Letter

Simulating time evolution on distributed quantum computers

Finn Lasse Buessen [1,2] Dvira Segal,[2,3] and Ilia Khait [1]
[1]*Entangled Networks Ltd., Toronto, Ontario M4R 2E4, Canada*
[2]*Department of Physics, University of Toronto, Toronto, Ontario M5S 1A7, Canada*
[3]*Department of Chemistry and Centre for Quantum Information and Quantum Control, University of Toronto, Toronto, Ontario M5S 3H6, Canada*

(Received 1 September 2022; accepted 9 February 2023; published 4 April 2023)

We study a variation of the Trotter-Suzuki decomposition, in which a Hamiltonian exponential is approximated by an ordered product of two-qubit operator exponentials such that the Trotter step size is enhanced for a small number of terms. Such decomposition directly reflects hardware constraints of distributed quantum computers, where operations on monolithic quantum devices are fast compared to entanglement distribution across separate nodes using interconnects. We simulate nonequilibrium dynamics of transverse-field Ising and XY spin chain models and investigate the impact of locally increased Trotter step sizes that are associated with an increasingly sparse use of the quantum interconnect. We find that the overall quality of the approximation depends smoothly on the local sparsity and that the proliferation of local errors is slow. As a consequence, we show that fast local operations on monolithic devices can be leveraged to obtain an overall improved result fidelity even on distributed quantum computers where the use of interconnects is costly.

DOI: 10.1103/PhysRevResearch.5.L022003

Introduction. The availability and rapid evolution of general purpose quantum computing hardware has lead to repeated claims of quantum supremacy [1–4]. It gives new impetus to Feynman's idea of using quantum computers as platforms to efficiently simulate the dynamics of quantum matter [5]. With quantum computers capable enough, the list of potential applications in physics and beyond is long; ranging from condensed matter physics [6–9] to the simulation of general quantum field theories [10,11] and from nuclear physics [12] to quantum chemistry [13–16] and drug discovery [17].

While quantum supremacy has indeed been claimed for select synthetic problems, current noisy intermediate-scale quantum (NISQ) devices do not yet exceed the performance of classical computers for purposeful algorithms like quantum simulation in the spirit of Feynman [18,19]. Applications on NISQ devices [20] are typically hindered by finite coherence times and insufficient gate precision, as well as by the overall small number of available qubits. A crucial step in surpassing the size limitation of current-generation quantum computers will be the transition to distributed quantum computers [21,22]. Similar to the limitations of monolithic NISQ devices, however, near-term interconnect hardware that is required for the facilitation of quantum gates across distributed quantum processing units (QPUs, cf. Fig. 1) is also facing challenges. Most notably, current interconnect hardware generates entanglement between remote qubits at a rate of approximately 182 Hz, which is over an order of magnitude slower than single-QPU gate operations [23]. Yet, time is a critical resource on NISQ architectures: Coherence times range from $\mathcal{O}(100\,\mu\mathrm{s})$ on superconducting qubit platforms to $\mathcal{O}(1\,\mathrm{s})$ on trapped ion systems [24,25]. Even on the latter systems, current interconnect rates appear slow, taking into account that each remote two-qubit unitary operation may consume up to three entangled qubit pairs [26]. In addition, further distillation of remote entanglement may be necessary to reach the desired fidelity levels [27]. It is therefore of great importance in quantum algorithm design to limit the number of required interconnect uses.

In this paper, we study the Trotter-Suzuki decomposition [28–30] from the perspective of potential implementations on distributed quantum hardware. The Trotter-Suzuki decomposition is a widely-used technique for approximating many-body time evolution operators by a product of two-qubit operators that can be implemented on NISQ devices. We propose a variation of the decomposition, in which the resulting two-qubit operators are no longer treated on equal footing: To limit the number of intercore operations, two-qubit gates that are facilitated by a quantum interconnect are treated at a coarser level of approximation, while a more accurate approximation, i.e., a shorter Trotter step size, is maintained for operations that are local within a single compute node. Our variation of the decomposition, which is illustrated in Fig. 1, allows us to systematically control the sparsity of interconnect usage.

