

Optimal compression of constrained quantum time evolutionMaurits S. J. Tepaske,¹ David J. Luitz,^{1,*} and Dominik Hahn²¹*Physikalisches Institut, Universität Bonn, Nussallee 12, 53115 Bonn, Germany*²*Max Planck Institute for the Physics of Complex Systems, Noethnitzer Strasse 38, 01167 Dresden, Germany*

(Received 1 February 2024; accepted 21 April 2024; published 15 May 2024)

The time evolution of quantum many-body systems is one of the most promising applications for near-term quantum computers. However, the utility of current quantum devices is strongly hampered by the proliferation of hardware errors. The minimization of the circuit depth for a given quantum algorithm is therefore highly desirable, since shallow circuits generally are less vulnerable to decoherence. Recently, it was shown that variational circuits are a promising approach to outperform current state-of-the-art methods such as Trotter decomposition, although the optimal choice of parameters is a computationally demanding task. In this work, we demonstrate a simplification of the variational optimization of circuits implementing the time evolution operator of local Hamiltonians by directly encoding constraints of the physical system under consideration. We study the expressibility of such constrained variational circuits for different models and constraints. Our results show that the encoding of constraints allows a reduction of optimization cost by more than one order of magnitude and scalability to arbitrary large system sizes, without losing accuracy in most systems. Furthermore, we discuss the exceptions in locally constrained systems and provide an explanation by means of an restricted lightcone width after incorporating the constraints into the circuits.

DOI: [10.1103/PhysRevB.109.205134](https://doi.org/10.1103/PhysRevB.109.205134)**I. INTRODUCTION**

Quantum computers are promising computational platforms which are expected to lead to breakthroughs for several difficult problems like integer factorization, optimization or machine learning algorithms by significant quantum speedups over their classical counterparts [1–4]. While most tasks will not become feasible before the realization of scalable quantum error correction, the simulation of quantum many-body systems is one of the most promising applications on near term devices [5]. The nonequilibrium behavior of different condensed matter systems [6–13], lattice gauge theories [14,15] or quantum interactive dynamics [16,17] are examples for applications which were studied recently. Current experimental platforms range from tens to more than a hundred qubits [18,19], with thousands of qubits projected for the near future. Such quantum processor sizes becomes increasingly challenging to simulate using the capabilities of classical computers [20–24].

While the physical system sizes are increasingly impressive, the reachable quantum volume [25] of present-day devices is still limited by the presence of noise. The coupling to the environment, gate imperfections, and measurement errors can destroy the coherence of the quantum state and limit the accuracy of the results of quantum computations. There are several attempts to quantify the amount of errors and to find protocols which allow one to average out the errors by increasing the number of required measurements. These techniques can be summarized by the term quantum error mitigation [26–39]. Recent work has shown that a combination of quantum simulation and quantum error mitigation

gives competitive results for the time evolution of quantum systems [23,24].

An alternative approach consists of reducing the circuit depth of the algorithm. The idea is simple: The proliferation of errors is suppressed when the possibilities where such an error can occur are reduced. In other words, one can try to squeeze a given quantum algorithm into a given quantum volume by finding a more efficient circuit representation of the algorithm. For simulations of time evolution in this era of noisy quantum computing, the state-of-the-art-method is the Trotter decomposition [40–43].

Here the unitary operator U generating the time evolution $U|\psi\rangle$ is expressed by a circuit of few-body gates, which can then be evaluated on a quantum computer. The error resulting from this mapping can be controlled by the circuit depth. Trotter circuits have the best known asymptotic error scaling for a large class of local quantum many-body systems and do not require any ancilla qubits [44]. They are projected to be the best iterative algorithms for gate and qubit counts that are expected in the early fault-tolerant era [45,46]. Nevertheless, at fixed time the circuit depth can potentially be made smaller by employing numerical optimization.

One promising approach are variational circuits [47–70]. The main idea is to describe the time-evolved state or the time evolution operator using a parametrized circuit. Choosing the parameters boils down to an optimization problem, which can be tackled using the gradient descent method. Several proposed algorithms try to optimize the circuit on a quantum computer or simulator [51–66]. However, this is currently unfeasible due to the high error rate on present-day devices. Other attempts, including our previous work, utilize a circuit optimization algorithm that is specifically designed for classical computers [67–70]. In some cases, we could reduce the circuit depth by almost 50% in comparison to standard

*david.luitz@uni-bonn.de

Trotter decomposition at fixed fidelity. While this approach is thus a promising strategy for the future, one major drawback is the large computational cost of the optimization. A possible solution for this issue is to exploit constraints of the system to reduce the number of independent parameters in the circuit. These can be encoded directly into the variational circuit ansatz.

In this work we study this approach for different constraints and analyze the resulting error scaling in comparison to generic variational circuits and Trotter decomposition.

This paper is organized as follows: In Secs. II A and II B, we describe our variational circuits and the implementation of constraints. Furthermore we describe our optimization procedure in detail in Sec. II C. In Sec. III, we present our results for optimizing the time evolution of three different models with different constraints: The Heisenberg XXZ chain, the PXP model and the quantum link model. Finally, we discuss the implications of our results for encoding constraints into variational circuits in Sec. IV.

II. METHODS

In this section, we explain in detail the different circuit architectures that we use to approximate the exact time evolution of a given Hamiltonian. In Sec. II A, we introduce a generic brickwall circuit architecture tailored to hardware-native gates. In the following discussions we will refer to these as unconstrained circuits. If the underlying time evolution has conserved charges or local constraints, these can be incorporated into our circuit template. This is described in Sec. II B. We refer to them as blocked circuits. In Sec. II C, we describe our cost functions and optimization strategies to determine the free parameters of our circuits. Finally, we explain in Sec. II D how the optimized circuits can be used to simulate time evolution on quantum computers for large systems and long time scales beyond the capability of classical devices.

A. Parameterized circuits

Since our main intention is to simplify the implementation of the time evolution of a many-body wave function using a quantum computer, it is sensible to use a set of native gates as building blocks. In our case, we use single-qubit gates parametrized as

$$u(\theta, \phi, \chi) = \begin{pmatrix} e^{i\phi} \cos(\theta) & e^{i\chi} \sin(\theta) \\ -e^{-i\chi} \sin(\theta) & e^{-i\phi} \cos(\theta) \end{pmatrix}, \quad (1)$$

and controlled NOT (CNOT) gates. These gates form a universal gate set [71].

The unconstrained circuit architecture C consists of M brickwall layers, as shown in Fig. 1(a). Each brickwall layer consists of two half-layers of CNOT gates connecting neighboring qubits, with general single-qubit layers interspersed. At the end of the circuit, we add a final layer of single-qubit unitaries. The advantage of the brickwall architecture is its minimal depth for a fixed amount of CNOTs. This is favorable in the presence of noncorrectable errors.

B. Symmetries and constraints

When the targeted unitary U has conserved charges because of an underlying symmetry or local constraint, we can incorporate these into the parameterized circuit by

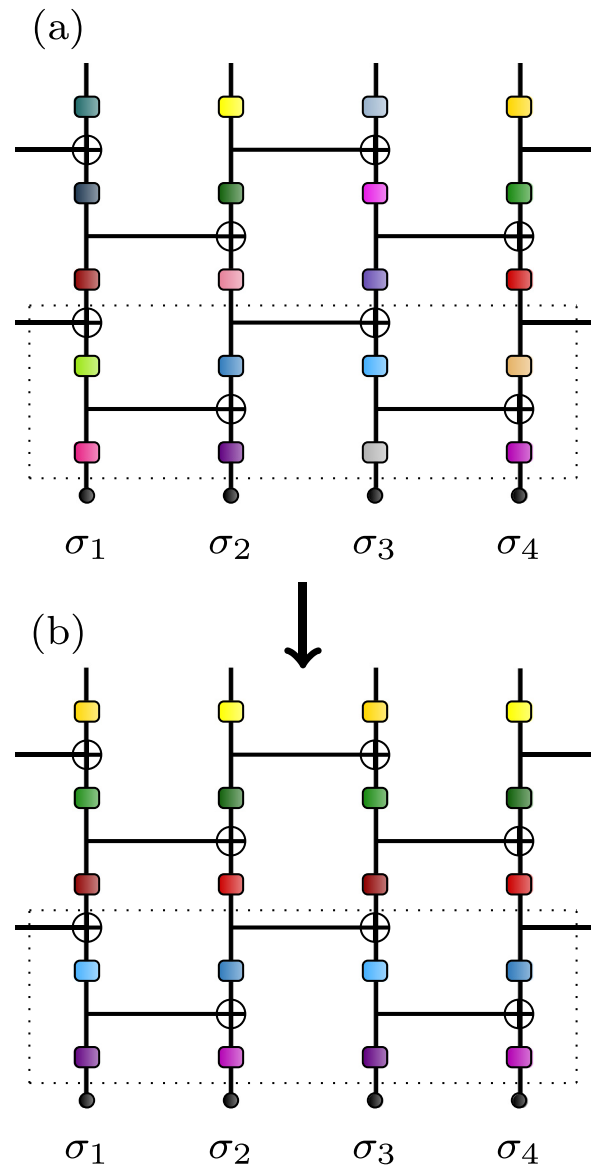


FIG. 1. (a) Template for an unconstrained circuit with $M = 2$ layers. Each layer of CNOT gates is interspersed with a single-qubit gate layer. At the end of the circuit, we add an extra layer of single-qubit unitaries. Equal colored single-qubit gates are identical. The dotted box encloses one layer of the circuit. (b) By adding translation symmetry, we reduce the number of different single-qubit unitaries to two independent single-qubits per half layer. Note that we cannot further reduce the number of different unitaries since the brickwall gate layout manifestly breaks site-inversion symmetry. For the rest of the work, we call this architecture translationally invariant variational brickwall (TIVB) circuits.

constraining the angles of the one-qubit unitaries or modifying the circuit structure.

