $n + m$ qubits in total. Using an additional $m - 1$ qubits in OAA allows for a linear scaling in $Q[S_m(\pi/3)]$. To achieve a fidelity $\langle\psi|R^\dagger\rho R|\psi\rangle = 1 - \epsilon$ for ancillae thermalization requires $N = \log(1/\epsilon)/\log[1/(1 - p_0)] - 1$ applications of V , where $\epsilon(1 - |\langle\psi|R\psi\rangle|^2)^{-1} \rightarrow \epsilon$. Insertions of m new ancillae and $m - 1$ NOR ancillae qubits are needed at each application. Therefore a total $O\{n + \log[1/(1 - p_0)]\log(1/\epsilon)m\}$ qubits are required for ancillae thermalization.

Operations: The most important resource in the comparison of ancillae thermalization and OAA is the number of gate operations required. This is computed in the case of post-selection, OAA and ancillae thermalization as follows.

(i) *Post-selection.* We define $Q(O_N)$ as the number of single qubit and C-NOT gates required to implement the N -qubit gate O . Therefore, U requires a total of $Q(U_{n+m})$ operations for n target and m ancillary qubits. No further coherence time nor operations are needed, since the ancillae becomes disentangled with the target qubits once measured and the circuit can be reset.

(ii) $\frac{\pi}{3}$ FP OAA. A detailed derivation of the resource costs can be found in Ref. [17]. We summarize the key results here. An upper bound on the error for the k th nested iteration is given by $(1 - p)^{3^k} \leq \epsilon$. Rearranging we find

$$k = \frac{1}{\log 3} \left[\log \log \left(\frac{1}{\epsilon} \right) - \log \log \left(\frac{1}{(1 - p)} \right) \right]. \quad (\text{A1})$$

The number of operations at each iteration can be computed using the recursive relation $Q[A_{n,m}(j)] = Q[A_{n,m}(j - 1)] + 2Q[S_m(\frac{\pi}{3})]$. Consequentially, its closed form can be written as

$$Q[A_{n,m}(k)] = \left\{ Q(U_{n+m}) + Q\left[S_m\left(\frac{\pi}{3}\right)\right] \right\} 3^k - Q\left[S_m\left(\frac{\pi}{3}\right)\right], \quad (\text{A2})$$

where the initial conditions are given in Eqs. (D2) and (D3).

The function $Q[S_m(\pi/3)]$ is the number of operations required to implement the controlled phase shift. Assuming access to additional ancillae, $S_m(\pi/3)$ can be constructed in a similar way to the NOR gate such that the number of operations scale linearly with m , for $m > 2$.

Using our expression for k from Eq. (A1) and $Q[S_m(\pi/3)]$, the result for the total number of operations in $\frac{\pi}{3}$ FP OAA follows.

(iii) *Ancillae thermalization.* Ancillae thermalization produces an overlap $\langle\psi|R\rho R|\psi\rangle = 1 - (1 - p)^{P+1}(1 - |\langle\psi|R\psi\rangle|^2)$ between the finalized target qubits and the desired state for P iterations of V . To achieve an overlap $1 - \epsilon$, $O[\log(1/\epsilon)]$ implementations of V are required. Each V consists of $Q(U_{n+m}) + Q(R_n)$ operations, of which controls need only be implemented on gates acting on the target qubits. The result for the number of operations required in ancillae thermalization follows.

2. Ground state preparation

In this paper we use ground state preparation as a demonstration to deterministically apply a nonunitary operation onto target qubits using ancillae thermalization. We also mention using OAA in addition to linear combination of unitaries to achieve the same task. Here we give a comparison of resources between the two schemes.

a. Phase estimation

The phase estimation algorithm (PEA) implements an exact projector onto the ground state of a Hamiltonian with $p = 1 - \epsilon$. Error ϵ is intrinsic to PEA and originates from the binary approximation of the eigenvalues. A higher precision is required to ensure these “imperfections” do not effect the computation. We assume a preprocessing shift has occurred such that the ground state energy value is 0 and initialize the target state in an equal superposition of all eigenstates.

Measurements: For an initial target state $|\psi\rangle$, given as an equal superposition of eigenstates, an average $O(p_0^{-1}) = 2^n$ measurements for $p_0 = 2^{-n}$ are required for ground state preparation. No measurements are required in ancillae thermalization nor block encoding + OAA.

