

This is the accepted manuscript made available via CHORUS. The article has been published as:

Recycling qubits in near-term quantum computers

Galit Anikeeva, Isaac H. Kim, and Patrick Hayden

Phys. Rev. A **103**, 042613 — Published 29 April 2021

DOI: [10.1103/PhysRevA.103.042613](https://doi.org/10.1103/PhysRevA.103.042613)

Recycling qubits in near-term quantum computers

Galit Anikeeva,¹ Isaac H. Kim,^{2,3} and Patrick Hayden³

¹*Department of Physics, Stanford University, Stanford, CA 94305, USA*

²*Centre for Engineered Quantum Systems, School of Physics,
University of Sydney, Sydney, NSW 2006, Australia*

³*Stanford Institute for Theoretical Physics, Stanford University, Stanford, CA 94305, USA*

(Dated: April 14, 2021)

Quantum computers are capable of efficiently contracting unitary tensor networks, a task that is likely to remain difficult for classical computers. For instance, networks based on matrix product states or the multi-scale entanglement renormalization ansatz (MERA) can be contracted on a small quantum computer to aid the simulation of a large quantum system. However, without the ability to selectively reset qubits, the associated spatial cost can be exorbitant. In this paper, we propose a protocol that can unitarily reset qubits when the circuit has a common convolutional form, thus dramatically reducing the spatial cost for implementing the contraction algorithm on general near-term quantum computers. This protocol generates fresh qubits from used ones by partially applying the time-reversed quantum circuit over qubits that are no longer in use. In the absence of noise, we prove that the state of a subset of these qubits becomes $|0 \dots 0\rangle$, up to an error exponentially small in the number of gates applied. We also provide a numerical evidence that the protocol works in the presence of noise, and formulate a condition under which the noise-resilience follows rigorously.

I. INTRODUCTION

A quarter century after the discovery that quantum computers could efficiently solve classically intractable problems [1, 2], there is now a worldwide academic and industrial effort to build machines realizing that vision. An important milestone was passed last year with Google’s experimental demonstration of quantum supremacy using a 54 qubit processor [3]. While the quantum supremacy experiment demonstrated that there are *some* classically intractable tasks that a quantum computer can efficiently perform, it is still unclear if a speedup can be realized in the near term for any practical problems of interest. One promising approach is to pursue variational quantum algorithms [4, 5]. These algorithms use a quantum computer as a device that can prepare quantum states, parametrized by a circuit of unitary operators. By variationally minimizing the energy of a physical system of interest, one can prepare a low-energy state of a quantum many-body system. If the states prepared in this minimization procedure are difficult to simulate on a classical computer, one may hope to obtain a quantum advantage in finding and characterizing the low-energy states.

However, to achieve a practical advantage, one needs to overcome several challenges. For generic circuits, the gradient decays exponentially with circuit depth [6], making gradient-based optimization challenging. Moreover, realistic devices are noisy and, as such, any intermediate or final outcomes obtained from these experiments will be noisy as well. These limitations suggest that, in order to use variational algorithms for practical purposes, one must use a structured circuit such as the ones introduced in Ref. [4, 5, 7–12].

The family of *convolutional quantum circuits* [13], in which a sliding active window moves over the qubits of the circuit, is particularly promising. The family includes

circuits referred to in previous literature as *holographic quantum circuits* [8, 9, 14–18]), which can be used to realize versions of the matrix product state [19], projected entangled pair states [20], and MERA [21]. Holographic circuits have a number of desirable properties. First, there are theoretical arguments for the faithfulness of the circuits in describing physical states of interest [8, 22]. Second, generically, the expectation values of the local observables obtained from these circuits are resilient to noise [9, 14]. Third, these circuits often lead to a reduced spatial cost. For instance, a quantum many-body system in D spatial dimensions can be simulated using a set of qubits arranged on a $(D - 1)$ -dimensional lattice. Therefore, the existing quantum computing architectures, which are either one- or two-dimensional, can simulate two- and three-dimensional quantum many-body systems. If the target system one intends to simulate has a volume of ℓ^D , the number of qubits one would need to carry out such a simulation would scale as $\mathcal{O}(\ell^{D-1})$. Indeed, a recent experiment demonstrates the success of this approach for $D = 1$, giving hope that these methods can become useful with a continued improvement in quantum computing technology [16, 18].

Despite their promise, convolutional quantum circuits have not been widely used on existing quantum computing platforms. This is mainly because they require a source of fresh $|0\rangle$ qubits to realize their advantages, but resetting qubits in the middle of a computation can be demanding. For example, the Sycamore processor which Google used in their quantum supremacy experiment does not have the ability to quickly reset a specific qubit [3]. Without that capability, the size of the system one can simulate is limited by the number of qubits present in the quantum computer. For instance, even if one uses convolutional quantum circuits for simulating a $\ell \times \ell$ system, without the ability to reset, one would still need a quantum computer consisting of $\mathcal{O}(\ell^2)$ qubits. In

contrast, with the ability to reset, one can reduce this qubit requirement to $\mathcal{O}(\ell)$. For instance, if $\ell = 10$, the requirement on the number of qubits can shrink from hundreds to tens.

Since having the reset capability can dramatically reduce the number of physical qubits required, it is clearly of interest to implement it in some form. This article provides a software-based reset protocol suitable for systems in which the hardware capability is either absent or too slow. Our procedure, which we refer to as the *rewinding protocol*, unitarily converts a set of qubits used in the computation to a fixed state, after which they can be reused in the remaining part of the computation. Because our protocol is unitary, the protocol can be implemented on any quantum computer, even if it is not equipped with a physical reset operation.

Importantly, we can provide a rigorous theoretical guarantee on the quality of the reset qubits. The infidelity of the reset qubits decays exponentially with the length of the protocol. Therefore, resetting a qubit up to a fixed target error tolerance takes time logarithmic in the inverse error.

While other strategies exist for reusing qubits, they have limitations that make them unsuitable for our purpose. As stated before, in certain experimental platforms, qubit reset is unavailable. Alternatively, one could resort to algorithmic cooling [23, 24]. However, algorithmic cooling can fail if the state of the qubits has off-diagonal components in the computational basis or if the qubits are correlated. Neither of these possibilities can be ruled out in our setup.

In contrast, our protocol enjoys the following properties. First, our protocol is unitary. Therefore, it can be implemented on any quantum computing platform. Second, it successfully produces fresh qubits even in a regime where algorithmic cooling fails. Third, our protocol is straightforward to implement. The protocol involves partially applying the time-reversed quantum circuit over the qubits that are no longer in use, a capability that we expect to be available in many near-term quantum computers. There is no need for measurement, reset or new unitary gates that are incompatible with the given hardware. Lastly, we numerically demonstrate that our protocol works in the presence of noise. While we do not have a theoretical understanding of this phenomenon, our numerical evidence seems compelling.

The rest of this paper is structured as follows. In Section II, we introduce our protocol. In Section III, we provide rigorous performance guarantee of our protocol, demonstrating its robustness to noise in Section V. In Section IV, we describe and analyze a method for using our protocol to effectively boost the number of qubits available for implementing a convolutional circuit. We conclude with a discussion in Section V.

II. REWINDING PROTOCOL

In this section, we introduce the rewinding protocol. Without loss of generality, a general quantum computation can be thought of as a sequence of gates applied to the $|0^n\rangle$ state, where n is the number of qubits. We shall refer to this sequence as

$$\mathcal{C} = (g_1, g_2, \dots, g_{N-1}, g_N), \quad (1)$$

where g_i belongs to some universal gate set for every i . This circuit implements a unitary

$$U(\mathcal{C}) := g_N g_{N-1} \cdots g_2 g_1. \quad (2)$$

In the quantum circuit, some of the qubits may no longer participate in the computation after a certain point. Specifically, we may be able to break up \mathcal{C} into

$$\begin{aligned} \mathcal{C}_1 &= (g_1, \dots, g_{k-1}), \\ \mathcal{C}_2 &= (g_k, \dots, g_N), \end{aligned} \quad (3)$$

such that \mathcal{C}_2 acts trivially on a subset of qubits. For such qubits, we will say that they are *idle* during \mathcal{C}_2 .