For select models of nonequilibrium quantum magnetism, we demonstrate that a significant enhancement of accuracy is achieved over the traditional Trotter-Suzuki decomposition when the Trotter step size is limited by the interconnect rate. Finally, we also benchmark a scenario in which the sparsity of interconnect usage is randomized. This emulates the operation of an interconnect at its latency limit, when due to the

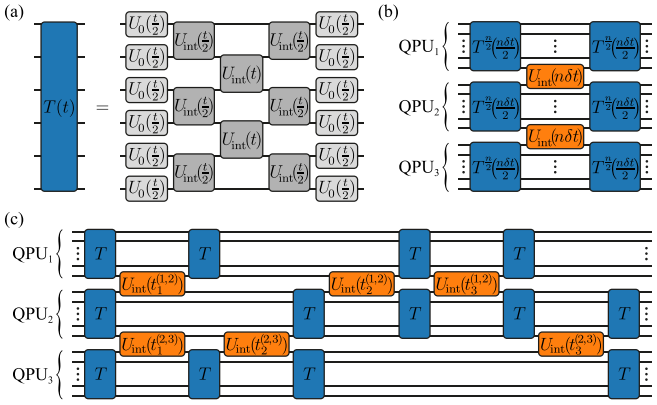

FIG. 1. Circuit models for uniform (conventional) and sparse Trotterization schemes on a distributed quantum computer. (a) Decomposition of a uniform Trotter step T into elementary operations U_0 and $U_{\rm int}$. (b) A single Trotter step on a distributed QPU with $k = 3$ compute nodes. The sparsity parameter n denotes the number of steps in the node-local Trotter operator (blue) that are performed before a single intercore operation (orange). (c) Stochastic sparse Trotterization on $k = 3$ distributed compute nodes with step sizes drawn from a random distribution (arguments of node-local Trotter steps suppressed).

nondeterministic nature of its dead-time after each use it cannot be guaranteed that the link is immediately available [31].

Models. We consider two distinct models of quantum magnetism that are defined on one-dimensional finite chains of L spin-1/2 operators $(\sigma_i^x, \sigma_i^y, \sigma_i^z)$, where i denotes the position on the chain. The first model is the XY model, governed by the Hamiltonian $H_{\rm XY} = -J\sum_{\langle i,j\rangle}(\sigma_i^x\sigma_j^x + \sigma_i^y\sigma_j^y)$, where $\langle,\rangle$ denotes nearest neighbor pairs on the chain and we fix the interaction energy scale $J = 1$. For the time evolution under this model, we assume that the system is initially prepared in a domain wall state $|\uparrow\ldots\uparrow\downarrow\ldots\downarrow\rangle$ in the eigenbasis of σ^z. The second model is the transverse-field Ising (TFI) model with Hamiltonian $H_{\rm TFI} = -J\sum_{\langle i,j\rangle}\sigma_i^z\sigma_j^z + h\sum_i\sigma_i^x$. We consider the case $h = 0.5$, which we refer to as a *slow* quench, and $h = 2.0$, which we call a *fast* quench [32,33]. In both cases for the TFI model, we assume the spin chain to initially be uniformly ordered with all spins in the $|\downarrow\rangle$ state. For quantum simulations, the spin chain is divided into k sections of equal size and the qubits are mapped onto k quantum compute nodes that are interconnected linearly, cf. Fig. 1.

Sparse Trotterization. The model Hamiltonians outlined in the previous section can be summarized in the generalized notation $H = \sum_i H_0^{(i)} + \sum_{\langle i,j\rangle} H_{\rm int}^{(i,j)}$, where $H_0^{(i)}$ denotes a local term on site i and $H_{\rm int}^{(i,j)}$ denotes an interaction term between neighboring sites i and j; the different terms generally do not commute. The time evolution operator $U(t) = e^{-iHt}$ can be approximated by a sequence $T^N(t)$ of one- and two-qubit operations $U_0^{(i)}(t) = e^{-iH_0^{(i)}t}$ and $U_{\rm int}^{(i,j)}(t) = e^{-iH_{\rm int}^{(i,j)}t}$ as

$$T^N(t) = \left(T_0\left(\tfrac{\delta t}{2}\right)T_{\rm even}\left(\tfrac{\delta t}{2}\right)T_{\rm odd}(\delta t)T_{\rm even}\left(\tfrac{\delta t}{2}\right)T_0\left(\tfrac{\delta t}{2}\right)\right)^N, \quad (1)$$

where $\delta t = \frac{t}{N}$, $T_0(t) = \prod_i U_0^{(i)}(t)$, $T_{\rm even}(t) = \prod_{\langle i,j\rangle_{\rm even}} U_{\rm int}^{(i,j)}(t)$ with the product running over all even pairs of nearest-neighbor sites (nearest-neighbor pairs on the spin chain are alternatingly labeled as even and odd) and $T_{\rm odd}(t)$ defined analogously. This approximation is known as the (N-step) second-order Trotter-Suzuki decomposition, depicted in Fig. 1(a) for a single step ($N = 1$) [28,29]. The approximation error scales $\sim N\delta t^3$ [34]. For the remainder of the paper, we shall also refer to this approximation as *uniform* Trotterization.