Number of qubits: Post-selection and the PEA act upon an n -qubit input state and m ancillae qubits. Here the number of ancillary qubits is dependent on the required precision of the eigenvalue $m \sim O[n/2 + \log(1/\epsilon) + \log(1/\Delta)]$, where Δ is the lower bound on the spectral gap [8]. Qubits can be reused upon failure. Ancillae thermalization with the PEA require a total $O\{m \log(1/\epsilon)/\log[1/(1 - p_0)]\}$ qubits, the result follows with $\log[1/(1 - p_0)] \sim O(2^n)$.

Operations:

(i) *Post-selection.* U_{n+m} requires 2^m applications of the controlled evolutionary unitary $A = e^{-iH}$. To implement A with error ϵ' requires $Q(A) \sim O[\Lambda \text{polylog}(2^n, 1/\epsilon')]$ operations where Λ is the number of operations required to simulate H [18]. We assume $\Lambda \sim O(d)$, where d is the sparsity of H , i.e., maximum number of nonzero entries in each row of H [18] and $Q(W_n) \sim O(1)$. Therefore, to prepare the ground state using the PEA requires $Q(U_{n+m}) \sim O[2^m Q(A)] = \tilde{O}(2^{n/2} \Delta^{-1} \epsilon^{-1} d)$.

(ii) *Ancillae thermalization.* The number of operations for the NOR gate scale linearly with m while two additional operations are needed per ancillae to implement a control. We require all qubits that have been acted on to remain coherent throughout the computation and assume $Q(W_n) \sim$

TABLE II. Comparison of computational resources for ground state preparation. Computational resources for post-selection and ancillae thermalization using phase estimation (PEA) as well ancillae thermalization using linear combination of unitaries (LCU) via the results found in Ref. [8]. We also show the results for most state-of-the-art quantum ground state preparation algorithm [10]. This algorithm uses block encoding in addition to amplitude amplification to approximately project the target qubits onto the ground state, where k is the number of qubits required in the block encoding. \tilde{O} is up to polylogarithmic factors. Here we assume the number of calls to H , $\Lambda \sim d$, where d is the sparsity of the Hamiltonian, i.e., maximum number of nonzero elements in each row of the Hamiltonian and Δ is the lower bound on the spectral gap. Note that we assume the input target state is given as an equal superposition of all eigenstates. Various schemes may exponentially reduce the number of ancillary qubits required for ancillae thermalization, e.g., by conditioning on only a subset of elements of a factorization of the unitary U . Indeed, a heat bath does not need to couple to all elements of a system to effectively cool.

Ground state preparation	Measurements	Qubits	Gates
Post-selection and PEA	$O(2^n)$	$O[n + \log(1/\Delta) + \log(1/\epsilon)]$	$\tilde{O}(2^{n/2} \Delta^{-1} \epsilon^{-1} \Lambda)$
Ancillae	0	$O[2^n \text{polylog}(1/\epsilon, 1/\Delta)]$	$\tilde{O}(2^{3n/2} \Delta^{-1} \epsilon^{-1} \Lambda)$
Thermalization and PEA	0	$O[2^n \text{polylog}(1/\epsilon, 1/\Delta)]$	$O[2^n \Delta^{-1} \Lambda \text{polylog}(2^n, \Delta^{-1}, \epsilon^{-1})]$
Ancillae	0	$O[2^n \text{polylog}(1/\epsilon, 1/\Delta)]$	$O[2^n \Delta^{-1} \Lambda \text{polylog}(2^n, \Delta^{-1}, \epsilon^{-1})]$
Thermalization and LCU	0	$O[2^n \text{polylog}(1/\epsilon, 1/\Delta)]$	$O[2^n \Delta^{-1} \Lambda \text{polylog}(2^n, \Delta^{-1}, \epsilon^{-1})]$
Block encoding and OAA	0	$O(n + k)$	$O[2^{n/2} \Delta^{-1} k \Lambda \log(1/\epsilon)]$

$O(1)$ or $\sim O[\log(n)]$ [19]. The total number of operations is $O[\log(1/\epsilon)/\log[1/(1 - p_0)]Q(U_{n+m})]$, the result follows.

b. Linear combination of unitaries

Linear combination of unitaries (LCU) can be used to construct a truncated Taylor series of the time-dependent evolutionary operator. This approximately projects the target qubits onto the ground state of the Hamiltonian. The implementation of this unitary is nondeterministic, thus either OAA or ancillae thermalization can be used to amplify the probability of success. A detailed derivation of the resource costs of LCU for ground state preparation can be found in [8]. We summarize the key results here.

Number of qubits: The number of ancillary qubits, i.e., the precision required for LCU is less than PEA, requiring $m \sim O[\log(1/\Delta) + \log\log(2^n/\epsilon)]$ qubits to achieve the same accuracy. The probability of correctly implementing the ground state projector by LCU is $p_0 \sim O(2^{-n})$, assuming $|\psi\rangle$ is in an equal superposition of eigenstates. We ignore any qubits required for the Hamiltonian simulation.