The main goal of our protocol is to convert some of these idle qubits to $|0^m\rangle$ for some $m > 0$. We propose to achieve this by implementing the following “rewinding” circuit. Let \mathcal{Q}_I be the set of idle qubits. Consider a subsequence of \mathcal{C}_1 consisting of gates that act exclusively on \mathcal{Q}_I . Let us refer to this circuit as

$$\mathcal{C}_{1,I} := (\tilde{g}_1, \dots, \tilde{g}_{N'}). \quad (4)$$

The rewinding circuit is defined as follows:

$$\mathcal{R}(\mathcal{C}_{1,I}) := (\tilde{g}_{N'}^\dagger, \dots, \tilde{g}_1^\dagger). \quad (5)$$

After applying the rewinding circuit (see Fig. 1 for an example), we obtain the following state:

$$\mathcal{R}(\mathcal{C}_{1,I})U(\mathcal{C})|0 \dots 0\rangle. \quad (6)$$

An important question is whether the state of the idle qubits becomes close to the $|0 \dots 0\rangle_I$ state in some sense. We formulate a sufficient condition under which this is possible. Of course, generally speaking, we should not expect this state to be close to $|0 \dots 0\rangle$. However, for a special class of circuits called *convolutional circuits* [8, 9, 14], one can guarantee this generically.

III. CONVOLUTIONAL CIRCUITS

In this section, we study how well the rewinding protocol works for convolutional circuits [13]. These circuits have a fixed-size active window that moves over the qubits of the circuit. More formally, a convolutional circuit is one which can be written in the following form:

$$U_{[n-k-1, n]} \cdots U_{[2, k+1]} U_{[1, k]}, \quad (7)$$

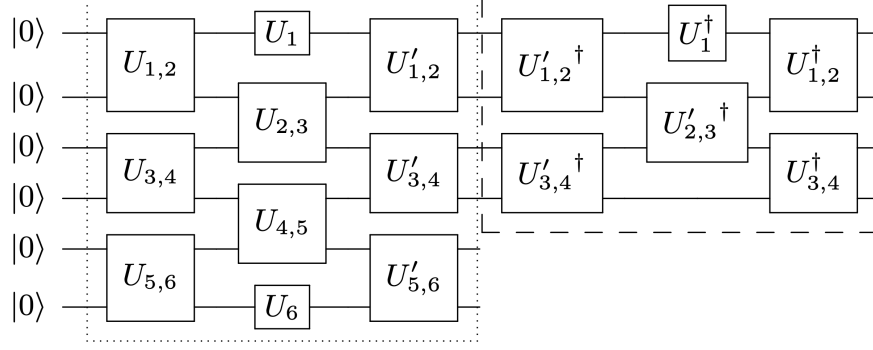


FIG. 1. An example of the rewinding circuit. Here the first 4 qubits are assumed to be idle. The subcircuit in the dashed box is the rewinding circuit of the circuit in dotted box.

over n qubits, where $U_{[i,j]}$ is a unitary operator acting on a set of qubits ranging from the i 'th to the j 'th index. Here k is a constant that is often chosen to be $\mathcal{O}(1)$.

Convolutional circuits play an important role in near-term quantum algorithms because they naturally appear in quantum circuits that can approximately prepare physical ground states of interest [8, 9, 14, 22]; see Appendix C for a review. Specifically, for the purpose of evaluating the expectation values of local observables, which is a key subroutine used in many near-term quantum algorithms, one can compress those circuits into a convolutional circuit. Without loss of generality, suppose our goal is to evaluate

$$\langle \psi | O | \psi \rangle \quad (8)$$

where $|\psi\rangle$ is the state prepared by the circuits in Ref. [8, 9, 14]. One can estimate this expectation value by preparing a new state $|\psi_O\rangle$, which can be expressed as

$$|\psi_O\rangle = U_{(n-1),n} \cdots U_{2,3} U_{1,2} |0 \dots 0\rangle, \quad (9)$$

where $U_{i,i+1}$ represents a unitary that acts on two groups of qubits, labeled by i and $i+1$ [15], and the support of the observable O lies in the union of the last two groups of qubits; see Fig. 2 for an illustration.

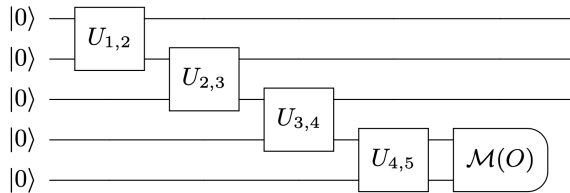


FIG. 2. A circuit diagram representing the state preparation of Eq. (9) for $n = 5$ followed by the measurement of an observable O .

An important property of convolutional circuits is that the state preparation of Eq. (9) can be implemented using a constant number of qubits (independent of n), provided

that one can reset the qubits at will; see Fig. 3. However, as explained above, the experimental implementation of reset remains challenging, which motivates the search for alternatives. By applying the rewinding protocol to the preparation circuit truncated after the gate $U_{m-1,m}$, we obtain the circuit diagram in Fig. 4. The rewind circuit prepares the same state on qubit m as the original. In addition, though, the first few qubits are returned to the $|0\rangle$ state up to an error exponentially small in m :

$$\langle 0^k | \rho_{[1,k]} | 0^k \rangle \geq 1 - c \exp(-\alpha(m-k)) \quad (10)$$

where c and α are numerical constants. For a sufficiently large m , the error becomes negligible. Therefore, one can inject these freshly reset qubits into the remainder of the preparation circuit.

The rest of this section will be devoted to explaining why Eq. (10) is true. A sketch of the proof will be presented in Section III A. We also report on a numerical experiment that corroborates the proof in Section III B, focusing on extracting a precise estimate on the average number of iterations one needs to bound the fidelity to a level that is available in realistic devices, namely 0.01.

A. Fidelity guarantee

In this section, we sketch the proof of Eq. (10), deferring the details to Appendix A. To explain the main idea, it will be convenient to replace the circuit diagram of Fig. 4 with a “process” diagram in which the unitary gate $U_{i,i+1}$ is replaced by a superoperator $\mathcal{U}_i(\cdot) = U_{i,i+1}(\cdot)U_{i,i+1}^\dagger$. For simplicity, suppose we trace out all but the first qubit; see Fig. 5.

Our argument will be directed at analyzing *transfer operators* associated to the circuit, which will describe how portions of the rewinding circuit affect the rest. Consider a sequence of quantum channels (Φ_i) for $i = 1, \dots, n$, whose action on a state ρ is defined as follows.

$$\begin{array}{c} \rho \\ |0\rangle\langle 0| \end{array} \begin{array}{c} \boxed{\mathcal{U}_i} \\ \boxed{\Phi_{i+1}} \end{array} \begin{array}{c} \boxed{\mathcal{U}_i^\dagger} \\ \text{Tr} \end{array} \Phi_i(\rho) \quad (11)$$

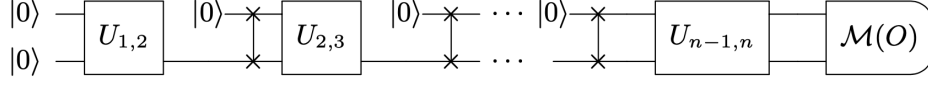


FIG. 3. A space-efficient implementation of Eq. (9). Compared to Fig. 2, the number of requisite qubits is reduced to a constant.

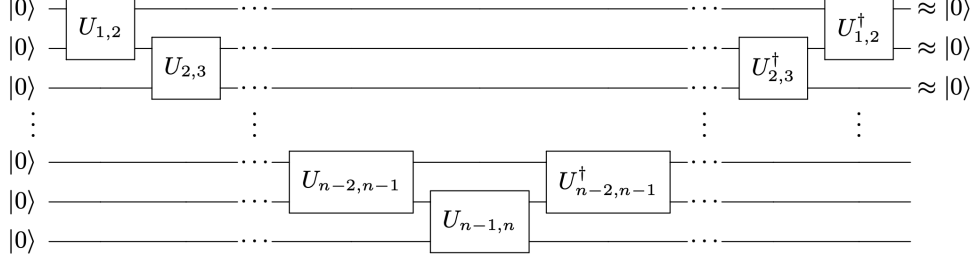


FIG. 4. Applying the rewinding protocol (see Section II) to Eq. (9).