Motivated by the expected constraints for near-term distributed quantum computing architectures—most importantly, the slow rate of entanglement generation in quantum interconnects—we define *sparse* Trotterization as follows. We assume a distributed quantum computer to consist of $k > 1$ compute nodes that are interconnected linearly, cf. Fig. 1. Interconnects may only be used at a fraction $1/n$ of the speed at which each individual compute node operates; we refer to n as the *sparsity*. The sparse Trotterization of the time evolution operator $U(t)$ is defined as

$$T_{\rm sparse}^{N,n}(t) = \left(\prod_{\kappa=1}^{k} T^{\frac{n}{2}}\Big|_\kappa\left(\frac{n\delta t}{2}\right)\prod_{\langle i,j\rangle_{\kappa\neq\kappa'}} U_{\rm int}^{(i,j)}(n\delta t)\prod_{\kappa=1}^{k} T^{\frac{n}{2}}\Big|_\kappa\left(\frac{n\delta t}{2}\right)\right)^{\frac{N}{n}}, \quad (2)$$

where $T|_\kappa(t)$ denotes the usual (uniformly Trotterized) time evolution within a single compute node κ and $\langle i,j\rangle_{\kappa\neq\kappa'}$ denotes nearest-neighbor pairs of qubits i and j on separate compute nodes κ and κ'. By virtue of this definition, the step size within each compute node remains $\delta t = \frac{t}{N}$, yet the interaction between qubits on different compute nodes is computed with larger step size $n\delta t$. The corresponding circuit model is depicted in Fig. 1(b) for one sparse Trotter step with a single interconnect use.

We further define a *stochastic* sparse Trotterization, in which the time steps for remote operations are randomized. In this case, time intervals $t_1^{(\kappa,\kappa')}, t_2^{(\kappa,\kappa')}, \ldots$ for the Trotter step size of remote operations on every interconnect are chosen randomly, and an appropriate number of node-local time evolution steps with the usual step size δt are inserted in between. An example of the resulting circuit model is displayed in Fig. 1(c), and a precise definition is given in Sec. SI in the Supplemental Material (SM) [35]. This variation of the Trotter-Suzuki decomposition is intended to reflect the limitations of interconnects, which generate entanglement nondeterministically. It allows to stretch the duration of the node-local time evolution—and thus the number of gates applied and the absolute computing time—until the required interconnect becomes available. We note that inhomogeneous variations of the Trotter-Suzuki decomposition at various orders [36,37] have been studied previously in the context of quantum chemistry, where different terms in the electronic Hamiltonian are separated by their energy scale to allow for a reduction of the Trotter step size within controlled error bounds [38–41]. Here, our motivation to consider inhomogeneous step sizes is rooted in hardware constraints of a distributed quantum computing platform and separation occurs according to qubit connectivity.

Results. We begin our analysis by investigating the role of sparsity for the example of the XY model and compare