Operations: LCU requires implementing a unitary B on the ancillae followed by $A_k = \exp(-iHk)$ on the target qubits, representing a segment of the Taylor series for A . B can be implemented with $Q(B) \sim O[\Delta^{-1} \log^{3/2}(2^n/\epsilon)]$ operations. The Hamiltonian simulation represented by A is implemented with $Q(A) \sim O[\Delta \text{polylog}(2^n, \tilde{\epsilon}^{-1})]$ operations to a required accuracy $\tilde{\epsilon} = \tilde{O}(2^{-n} \Delta \epsilon)$. Thus the number of operations needed to perform LCU, $Q(\text{LCU}) \sim O[\Delta^{-1} \Lambda \text{polylog}(2^n, \Delta^{-1}, \epsilon^{-1})]$. Assuming $Q(W_n)$ and $Q(\text{C-LCU}) - Q(\text{LCU}) \sim O(1)$ leads to the result shown in Table II for ancillae thermalization.

APPENDIX B: METHODS

We provide a detailed description of the methods and results found in the main text.

1. Ancillae coupling and reset gates

To construct the ancillae thermalization circuit we use the methods discussed in Sec. II A to initially apply U followed by iterations of the controlled unitary V . For ground state preparation, V includes a reset gate and unitary U of the phase estimation circuit acting on excited states in the wave function of the target qubits. It is constructed as follows: First, we only apply V if the ancillae from the previous iteration are not in the state $|0\rangle^{\otimes m}$. This state corresponds to the preparation of the target qubits in the ground state. The condition is checked through use of a NOT-OR (NOR) gate. This logic gate acts upon all ancillae from the previous iteration and an additional $m - 1$ NOR ancillae qubits. The result of whether the previous ancillae correspond to the preparation of the ground state is outputted onto the last NOR ancilla. Second, a reset gate consisting of a scrambling operation acts upon the target qubits conditioned by the last NOR ancilla. The purpose of the scrambling operation is to redistribute the probability of an eigenstate to all other eigenstates equally. A full description of the NOR and scrambling gate can be found in Appendices B 2 and B 3, respectively. Finally, another application of the phase estimation unitary U —controlled by the last NOR ancillary qubit—acts upon the target qubits and a batch of m fresh ancillae.

2. NOR gate

The quantum NOT-OR (NOR) gate shown in Fig. 5 is an quantum logic gate. Its purpose is to compile the controls on all excited states represented by the ancillae onto a single qubit. The gate acts on the m precision ancillae from each iteration in addition to $m - 1$ NOR ancillae in initial state $|0\rangle^{\otimes m-1}$. The state of the last NOR ancilla is $|0\rangle$ if and only if the precision ancillae are in $|0\rangle^{\otimes m}$, otherwise the NOR qubit is in state $|1\rangle$. The scrambling gate W and U in the next iteration of V are both controlled by the last NOR ancilla. If the ground state energy $E_G \neq 0$, then a preprocessing procedure needs to be implemented to shift all the energies by a constant such that