Because Φ_{i+1} is a composition of quantum channels, so is Φ_i . The transfer operator $\mathcal{T}_{i+1 \rightarrow i}$ is the transformation taking the channel Φ_{i+1} to Φ_i .

The transfer operator has several useful properties. First, it maps the identity channel to the identity channel. In particular, viewing the transfer operator as a linear operator, the identity channel becomes an eigenvector with an eigenvalue 1. Moreover, the transfer operator has an eigenvalue with modulus less or equal to 1. Therefore, generically, by applying the transfer operator sufficiently many times, all the non-identity components decay away. What is left is a quantum channel that is exponentially close to the identity channel. We make this argument more precise in Appendix A, proving that

$$\|\Phi_1 - \mathcal{I}\|_\diamond = \mathcal{O}(\lambda^n), \quad (12)$$

where λ is the modulus of the second largest eigenvalue of $\mathcal{T}_{i+1,i}$, maximized over all i , and $\|\cdot\|_\diamond$ is the diamond norm. In particular, we obtain

$$\|\rho - \Phi_1(\rho)\|_1 = \mathcal{O}(\lambda^n). \quad (13)$$

In practice, it could be difficult to estimate the precise value of λ . Without that knowledge, one cannot provide a quantitative bound on the fidelity of the rewinding process. However, there is a simple method that can overcome this problem. One can simply estimate the fidelity of the recycled state by measuring those qubits in the computational basis. Assuming that we expect the fidelity to be bounded from below by $1 - \delta$, to estimate the fidelity up to an error of ϵ , it suffices to take $\sim \delta(1 - \delta)/\epsilon^2$ samples.

Of course, Eq. (13) will yield a nontrivial bound only if $\lambda < 1$. There are examples for which λ is equal to 1, for instance, when U is chosen to be the SWAP operation. However, there is an evidence suggesting that $\lambda < 1$ generically. As we shall briefly mention in Section III B, if we choose $U_{i,i+1}$ to be identical for all i and choose the

underlying unitary randomly, we generically observe an exponential decay in n for $\|\rho - \Phi_1(\rho)\|_1$.

B. Numerical experiment

While our proof firmly establishes that the rewinding protocol works for convolutional circuits, Eq. (12) is a conservative bound. To get a more precise quantitative estimate of the actual fidelity we can typically achieve, we have performed a numerical experiment using the Tensor Network library [25].

Consider Haar-random gates $U, V \in SU(4)$. We choose the gates in the circuit to be $U_{1,2} = U_{2,3} = \dots = U_{n-2,n-1} = U$ and $U_{n-1,n} = V$. We applied the rewinding protocol to this circuit, making the length $n = 150$. Fig. 6 presents the median (over 2000 trials) of the overlap of the i -th qubit with $|0\rangle$ for different values of i .

As shown in Fig. 6, our protocol yields a fidelity converging exponentially to 1 with an average decay rate of 0.22. Typically, to reach a single-qubit preparation error of 0.01 and 0.001, we needed to apply the protocol of length 19 and 29 respectively.

In all the randomly generated instances we studied, we observed that the infidelities decay exponentially, albeit with varying decay constants. For instance, there were rare instances for which the exponential decay became evident only at around $x \approx 10^3$. This suggests that λ is likely to be generically strictly smaller than 1 and the probability of λ being close to 1 is small. Making these numerical observations mathematically rigorous is an important open problem left for future work.

We have also performed a similar experiment in which $U_{i,i+1}$ were chosen randomly and independently from each other. The result was similar to Fig. 6. An interested reader may access the code in Ref. [26].

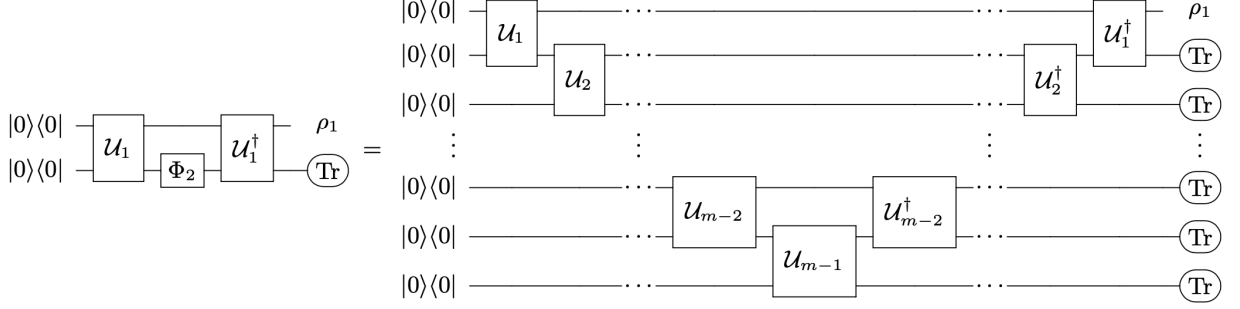


FIG. 5. Circuit in Fig. 4 in the density matrix picture. Here \mathcal{U}_i is a quantum channel that applies a unitary gate $U_{i,i+1}$. After tracing out all but the first qubit, we obtain a density matrix ρ_1 . One can view the right hand side as a recursive application of Eq. (11).

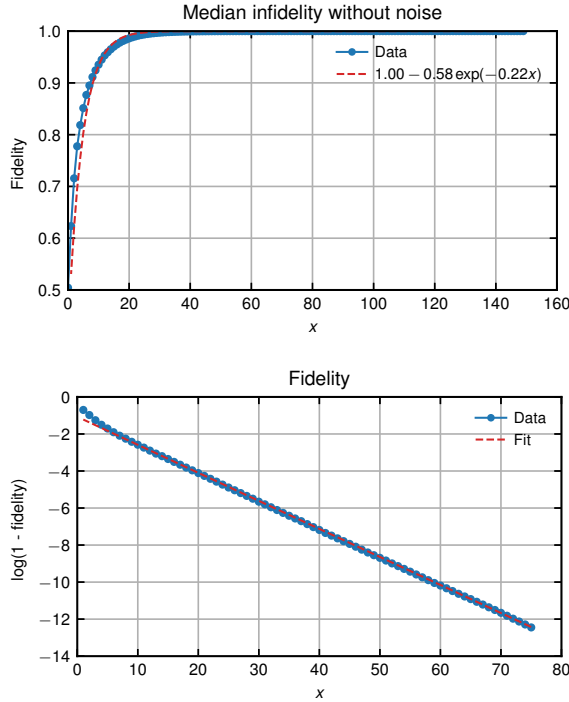


FIG. 6. Median fidelity plots. Here x is $150 - i$, where i is the index of the qubit as defined above. The base of the logarithm is e .

C. Robustness to noise

Numerical evidence indicates that the rewinding protocol is also remarkably robust to noise. To model those imperfections, we redid the numerical experiment of Section III B but followed every gate with noise. Specifically, each qubit the gate acts on was completely depolarized independently with probability 1%. The results are summarized in Fig. 7. The fidelity of the qubits being reset converges to a quantity that is close to 1 even in the presence of noise. In fact, the fidelity converges to a sta-

tionary value close to 0.99, which suggests that the stationary value is $1 - \mathcal{O}(p)$, where p is the error probability per gate. Crucially, there is no indication that noise accumulates with increasing circuit depth. Moreover, up to the precision in our numerical experiment, the decay rate is 0.22, which is the same value obtained *in the absence of noise*. In Appendix B, we have formulated a condition under which this robustness can be explained rigorously.