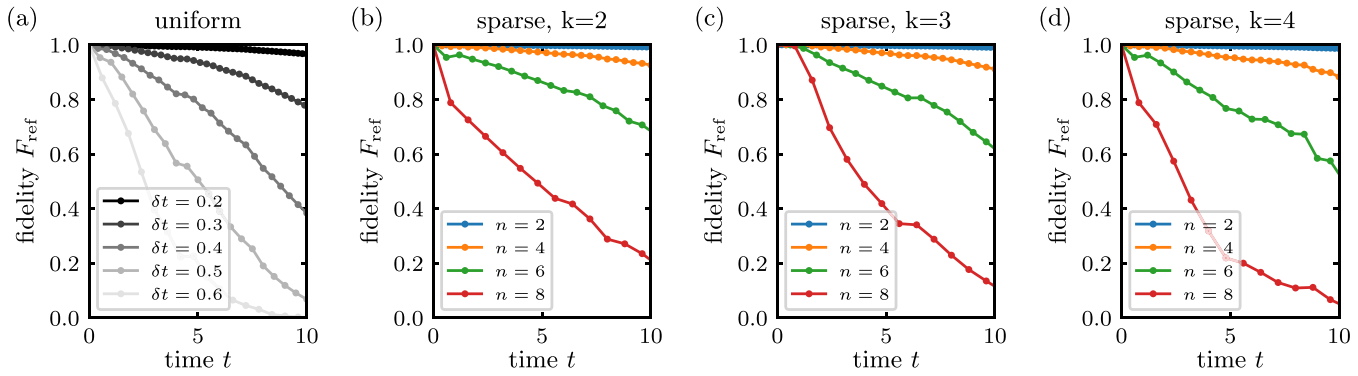

FIG. 2. Sparse Trotterization for the XY model with an initial domain wall configuration at the chain center. (a) Fidelity F_{ref} obtained for the time evolution under uniform Trotterization with various step sizes δt. The fidelity is computed with respect to the reference state obtained from step size $\delta t_{\text{ref}} = 0.1$. [(b)–(d)] Fidelity of states obtained from sparse Trotterization across $k = 2$, 3, and 4 distributed quantum compute nodes at different levels of sparsity n. The node-local step size is $\delta t_{\text{ref}} = 0.1$. Data is obtained for $L = 24$ spins.

its performance to the conventional uniform Trotterization. In the following, we shall assume a reference step size of $\delta t_{\text{ref}} = 0.1$ for node-local operations. The reference step δt_{ref} is chosen large enough such that relevant time scales could realistically be reached on near-term quantum hardware, and at the same time it is fine enough to closely reproduce the exact solution, see Sec. SII of the SM [35]. Note that for sparse Trotterization, a sparsity of n would then entail that remote operations are performed with a step size of $n\delta t_{\text{ref}}$. We evaluate all approximations by computing the associated wave functions and benchmarking against a reference state $|\psi_{\text{ref}}(t)\rangle$ that is obtained from uniform Trotterization with time step δt_{ref}. The quality of any state $|\psi(t)\rangle$ can then be quantified by the reference fidelity $F_{\text{ref}} = |\langle\psi_{\text{ref}}(t)|\psi(t)\rangle|^2$. We find that for increased step size $\delta t = 0.2$ in a uniform Trotterization the fidelity remains acceptable with $F_{\text{ref}} = 0.97$ after evolution to $t = 10$ but it diminishes quickly for larger step sizes, see Fig. 2(a). In contrast, on a distributed architecture with $k = 2$ compute nodes, introducing a sparsity of $n = 2$ still yields a fidelity of $F_{\text{ref}} = 0.99$ after evolution to $t = 10$ and the decay in result quality for increasing n is reduced significantly [Fig. 2(b)]: For example, sparse Trotterization with $n = 4$ ($F_{\text{ref}} = 0.93$) still performs better than uniform Trotterization with $\delta t = 0.3$ ($F_{\text{ref}} = 0.77$), despite the larger step size for the interconnect-mediated interaction. The trend not only holds for larger n, but also as the number of compute nodes k is increased moderately, see Figs. 2(c) and 2(d).

We make similar observations for the slow TFI quench model, where the fidelity is significantly more robust against local sparsity n than against a globally increased step size [Figs. 3(a) and 3(b)]. The robustness may be related to the finite magnetization in the initial state, which persists well beyond $t = 10$ and can act as a self-stabilizing mechanism against local perturbations, see Sec. SIII of the SM [35]. In the fast TFI quench model, the fidelity remains robust for $n < 6$ before decreasing substantially for $n \geqslant 6$. This still marks a substantial improvement over a global increase of the step size δt, especially for small $n = 2$ and $n = 4$ [Figs. 3(c) and 3(d)]. We note that the sudden drop in fidelity that occurs upon crossing the threshold of $n \geqslant 6$ is reminiscent of an abrupt transition into a quantum chaotic regime that has previously been identified under a uniform Trotterization when the uniform Trotter step size was chosen too large [42]. While the threshold into the chaotic regime is near $\delta t \geqslant 0.2$ for the uniform Trotterization (cf. Sec. SII of the SM [35]), it is $n \geqslant 6$ in the sparse Trotterization scheme, which corresponds to a step size of $\delta t \geqslant 0.6$ for interconnect-mediated interactions only. We thus observe a significant enhancement of the quantum chaos threshold in the sparse Trotterization scheme.