1. Lattice symmetries

We are often interested in translationally invariant systems. In this case we choose the one-qubit unitaries as in Fig. 1(b), where we illustrate a translation-invariant circuit with $M = 2$ brickwall layers. Here each horizontal layer of one-qubit unitaries contains two independent unitaries (equal colors have equal parameters), giving a circuit with $12M + 6$

real parameters. We note that the translation invariance in a brickwork circuit is necessarily reduced: The minimal unit cell for translations consists of two qubits. This implies that we require two different single-qubit unitaries per half layer.

Apart from reducing the number of parameters to optimize, the implementation of translation symmetry has additional advantageous implications for the optimization: A fundamental property that C should reproduce is the Lieb-Robinson bound of correlation spreading for local many-body quantum systems [72,73]. As a result, if we want to approximate the unitary $U(t)$ for system size L^* , and the corresponding light-cone of correlations has width $W(t) < L^*$, then optimizing C at a smaller size $L > W(t)$ is sufficient.

This allows us to perform a difficult optimization procedure for a small amount of qubits and use the result for larger system sizes. We find that this extrapolation works extremely well, sometimes even when $L < W(t)$. For the time scales we have considered, we find that optimization for $L = 8$ qubits is sufficient, with additional optimization at larger sizes giving only insignificant improvements. For the rest of this work, we have implemented this translation symmetry in our most generic variational circuits, denoting them as TIVB circuits.

2. Local constraints

Apart from lattice symmetries, many systems of interest exhibit symmetric couplings that induce a global symmetry. Here each coupling only connects states with equal global charges. These charges can also be local, as it is the case for gauge theories, or do not need to be associated to a symmetry or gauge freedom at all. Instead, a local constraint is sufficient, as for example in the PXP model [74].

The constraints or conserved quantities imposes a block-diagonal time evolution operator $U(t)$. Each block corresponds to a different charged sector. In the following, we want to encode such symmetries or constraints directly into the circuit architecture, which we will then denote as a blocked circuit architecture. By incorporating such special properties into the circuit architecture, we get a blocked circuit that manifestly has the corresponding block-diagonal structure. This restricts the space of variational circuits to a subspace in which the targeted time-evolution operator U lives, and reduces the parameter count per number of CNOTs. However, it does not guarantee an increased accuracy. In fact, such a restriction can reduce the expressibility, e.g., restricting the maximum possible distance that correlations can travel among the qubits in M layers of CNOTs.

C. Optimization

As a cost function for the optimization of circuits C to faithfully represent a given unitary transformation U , we use the normalized Frobenius distance between U and C ,

$$\epsilon = \frac{\|C - U\|_F^2}{2^{L+1}} = \frac{\sum_{ij} |C_{ij} - U_{ij}|^2}{2^{L+1}} = 1 - \frac{\text{Tr}[C^\dagger U]}{2^L}. \quad (2)$$

The last equality holds because both C and U are unitary. For small system sizes up to $L \approx 16$ this expression can be evaluated exactly, representing U and C as matrices with dimension 2^L . Otherwise, we represent U and C as matrix-product operators (MPOs), where we inevitably lose information at large t due to entanglement truncation [70]. To represent the exact evolution operator U as a MPO, we Trotterize the exact time

evolution using a sufficiently small step size Δt such that the discretization error is smaller than machine precision or truncation errors due to the entanglement barrier.

To determine the parameters of the single-qubit unitaries in the specific architecture, we minimize the distance (2), which we do with first-order gradient descent, using the Adam optimizer [75] (see Algorithm 1). This optimizer uses exponentially decaying averages of the first and second moments of previous parameter updates to modulate the next update.

Our goal is to compress the time-evolution operator $U(t)$ for a sequence of times $t = 1, 2, \dots, 10$ into parameterized circuits with M layers. The minimization of ϵ is performed for each pair (t, M) separately. For a fixed layer count M , we first optimize the circuit for the smallest timestep $t = 1$. We initialize C as the identity circuit and choose a set of Adam hyperparameters $(\lambda, \beta_1, \beta_2, \epsilon_{\text{reg}})$. We perform $\mathcal{O}(10^5)$ iterations of gradient descent to reduce the possibility of getting stuck in local minima. For $t > 1$, we initialize with the same set of hyperparameters and the optimal result of $t - 1$.

We perform this optimization simultaneously for a large grid of hyperparameters. In particular, we choose

$$\begin{aligned} \beta_1, \beta_2 &\in \{0.9, 0.99, 0.999, 0.9999\}, \\ \epsilon_{\text{reg}} &\in \{10^{-2}, 10^{-4}, 10^{-8}, 10^{-12}\}, \\ \lambda &\in \left\{ \begin{array}{l} \{0.5, 0.2, 0.1, 0.01, 0.001, 0.0001\}, \text{ blocked} \\ \{10^{-1}, 10^{-2}, \dots, 10^{-6}\}, \text{ TIVB circuit.} \end{array} \right. \end{aligned} \quad (3)$$

ALGORITHM 1. *The Adam optimizer* [75]. We want to minimize the distance ϵ as a function of the the circuit parameters $\bar{\theta}$. Adam uses first-order gradient descent to update the parameters. With Adam, we calculate exponentially decaying running averages of the first moment m and the second moment v of the gradient, which are then used to update the parameters as $\delta\bar{\theta} \propto m/\sqrt{v}$. This ensures that the updates are steered away from tiny or huge values to improve convergence to the global minimum, by providing a mechanism to escape local minima without excessively large updates (which are likely to be inaccurate since we are using only local gradient information).

Hyperparameters:

λ : Base learning-rate
 β_1 : First moment decay rate
 β_2 : Second moment decay rate
 δ : Regularization
 N_{iters} : Amount of iterations

Initial conditions:

$m_0 \leftarrow 0$ (First moment initially set to zero)
 $v_0 \leftarrow 0$ (Second moment initially set to zero)

for ($i = 0; i < N_{\text{iters}}; i = i + 1$) do

$g_i \leftarrow \nabla_{\bar{\theta}_{i-1}} \epsilon(\bar{\theta}_{i-1})$ (Calculate gradient at current parameters)
 $m_i \leftarrow \beta_1 m_{i-1} + (1 - \beta_1) g_i$ (Update running average of first moment)
 $m_i^* \leftarrow m_i / (1 - \beta_1^i)$ (Bias correction)
 $v_i \leftarrow \beta_2 v_{i-1} + (1 - \beta_2) g_i^2$ (Update running average of second moment)
 $v_i^* \leftarrow v_i / (1 - \beta_2^i)$ (Bias correction)
 $\bar{\theta}_i \leftarrow \bar{\theta}_{i-1} - \lambda m_i^* / (\sqrt{v_i^*} + \delta)$ (Update parameters)

end for

return $\bar{\theta}_i$ (Final circuit parameters)

The difference in the choice of λ accounts for the fact that the blocked circuits have less parameters per number of CNOTs. After this sequential optimization, we determine the best set of hyperparameters at each (t, M) and use the corresponding parameters to initialize another round of optimization at $L = 6$, with the same set of hyperparameters and amount of iterations as before. These results are used to initialize the last round of optimization, now for $L = 8$, with $\mathcal{O}(10^4)$ iterations.

Removing any of the sequential optimization steps leads to results that are orders of magnitude worse in accuracy. Furthermore, we find that using a finer sequential optimization, e.g., starting at $t = 10^{-2}$ and progressing in steps of 10^{-2} , decreases the accuracy of the circuits after every round of optimization.

D. Stacking circuits

The circuits shown in this work are optimized on small system sizes and for short evolution times. As mentioned in Sec. II B 1, the translation symmetry of the circuit template allows to extend the optimized circuits to arbitrarily large system sizes. Furthermore, stacking the circuits [70] allows to reach time scales on a quantum device beyond classical capabilities: Consider a circuit, optimized for a given time t^* . By repeatedly applying this circuit n times on a quantum computer, this allows to implement the time step nt^* on the quantum device. It was analyzed in Ref. [70] that the advantageous scaling in terms of resource cost for the compressed circuit in comparison to Trotter decomposition persists while stacking a circuit multiple times.

III. MODELS AND RESULTS

In this section, we consider specific models and investigate whether shallow depth circuits benefit from incorporating conserved quantities or local constraints into the compressed circuit ansatz C .