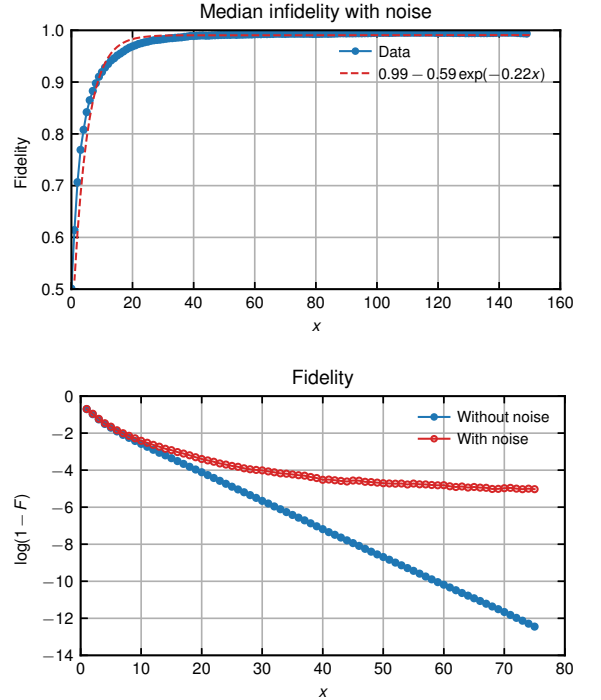


FIG. 7. Median fidelity in the presence of noise, with the convention $x = 150 - i$, where i is the index of the qubit as defined above. The base of the logarithm is e .

IV. BOOSTING THE NUMBER OF AVAILABLE QUBITS

By repeatedly using the rewinding protocol, we can obtain an *asymptotic* reduction in the number of qubits needed. Suppose, for concreteness, that we are equipped with an n_q qubit quantum computer but somehow the circuit of interest consists of $n > n_q$ qubits. What should we do? Often, one is interested in estimating the expectation value of the observable up to some error, say ϵ . The main question we address in this section is the size of the circuit one can simulate using n qubits while ensuring that the error is below ϵ .

The idea is to use the rewinding protocol recursively. Suppose, for instance, that we run the circuit only up to the first $n_1 = n_q$ qubits and then apply the rewinding protocol, resetting the first k_1 qubits. We can feed those k_1 qubits in as new ancillae to be used for the remaining part of the circuit, incurring an error bounded by $\epsilon_1 = ce^{-\alpha(n_1-k_1)}$ in doing so. We are once again constrained by space so we apply a circuit over $n_2 = k_1 + 1$ qubits, namely the output of the first portion in addition to the new ancillae, and then apply the rewinding protocol. If we keep the first k_2 qubits, we incur an error on these qubits which can be bounded by

$$\epsilon_2 \leq \epsilon_1 + ce^{-\alpha(n_2-k_2)}. \quad (14)$$

By repeating this t times, the total error is bounded by:

$$\epsilon_{\text{total}} \leq \sum_{i=1}^t c(t-i+1)e^{-\alpha(k_{i-1}-k_i)}, \quad (15)$$

where $k_0 = n_q$. The number of qubits effectively available for the original convolutional circuit is therefore $\sum_{i=0}^t k_i$. Since the size of a convolutional quantum circuit is proportional to the number of available qubits, the size is itself $\Theta(\sum_{i=0}^t k_i)$.

Let us choose $k_{i-1} - k_i$ to be a constant, denoted Δk . To leading order, the Δk that ensures ϵ_{total} to be below ϵ is

$$\Delta k \approx \frac{\ln(t^2/\epsilon)}{\alpha}. \quad (16)$$

We can then choose $t \sim \frac{\alpha n_q}{\ln(1/\epsilon) \ln(n_q)}$ to get, to leading order,

$$n_{\text{circuit}} \approx \frac{\alpha n_q^2}{2 \ln(1/\epsilon) \ln n_q}. \quad (17)$$

Therefore, we get a nearly quadratic boost in the size of the circuit we can implement. For $\epsilon = 0.01$ and the numerically observed average value of $\alpha = 0.22$, we obtain $n_{\text{circuit}} \approx 0.024 n_q^2 / \ln n_q$. The rewinding protocol always increases the number of available qubits but this bound only detects the advantage at $n_q \sim 90$ qubits. A more optimized choice of $k_{i-1} - k_i$ could potentially yield a better bound.

The preceding analysis allowed for the possibility that the induced error on the recycled qubits could be highly correlated. In practice, the error model might be close to being independent. While we do not have a general proof for such a statement, we have performed a preliminary numerical experiment that supports the hypothesis. If that is really the case, an *extensive* error occurring on every qubit, say of fixed strength ϵ , generically only gives rise to an error of order ϵ to the final answer [8, 9, 14]. Understanding to what degree this is true and what the ultimate limit of our protocol is an open problem left for future work.

As a practical matter, it is important to note that the procedure described above is highly parallelizable. There is no reason to wait until a given iteration of the rewinding protocol is complete before recycling qubits. As soon as a sufficiently high-fidelity qubit is prepared, one can use this qubit while running the remainder of the recycling protocol concurrently. The procedure discards the first Δk qubits processed by the rewinding protocol, which will impose a corresponding delay at each level of recursion, but the delay is independent of the number of qubits n_q , depending only on the quality of the desired recycled qubits. In fact, the delay can even be eliminated by maintaining a buffer of Δk fresh qubits. Therefore, the depth of the convolutional circuit implemented using the rewinding protocol to achieve a nearly quadratic savings in spatial cost can be made to be the same as the circuit executed with a physical supply of fresh qubits.

V. DISCUSSION

In this paper, we proposed a simple protocol for unitarily implementing a reset operation on a qubit. This protocol is particularly useful for reducing the spatial cost in implementing convolutional quantum circuits, which frequently appear in a large family of promising variational quantum circuits that can efficiently prepare physical ground states of interest [8, 9, 14–18]. While implementation of such circuits using a physical reset operation can be realized directly in ion qubits [18] and indirectly via measurements and conditional correction in superconducting qubits [27, 28], the reported experimental system consists of a small number of qubits. By using our method, it should be possible to run these circuits on the existing larger quantum computers which are not equipped with physical reset operations.

Recently, Borregaard et al. [15] pointed out that convolutional circuits can be truncated to shorter circuits, leading to a reduction in spatial (and temporal) cost. Interestingly, the key parameter in their work that determines the savings is different from ours. In their case, the modulus of the second largest eigenvalue of the transfer matrix acting on the space of density matrices determines the optimal depth of the circuit beyond which it does not make sense to run the circuit any further. In our case, the savings are determined by the second largest eigenvalue

of the transfer matrix acting on the *space of completely positive trace-preserving maps*; see Appendix A. Understanding which of these two methods is better in practice is an open problem left for future work.

There are several directions to improve our method. First, one may be able to devise a method to efficiently certify that the newly prepared ancilla states have large overlaps with the $|0\rangle$ state. One straightforward way to do this is to measure the state of the qubit directly, thus gathering the statistics. However, this process can predict a lower fidelity because of the measurement error. A method to remove such measurement error in the spirit of randomized benchmarking [29, 30] would be desirable. Second, one may be able to concatenate our method with the well-known algorithmic cooling method to further improve the purity of the ancilla qubits [23, 24].

Our work focused on convolutional quantum circuits for which the local Hilbert space dimension is bounded. This is an important assumption because the bounds we derived in this paper were in terms of 2-norm in some vector space. However, in order to convert a bound on this norm to a statement about the fidelity, one generally introduces a dimension-dependent factor. Therefore, if the local Hilbert space dimension is large, this bound may not be so useful. A similar problem was tackled in Ref. [14], by introducing a more refined notion of mixing called *local mixing*. It will be interesting to understand if there is an analogous physically plausible local mixing

condition in our context.