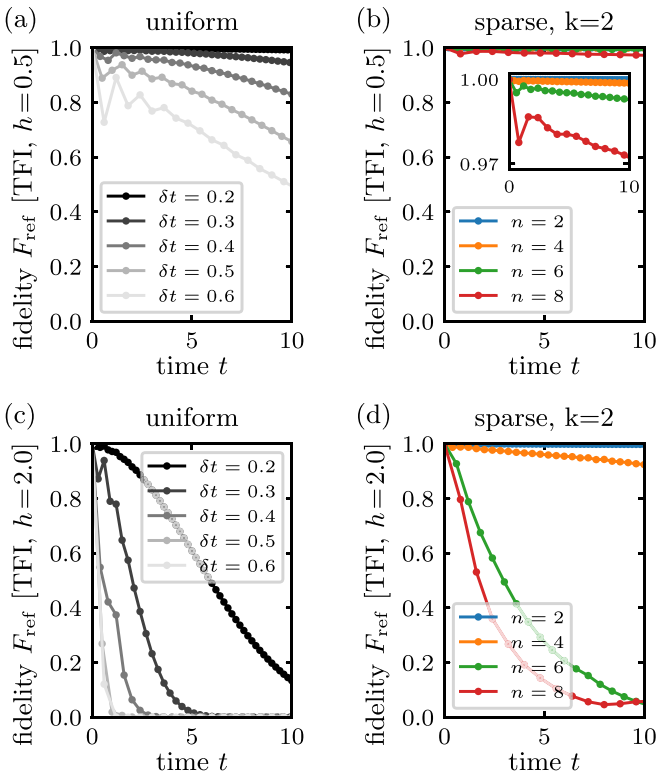

FIG. 3. Sparse Trotterization for TFI quench models. (a) Fidelity F_{ref} for uniform Trotterization with various step sizes δt, obtained for the slow TFI quench model. The fidelity is computed with respect to the reference state obtained from step size $\delta t_{\text{ref}} = 0.1$. (b) Fidelity obtained from sparse Trotterization at varying sparsity n on $k = 2$ distributed quantum compute nodes. The node-local step size is $\delta t_{\text{ref}} = 0.1$. [(c), (d)] Same as panels (a) and (b) but for the fast TFI quench model. Data is obtained for $L = 24$ spins.

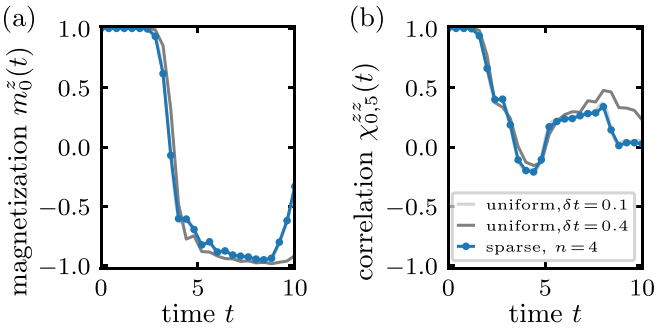

FIG. 4. Magnetization and magnetic correlations in the XY model. (a) Local magnetization at the chain boundary obtained from sparse Trotterization with $n = 4$ and uniform Trotterization with $\delta t = 0.1, 0.4$. The curve for sparse Trotterization coincides with the curve for uniform Trotterization at $\delta t = 0.1$. (b) Same as panel (a) but for the magnetic correlation between chain boundary and bulk. Data is computed for $L = 24$ spins across $k = 2$ compute nodes and a node-local step size $\delta t_{\text{ref}} = 0.1$.