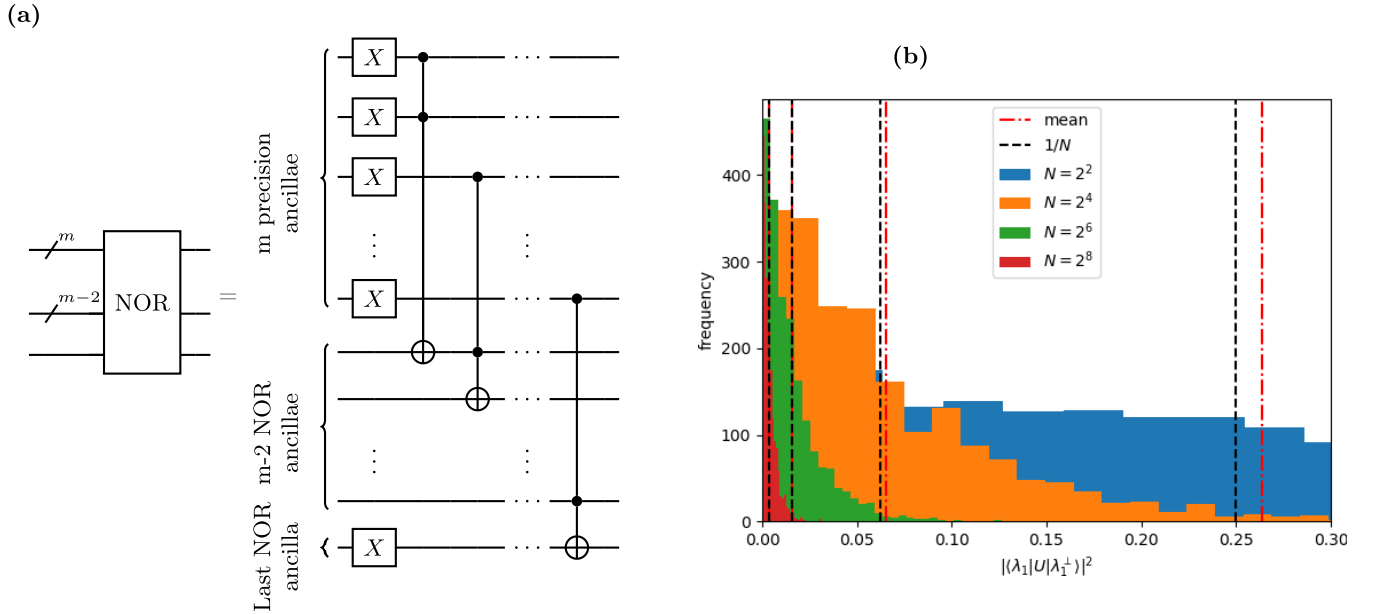


FIG. 5. Operations required in ancillae thermalization for ground state preparation. (a) The NOR gate quantum circuit used to compile the precision ancillae conditions into a single qubit. The gate acts on all ancillae from the iteration and $m - 1$ fresh NOR ancillae. The output state of the last NOR ancilla is $|0\rangle$ if and only if the state of the precision ancillae is $|0\rangle^{\otimes r}$, otherwise the last NOR qubit state is $|1\rangle$. The scrambling gate and next iteration of U are both controlled by the last NOR qubit. (b) Histogram showing the overlap between a randomly chosen state and a scrambled state which has been transformed by an arbitrary local unitary. We can see numerically that the unitary increases the overlap between the states to $\sim \frac{1}{N}$, becoming more accurate as the system size increases. This proves that the application of a scrambling gate gives the desired result of redistributing the probability weights of each state equally.

$|0\rangle^{\otimes m}$ corresponds to the preparation of the ground state of an arbitrary Hamiltonian on the target qubits.

3. Scrambling gate

The purpose of the scrambling gate is to redistribute the weight of each eigenstate equally among all other eigenstates on the target qubits. We assume that an equal distribution of eigenstates corresponds to a maximally mixed set of bit strings. This approximation is discussed further in this section below. Figure 5(b) shows that for a randomly chosen eigenstate of U , applications of arbitrary local unitaries will increase its overlap with another eigenstate. As N increases, the mean overlap between the transformed eigenstate and a perpendicular state converges to $1/N$. To numerically prove this result, N -qubit state $|\lambda_1\rangle$ and its perpendicular state $|\lambda_1^\perp\rangle$ were chosen uniformly with the Haar measure. The overlap between $|\lambda_1^\perp\rangle$ and $|\lambda_1\rangle$ acted on by local Hadamard gates, i.e., $H^{\otimes N}$ was computed and the process was repeated 1000 times. The simulation was also repeated with local X gates which gave the same result.

Scrambling approximation: The eigenstates of a diagonal unitary correspond to single bit strings. Consequently, a maximally mixed state of bit strings corresponds to an equal distribution of eigenstates, where application of a Hadamard gate on a bit string scrambles the state entirely. This is not true for nondiagonal unitaries, where eigenstates correspond to a linear combination of bit strings. However, as N increases so does the number of eigenstates which have an average overlap $1/N$ with each bit string. Therefore, the assumption that a maximally mixed set of bit strings can approximate an equal

distribution of eigenstates for an arbitrary unitary becomes more accurate as the number of target qubits increases.

4. Fidelity calculations

In this section we derive the fidelity between the finalized target qubits from the tracing out method and the desired state for the quantum perception and ground state preparation algorithms.

a. Quantum perceptron

Applying the quantum perceptron post-select unitary U onto the state $|0\rangle|\psi\rangle$ produces

$$|\phi_0\rangle = \sqrt{p(\theta)}|0\rangle|\psi'\rangle + \sqrt{1-p(\theta)}|1\rangle E|\psi\rangle, \quad (\text{B1})$$

where $|\psi'\rangle = R(\theta)|\psi\rangle$. The reset gate W transforms $E|\psi\rangle \rightarrow |\psi\rangle$ and a new ancilla is inserted. U , conditioned by the state of the previous ancilla, acts on the new ancilla and target qubits to give

$$|\phi_1\rangle = (\sqrt{p(\theta)}|0\rangle + \sqrt{1-p(\theta)}\sqrt{p(\theta)}|1\rangle)|0\rangle|\psi'\rangle + [1-p(\theta)]|1\rangle|1\rangle|\psi\rangle. \quad (\text{B2})$$