Lastly, the noise-resilience of our protocol warrants a further study. Specifically, if every gate in the circuit is replaced by an ideal gate followed by a local noise channel, would the fidelity of the prepared qubits be still $1 - \mathcal{O}(\epsilon)$ where ϵ is the noise strength? Alternatively, would the fidelity scale as $1 - \mathcal{O}(n\epsilon)$, where n is the circuit depth? Our numerical experiment suggests the former, and our analysis in Appendix B definitively shows that there is a sufficient condition under which one can rigorously prove the noise-resilience of our protocol. Whether this condition holds generically for convolutional quantum circuits is an important open problem that is left for future work.

ACKNOWLEDGMENTS

We thank Dave Bacon and Stefan Leichenauer for helpful conversations. GA acknowledges the support of the Stanford Institute for Theoretical Physics and her coauthors. IK was supported by Google, the Simons Foundation It from Qubit Collaboration and by the Australian Research Council via the Centre of Excellence in Engineered Quantum Systems (EQUS) project number CE170100009. PH acknowledges the support of Google, as well as AFOSR (award FA9550-19-1-0369), CIFAR, and the Simons Foundation.

-
- [1] P. W. Shor, SIAM Journal on Computing **26**, 1484 (1997).
 - [2] S. Lloyd, Science **273**, 1073 (1996).
 - [3] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Nature **574**, 505 (2019).
 - [4] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, Nat. Commun **5**, 4213 (2014).
 - [5] E. Farhi, J. Goldstone, and S. Gutmann, (2014), arXiv:1411.4028.
 - [6] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Nat. Commun **9**, 1 (2018).
 - [7] D. Wecker, M. B. Hastings, and M. Troyer, Phys. Rev. A **92**, 042303 (2015).
 - [8] I. H. Kim, (2017), arXiv:1702.02093.
 - [9] I. H. Kim and B. Swingle, (2017), arXiv:1711.07500.
 - [10] I. Cong, S. Choi, and M. D. Lukin, Nat. Phys. **15**, 1273 (2019).
 - [11] J. Bausch, (2020), arXiv:2006.14619.
 - [12] S. Hadfield, Z. Wang, B. O’Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, Algorithms **12** (2019), 10.3390/a12020034.
 - [13] D. Poulin, J.-P. Tillich, and H. Ollivier, “Quantum serial turbo-codes,” (2009), arXiv:0712.2888 [quant-ph].
 - [14] I. H. Kim, (2017), arXiv:1703.00032.
 - [15] J. Borregaard, M. Christandl, and D. S. França, (2019), arXiv:1909.04786.
 - [16] F. Barratt, J. Dborin, M. Bal, V. Stojevic, F. Pollmann, and A. G. Green, (2020), arXiv:2003.12087.
 - [17] C. T. Olund, M. Block, S. Gazit, J. McGreevy, and N. Y. Yao, Phys. Rev. B **101**, 155152 (2020).
 - [18] M. Foss-Feig, D. Hayes, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, J. M. Pino, and A. C. Potter, (2020), arXiv:2005.03023.
 - [19] I. Cirac, D. Perez-Garcia, N. Schuch, and F. Verstraete, arXiv preprint arXiv:2011.12127 (2020).
 - [20] F. Verstraete and J. I. Cirac, (2004), arXiv:cond-mat/0407066.
 - [21] G. Vidal, Phys. Rev. Lett. **99**, 220405 (2007).
 - [22] J. Haegeman, B. Swingle, M. Walter, J. Cotler, G. Evenbly, and V. B. Scholz, Phys. Rev. X **8**, 011003 (2018).

- [23] J. M. Fernandez, S. Lloyd, T. Mor, and V. Roychowdhury, (2004), arXiv:quant-ph/0401135.
- [24] L. J. Schulman, T. Mor, and Y. Weinstein, Phys. Rev. Lett. **94**, 120501 (2005).
- [25] C. Roberts, A. Milsted, M. Ganahl, A. Zalcman, B. Fontaine, Y. Zou, J. Hidary, G. Vidal, and S. Leichenauer, (2019), arXiv:1905.01330.
- [26] “Recycling Qubits: Repository,” <https://github.com/ikim-quantum/cooling>.
- [27] D. Ristè, L. C. G. Govia, B. Donovan, S. D. Fallek, W. D. Kalfus, M. Brink, N. T. Bronn, and T. A. Ohki, npj Quantum Information **6**, 71 (2020).
- [28] A. D. Corcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Mineev, J. M. Chow, and J. M. Gambetta, (2021), arXiv:2102.01682.
- [29] J. Emerson, R. Alicki, and K. Życzkowski, J. Opt. B **7**, S347 (2005).
- [30] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Phys. Rev. A **77**, 012307 (2008).
- [31] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf, Phys. Rev. Lett. **95**, 110503 (2005).
- [32] G. Evenbly and G. Vidal, Phys. Rev. B **79**, 144108 (2009).
- [33] W. F. Stinespring, Proceedings of the American Mathematical Society **6**, 211 (1955).

Appendix A: Transfer operator

In this appendix, we will prove properties of the transfer operator introduced in Section III A. Let us first recall the definition of the transfer operator. Without loss of generality, let $U : \mathcal{H}_A \otimes \mathcal{H}_B \rightarrow \mathcal{H}_A \otimes \mathcal{H}_B$ be a unitary on two quantum systems, A and B . Let $\mathcal{U}(\rho) = U\rho U^\dagger$ be the corresponding channel. The transfer operator is defined as follows:

$$\mathcal{T}_U[\Phi] = \text{Tr}_B \circ \mathcal{U}^\dagger \circ (\mathcal{I}_A \otimes \Phi) \circ \mathcal{U} \circ \mathcal{E} \quad (\text{A1})$$

where $\mathcal{E}(\rho) = \rho \otimes |0\rangle_B \langle 0|_B$.

Clearly, \mathcal{T}_U maps quantum channels to quantum channels. It will be helpful to be precise about the domain and range \mathcal{T}_U . In order to be able to introduce eigenvectors, it will be convenient to view \mathcal{T}_U as an operator acting on the complex vector space $\langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle$ spanned by the CPTP maps $\mathcal{CP}_{\mathcal{H}}^{TP}$. There is a natural inner product $\langle \cdot, \cdot \rangle$ on $\langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle$ induced by the Hilbert space inner product on \mathcal{H} .

In order to study properties of \mathcal{T}_U , it will be convenient to define two norms: the “Hilbert-Schmidt” norm over $\langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle$, which we shall denote as $\|\cdot\|_2$, and an operator norm for the operators that act on $\langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle$, which we shall denote as $\|\cdot\|$. Formally, these norms are defined as follows.

$$\|\Phi\|_2 := \sqrt{\langle \Phi, \Phi \rangle} \quad (\text{A2})$$

and

$$\|\mathcal{T}\| = \sup_{\Phi \in \langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle} \frac{\|\mathcal{T}[\Phi]\|_2}{\|\Phi\|_2}. \quad (\text{A3})$$

As a first application, we first show that \mathcal{T}_U has eigenvalues with modulus less or equal to 1.

Theorem A.1. *Suppose U is a unitary and \mathcal{T}_U is defined as in Eq. (A1). Let $\Phi \in \langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle$ be a right eigenvector of \mathcal{T}_U with eigenvalue λ . Then $|\lambda| \leq 1$.*

Proof. Without loss of generality, let $\Phi = \sum_{i \in I} \alpha_i \Psi_i$, for some $\alpha_i \in \mathbb{C}$ and $\Psi_i \in \mathcal{CP}_{\mathcal{H}}^{TP}$ for some finite index set I . The index can be made to be finite because the vector space $\langle \mathcal{CP}_{\mathcal{H}}^{TP} \rangle$ is finite-dimensional. Thus $\mathcal{T}_U^k[\Phi] = \sum_{i \in I} \alpha_i \mathcal{T}_U^k[\Psi_i]$ by linearity, where \mathcal{T}_U^k is the k -th repeated application of \mathcal{T}_U .