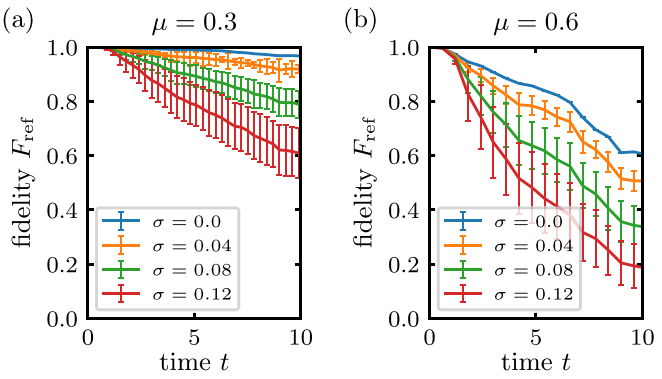

FIG. 5. Stochastic sparse Trotterization of the time evolution for the XY model at $k = 3$. The fidelity is shown for different levels of standard deviation σ. Average sparsity is set to (a) $\mu = 0.3$ and (b) $\mu = 0.6$, respectively. Data is obtained for $L = 18$ spins and node-local step size $\delta t_{\text{ref}} = 0.1$.

For practical applications, usually the goal is to accurately predict physical observables like the time-dependent magnetization $m_i^z(t) = \langle\psi(t)|\sigma_i^z|\psi(t)\rangle$ or the magnetic correlation function $\chi_{i,j}^{zz}(t) = \langle\psi(t)|\sigma_i^z\sigma_j^z|\psi(t)\rangle$. We find that deviations from the magnetization and correlations of the reference state $|\psi_{\text{ref}}(t)\rangle$ are generally small, especially when compared to the error that accumulates when the step size is globally increased in a uniform Trotterization. This is illustrated for the magnetization of the XY model at the chain boundary $m_0^z(t)$ in Fig. 4(a) and for the correlation between the chain boundary and the bulk $\chi_{0,5}^{zz}(t)$ in Fig. 4(b). More systematically, we compute deviations from the reference magnetization as $\Delta m_i^z(t) = m_i^z(t) - \langle\psi_{\text{ref}}(t)|\sigma_i^z|\psi_{\text{ref}}(t)\rangle$. The maximum deviation $\max_{i,t}(|\Delta m_i^z(t)|)$ obtained with sparsity $n = (2, 4, 6)$ is $(0.03, 0.10, 0.22)$. In contrast, the maximum deviation for states obtained from uniform Trotterization with analogous uniform step sizes $\delta t = (0.2, 0.4, 0.6)$ is significantly larger, yielding $(0.15, 0.58, 0.94)$. Differences of similar magnitude are also observed for the deviation of the magnetic correlations. The full space- and time-resolved data for the XY model and the TFI quench models is shown in Sec. SIII of the SM [35].

The above results indicate that the quality of the sparse Trotterization, for the Hamiltonians with two-spin interactions considered in this paper, smoothly depends on the sparsity parameter n and on the number of sparse qubit pairs $k - 1$. Increasing the Trotter step size only between a small number of $k - 1$ qubit pairs does not immediately lead to a proliferation of the error to levels that are associated with a uniform increase of δt between all qubit pairs. Such inhibition of error proliferation is intimately tied to the way that information spreads in quantum mechanical systems; it has previously been demonstrated that for systems with short ranged interactions information propagates within a light-cone-like horizon [43,44]. Whereas the propagation speed is model dependent, its confining effect on the proliferation of local errors is universal. As a consequence, the results of a conventional Trotter-Suzuki decomposition with a given step size $\delta t_{\text{uniform}}$ can be matched or improved by an inhomogeneous decomposition for which the step size is more fine-grained between most of the qubit pairs ($\delta t < \delta t_{\text{uniform}}$) but coarser between a small number of qubit pairs ($\delta t > \delta t_{\text{uniform}}$). For the models studied here, a reduction of the number of Trotter steps between a small number of qubit pairs by about 50% seems feasible.

Finally, we explore the effect of randomness in the stochastic sparse Trotterization. For this purpose, the time steps on the sparse qubit pairs are drawn from a normal distribution with mean μ and standard deviation σ. We then calculate F_{ref} and average it over 1000 instances of randomized configurations for each parameter set. In practical applications, stochastic averaging occurs naturally when measurement results are sampled statistically and interconnect uses on distributed quantum computing systems are inherently randomized. We find that despite relatively large variation across the different randomized configurations, the mean value for the fidelity remains smooth and systematically depends on μ and σ, see Fig. 5 for data on the XY model on $k = 3$ distributed compute nodes. Data for the TFI quench models is shown in Sec. SI of the SM [35]. Our data also indicates that randomness σ can have a more adverse effect than a systematic increase of the mean step size μ. For example, F_{ref} at $(\mu, \sigma) = (0.3, 0.12)$ is comparable to the fidelity achieved for much larger mean step size in the absence of randomness, $(\mu, \sigma) = (0.6, 0.0)$, see Fig. 5. Note that in the former case a step size of less than 0.6 would effectively occur with probability greater than 99% and a fidelity enhancement would therefore naively be expected. We speculate that the randomness leads to a reduction in the cancellation of Trotterization error terms that has been observed to be relevant for practical models in condensed matter physics [42,45,46].