A. XXZ model

We start with the following XXZ model on a periodic chain:

$$H = \sum_{(i,j)} S_i^x S_j^x + S_i^y S_j^y + \frac{1}{2} S_i^z S_j^z, \quad (4)$$

which satisfies $[H, Q] = 0$ with $Q = \sum_{i=1}^L S_i^z$. The conservation of the total z spin can be associated with the global $U(1)$ symmetry corresponding to the invariance of (4) under a simultaneous rotation of all spins in the XY plane. This conserved global charge Q induces a splitting of the time-evolution operator $U(t) = \exp(-itH)$ into $\mathcal{O}(L)$ blocks with fixed Q . Specifically, because Q has $L + 1$ possible values, i.e., $Q = -L/2, -L/2 + 1, \dots, L/2$, there is an equal amount of diagonal blocks. The dimension of the largest block, i.e., the zero-magnetization sector, is equal to the binomial coefficient $\binom{L}{L/2}$. The second-largest sector has $\binom{L}{L/2-1}$, such that their ratio is $((L/2)!)^2 / ((L/2 + 1)!(L/2 - 1)!)$. This approaches 1 for large L . More broadly, for the XXZ model there is never a single block that significantly outsizes the rest.




FIG. 2. The blocked architecture that we use to compress the XXZ time-evolution operator. Equal-colored gates are identical, with the architecture consisting of $U(1)$ -symmetric two-qubit unitaries, Eq. (5). The dotted box encloses one layer of circuits. The colors display the manifest translational and bond-inversion symmetries of the XXZ model.

A TIVB circuit architecture from Sec. II A does not respect this conservation by default. We can modify this circuit, however, by replacing the elementary CNOT gate with the $U(1)$ symmetric gate

$$U_{XXZ}(\theta, \phi) = e^{i\theta(\sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y) + i\phi \sigma^z \otimes \sigma^z}. \quad (5)$$

Furthermore, we remove the one-qubit unitaries. This is illustrated in Fig. 2.

The gate (5) can be decomposed into three CNOTs with one-qubit unitaries, see, e.g., Eq. (6) in Ref. [76]. This yields a blocked circuit of \tilde{M} $U(1)$ -symmetric brickwall layers with $2\tilde{M}$ parameters. The number of CNOT gates is the same as for the previous TIVB circuit with $M = 3\tilde{M}$ layers. The lightcone width for a fixed amount of CNOTs is reduced by a factor three, since a minimum of three CNOTs is required to entangle neighboring spins, instead of only one CNOT for the TIVB architecture.

Choosing $\theta = t/2$ and $\phi = t/4$, we obtain the local evolution operator

$$U_l = e^{-it(S_i^x S_j^x + S_i^y S_j^y + \frac{1}{2} S_i^z S_j^z)}, \quad (6)$$

so the blocked circuit is a generalization of first-order Trotter decomposition with variational time steps [56,64,65,68]. We find that compressing into the second-order Trotter layout, i.e., adding an extra half-layer on top of the circuit in Fig. 2, provides no advantage over the first-order layout. To be more concrete, the optimized ϵ decreases smoothly as we add half-layers to the architecture. We checked that in both cases we get the Trotter scaling from Ref. [77] when stacking the optimized circuits.

In order to see the advantages of the blocked architecture, we benchmark it against the TIVB architecture. We compare blocked circuits with $\tilde{M} = 1, 3, 4, 5, 7, 8$ layers against TIVB circuits with $M = 4, 8, 12, 16, 20, 24$ layers. As we discussed before, each $U(1)$ -conserving two-qubit gate can be expressed by a combination of unitaries and at most three CNOT layers. The number of CNOT gates for a blocked circuit with \tilde{M} layers is thus the same as for the previous TIVB circuit with $M = 3\tilde{M}$ layers. Therefore we have a slight mismatch in our comparison by means of circuit depth, since only for $\tilde{M} = 4, 8$ can we exactly match the CNOT count with $M = 12, 24$. The




FIG. 3. Results for the compression of the XXZ time-evolution operator for system size $L = 16$ and times up to $t = 10$ for brickwall circuits for a different number of layers M . We compare TIVB circuits with $M = 4, 8, 12, 16, 20, 24$ with blocked circuits with $\tilde{M} = 1, 3, 4, 5, 7, 8$. (a) The normalized distance ϵ of the optimized circuits as a function of the evolution time t , for the TIVB circuits (solid lines), blocked circuits (dashed lines) and second-order Trotter decomposition with $\tilde{M} + 1/2$ steps (dotted lines). The blocked architecture outperforms the TIVB architecture in ϵ per number of CNOTs for all t . Both parameterized architectures outperform the second-order Trotter circuits. (b) The z -magnetization imbalance P of the propagated Néel state $|\mathbb{Z}_2\rangle$ for the different circuits architectures. As before, the blocked architecture outperforms the other approaches, as long as ϵ is on the order of 1% at most. The observable is calculated on a time grid with spacing $\delta t = 1$ and lines are guides to the eye.

other compared circuits differ at most by one additional brick-wall layer of CNOTs.

We minimize Eq. (2) for the time-evolution operator of (4) for times $t = 1, 2, \dots, 10$. The results are shown in Fig. 3 for system size $L = 16$. We optimized the parameters of the circuit for system sizes $L = 6$ and $L = 8$ and interpolated to larger system sizes by imposing translation symmetry as described above. We compare the error for TIVB (solid lines), blocked (dashed), and second-order Trotter circuits with $\tilde{M} + 1/2$ layers (dotted).

In Fig. 3(a), we show the normalized distance ϵ of the optimized circuits. The blocked circuits outperform the TIVB circuits for all investigated timesteps t . This difference is most prominent at short times and vanishes for large times. For small time steps t , the second-order Trotter circuits perform worse than the optimized circuits. For large t , the

performance is similar to the optimized blocked circuit. Even if this observation is tight only for $M = 12, 24$ due to the equal number of CNOT gates, it extends to other numbers of layers M . However, in this case the comparison is more difficult due to the mismatch in CNOTs.

We are not only interested in the distance between the circuit and the targeted unitary, but also in the resulting error for measurable observables. Moreover, we want to know if optimizing the full distance also yields systematically increasing accuracy on the dynamics of the biggest block. As a check for that, we show in Fig. 3(b) the z -magnetization imbalance

$$P = \sum_j (-1)^j \langle \sigma_j^z \rangle \quad (7)$$

for the Néel-state

$$|\mathbb{Z}_2\rangle = |\uparrow\downarrow\uparrow\dots\rangle \quad (8)$$

propagated to time t by the different circuit architectures. The time evolution of $|\mathbb{Z}_2\rangle$ is constrained to the largest block when the global $U(1)$ constraint is satisfied. We compare with the exact time evolution obtained from exact diagonalization (red line). The blocked circuits reproduce P with the highest accuracy as long as at least $\epsilon \sim \mathcal{O}(10^{-2})$. As an example, for $M = 12$ the blocked circuit is best until $t = 5$, where its error is a few percent. The blocked circuit with $M = 24$ is best until $t = 8$, after which the Trotter circuit is best. For larger distance ϵ , the TIVB and Trotter circuits perform better, with almost all Trotter curves lying around the exact curve until $t = 10$. As long as at most $\epsilon \sim \mathcal{O}(10^{-2})$, the accuracy on P increases systematically, i.e., determining a deeper circuit with lower ϵ generally also improves the accuracy on P .

It is unclear whether this situation will change significantly when performing a simulation on noisy quantum computers. In that case, we can perform classical postprocessing based on the constraint, discarding output states that violate it, in an effort to mitigate errors [29]. Furthermore, it should be noted that the TIVB architecture has more than eighteen times as many parameters as the blocked architecture for an equivalent circuit depth. This makes the backpropagation computations during the optimization process at least eighteen times as expensive.

In Appendix B we show that the relative performance of the blocked and TIVB circuits cannot be improved by constraining the cost function to one of the largest blocks and its off-diagonal.

B. PXP model

The PXP model on a chain of L qubits with periodic boundary conditions is given by [74]

$$H = \sum_{j=1}^L P_{j-1}^z \sigma_j^x P_{j+1}^z, \quad (9)$$

where σ_j^α denote the Pauli operators acting on spin j and $P_j^z = (1 - \sigma_j^z)/2$ is a projector of the j -th qubit into its ground state. The three-body term introduces a local constraint: neighboring excited states are immobile and cannot be created or annihilated. This can be associated with an extensive number of conserved local charges $Q_j = (1 + \sigma_j^z)(1 + \sigma_{j+1}^z)$ that




FIG. 4. The three-qubit circuit $U_3(\theta)$ that implements a PXP-blocked rotation on the middle site. The blue unitaries correspond to $u(\theta, 0, 0)$ and the gray to $u(0, \phi, 0)$. The gray and CNOT gates ensure that $u(4\theta, 0, 0)$ is applied to the inner qubit only when its neighboring qubits are in the ground state.

encode the absence ($Q_j = 0$) or existence ($Q_j = 1$) of a frozen pair on the bond between sites j and $j + 1$. The time-evolution operator splits into $\mathcal{O}(2^L)$ blocks, each block being labeled by a particular charge configuration $\{Q_j\}$. The largest block has $Q_j = 0 \forall j$ and is well known for its weak ergodicity breaking: its Hilbert space has an exponentially small fraction of ETH violating states, as revealed by the revivals in return probability and z -magnetization imbalance that occur when time-evolving Néel product states [74].