After applying \mathcal{T}_U k times, we obtain:

$$\|\mathcal{T}_U^k[\Phi]\|_2 = |\lambda|^k \|\Phi\|_2. \quad (\text{A4})$$

The key point is that this quantity becomes unbounded if $|\lambda| > 1$ in the $k \rightarrow \infty$ limit. This is a contradiction because, using triangle inequality, we get

$$\begin{aligned} \|\mathcal{T}_U^k[\Phi]\|_2 &\leq \sum_{i \in I} \|\alpha_i \mathcal{T}_U^k[\Psi_i]\|_2 \\ &\leq \sum_{i \in I} |\alpha_i| \|\tilde{\Psi}_i\|_2, \end{aligned} \quad (\text{A5})$$

where $\tilde{\Psi}_i \in \mathcal{CP}_{\mathcal{H}}^{TP}$ because \mathcal{T}_U^k maps quantum channels to quantum channels. Importantly, the last line of Eq. (A5) is bounded for all k because the set of quantum channels itself is bounded. If $|\lambda| > 1$, there exists a finite $k \in \mathbb{N}$ such that

$$|\lambda|^k \|\Phi\|_2 > \sum_{i \in I} |\alpha_i| \|\tilde{\Psi}_i\|_2, \quad (\text{A6})$$

which is a contradiction. Therefore, we conclude that $|\lambda| \leq 1$, completing the proof. \square

From Theorem A.1, we found that the eigenvalues of the right eigenvectors of \mathcal{T}_U must lie on $|\lambda| \leq 1$. Moreover, one can verify straightforwardly that the identity channel, denoted as \mathcal{I} , satisfies the following relation:

$$\mathcal{T}_U[\mathcal{I}] = \mathcal{I}. \quad (\text{A7})$$

This means that \mathcal{I} is a right eigenvector of \mathcal{T}_U with an eigenvalue 1. Provided that this is the only right eigenvector whose eigenvalue has a modulus of 1, the norm of the other right eigenvectors must be strictly less than 1. In particular, after a repeated application of \mathcal{T}_U , the norm in fact decays exponentially. Let us note in passing that the eigenvectors of \mathcal{T}_U need not be CPTP; the vector space on which \mathcal{T}_U is defined consists of linear combination of CPTP maps, which are generally not CPTP.

Corollary A.2. *Let \mathcal{T} be a transfer operator. Suppose there is only one right eigenvector of \mathcal{T} that has an eigenvalue with modulus 1. Let Φ be a CPTP map and define a sequence of CPTP maps as $\Phi_0 := \Phi$ and $\Phi_i := \mathcal{T}[\Phi_{i-1}]$. Then*

$$\|\Phi_k - \mathcal{I}\|_2 \leq \mathcal{O}(\Delta^k), \quad (\text{A8})$$

where $0 < \Delta < 1$. Specifically, Δ is the modulus of the largest non-unit eigenvalue of \mathcal{T} .

Proof. Recall that $\mathcal{T}[\mathcal{I}] = \mathcal{I}$. Therefore, if \mathcal{T} has only one right eigenvector, it must be \mathcal{I} . Without loss of generality,

$$\mathcal{T} = P + D, \quad (\text{A9})$$

where P is a projector onto the one-dimensional subspace spanned by \mathcal{I} and $D = \mathcal{T} - P$. Let us first prove that $P[\Phi] = \mathcal{I}$. Since \mathcal{T}_i maps a CPTP map to a CPTP map,

$$\lim_{n \rightarrow \infty} \mathcal{T}^n[\Phi] = P[\Phi] \quad (\text{A10})$$

is also a CPTP map. However, the only multiple of the identity that is a CPTP map is the identity map itself. Therefore, $P[\Phi] = \mathcal{I}$.

Without loss of generality, consider the decomposition of Φ into the linear combination of right eigenvectors of \mathcal{T} :

$$\Phi = \sum_{i \in I} \alpha_i \Psi_i, \quad (\text{A11})$$

where $\mathcal{T}[\Psi_i] = \lambda_i \Psi_i$. Since $P[\Phi] = \mathcal{I}$,

$$\Phi - \mathcal{I} = \sum_{\substack{i \in I, \\ |\lambda_i| < 1}} \alpha_i \Psi_i. \quad (\text{A12})$$

Using the fact that $\mathcal{T}[\mathcal{I}] = \mathcal{I}$,

$$\begin{aligned} \|\Phi_k - \mathcal{I}\|_2 &= \|\mathcal{T}^k[\Phi - \mathcal{I}]\|_2 \\ &= \left\| \sum_{\substack{i \in I, \\ |\lambda_i| < 1}} \alpha_i \Psi_i \lambda_i^k \right\|_2 \\ &\leq \Delta^k \sum_{\substack{i \in I, \\ |\lambda_i| < 1}} |\alpha_i| \|\Psi_i\|_2 \end{aligned} \quad (\text{A13})$$

Since we are working in a finite-dimensional vector space, each of the terms in the last line is bounded. \square

Let us remark that the distance $\|\Phi_k - \mathcal{I}\|_2$ upper bounds a more operationally meaningful distance of

$\|\Phi_k - \mathcal{I}\|_\diamond$ up to a dimension-dependent factor. Because the underlying Hilbert space is finite, the fact that this distance decays exponentially in k remains the same. This completes the main claim in Section III A.

Appendix B: Noise resilience

In this appendix, we formulate a sufficient condition under which the rewinding protocol becomes resilient to noise. The central concept is the notion of *contractivity* of the transfer matrix.

Definition 1. Let \mathcal{T} be a linear map which maps a channel to another channel. \mathcal{T} is contractive if

$$\|\mathcal{T}[\Phi_1 - \Phi_2]\|_2 \leq \gamma(\mathcal{T}) \|\Phi_1 - \Phi_2\|_2 \quad (\text{B1})$$

for some real number $\gamma(\mathcal{T}) < 1$. If such a $\gamma(\mathcal{T})$ exists, the smallest such number will be defined as the contraction coefficient of \mathcal{T} .

Below, we will show that if the transfer matrix in Eq. (A1) is contractive, the rewinding protocol is resilient to noise. Specifically, without loss of generality, consider the noisy version of the transfer matrix (cf. Eq. (A1)):

$$\mathcal{T}_{U,\delta}[\Phi] = \text{Tr}_B \circ \mathcal{U}_\delta^\dagger \circ (\mathcal{I}_A \otimes \Phi) \circ \mathcal{U}_\delta \circ \mathcal{E}, \quad (\text{B2})$$

where \mathcal{U} and $\mathcal{U}_\delta^\dagger$ are noisy versions of \mathcal{U} and \mathcal{U}^\dagger :

$$\begin{aligned} \|\mathcal{U}_\delta - \mathcal{U}\| &\leq \delta, \\ \|\mathcal{U}_\delta^\dagger - \mathcal{U}^\dagger\| &\leq \delta. \end{aligned} \quad (\text{B3})$$

Our goal is to show that the fixed point of $\mathcal{T}_{U,\delta}$ is close to that of \mathcal{T}_U up to a distance $\mathcal{O}(\delta/(1-\gamma))$.

For small δ , $\mathcal{T}_{U,\delta}$ and \mathcal{T}_U are close to each other, as we prove below.

Lemma 1.

$$\|\mathcal{T}_{U,\delta} - \mathcal{T}_U\| \leq C\delta \quad (\text{B4})$$

for some constant C .