Conclusions. We have demonstrated that issues in the implementation of quantum simulations on distributed quantum computers that arise from slow interconnect hardware can be mitigated by modifying the Trotter-Suzuki decomposition to allow for a nonuniform variation of the Trotter step size. For the XY model and the TFI quench models studied in this letter, a coarsening of the step size between a small number of qubit pairs could be compensated by a refinement of the

step size across the remaining qubit pairs. Notably, the sparse Trotterization could be applied successfully despite the fact that the underlying models only have a single principal energy scale and therefore ruling out Trotter constructions that rely on scale separation [38–41]. This observation has important consequences for the implementation of quantum simulations on distributed quantum computers. Gate operations that are facilitated by an interconnect are expected to remain slower than operations within a single node, even as future hardware generations evolve. To remedy the speed deficiency, instead of using Trotterization with a uniform time step that is bounded by the interconnect speed, results of similar or better fidelity can be obtained by maintaining fine-grained time stepping within each compute node. For the examples considered here, we find that a reduction in the number of interconnect uses by as much as 50% can be viable.

Further, we explored the possibility of exploiting the nondeterministic dead time after every interconnect use to perform additional node-local Trotter steps until the interconnect becomes available. Our data suggests that the randomized execution of additional Trotter steps quickly degrades the result quality. Unless the overall reduction in total compute time can compensate for the randomness-induced fidelity loss, it remains more beneficial to delay the execution of additional Trotter steps. Further calculations with hardware specific error models are required to find the optimal tradeoff.

In this paper, we focused on one-dimensional models. To explore more general applications in the future, it would be interesting to benchmark the performance of sparse Trotterization for quantum spin models with next-nearest-neighbor interactions or beyond, as well as for models in higher dimensions. In such generalizations, increased internode communication is expected and additional optimization of the interconnect usage may be necessary [47,48].

Further practical applications could also include the sparse Trotterization of imaginary time evolution [49] or the integration with variational algorithms [9,44,50] on distributed quantum computers.

Acknowledgments. This work was supported by Mitacs through the Mitacs Elevate program. D.S. acknowledges the NSERC Discovery Grant and the Canada Research Chairs Program.

[1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).

[2] H. S. Zhong, H. Wang, Y. H. Deng, M. C. Chen, L. C. Peng, Y. H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[3] Y. Wu, W. S. Bao, S. Cao, F. Chen, M. C. Chen, X. Chen, T. H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H. L. Huang, Y. H. Huo, L. Li, N. Li *et al.*, Strong Quantum Computational Advantage Using a Superconducting Quantum Processor, Phys. Rev. Lett. **127**, 180501 (2021).

[4] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins *et al.*, Quantum computational advantage with a programmable photonic processor, Nature (London) **606**, 75 (2022).

[5] R. P. Feynman, Simulating physics with computers, Int. J. Theor. Phys. **21**, 467 (1982).

[6] S. Raeisi, N. Wiebe, and B. C. Sanders, Quantum-circuit design for efficient simulations of many-body quantum dynamics, New J. Phys. **14**, 103017 (2012).

[7] A. Macridin, P. Spentzouris, J. Amundson, and R. Harnik, Electron-Phonon Systems on a Universal Quantum Computer, Phys. Rev. Lett. **121**, 110504 (2018).

[8] 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).

[9] 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).

[10] S. P. Jordan, K. S. Lee, and J. Preskill, Quantum algorithms for quantum field theories, Science **336**, 1130 (2012).

[11] S. P. Jordan, K. S. Lee, and J. Preskill, Quantum computation of scattering in scalar quantum field theories, Quantum Inf. Comput. **14**, 1014 (2014).

[12] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, Cloud Quantum Computing of an Atomic Nucleus, Phys. Rev. Lett. **120**, 210501 (2018).

[13] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, Gate-count estimates for performing quantum chemistry on small quantum computers, Phys. Rev. A **90**, 022305 (2014).

[14] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation, Phys. Rev. A **91**, 022311 (2015).