Resetting, inserting new ancillae, and applying U conditioned by the previous iteration's ancilla P times leads to the state

$$|\phi_P\rangle = \left(\sum_{k=0}^{P-1} \sqrt{p(\theta)[1-p(\theta)]^k} |k\rangle \right) |\psi'\rangle + [1-p(\theta)]^{\frac{P}{2}} |P\rangle |\psi\rangle, \quad (\text{B3})$$

where $|k\rangle = |1\rangle^{\otimes k}|0\rangle^{\otimes(P-k)}$ and $|P\rangle = |1\rangle^{\otimes P}$. The density matrix after P steps is given by

$$\rho^P = |A\rangle|\psi'\rangle\langle\psi'| + |B\rangle|\psi\rangle\langle\psi| + \dots, \quad (\text{B4})$$

where $|A\rangle = \sum_{k=0}^{P-1} \sqrt{p(\theta)[1-p(\theta)]^k} |k\rangle$ and $|B\rangle = [1-p(\theta)]^{\frac{P}{2}} |P\rangle$. Using $\langle k'|k\rangle = 0$ for all $k \neq k'$, a partial trace is performed on the ancillae to obtain

$$\rho_{\text{target}}^P = \{1 - [1 - p(\theta)]^P\} |\psi'\rangle\langle\psi'| + [1 - p(\theta)]^P |\psi\rangle\langle\psi|. \quad (\text{B5})$$

Using the equation above, the fidelity between the finalized target qubits and desired state $F = \text{Tr}(\rho_{\text{target}}^P |\psi'\rangle\langle\psi'|)$ can be written as

$$F(\rho_{\text{target}}^P, |\psi'\rangle) = 1 - (1 - p)^P (1 - \langle\psi|\psi'\rangle). \quad (\text{B6})$$

As the quantum perceptron is an example of a single ancilla repeat-until-success circuit, the result can easily be generalized to the m ancillae case.

b. Ground state preparation

To prepare the target qubits in the ground state of a specified Hamiltonian we utilize quantum phase estimation. This algorithm computes $0 \leq \theta \leq 1$ which satisfies $A|\lambda\rangle = e^{2\pi i\theta} |\lambda\rangle$ up to a finite precision for m ancillae qubits in the first register. In other words, it computes the binary value $\theta = 0, \theta_1, \theta_2, \dots, \theta_m$ with $\theta_i \in [0, 1]$. The unitary A can always be constructed from Hermitian matrix H such that $A|\lambda\rangle = e^{-iE\tau} |\lambda\rangle \Rightarrow E = -\frac{2\pi\theta}{\tau}$ where E is the energy corresponding to eigenstate $|\lambda\rangle$. Note that in all of our experiments $\tau = 1$.

The n -target qubits are initialized to an equal superposition $|\psi_n\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N |\lambda_i\rangle$ of all eigenstates. Implementing the phase estimation unitary U onto $|0^{\otimes m}\rangle |\psi_n\rangle$ gives

$$|\phi_1\rangle = \frac{|\theta_G\rangle}{\sqrt{N}} \sum_{i=1}^{N^*} |\lambda_i\rangle + \sum_{i=N^*+1}^N \frac{|\theta_i\rangle |\lambda_i\rangle}{\sqrt{N}}. \quad (\text{B7})$$

The false positive in the prepared ground state originates from the finite precision on the eigenvalues. A scrambling operation S is performed on the incorrectly prepared eigenstates by placing a condition on the ancillae. This operation produces an overlap $\langle \lambda_i | S \lambda \rangle \sim \frac{1}{N}$ for all $i = 1, \dots, N$. Details of S and the conditioning on the ancillae can be found in Appendix B. A batch of m new ancillae are inserted. U acts on the new ancillae and target qubits conditioned by the previous ancillae to transform the state as

$$|\phi_2\rangle = \left(\frac{|\theta_G\rangle |\theta_G\rangle}{\sqrt{N}} + \frac{|\theta_G\rangle}{N} |k\rangle \right) \sum_{j=1}^{N^*} |\lambda_j\rangle + \frac{|k\rangle |k\rangle |\psi_n\rangle}{N}, \quad (\text{B8})$$

where $|k\rangle = \sum_{i=N^*+1}^N |\theta_i\rangle$ and a conditioned S has been applied to the target qubits.