Besides the PXP model having exponentially many more blocks than the XXZ model (4), another qualitative difference is the relative size between the largest and second largest blocks: The largest block has dimension $F_{L-1} + F_{L+1}$, where F_j is the Fibonacci sequence $F_0 = 0, F_1 = 1, \dots$ [74]. The second-largest block has a dimension equal to that of the largest block of the $L - 4$ PXP model with OBC, i.e., F_{L-2} . Their ratio approaches $5.8541\dots$ for $L \rightarrow \infty$. The PXP model thus has a dominant block in the Hilbert space, in contrast to the XXZ model where this ratio approaches 1.

1. Circuit architectures

We compare again the compression of the PXP time-evolution operator into TIVB and blocked circuit architectures. The TIVB architecture is the same as in Fig. 1, which breaks the site-inversion symmetry of the PXP model (9).

In order to incorporate the local constraint of the PXP model, we have to fix the relation between different one-qubit angles. To do so, we consider a blocked circuit $U_3(\theta)$ acting on three qubits that satisfies the local constraint. It has a single parameter and is displayed in Fig. 4. In this figure, the blue unitaries denote $u(\theta, 0, 0) = \exp(i\theta\sigma^y)$. The gray

unitary gates correspond to $u(0, \pm\pi/4, 0)$. The interplay of $u(0, \pm\pi/4, 0)$ and the CNOT gates ensures that the inner qubit is only rotated when the neighboring qubits are in the ground state. With the choice $\theta = t/4$, it is an implementation of the local time-evolution operator $U_j = \exp(-itP_{j-1}\sigma_j^x P_{j+1})$.

It is important to note that we can generalize $U_3(\theta)$ to a full circuit which implements the constraint on all even qubits (and then on all odd qubits) simultaneously. This allows us to construct a blocked circuit as shown in Fig. 5, implementing the PXP constraint for the entire system. From this circuit diagram it is immediately clear that information can only travel one site after four brickwall layers of CNOT gates.

We find that we can implement a time-inversion symmetry without losing accuracy, i.e., choosing the angles θ_i symmetric as $\theta_i = \theta_{M/2-i}$. The resulting brickwall circuit with M layers has $M/4$ free parameters.

Similarly to the U(1)-symmetric circuits in Sec. III A, the blocked ansatz is a first-order Trotter circuit with variational time steps, and it always satisfies $M \bmod 4 = 0$ due to its construction. It gives the identity operator when all angles are set to zero.

We compress the PXP time-evolution operator for times $t = 1, 2, \dots, 10$ into TIVB and blocked circuits with $M = 4, 8, \dots, 24$ brickwall layers of CNOTs. The results are shown in Fig. 6, evaluated at $L = 16$. The TIVB circuits are shown as solid lines, the blocked circuits as dashed lines, and the second-order Trotter circuits with $M + 2$ layers as dotted lines.

2. Distance and imbalance

In Fig. 6(a) we show the normalized distance ϵ . The accuracy of the TIVB circuits is at least one order of magnitude better than the blocked circuits for fixed circuit depth. This is in contrast to the case for the XXZ chain. As before, both architectures outperform the Trotter circuits in terms of distance.

We evolve the Néel state (8) up to a time t and measure the imbalance P defined in Eq. (7). In the case of the PXP model, the dynamics gives rise to prominent revivals in the imbalance P , since the initial $|\mathbb{Z}_2\rangle$ state has large overlap with scarred eigenstates of the model [29, 74]. In Fig. 6(b) we show the results for the optimized circuits from Fig. 6(a). The red line displays the exact values P_{exact} obtained from exact diagonalization.

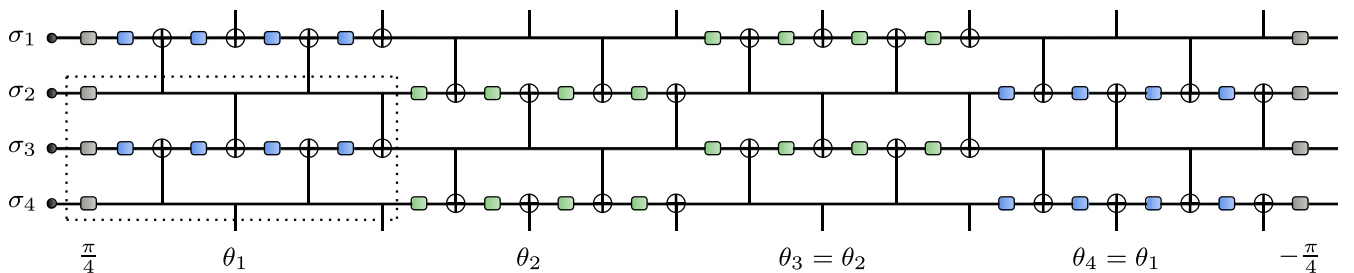


FIG. 5. The blocked architecture with $M = 8$ brickwall layers. The building blocks is the sub-circuit $U_3(\theta)$ defined in Fig. 4 and indicated by a black dashed box. The implemented circuit has a time-inversion symmetry, resulting in $M/4$ parameters for M brickwall layers. The gray unitaries correspond to $u(0, \pm\pi/4, 0)$, and the blue and green unitaries to $u(\theta_j, 0, 0)$, with equal-colored gates being equivalent.




FIG. 6. Results for the compression of the PXP time-evolution operator for system size $L = 16$ and times up to $t = 10$ for brickwall circuits with up to $M = 24$ layers. (a) The normalized distance ϵ as a function of the timestep t , results for the TIVB circuits (solid lines), blocked circuits (dashed) and second-order Trotter circuits with $M + 2$ layers (dotted). The TIVB architecture outperforms the blocked architecture with the same circuit depth. Both architectures outperform the second-order Trotter circuits, although the blocked architecture is only marginally better at large t . (b) The imbalance P (7) as a function of time, starting from the Néel state. The exact result is shown as a solid red line. The same qualitative picture emerges as in panel (a).

Although we focus here on a state that evolves within a single block of the Hilbert space, the same picture as in Fig. 6 carries over: The TIVB circuits outperform the blocked circuits with the same circuit depth. For a fixed amount of gates and a fixed accuracy threshold ϵ , the TIVB circuits can reach times around $2t$ when the blocked and Trotter circuits reach t . Similar to the situation for the XXZ model, the blocked circuits outperform the Trotter circuits when ϵ is of order $\mathcal{O}(10^{-2})$ at most. Afterwards, there is a short region where the Trotter circuit is most accurate, before both the Trotter and constrained circuits are inaccurate. This is in contrast to our results for the XXZ model, where the Trotter circuits remained fairly accurate for the magnetization even when ϵ was highly inaccurate.

3. Detailed error analysis

To understand the origin of the discrepancy between the architectures, we consider the component-wise absolute error




FIG. 7. The component-wise absolute error $|C_{ij} - U_{ij}|^2$ of a TIVB circuit (a) and a blocked circuit (b) with $M = 20$ brickwall layers for $L = 6$ at $t = 5$. The blocks are shown as turquoise boxes. The TIVB circuit connects different subspaces: It hybridizes blocks that have a fixed amount of frozen pairs but at slightly shifted locations. This hybridization is accompanied with a higher approximation accuracy inside of the blocks, where the errors are more than an order of magnitude smaller than for the blocked circuit architecture with the same circuit depth.

$|U_{ij} - C_{ij}|^2$ for $L = 6$ at time $t = 5$ with $M = 20$ in Fig. 7. We compare the TIVB architecture (left), with a total distance $\epsilon \approx 10^{-2}$, with the blocked architecture (right), which is significantly worse with $\epsilon \approx 10^{-1}$. The block structure of the Hilbert space is indicated by turquoise lines. The main contributions of errors in the TIVB circuit stem from violation of the block structure. In particular, we find that the optimization systematically leads to hybridization of blocks that have an equal amount of frozen spin pairs, i.e., it allows for small moves of the frozen spin-up pairs.

As a compensation, this allows a significant increase of expressibility within the diagonal blocks in comparison to




FIG. 8. The out-of-time-ordered correlators $C_{L/2,j}(t)$ (10) as a function of time t and position j , traced over the largest subspace. Results are shown for the exact $L = 16$ PXP time-evolution operator (left panel), the TIVB brickwall circuit approximation (center), and the blocked brickwall circuit approximation (right). Both circuits have a fixed depth of $M = 16$ layers for all times t . The architecture of the blocked circuit constraints its correlation range. As a consequence, it can reproduce $C_{L/2,j}(t)$ only for times $t < 4$. In contrast, the TIVB architecture is only limited by the expressibility due to its given circuit depth.

the blocked architecture: The errors within a diagonal block are more than one order of magnitude lower on average. In case the violation of the constraint is not severe, this still allows the use of post-processing quantum simulation data in real experiments, e.g., discarding measurement output which violates the constraint. This improves the mitigation of errors in a quantum simulation. The downside of the TIVB architecture is the optimization cost: This architecture has 48 times more parameters than the blocked circuits for fixed circuit depth. This increases the cost of backpropagation during the optimization process by at least the same amount.