Proof. Note the following chain of inequalities.

$$\begin{aligned}
\|\mathcal{T}_{U,\delta} - \mathcal{T}_U\| &= \sup_{\Phi \in \langle \mathcal{CP}_{\mathcal{H}}^{T,P} \rangle} \frac{\|(\mathcal{T}_{U,\delta} - \mathcal{T}_U)[\Phi]\|_2}{\|\Phi\|_2} \\
&= \sup_{\Phi \in \langle \mathcal{CP}_{\mathcal{H}}^{T,P} \rangle} \frac{\|\text{Tr}_B \circ (\mathcal{U}_\delta^\dagger - \mathcal{U}^\dagger) \circ (\mathcal{I}_A \otimes \Phi) \circ \mathcal{U}_\delta \circ \mathcal{E} + \text{Tr}_B \circ \mathcal{U}^\dagger \circ (\mathcal{I}_A \otimes \Phi) \circ (\mathcal{U}_\delta - \mathcal{U}) \circ \mathcal{E}\|_2}{\|\Phi\|_2} \\
&\leq \sup_{\Phi \in \langle \mathcal{CP}_{\mathcal{H}}^{T,P} \rangle} \frac{\|\text{Tr}_B \circ (\mathcal{U}_\delta^\dagger - \mathcal{U}^\dagger) \circ (\mathcal{I}_A \otimes \Phi) \circ \mathcal{U}_\delta \circ \mathcal{E}\|_2}{\|\Phi\|_2} + \sup_{\Phi \in \langle \mathcal{CP}_{\mathcal{H}}^{T,P} \rangle} \frac{\|\text{Tr}_B \circ \mathcal{U}^\dagger \circ (\mathcal{I}_A \otimes \Phi) \circ (\mathcal{U}_\delta - \mathcal{U}) \circ \mathcal{E}\|_2}{\|\Phi\|_2} \\
&\leq \|\text{Tr}_B\| \|\mathcal{I}_A\| \|\mathcal{E}\| \left(\|\mathcal{U}_\delta^\dagger - \mathcal{U}^\dagger\| \|\mathcal{U}_\delta\| + \|\mathcal{U}_\delta - \mathcal{U}\| \|\mathcal{U}^\dagger\| \right) \\
&\leq \delta \|\text{Tr}_B\| \|\mathcal{I}_A\| \|\mathcal{E}\| (\|\mathcal{U}_\delta\| + \|\mathcal{U}^\dagger\|).
\end{aligned} \tag{B5}$$

Since the operators placed in the norm in the last line are bounded operators acting on a finite-dimensional vector space, their norms are bounded. \square

Proposition 1. Consider a sequence $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_k$ of transfer operators, each potentially defined by a different unitary transformation. Suppose \mathcal{T}_i has a contraction coefficient of γ_i and let $\gamma = \max_i \gamma_i$.

For any sequence of operators $\tilde{\mathcal{T}}_1, \tilde{\mathcal{T}}_2, \dots, \tilde{\mathcal{T}}_k$ which map channels to channels such that $\|\mathcal{T}_i - \tilde{\mathcal{T}}_i\| \leq \delta$ for all $i \in \{1, \dots, k\}$,

$$\|(\tilde{\mathcal{T}}_k \circ \dots \circ \tilde{\mathcal{T}}_1 - \mathcal{T}_k \circ \dots \circ \mathcal{T}_1)[\Phi]\|_2 \leq \frac{C'\delta}{1-\gamma} \tag{B6}$$

for some constant $C' > 0$.

Proof. Consider the telescopic decomposition of $\tilde{\mathcal{T}}_k \circ \dots \circ \tilde{\mathcal{T}}_1 - \mathcal{T}_k \circ \dots \circ \mathcal{T}_1$:

$$\begin{aligned}
&\tilde{\mathcal{T}}_k \circ \dots \circ \tilde{\mathcal{T}}_1 - \mathcal{T}_k \circ \dots \circ \mathcal{T}_1 \\
&= \sum_{i=1}^{k-1} \mathcal{T}_{[i+1,k]} \circ (\tilde{\mathcal{T}}_i - \mathcal{T}_i) \circ \tilde{\mathcal{T}}_{[0,i-1]},
\end{aligned} \tag{B7}$$

where $\mathcal{T}_{[i,j]}$ and $\tilde{\mathcal{T}}_{[i,j]}$ is a sequential application of \mathcal{T}_k for k from i to j . Here we assume that $j \geq i$ and also that \mathcal{T}_0 is an identity operation, returning every channel to itself.

Therefore,

$$\begin{aligned}
&\left\| \sum_{i=1}^{k-1} \mathcal{T}_{[i+1,k]} \circ (\tilde{\mathcal{T}}_i - \mathcal{T}_i) \circ \tilde{\mathcal{T}}_{[0,i-1]}[\Phi] \right\|_2 \\
&\leq \sum_{i=1}^{k-1} \|\mathcal{T}_{[i+1,k]} \circ (\tilde{\mathcal{T}}_i - \mathcal{T}_i)[\Phi_i]\|_2
\end{aligned} \tag{B8}$$

for some channel Φ_i . Note that

$$\|\tilde{\mathcal{T}}_i[\Phi_i] - \mathcal{T}_i[\Phi_i]\| \leq C'\delta \tag{B9}$$

for some constant C' and that both $\tilde{\mathcal{T}}_i[\Phi_i]$ and $\mathcal{T}_i[\Phi_i]$ are channels. Using the definition of the contraction coefficient, (see Definition 1) we obtain the following bound:

$$\begin{aligned}
&\left\| \sum_{i=1}^{k-1} \mathcal{T}_{[i+1,k]} \circ (\tilde{\mathcal{T}}_i - \mathcal{T}_i) \circ \tilde{\mathcal{T}}_{[0,i-1]}[\Phi] \right\|_2 \\
&\leq C'\delta \sum_{i=1}^{k-1} \gamma^{k-i-1} \\
&\leq \frac{C'\delta}{1-\gamma}
\end{aligned} \tag{B10}$$

\square

Appendix C: Examples of convolutional circuits

Convolutional circuits appear in many situations, including the encoding circuits for convolutional quantum codes [13]. We focus on three examples: matrix product states [16, 31], holographic quantum circuits [8, 14], and deep multi-scale entanglement renormalization ansatz (DMERA) [9, 21, 32].

1. Matrix product states

Matrix product state (MPS) is a many-body quantum state which can be expressed in the following form:

$$|\Psi\rangle = \sum_{\{s\}} \text{Tr}[A_1^{(s_1)} A_2^{(s_2)} \dots A_n^{(s_n)}] |s_1 s_2 \dots s_n\rangle, \tag{C1}$$

where $\{A_i^{(s_i)}\}$ is a set of matrices of dimension χ and $s_i \in \{0, \dots, d-1\}$. Here χ is the dimension of the *virtual space* \mathcal{H}_v of the MPS and d is the local Hilbert space dimension of the constituent particles, which has a tensor product form $\otimes_i \mathcal{H}_i$ where $\dim(\mathcal{H}_i) = d$. Formally, one can view $A_i^{(s_i)}$ as a matrix A_i whose domain and image

is \mathcal{H}_v and $\mathcal{H}_v \otimes \mathcal{H}_i$. Specifically,

$$A_i|\alpha\rangle = \sum_{\beta, s_i} (A_i)_{\beta\alpha}^{(s_i)} |\beta\rangle |s_i\rangle, \quad (\text{C2})$$

where $\alpha, \beta \in \{0, \dots, \chi-1\}$ and $s_i \in \{0, \dots, d-1\}$.

One can consider the following special form of MPS:

$$|\Psi\rangle = \sum_{\{s\}} \langle \phi_F | A_1^{(s_1)} A_2^{(s_2)} \dots A_n^{(s_n)} | \phi_I \rangle |s_1 s_2 \dots s_n\rangle, \quad (\text{C3})$$

where $|\phi_I\rangle$ and $|\phi_F\rangle$ are χ -dimensional vectors. While a class of states in the form of Eq. (C3) forms a subclass of Eq. (C1), this subclass is still capable of describing many physical states of interest, such as the W-state, cluster state, and GHZ state [31].