After P iterations of inserting ancillae, applying U and S , the state is given by

$$|\phi_{P1}\rangle = \left(\sum_{i=0}^{P-1} \frac{1}{N^{\frac{i+1}{2}}} |\theta_G\rangle^{\otimes(P-i)} |k\rangle^{\otimes i} \right) \sum_{j=1}^{N^*} |\lambda_j\rangle + \frac{1}{N^{\frac{P}{2}}} |k\rangle^{\otimes P} |\psi_n\rangle. \quad (\text{B9})$$

By expanding $|\psi_n\rangle$ it can be shown that scrambling the state increases the overlap with the ground state

$$|\phi_{P2}\rangle = \left(\sum_{i=0}^{P-1} \frac{1}{N^{\frac{i+1}{2}}} |\theta_G\rangle^{\otimes(P-i)} |k\rangle^{\otimes i} + \frac{1}{N^{\frac{P+1}{2}}} |k\rangle^{\otimes P} \right) |\tilde{\lambda}_G\rangle + \frac{1}{N^{\frac{P+1}{2}}} |k\rangle^{\otimes P} |\lambda_E\rangle, \quad (\text{B10})$$

where $|\tilde{\lambda}_G\rangle = \sum_{j=1}^{N^*} |\lambda_j\rangle$ and $|\lambda_E\rangle = \sum_{j=N^*+1}^N |\lambda_j\rangle$. The density matrix after P steps is given by

$$\rho^P = |A\rangle |\tilde{\lambda}_G\rangle \langle \tilde{\lambda}_G| + |B\rangle |\lambda_E\rangle \langle \lambda_E| + \dots, \quad (\text{B11})$$

where

$$|A\rangle = \left(\sum_{i=0}^{P-1} \frac{1}{N^{\frac{i+1}{2}}} |\theta_G\rangle^{\otimes(P-i)} |k\rangle^{\otimes i} + \frac{1}{N^{\frac{P+1}{2}}} |k\rangle^{\otimes P} \right), \quad (\text{B12})$$

$$|B\rangle = \frac{1}{N^{\frac{P+1}{2}}} |k\rangle^{\otimes P}. \quad (\text{B13})$$

Performing a partial trace on the ancillae the density matrix of the target qubits is given by

$$\rho_{\text{target}}^P = \frac{1}{N^*} \left[1 - \left(\frac{N - N^*}{N} \right)^{P+1} \right] |\tilde{\lambda}_G\rangle \langle \tilde{\lambda}_G| + \frac{(N - N^*)^P}{N^{P+1}} |\lambda_E\rangle \langle \lambda_E| + \dots \quad (\text{B14})$$

Using the equation above, the fidelity between the finalized target qubits and ground state can be written as

$$F(\rho_{\text{target}}^P, |\lambda_G\rangle) = \frac{1}{N^*} \left[1 - \left(\frac{N - N^*}{N} \right)^{P+1} \right]. \quad (\text{B15})$$

The fidelity is bound by the distinguishability of the eigenstates. In the main paper we choose m such that $N^* = 1$.

APPENDIX C: NUMERICAL SIMULATION PARAMETERS

Here we provide the parameters used in the numerical simulations shown in the main text.

1. Qiskit code

We construct the example circuits discussed in this paper using the python qiskit API [20]. This allows the decomposition of our circuits into the universal gate set consisting of arbitrary single qubit rotations and C-NOT gates. We use a total of 8192 shots for each data point in both examples and chose not to display error bars since they are statistically negligible. Figure 6 shows the $\pi/3$ FP OAA circuit for the quantum gearbox simulation, implemented using qiskit. The code used for the experiments in this paper can be found at [21].