4. Growth of correlations

To gain more insight, we analyze the results for the out-of-time-ordered correlator (OTOC)

$$C_{L/2,j}(t) = \left\| \left[\sigma_{L/2}^z(t), \sigma_j^z \right] \right\|_F^2 \quad (10)$$

as a function of j and t . The average is taken only over the largest block without frozen spin-up pairs. The results are shown in Fig. 8 for a system with $L = 16$ sites, up to time $t = 8$. In the left panel we show the exact results. In the middle and right panels, we look to reproduce this using the TIVB and blocked architecture with $M = 16$ layers.

Due to the specific construction of U_3 shown in Figs. 4 and 5, the blocked circuit can only reproduce correlations within $M/4$ sites. As a result, the blocked architecture has a reduced maximum-velocity lightcone width, which limits its expressibility. As is shown in Fig. 8, this restricts the reachable times of the blocked circuit: In case of $M = 16$, the maximum reachable time is thus $t = 4$. In contrast, the TIVB architecture has no such limitation and is capable to reproduce the lightcone structure of the OTOC up to $t = 8$.

5. Scaling with system size

Now we will check the scaling of accuracy with system size. We use the $M = 24$ circuits optimized at $t = 1, 2, \dots, 10$




FIG. 9. The system-size extrapolation of the normalized distance ϵ for the TIVB (solid), blocked (dashed) and second-order Trotter circuits with $M = 24$ layers (dotted) at times $t = 1, 2, \dots, 10$. The circuits are optimized at $L = 8$ and then used to evaluate the quantities at larger sizes $L = 10, 12, 14, 16$ by exploiting translational invariance. All optimized circuits extrapolate as good as the Trotter circuits, with errors that show only weak dependence on system size. Thus the optimization of circuits is scalable to larger system sizes without losing accuracy.

with $L = 8$ to calculate their distances to the exact time-evolution operator at larger sizes $L = 10, 12, 14, 16$. In Fig. 9 we present the results. All architectures have a weak dependence on system size L . This shows the scalability of our approach to system sizes larger than investigated in this work.

6. Comparison with Trotter decomposition

In order to compare the optimized blocked circuits with the Trotter circuits, we analyze the optimized circuit parameters θ_l . As mentioned above, the Trotter circuits have the same structure as the blocked circuits, but with fixed angles $\theta_1 = t/(2M)$ and $\theta_l = t/M$ for $l > 1$.

In Fig. 10(a), we compare the six angles θ_l of the blocked circuit with $M = 24$ layers at various times and compare it with the angles of the Trotter circuit. In Fig. 10(b), we compare the sum of the optimized angles $\sum_j \theta_j$ as a blue solid line, with the sum of Trotter angles $t/4$ shown as a dotted line.

The optimized angles are far off the Trotter parameters in most cases. The blocked circuits implement a sequence of forward and backward time evolutions. In contrast their sums $\sum_j \theta_j$ are close as long as $\epsilon \sim \mathcal{O}(10^{-2})$. Higher-order Trotter circuits also contain backward time evolution [43], but we did not find any symmetrical Trotter decomposition that matches our optimized angles. Specifically, the angles of the often-used decompositions from Ref. [43] are significantly smaller.

In Appendix A we consider the scar states of the PXP model in more detail. In Appendix B we consider the distance (2) restricted to various blocks. There we also show that the distance of the largest block is not improved by restricting the cost function to it. We show this for the PXP and XXZ models.




FIG. 10. (a) The six optimized angles θ_i of the blocked architecture with $M = 24$ brickwall layers. The angles of the second-order Trotter circuit with two extra brickwall layers are shown as dotted lines. While the angles of the Trotter circuit θ_i are positive, the optimized blocked circuit relies on the cancellation of multiple forward and backward evolutions. (b) The sum of the optimized angles from panel (a) shown as a solid line, and the sum of the Trotter angles $t/4$ shown as a dotted line. Both angles sum up to the same value of $t/4$, as long as the blocked circuit has an optimized distance below $\epsilon \approx 10^{-1}$. The sequences of forward and backward evolutions observed in panel (a) always accumulate to $t/4$ when it is sufficiently accurate.

C. Quantum link model

As a third model, we consider the massless spin-1/2 quantum link model (QLM)

$$H = - \sum_{j=1}^L \sigma_j^+ s_{j,j+1}^+ \sigma_{j+1}^- + \text{h.c.}, \quad (11)$$

which describes matter spins σ_j that are coupled by gauge spins $s_{j,j+1}$ that live on the links. It is equivalent to a discrete massless Schwinger model with staggered fermions [78]. The time-evolution operator splits into exponentially many blocks, to account for the gauge freedom that is generated by the conserved charges

$$Q_j = (\sigma_j^z + s_{j-1,j}^z - s_{j,j+1}^z + (-1)^j)/2, \quad (12)$$

which act on a matter spin and its neighboring gauge spins. It takes on the values $Q_j = 0, \pm 1, \pm 2$.




FIG. 11. One layer of the TIVB architecture used for the quantum link model (11). Each horizontal layer of one-qubit unitaries contains four distinct one-qubit unitaries.

One of the two largest blocks corresponds to the gauge-invariant sector $Q_j|\psi\rangle = 0 \forall j$. When tracing out the matter spins in this sector, we get a PXP model for the gauge spins [79]. As such, the gauge-invariant block is simply the largest block of a PXP model with L sites. The other largest block corresponds to the charged sector with $Q_j|\psi\rangle = (-1)^j|\psi\rangle$, which is the largest block of a dual PXP model where neighboring gauge spins are frozen when they both point down. The next-largest blocks correspond to the charge configurations Q_j that can be obtained from the gauge-invariant sector by replacing pairs of neighboring $Q_j = 0$ with $Q_j = (-1)^j$, leaving at least one pair of zeros. The ratio of the dimensions of the largest and second-largest blocks is 1.206... Hence there are many similarly sized blocks, unlike the PXP model where the largest block is multiple times larger than the rest. The situation is analogous to the XXZ model, where at finite L the second-largest block is only marginally smaller than the largest block. However, we now have L conserved charges instead of one, such that there are exponentially more blocks (that are therefore exponentially smaller).

To account for the doubled unit cell of the QLM, we compress its time-evolution operator into TIVB brickwall circuits that contain four unique one-qubit unitaries per half-brickwall layer. This is illustrated in Fig. 11 for $M = 1$. This TIVB architecture has $24M + 12$ parameters. Due to the chosen layout of matter and gauge spins, each elementary nearest-neighbor gate always acts on a matter and gauge spin. To compress while respecting the local constraint, we start by decomposing the local time-evolution operator into nearest-neighbor CNOTs and one-qubit unitaries. This yields the circuit $G_3(\theta)$ shown in Fig. 12. The local time-evolution operator is obtained by setting $\theta = t/8$. $G_3(\theta)$ is our building block for a charge-conserving circuit, as is shown in Fig. 13 for $M = 2$ layers. Each three-qubit gate represents a two-local subcircuit $G_3(\theta)$,




FIG. 12. CNOT implementation of the parameterized constrained coupling $G_3(\theta)$ between a gauge field and its neighboring matter sites. It contains a single parameter θ . The gray one-qubit unitaries are $u(\pm\pi/4, 0, 0)$, the green are $u(\pm\theta, -\pi/2, 0)$ and the yellow are $u(\pm\theta, \pi/2, 0)$.




FIG. 13. The translationally invariant blocked architecture with $M = 2$ layers, where every three-qubit “gate” is a $G_3(\theta)$ subcircuit. The circuit has a time-inversion symmetry, as indicated by the colors, such that a blocked circuit with \tilde{M} layers has only \tilde{M} parameters.

and equal-colored gates have equal angles in order to enforce translation and time-inversion symmetry. Consequently, the blocked circuit with \tilde{M} layers has only \tilde{M} parameters, and has an equal amount of CNOTs as a brickwall circuit with $M = 6\tilde{M}$. However, because a blocked circuit cannot be cast into a two-local brickwall circuit, the comparison at fixed gate count is not entirely straightforward.

The performance of the circuit compressions for the QLM model is shown in Fig. 14. The normalized distance ϵ for the different circuit architectures is shown in Fig. 14(a). The situation is analogous to that for the XXZ model from Fig. 3. The blocked circuits outperform the other approaches for fixed gate count when $\epsilon \leq 10^{-2}$, and afterwards they all perform similarly. For the circuit sizes that we have considered, any significant advantage over the Trotter circuits is restricted to $t < 3$.

The curves for $M = 6, 12$ are an exception, since here the second-order Trotter circuits appear to outperform the optimized circuits. However, the Trotter architecture gets an extra half-layer in Fig. 13, such that in Fig. 14 we, e.g., compare a blocked circuit with $\tilde{M} = 1$ with a Trotter circuit with $\tilde{M} = 1 + 1/2$. The difference is three brickwall layers worth of CNOT gates. This difference becomes negligible at large \tilde{M} , but for small \tilde{M} it skews the comparison, as we see for $M = 6, 12$ in Fig. 14.