While the matrices appearing in Eq. (C3) can be arbitrary, up to normalization, such states can be reexpressed as

$$|\Psi\rangle \otimes |\phi_F\rangle = \sum_{\{s\}} V_1^{(s_1)} V_2^{(s_2)} \dots V_n^{(s_n)} |\phi_I\rangle |s_1 s_2 \dots s_n\rangle, \quad (\text{C4})$$

where $V_i^{(s_i)}$ is an isometry from the virtual Hilbert space to the tensor product of the virtual Hilbert space and the local Hilbert space. Specifically, the action of this isometry, acting on a basis state $|\alpha\rangle$ in the virtual Hilbert space and $|s_i\rangle$ in the local Hilbert space, acts as

$$V_i|\alpha\rangle |s_i\rangle = \sum_{\beta, s_i} (V_i)_{\beta\alpha}^{(s_i)} |\beta\rangle |s_i\rangle. \quad (\text{C5})$$

Since V_i is an isometry from \mathcal{H}_v to $\mathcal{H}_v \otimes \mathcal{H}_i$, one can rewrite it as

$$V_i = U_i(|0\rangle \otimes I_v), \quad (\text{C6})$$

for some unitary acting on $\mathcal{H}_v \otimes \mathcal{H}_i$, where $|0\rangle \in \mathcal{H}_i$ and I_v is the identity operator acting on \mathcal{H}_v . Viewed this way, the state in Eq. (C3) can be created by first preparing the following state:

$$|\phi_I\rangle \otimes \underbrace{|0\rangle \otimes \dots \otimes |0\rangle}_{n \text{ copies}} \quad (\text{C7})$$

and applying U_n to U_1 sequentially.

The resulting circuit, up to a relabeling, has a convolutional structure. To see why, it is convenient to view the virtual Hilbert space as a $\lceil \log_d \chi \rceil$ -qudit system. The key idea is to convert the following tensor representation of U_i

$$\begin{array}{ccc} \mathcal{H}_i & \boxed{U_i} & \mathcal{H}_i \\ \mathcal{H}_v & & \mathcal{H}_v \end{array} \quad (\text{C8})$$

to the following form.

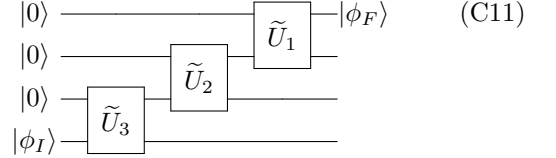
$$\begin{array}{ccc} \mathcal{H}_i & \boxed{\tilde{U}_i} & \mathcal{H}_v \\ \mathcal{H}_v & & \mathcal{H}_i \end{array} \quad (\text{C9})$$

Formally, \tilde{U}_i is U_i followed by this map:

$$|\alpha\rangle |s_i\rangle \rightarrow |s_i\rangle |\alpha\rangle \quad (\text{C10})$$

for all $\alpha \in \{0, \dots, \chi-1\}$ and $s_i \in \{0, \dots, d-1\}$. Since this is a permutation of qudits, it can be realized unitarily. Therefore, \tilde{U}_i is simply a unitary acting on $\lceil \log_d \chi \rceil + 1$ qudits.

Using \tilde{U}_i , the convolutional structure of MPS becomes evident. We provide the following example, which is a MPS over 3 qudits.



Generalization to n -qudit MPS is straightforward.

2. Holographic quantum circuits

Holographic quantum circuits [8, 14, 15, 18] generalize MPS to higher dimensions by replacing a qudit with a *collection* of qudits and imposing a spatial structure within U_i . These circuits can be used as ansatz to simulate a \mathcal{D} -dimensional quantum many-body systems by applying a sequence of local gates on a set of qudits arranged on a $(\mathcal{D}-1)$ -dimensional lattice.

For concreteness, let us describe a $\mathcal{D} = 2$ -dimensional example. Without loss of generality, suppose the physical system of interest consists of qudits with local Hilbert space dimension d , arranged on a square lattice of size $\ell_x \times \ell_y$. Each qudit shall be specified by its coordinate (x, y) , where $x \in [\ell_x] = \{0, \dots, \ell_x - 1\}$ and $y \in [\ell_y] = \{0, \dots, \ell_y - 1\}$. To describe our circuit, we will also need to introduce a total of ℓ_y “virtual” qudits, each of which has local Hilbert space dimension of χ . We shall specify them with an index $y \in [\ell_y]$.

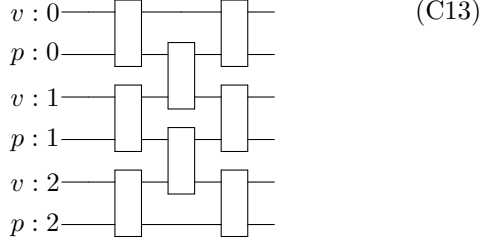
Just like in the MPS example, a holographic quantum circuit consists of a sequence of unitaries $\{U_i\}$ that act on the virtual qudits and a set of qudits at $x = i$. The state is prepared by first initializing all the qudits as follows

$$\underbrace{\left(\bigotimes_{y=0}^{\ell_y-1} |0\rangle_y \right)}_{\text{Virtual qudits}} \otimes \underbrace{\left(\bigotimes_{y=0}^{\ell_y-1} |0\rangle_{x,y} \right) \otimes \dots \otimes \left(\bigotimes_{y=0}^{\ell_y-1} |0\rangle_{x,y} \right)}_{\ell_x \text{ copies Physical qudits}} \quad (\text{C12})$$

and then applying U_i from $i = \ell_x - 1$ to 0.

Importantly, U_i is a *finite-depth local quantum circuit*, meaning that it is a $\mathcal{O}(1)$ -depth quantum circuit consisting of gates that act on qudits whose y -coordinates are close to each other. An example of such U_i for a depth

of 3 and $\ell_y = 3$ is described below.



Here $v : y$ is the virtual qudit labeled by y and $p : y$ is the physical qudit labeled by (i, y) .

To restore the convolutional structure, we can swap each virtual qudit with a physical qudit with the same y -coordinate, similar to how we did in Eqs. (C8) and (C9).

3. DMERA

DMERA is a multi-scale ansatz that can well-approximate many-body wavefunctions of scale-invariant quantum many-body systems [9, 21]. This ansatz can be prepared in a recursive manner. Without loss of generality, suppose we are interested in simulating a scale-invariant system in \mathcal{D} spatial dimensions. The DMERA ansatz over $2^{n\mathcal{D}}$ qudits can be constructed by first preparing an ansatz over the qudits arranged on a lattice of size $\underbrace{2^{n-1} \times \dots \times 2^{n-1}}_{\mathcal{D}}$, embedding the constituent qudits into a lattice of size $\underbrace{2^n \times \dots \times 2^n}_{\mathcal{D}}$ with a lattice spacing of $\underbrace{(2, \dots, 2)}_{\mathcal{D}}$, filling the remaining lattice sites by qudits initialized in the $|0\rangle$ state, and then applying a finite-

depth local quantum circuit. Without loss of generality, we will assume that the depths of these local circuits are bounded by D .

For the purpose of estimating the expectation value of a local observable O , one can remove some of the spurious gates that do not affect its expectation value. For instance, suppose the DMERA ansatz $|\Psi\rangle$ can be represented as $u|\Psi'\rangle$ for some gate u . If u commutes with O , then the following identity holds:

$$\langle \Psi | O | \Psi \rangle = \langle \Psi' | O | \Psi' \rangle. \quad (\text{C14})$$

By applying this identity recursively, one can remove gates used in preparing $|\Psi\rangle$ until this procedure no longer yields a smaller circuit. The resulting minimal circuit is referred to as the *past causal cone* of the observable O [32].

Once the past causal cone is obtained, the expectation value of O can be expressed as follows:

$$\langle O \rangle = \text{Tr}[\Phi_n \circ \dots \circ \Phi_1(\rho)O], \quad (\text{C15})$$

where Φ_i is a completely-positive trace preserving (CPTP) map called the *transfer operator* acting on $\mathcal{O}(D^{\mathcal{D}})$ qudits, whose exact form can be deduced from the DMERA circuit; see Ref. [9, 32] for details.

By the Stinespring dilation theorem [33], every CPTP map can be viewed as an isometry followed by a partial trace. Moreover, this isometry can be directly deduced from the DMERA circuit, as explained in Ref. [9, 32]. Therefore, Eq. (C15) can be viewed as an expectation value of O evaluated with respect to a state created by applying a sequence of isometries $\{V_i\}$ which are the dilations of the CPTP maps $\{\Phi_i\}$. Then, by the analysis of Appendix C 1, the expectation value of O can be viewed as an expectation value evaluated with respect to a state prepared by a convolutional circuit.