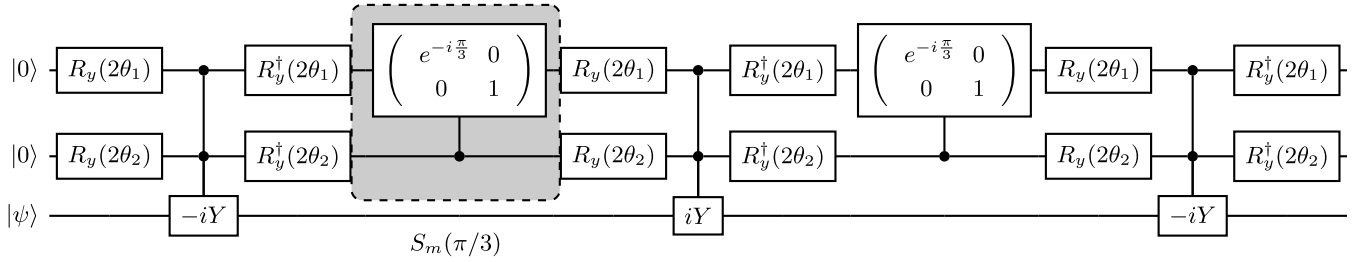


FIG. 6. $\frac{\pi}{3}$ FP OAA circuit for the quantum gearbox. First iteration of $\frac{\pi}{3}$ FP OAA for an equivalent $O(N)$ attempts of post-selection with $N = 3$ applications of the circuit in Fig. 4(a). The circuits consists of repetitions of the post-select unitary U and controlled phase gate $S_m(\pi/3)$ for $m = 2$, which performs a phase shift on the $|0\rangle$ state of the top ancillary qubit. The quantum gearbox with $\frac{\pi}{3}$ FP OAA was simulated with depolarizing and thermal relaxation noise in addition to being run on IBMQ's quantum device. Fidelity between the finalized target and desired state was computed for different numbers of nested iterations and compared against ancillae thermalization. In all three noise models it is shown in Figs. 4(c) and 4(d) that ancillae thermalization has an increased robustness to gate noise as a function of circuit depth.

2. Hamiltonians

The Hamiltonians used in the experiments are given by

$$H_1 = \begin{pmatrix} 0 & 0 \\ 0 & -\frac{3\pi}{2} \end{pmatrix},$$

$$H_2 = \begin{pmatrix} -0.08609 & -0.22467 & -0.41822 & -0.10511 \\ -0.22467 & -1.40667 & -0.16506 & -0.67003 \\ -0.41822 & -0.16506 & -3.06202 & 0.09996 \\ -0.10511 & -0.67003 & 0.09996 & 1.41319 \end{pmatrix}. \quad (C1)$$

H_2 was chosen such that $A = e^{-iH_2} = VDV^\dagger$ where $D = \sum_{k=0}^3 e^{\frac{k\pi i}{2}} |k\rangle\langle k|$. V contains the eigenvectors of A and has the form $V = I + \epsilon B$ where the elements $B_{ij} \sim \mathcal{N}(0, 0.5)$ and perturbation $\epsilon = 0.5$. Orthonormalization of V is ensured by the Gram-Schmidt process. The motivation behind H_2 comes from the accurate approximation on the scrambling gate discussed in Appendix B 3. As the number of target qubits increase, the space of applicable Hamiltonians increases and this approximation becomes more accurate.

3. Simulated noise

The simulated noise for ground state preparation represents the thermal relaxation between each qubit and their environment. This was parametrized by the thermal relaxation time T_1 , the dephasing constant T_2 , and the implementation time of each gate. The thermal relaxation noise model provided by Qiskit was used in the ground state preparation experiment of H_2 . This model is parametrized by the thermal relaxation time T_1 , dephasing constant T_2 , and implementation time of CC-A, C-A, C-NOT, and single qubit gates. Table III shows the range of parameter values used in the ground state preparation experiments with different levels of noise. The noise was computed by decomposing each gate into C-NOT and single qubit gates where the noises are given explicitly. The gates C-A and CC-A were an exception to this decomposition and custom gate noises were computed, respectively, using the values below. T_1 and T_2 were sampled for each qubit from a normal distribution with means μ_1 and μ_2 , respectively, and shared variance σ .

APPENDIX D: BACKGROUND

Here we provide a brief explanation of the other techniques mentioned in the main text.

1. Classical perceptron

The classical perceptron whose quantum analog is discussed in the paper consists of two parts: The first part takes n -inputs x_1, \dots, x_n and performs linear regression with synaptic weights w_1, \dots, w_n plus a bias b . This computes the input signal to the perceptron $\theta = x_1 w_1 + x_2 w_2 + \dots + b$. The second part maps θ onto the activation function $a(\theta) \in [0, 1]$. This is known as the state of a perceptron and is used either as an input for a next perceptron or an output for a neural network. Within the quantum perceptron the latter of the two processes is represented by an angle of rotation upon a target qubit as a function of θ . The challenge is to overcome the innate linearity of quantum dynamics to find a realization of this nonlinear function.

2. Oblivious amplitude amplification

Oblivious amplitude amplification (OAA) replaces post-selection with tracing out ancillary qubits to guarantee a specified unitary transformation, without knowledge of the target state [5]. The allocation of resources for OAA differ to that of our proposed method. We focus upon an implementation that monotonically decreases the error of implementing

TABLE III. Table of parameter values used in the simulation of thermal relaxation noise for the two-qubit ground state preparation experiment.