As a further test, we simulate the dynamics of the gauge-invariant block, again in order to determine whether optimizing the global distance ϵ also yields systematic improvement on a block-restricted quantity. Specifically, because the QLM restricted to the gauge-invariant sector is a PXP model on the gauge spins, the propagation of states in this block can again yield revivals [79]. For this correspondence, the gauge z -spin operator is mapped onto a staggered z -spin operator. Consequently, revivals now occur for uniform up or down states of the gauge spins, as captured by the magnetization

$$\mathcal{P} = \sum_{i=1}^L \langle s_{i,i+1}^z \rangle / L. \quad (13)$$

We consider the initial state

$$|\mathbb{Z}_2^{\text{QLM}}\rangle = |\downarrow_{\sigma_1} \downarrow_{s_1} \uparrow_{\sigma_2} \downarrow_{s_2} \downarrow_{\sigma_3} \downarrow_{s_3} \dots\rangle, \quad (14)$$

which is the product of a staggered state for the matter spins $|\downarrow_{\sigma_1} \uparrow_{\sigma_2} \dots\rangle$ and a uniform down state for the gauge spins




FIG. 14. Results for the compression of the massless spin-1/2 quantum link model time-evolution operator for system size $L = 16$, counting both matter and gauge spins, and times up to $t = 7$ for brickwall circuits with up to $M = 36$ layers. The blocked circuit with \tilde{M} layers has an equal amount of CNOTs as the TIVB circuit with $M = 6\tilde{M}$, but the CNOT architecture is deeper in the blocked case. For convenience we label \tilde{M} as $6M$ in the plots. (a) The normalized distance ϵ of the optimized circuits. We compare TIVB circuits (solid lines), blocked circuits (dashed) and second-order Trotter circuits with $\tilde{M} + 1/2$ layers (dotted). The blocked circuits outperform the others for fixed gate count, most prominently at low t . (b) The string-order parameter \mathcal{P} of the gauge spins, evaluated for the circuits from (a). The exact result is shown as a red line. As for the PXP model, this quantity shows revivals. The relative performance observed in panel (a) carries over for the circuits with ϵ below some threshold of order $\mathcal{O}(10^{-2})$.

$|\downarrow_{s_1} \downarrow_{s_2} \dots\rangle$. The state (14) is part of the gauge-invariant sector, in which it remains as long as it is evolved with a circuit that obeys the local constraint. We show the results in Fig. 14(b), where the exact curve is shown as a red line. As for the other models, the picture from panel (a) carries over for the circuits with distance smaller than $\epsilon \lesssim \mathcal{O}(10^{-2})$. For larger compression errors, all circuits show large deviations from the exact curve.

IV. DISCUSSION

In this work, we have studied the performance of different circuit architectures to simulate time evolution in the presence

of conserved charges. In particular, we put an emphasis on the accuracy that the compressed circuits can achieve when we directly encode constraints into the circuit architecture.

In the presence of global charges and also for lattice gauge theory, we can decrease the optimization cost by more than two orders of magnitude and simultaneously increase the accuracy. At the same time, translation invariance allows for a scalability of our circuits to arbitrary large system sizes while remaining accurate. The only exception arises for systems with local constraints such as the PXP model. In this case, the expressibility of the blocked circuits is reduced to the extent that lifting the constraint gives rise to substantial improvement. It manifests itself as a severely restricted light-cone of correlation spreading (see Fig. 8) and thus imposes fundamental limitations on the accuracy that can be reached with shallow circuits.

Furthermore, we believe that the ability of the TIVB architecture to outperform the blocked architecture on only one of the three considered models is related to the different symmetry block structures of the Hamiltonians. As shown in Fig. 7, lifting the constraint allows for a hybridization of different symmetry blocks in the PXP model. While this induces additional errors violating the constraint, the hybridization of different blocks allows a larger expressibility and thus error reduction within a symmetry block.

Another interesting observation arises for the blocked circuits: In many cases, the constraint puts such heavy constraints on the final architecture that they automatically turn out as a Trotter circuit with variational time steps. However, our optimized circuits did not coincide with any conventional Trotter decompositions, relying on the cancellation of relatively large forward and backward local time evolutions. In the case of the PXP model they achieved an increase in accuracy of more than an order of magnitude in comparison to standard second-order Trotter decomposition with the same gate count, as measured in the normalized distance. It is desirable to get a more rigorous understanding of the origins of this improved accuracy [44].

Another interesting direction for future studies is the use of different cost functions to simulate quantum dynamics. In the case of adaptive Trotter decomposition, it was shown that bounding the errors in the mean and the variance of the energy is sufficient to obtain very precise expectation values for local observables at intermediate to long times [64,65]. This raises the question whether we can use shallower circuits if we only care about reproducing the time evolution of local observables.

ACKNOWLEDGMENTS

We thank Z. Mineev, A. Seif, O. Shtanko, and D. Wang for inspiring discussions. This work was supported by the Deutsche Forschungsgemeinschaft through the cluster of excellence ML4Q (EXC2004, Project-id No. 390534769). We also acknowledge support from the QuantERA II Programme, which has received funding from the European Union's Horizon 2020 research innovation programme (GA No. 101017733), and from the DFG through the project DQUANT (Project-id No. 499347025).

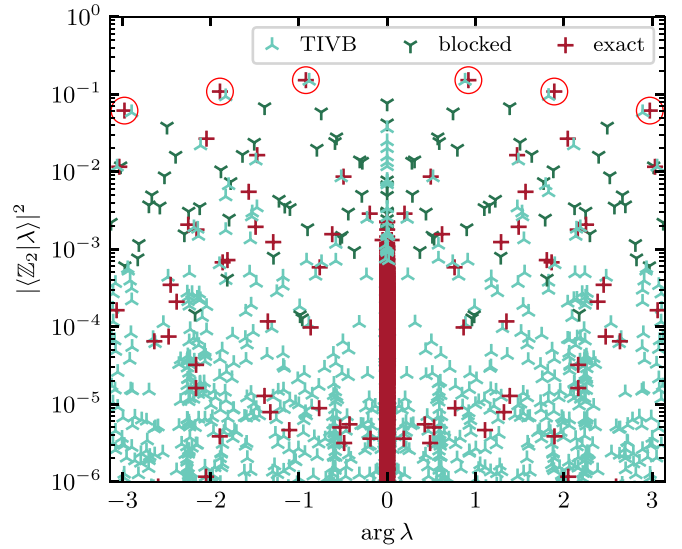


FIG. 15. The absolute overlap of the eigenstates $|\lambda\rangle$ with the Néel state $|\mathbb{Z}_2\rangle$, versus the argument of the complex eigenvalues λ , for the exact PXP time-evolution operator and compressed circuits with $M = 16$ layers at time $t = 4$ for system size $L = 14$. Note that these are the eigenstates of the entire unitary, not only the first block. Even though the TIVB circuit simulates the time evolution most accurately, the breaking of the constraint leads to multiple additional non-degenerate eigenstates with small overlap. The six exact eigenstates with the largest overlap are encircled in red. The TIVB architecture can reproduce them to high precision.

APPENDIX A: QUANTUM MANY-BODY SCARS IN THE PXP MODEL

The PXP model is known to host so-called quantum many-body scars [6,74]. In this section we probe the accuracy with which these states are reproduced by the different architectures. To do so, we show in Fig. 15 the absolute overlap between the eigenstates $|\lambda\rangle$ of the circuits with eigenvalues λ and the \mathbb{Z}_2 Néel state. The calculations were performed at $L = 14$ for $t = 4$. The overlap with eigenstates of the blocked circuit with $M = 16$ are shown in dark blue, and of the TIVB circuit with $M = 16$ in cyan. As a reference we show the overlap for the exact time evolution with red crosses. The TIVB circuit reproduces the uppermost scarred eigenstates to high accuracy, while the blocked circuit is visibly off. This is in line with their distance, namely $\epsilon \sim \mathcal{O}(10^{-3})$ for the TIVB circuit and $\epsilon \sim \mathcal{O}(10^{-2})$ for the blocked circuit, as can be seen in Fig. 6(a). On account of breaking the constraint, there are many spurious states for the TIVB circuit. It is shown in Fig. 6(b) that these do not significantly affect the imbalance revivals.

APPENDIX B: RESTRICTED DISTANCES

To better understand how the TIVB circuits can outperform the blocked circuits in compressing the PXP time-evolution operator, we plot in Fig. 16 the Frobenius distance (2) restricted to various parts of the unitary. We do this for a TIVB and a blocked circuit, both with $M = 24$ brickwall layers, shown as the solid and dashed light blue lines. The second-order Trotter circuit with $M = 26$ is shown as a dashed green




FIG. 16. The normalized Frobenius distance for various parts of the PXP time-evolution operator with size $L = 12$. We consider a TIVB and a blocked circuit with $M = 24$, shown as solid and dashed lines, respectively. The second-order Trotter circuit with $M = 26$ is shown as a dashed green line. The light blue curves correspond to circuits that were optimized with the total distance ϵ as cost function, which for the dark blue curves was restricted to the largest diagonal block and its off-diagonal rectangles. In (a) we show the distance of the largest diagonal block ϵ_{d1} , in (b) the distance of the remaining diagonal blocks ϵ_{dr} , in (c) the distance of the off-diagonal rectangles adjacent to the largest diagonal block ϵ_{o1} , and in (d) the distance of the remaining off-diagonal rectangles ϵ_{or} . The block cost function provides a marginal increase in accuracy for the first diagonal block, which is outweighed by the exponential complexity of determining the block. Interestingly, for $t > 1$ the restricted cost function achieves equally high accuracy outside of the first block.

line. In (a) we show the distance of the largest diagonal block ϵ_{d1} , in (b) the distance of the remaining diagonal blocks ϵ_{dr} , in (c) the distance of the off-diagonal rectangles adjacent to the largest block ϵ_{o1} , and in panel (d) the distance of the remaining distances. To normalize the Frobenius norm we divide it by $2\sqrt{N_e}$, with N_e being the amount of matrix entries in the average. This reduces to the normalized distance (2) when the average is over one or multiple diagonal blocks.

We see that all distances are roughly equal, indicating that optimizing the total ϵ leads to systematic improvement of all sectors. The TIVB circuit has almost two orders of magnitude higher accuracy on the diagonal blocks than the blocked circuit. We know from Sec. III B 2 that this gain carries over to the simulation accuracy of the imbalance revivals, even at times when the largest diagonal block couples to the other diagonal blocks with $\epsilon_{o1} \sim \epsilon_{or} \sim \mathcal{O}(10^{-2})$.

Given these results, it is sensible to ask whether we can increase the accuracy of the largest block by restricting the cost function to the largest block and the adjacent off-diagonal rectangles. The results of this optimization are shown as the dark blue lines in Fig. 16. First we consider the TIVB




FIG. 17. The normalized Frobenius distance for various parts of the XXZ time-evolution operator with size $L = 12$. We consider a TIVB circuit with $M = 24$ and a blocked circuit with $\tilde{M} = 8$, shown as solid and dashed lines, respectively. The second-order Trotter circuit with $\tilde{M} = 8 + 1/2$ is shown as a dashed green line. The light blue curves correspond to circuits that were optimized with the total distance ϵ as cost function, which for the dark blue curves was restricted to the largest diagonal block and its off-diagonal rectangles. In panel (a) we show the distance of the largest diagonal block ϵ_{d1} , in (b) the distance of the remaining diagonal blocks ϵ_{dr} , in (c) the distance of the off-diagonal rectangles adjacent to the largest diagonal block ϵ_{o1} , and in (d) the distance of the remaining off-diagonal rectangles ϵ_{or} . The block cost function does not improve the accuracy of the first diagonal block, and it no longer automatically yields high accuracy on the smaller diagonal blocks.

circuit. There is a marginal improvement in the distance of the largest diagonal block ϵ_{d1} and the distance of its off-diagonal rectangles ϵ_{o1} . Furthermore, the distance of the other diagonal blocks ϵ_{dr} and the distance of its off-diagonals ϵ_{or} are automatically reproduced to almost the same extent as when using the full cost function. Only at $t = 1$ the accuracy of the distances ϵ_{dr} and ϵ_{or} is diminished, which does not affect the distances ϵ_{d1} and ϵ_{o1} . We have checked that the imbalance revivals are reproduced equally accurate with both cost functions. For the blocked circuit there is no gain at all, with the differences stemming from the optimization procedure. In summary, there seems to be little merit to a restricted cost function, especially because determining the blocks is exponentially complex in L and requires *a priori* knowledge of the particular structure of U .

In Fig. 17 we repeat this analysis for the XXZ model, now for a TIVB circuit with $M = 24$ brickwall layers of CNOT gates and a blocked circuit with $\tilde{M} = 8$ brickwall layers of $U(1)$ -symmetric gates. The second-order Trotter circuit has $\tilde{M} = 8 + 1/2$ brickwall layers of $U(1)$ -symmetric gates, with the $+1/2$ indicating an extra half brickwall layer. Clearly, the restricted cost function provides no benefit.

- [1] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Comput.* **26**, 1484 (1997).
- [2] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [3] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini *et al.*, Quantum optimization of maximum independent set using rydberg atom arrays, *Science* **376**, 1209 (2022).
- [4] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, *Science* **376**, 1182 (2022).
- [5] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [6] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature (London)* **551**, 579 (2017).
- [7] C. Gross and I. Bloch, Quantum simulations with ultracold atoms in optical lattices, *Science* **357**, 995 (2017).
- [8] Google AI Quantum and Collaborators, F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. J. Huggins, L. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, O. Martin, J. M. Martinis, J. R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, H. Neven, M. Y. Niu, T. E. O'Brien, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, D. Strain, K. J. Sung, M. Szalay, T. Y. Takeshita, A. Vainsencher, T. White, N. Wiebe, Z. J. Yao, P. Yeh, and A. Zalcman, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [9] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard *et al.*, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, *Nature (London)* **595**, 233 (2021).
- [10] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature (London)* **595**, 227 (2021).
- [11] P. Frey and S. Rachel, Realization of a discrete time crystal on 57 qubits of a quantum computer, *Sci. Adv.* **8**, eabm7652 (2022).
- [12] K. J. Satzinger, Y.-J. Liu, A. Smith, C. Knapp, M. Newman, C. Jones, Z. Chen, C. Quintana, X. Mi, A. Dunsworth, C. Gidney, I. Aleiner, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, R. Barends, J. Basso, A. Bengtsson, A. Bilmes, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, D. Eppens, C. Erickson, L. Faoro, E. Farhi, A. G. Fowler, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, S. Hong, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, T. Khattar, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, A. Locharla, E. Lucero, O. Martin, J. R. McClean, M. McEwen, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, T. E. O'Brien, A. Opremcak, B. Pató, A. Petukhov, N. C. Rubin, D. Sank, V. Shvarts, D. Strain, M. Szalay, B. Villalonga, T. C. White, Z. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, A. Megrant, Y. Chen, J. Kelly, V. Smelyanskiy, A. Kitaev, M. Knap, F. Pollmann, and P. Roushan, Realizing topologically ordered states on a quantum processor, *Science* **374**, 1237 (2021).
- [13] A. Smith, M. S. Kim, F. Pollmann, and J. Knolle, Simulating quantum many-body dynamics on a current digital quantum computer, *npj Quantum Inf.* **5**, 106 (2019).
- [14] B. Yang, H. Sun, R. Ott, H.-Y. Wang, T. V. Zache, J. C. Halimeh, Z.-S. Yuan, P. Hauke, and J.-W. Pan, Observation of gauge invariance in a 71-site Bose-Hubbard quantum simulator, *Nature (London)* **587**, 392 (2020).
- [15] M. Meth, J. F. Haase, J. Zhang, C. Edmunds, L. Postler, A. Steiner, A. J. Jena, L. Dellantonio, R. Blatt, P. Zoller, T. Monz, P. Schindler, C. Muschik, and M. Ringbauer, Simulating 2D lattice gauge theories on a qudit quantum computer, [arXiv:2310.12110](https://arxiv.org/abs/2310.12110).
- [16] Google AI Quantum and Collaborators, Measurement-induced entanglement and teleportation on a noisy quantum processor, *Nature (London)* **622**, 481 (2023).
- [17] C. Noel, P. Niroula, D. Zhu, A. Risinger, L. Egan, D. Biswas, M. Cetina, A. V. Gorshkov, M. J. Gullans, D. A. Huse *et al.*, Measurement-induced quantum phases realized in a trapped-ion quantum computer, *Nat. Phys.* **18**, 760 (2022).
- [18] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner *et al.*, A quantum processor based on coherent transport of entangled atom arrays, *Nature (London)* **604**, 451 (2022).
- [19] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, J. G. Bohnet, N. C. Brown, N. Q. Burdick, W. C. Burton, S. L. Campbell, J. P. C. III au2, C. Carron, J. Chambers, J. W. Chan, Y. H. Chen, A. Chernoguzov, E. Chertkov, J. Colina, J. P. Curtis, R. Daniel, M. DeCross, D. Deen, C. Delaney, J. M. Dreiling, C. T. Ertsgaard, J. Esposito, B. Estey, M. Fabrikant, C. Figgatt, C. Foltz, M. Foss-Feig, D. Francois, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, E. Glynn, A. Hall, A. M. Hankin, A. Hansen, D. Hayes, B. Higashi, I. M. Hoffman, B. Horning, J. J. Hout, R. Jacobs, J. Johansen, L. Jones, J. Karcz, T. Klein, P. Lauria, P. Lee, D. Liefer, C. Lytle, S. T. Lu, D. Lucchetti, A. Malm, M. Matheny, B. Mathewson, K. Mayer, D. B. Miller, M. Mills, B. Neyenhuis, L. Nugent, S. Olson, J. Parks, G. N. Price, Z. Price, M. Pugh, A. Ransford, A. P. Reed, C. Roman, M. Rowe, C. Ryan-Anderson, S. Sanders, J. Sedlacek, P. Shevchuk, P. Siegfried, T. Skripka, B. Spaun, R. T. Sprenkle, R. P. Stutz, M. Swallows, R. I. Tobey, A. Tran, T. Tran, E. Vogt, C. Volin, J. Walker, A. M. Zolot, and J. M. Pino, A race track trapped-ion quantum processor, *Phys. Rev. X* **13**, 041052 (2023).
- [20] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).