Noise level	Implementation time (ns)						Statistical params. (μs)		
	U_1	U_2	U_3	C-NOT	C-A	CC-A	μ_1	μ_2	σ
Low	0	50	100	300	1600	3000	1800	2000	10
Medium	0	50	100	300	1600	3000	180	200	10
High	0	50	100	300	1600	3000	50	70	10

the specified unitary transformation in the regime of large initial success probabilities, $\frac{\pi}{3}$ FP OAA [3]. See Fig. 6.

Repeating Eq. (1) of the main text here for clarity, we seek a transformation U that achieves a desired unitary transformation R of a target set of qubits with some probability p_0 :

$$U|0\rangle^{\otimes m}|\psi\rangle = \sqrt{p_0}|0\rangle^{\otimes m}R|\psi\rangle + \sum_{k=1}^{2^m-1} \sqrt{p_k}|k\rangle E_k|\psi\rangle. \quad (\text{D1})$$

In essence, $\frac{\pi}{3}$ FP OAA “boosts” the final success probability from $p_0 = 1 - \epsilon$ to $1 - \epsilon^3$ using the equality $(1 - e^{i\pi/3}) = e^{-i\pi/3}$. This is done by replacing U with A_1

given by

$$A_0 = U, \quad (\text{D2})$$

$$A_k = -A_{k-1}S(\pi/3)A_{k-1}^\dagger S(\pi/3)A_{k-1}, \quad (\text{D3})$$

where $S(\pi/3) = \mathbb{I}^m - (1 - e^{i\pi/3})|0^m\rangle\langle 0^m|$ is a controlled $\frac{\pi}{3}$ phase shift applied to the ancillary qubits. A_k concatenates this procedure k times to obtain a final success probability $p_{\text{final}} = 1 - \epsilon^{3^k}$. Each recursion increases p_{final} superexponentially at the cost of an exponential number of operators. The larger number of gate operations acting on each ancillary qubit in $\frac{\pi}{3}$ FP OAA, as discussed in Sec. III, may lead to a reduced robustness to noise when compared with ancillae thermalization.

-
- [1] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (ACM, New York, 1996), pp. 212–219.
 - [2] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, *AMS Contemp. Math. Ser.* **305**, 06 (2002).
 - [3] L. K. Grover, Fixed-Point Quantum Search, *Phys. Rev. Lett.* **95**, 150501 (2005).
 - [4] T. J. Yoder, G. H. Low, and I. L. Chuang, Fixed-Point Quantum Search with an Optimal Number of Queries, *Phys. Rev. Lett.* **113**, 210501 (2014).
 - [5] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Exponential improvement in precision for simulating sparse Hamiltonians, in *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing* (ACM, New York, 2014), pp. 283–292.
 - [6] Y. Cao, G. G. Guerreschi, and A. Aspuru-Guzik, Quantum neuron: An elementary building block for machine learning on quantum computers, [arXiv:1711.11240](https://arxiv.org/abs/1711.11240) (2017).
 - [7] M. Schuld, I. Sinayskiy, and F. Petruccione, The quest for a quantum neural network, *Quant. Info. Proc.* **13**, 2567 (2014).
 - [8] Y. Ge, J. T. Brugués, and J. Cirac, Faster ground state preparation and high-precision ground energy estimation with fewer qubits, *J. Math. Phys.* **60**, 022202 (2019).
 - [9] N. Wiebe and C. Granade, Efficient Bayesian Phase Estimation, *Phys. Rev. Lett.* **117**, 010503 (2016).
 - [10] L. Lin and Y. Tong, Near-optimal ground state preparation, *Quantum* **4**, 372 (2020).
 - [11] D. W. Berry, A. M. Childs, and R. Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (IEEE, New York, 2015), pp. 792–809.
 - [12] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2019), pp. 193–204.
 - [13] N. Wiebe and V. Kliuchnikov, Floating point representations in quantum circuit synthesis, *New J. Phys.* **15**, 093041 (2013).
 - [14] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
 - [15] D. Maslov, Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization, *Phys. Rev. A* **93**, 022311 (2016).
 - [16] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, Open quantum assembly language, [arXiv:1707.03429](https://arxiv.org/abs/1707.03429).
 - [17] G. G. Guerreschi, Repeat-until-success circuits with fixed-point oblivious amplitude amplification, *Phys. Rev. A* **99**, 022306 (2019).
 - [18] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
 - [19] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, *Phys. Rev. A* **80**, 012304 (2009).
 - [20] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, and C.-F. Chen, *Qiskit: An Open-source Framework for Quantum Computing* (Zenodo, 2019).
 - [21] https://github.com/Lewis1304/Ancillae_Thermalisation