- [21] Y. Zhou, E. M. Stoudenmire, and X. Waintal, What limits the simulation of quantum computers? *Phys. Rev. X* **10**, 041038 (2020).
- [22] J. Tindall, M. Fishman, M. Stoudenmire, and D. Sels, Efficient tensor network simulation of ibm's eagle kicked ising experiment, *PRX Quantum* **5**, 010308 (2024).
- [23] S. Anand, K. Temme, A. Kandala, and M. Zaletel, Classical benchmarking of zero noise extrapolation beyond the exactly-verifiable regime, [arXiv:2306.17839](https://arxiv.org/abs/2306.17839).
- [24] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme *et al.*, Evidence for the utility of quantum computing before fault tolerance, *Nature (London)* **618**, 500 (2023).
- [25] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Müller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, and K. Temme, Quantum optimization using variational algorithms on near-term quantum devices, *Quantum Sci. Technol.* **3**, 030503 (2018).
- [26] Y. Kim, C. J. Wood, T. J. Yoder, S. T. Merkel, J. M. Gambetta, K. Temme, and A. Kandala, Scalable error mitigation for noisy quantum circuits produces competitive expectation values, *Nat. Phys.* **19**, 752 (2023).
- [27] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [28] Y. Guo and S. Yang, Quantum error mitigation via matrix product operators, *PRX Quantum* **3**, 040313 (2022).
- [29] I.-C. Chen, B. Burdick, Y. Yao, P. P. Orth, and T. Iadecola, Error-mitigated simulation of quantum many-body scars on quantum computers with pulse-level control, *Phys. Rev. Res.* **4**, 043027 (2022).
- [30] S. Filippov, B. Sokolov, M. A. C. Rossi, J. Malmi, E.-M. Borrelli, D. Cavalcanti, S. Maniscalco, and G. García-Pérez, Matrix product channel: Variationally optimized quantum tensor network to mitigate noise and reduce errors for the variational quantum eigensolver, [arXiv:2212.10225](https://arxiv.org/abs/2212.10225).
- [31] J. Vovrosh, K. E. Khosla, S. Greenaway, C. Self, M. S. Kim, and J. Knolle, Simple mitigation of global depolarizing errors in quantum simulations, *Phys. Rev. E* **104**, 035309 (2021).
- [32] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, *Phys. Rev. X* **8**, 031027 (2018).
- [33] E. van den Berg, Z. K. Mineev, and K. Temme, Model-free readout-error mitigation for quantum expectation values, *Phys. Rev. A* **105**, 032620 (2022).
- [34] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, *Rev. Mod. Phys.* **95**, 045005 (2023).
- [35] M. S. J. Tepaske and D. J. Luitz, Compressed quantum error mitigation, *Phys. Rev. B* **107**, L201114 (2023).
- [36] Y. Li and S. C. Benjamin, Efficient variational quantum simulator incorporating active error minimization, *Phys. Rev. X* **7**, 021050 (2017).
- [37] E. van den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
- [38] S. Endo, Q. Zhao, Y. Li, S. Benjamin, and X. Yuan, Mitigating algorithmic errors in a hamiltonian simulation, *Phys. Rev. A* **99**, 012334 (2019).
- [39] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature (London)* **567**, 491 (2019).
- [40] S. Lie, Theorie der transformationsgruppen i, *Math. Ann.* **16**, 441 (1880).
- [41] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [42] M. Suzuki, Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, *Commun. Math. Phys.* **51**, 183 (1976).
- [43] M. Suzuki, Fractal decomposition of exponential operators with applications to many-body theories and monte carlo simulations, *Phys. Lett. A* **146**, 319 (1990).
- [44] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [45] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proc. Natl. Acad. Sci.* **115**, 9456 (2018).
- [46] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, Alá n Aspuru-Guzik, H. Neven, and R. Babbush, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, *Quantum* **4**, 296 (2020).
- [47] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [48] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [49] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [50] R. Haghshenas, J. Gray, A. C. Potter, and G. K.-L. Chan, Variational power of quantum circuit tensor networks, *Phys. Rev. X* **12**, 011047 (2022).
- [51] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, Variational fast forwarding for quantum simulation beyond the coherence time, *npj Quantum Inf.* **6**, 82 (2020).
- [52] K. Mizuta, Y. O. Nakagawa, K. Mitarai, and K. Fujii, Local variational quantum compilation of large-scale Hamiltonian dynamics, *PRX Quantum* **3**, 040302 (2022).
- [53] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, Quantum-assisted quantum compiling, *Quantum* **3**, 140 (2019).
- [54] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, *New J. Phys.* **22**, 043006 (2020).
- [55] B. Commeau, M. Cerezo, Z. Holmes, L. Cincio, P. J. Coles, and A. Sornborger, Variational hamiltonian diagonalization for dynamical quantum simulation, [arXiv:2009.02559](https://arxiv.org/abs/2009.02559).
- [56] N. F. Berthussen, T. V. Trevisan, T. Iadecola, and P. P. Orth, "Quantum dynamics simulations beyond the coherence time on nisq hardware by variational trotter compression, *Phys. Rev. Res.* **4**, 023097 (2022).
- [57] S.-H. Lin, R. Dilip, A. G. Green, A. Smith, and F. Pollmann, Real- and imaginary-time evolution with compressed quantum circuits, *PRX Quantum* **2**, 010342 (2021).

- [58] M. Benedetti, M. Fiorentini, and M. Lubasch, Hardware-efficient variational quantum algorithms for time evolution, *Phys. Rev. Res.* **3**, 033083 (2021).
- [59] S. Barison, F. Vicentini, and G. Carleo, An efficient quantum algorithm for the time evolution of parameterized circuits, *Quantum* **5**, 512 (2021).
- [60] T. Jones and S. C. Benjamin, Robust quantum compilation and circuit optimisation via energy minimisation, *Quantum* **6**, 628 (2022).
- [61] K. Heya, Y. Suzuki, Y. Nakamura, and K. Fujii, Variational quantum gate optimization, [arXiv:1810.12745](https://arxiv.org/abs/1810.12745).
- [62] J. Gibbs, K. Gili, Z. Holmes, B. Commeau, A. Arrasmith, L. Cincio, P. J. Coles, and A. Sornborger, Long-time simulations for fixed input states on quantum hardware, *npj Quantum Inf.* **8**, 135 (2022).
- [63] K. Heya, K. M. Nakanishi, K. Mitarai, Z. Yan, K. Zuo, Y. Suzuki, T. Sugiyama, S. Tamate, Y. Tabuchi, K. Fujii, and Y. Nakamura, Subspace variational quantum simulator, *Phys. Rev. Res.* **5**, 023078 (2023).
- [64] H. Zhao, M. Bukov, M. Heyl, and R. Moessner, Making trotterization adaptive and energy-self-correcting for NISQ devices and beyond, *PRX Quantum* **4**, 030319 (2023).
- [65] H. Zhao, M. Bukov, M. Heyl, and R. Moessner, Adaptive trotterization for time-dependent Hamiltonian quantum dynamics using instantaneous conservation laws, [arXiv:2307.10327](https://arxiv.org/abs/2307.10327).
- [66] K. Wada, R. Raymond, Y.-Y. Ohnishi, E. Kaminishi, M. Sugawara, N. Yamamoto, and H. C. Watanabe, Simulating time evolution with fully optimized single-qubit gates on parametrized quantum circuits, *Phys. Rev. A* **105**, 062421 (2022).
- [67] C. Mc Keever and M. Lubasch, Classically optimized hamiltonian simulation, *Phys. Rev. Res.* **5**, 023146 (2023).
- [68] R. Mansuroglu, T. Eckstein, L. Nützel, S. A. Wilkinson, and M. J. Hartmann, Variational Hamiltonian simulation for translational invariant systems via classical pre-processing, *Quantum Sci. Technol.* **8**, 025006 (2023).
- [69] R. Mansuroglu, F. Fischer, and M. J. Hartmann, Problem-specific classical optimization of hamiltonian simulation, *Phys. Rev. Res.* **5**, 043035 (2023).
- [70] M. S. J. Tepaske, D. Hahn, and D. J. Luitz, Optimal compression of quantum many-body time evolution operators into brickwall circuits, *SciPost Phys.* **14**, 073 (2023).
- [71] 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).
- [72] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, *Commun. Math. Phys.* **28**, 251 (1972).
- [73] S. Bravyi, M. B. Hastings, and F. Verstraete, Lieb-robinson bounds and the generation of correlations and topological quantum order, *Phys. Rev. Lett.* **97**, 050401 (2006).
- [74] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, Quantum scarred eigenstates in a rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations, *Phys. Rev. B* **98**, 155134 (2018).
- [75] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [76] G. Vidal and C. M. Dawson, Universal quantum circuit for two-qubit transformations with three controlled-not gates, *Phys. Rev. A* **69**, 010301(R) (2004).
- [77] D. Layden, First-order trotter error from a second-order perspective, *Phys. Rev. Lett.* **128**, 210501 (2022).
- [78] P. Hauke, D. Marcos, M. Dalmonte, and P. Zoller, Quantum simulation of a lattice Schwinger model in a chain of trapped ions, *Phys. Rev. X* **3**, 041018 (2013).
- [79] F. M. Surace, P. P. Mazza, G. Giudici, A. Lerose, A. Gambassi, and M. Dalmonte, Lattice gauge theories and string dynamics in Rydberg atom quantum simulators, *Phys. Rev. X* **10**, 021041 (2020).