[15] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, The trotter step size required for accurate quantum simulation of quantum chemistry, Quantum Inf. Comput. **15**, 361 (2015).

[16] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya *et al.*, Quantum chemistry in the age of quantum computing, Chem. Rev. **119**, 10856 (2019).

[17] Y. Cao, J. Romero, and A. Aspuru-Guzik, Potential of quantum computing for drug discovery, IBM J. Res. Dev. **62**, 6 (2018).

[18] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[19] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, Nature (London) **607**, 667 (2022).

[20] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke *et al.*, Noisy intermediate-scale quantum algorithms, Rev. Mod. Phys. **94**, 015004 (2022).

[21] D. Cuomo, M. Caleffi, and A. S. Cacciapuoti, Towards a distributed quantum computing ecosystem, IET Quantum Commun. **1**, 3 (2020).
[22] R. Van Meter and S. J. Devitt, The path to scalable distributed quantum computing, Computer **49**, 31 (2016).
[23] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network, Phys. Rev. Lett. **124**, 110501 (2020).
[24] Eagle's Quantum Performance Progress, https://research.ibm.com/blog/eagle-quantum-processor-performance. Accessed: 02-12-2022
[25] IonQ Aria: Practical Performance (Part One), https://ionq.com/posts/july-25-2022-ionq-aria-part-one-practical-performance. Accessed: 02-12-2022.
[26] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, Phys. Rev. A **69**, 032315 (2004).
[27] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels, Phys. Rev. Lett. **76**, 722 (1996).
[28] M. Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, J. Math. Phys. **32**, 400 (1991).
[29] H. F. Trotter, On the product of semi-groups of operators, Proc. Amer. Math. Soc. **10**, 545 (1959).
[30] 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).
[31] S. Olmschenk, D. N. Matsukevich, P. Maunz, D. Hayes, and C. Monroe, Distant matter qubits, Science **323**, 486 (2009).
[32] P. Pfeuty, The one-dimensional Ising model with a transverse field, Ann. Phys. **57**, 79 (1970).
[33] P. Calabrese, F. H. Essler, and M. Fagotti, Quantum quench in the transverse field Ising chain: I. Time evolution of order parameter correlators, J. Stat. Mech. (2012) P07016.
[34] In some cases, tighter bounds can be formulated, depending on the structure of the underlying Hamiltonian [42,45,46,51].
[35] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.5.L022003 for additional data and details on the stochastic sparse Trotterization.
[36] A. Papageorgiou and C. Zhang, On the efficiency of quantum algorithms for Hamiltonian simulation, Quantum. Inf. Proc. **11**, 541 (2012).
[37] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Efficient quantum algorithms for simulating sparse Hamiltonians, Commun. Math. Phys. **270**, 359 (2007).
[38] S. Hadfield and A. Papageorgiou, Divide and conquer approach to quantum Hamiltonian simulation, New J. Phys. **20**, 043003 (2018).
[39] A. M. Childs, A. Ostrander, and Y. Su, Faster quantum simulation by randomization, Quantum **3**, 182 (2019).
[40] E. Campbell, Random Compiler for Fast Hamiltonian Simulation, Phys. Rev. Lett. **123**, 070503 (2019).
[41] Y. Ouyang, D. R. White, and E. T. Campbell, Compilation by stochastic hamiltonian sparsification, Quantum **4**, 235 (2020).
[42] M. Heyl, P. Hauke, and P. Zoller, Quantum localization bounds Trotter errors in digital quantum simulation, Sci. Adv. **5**, eaau8342 (2019).
[43] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, Commun. Math. Phys. **28**, 251 (1972).
[44] M. Benedetti, M. Fiorentini, and M. Lubasch, Hardware-efficient variational quantum algorithms for time evolution, Phys. Rev. Res. **3**, 033083 (2021).
[45] D. Layden, First-Order Trotter Error from a Second-Order Perspective, Phys. Rev. Lett. **128**, 210501 (2022).
[46] L. M. Sieberer, T. Olsacher, A. Elben, M. Heyl, P. Hauke, F. Haake, and P. Zoller, Digital quantum simulation, Trotter errors, and quantum chaos of the kicked top, npj Quantum Inf. **5**, 78 (2019).
[47] E. Tham, I. Khait, and A. Brodutch, Quantum circuit optimization for multiple QPUs using local structure, in *Proceedings of the 2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Broomfield, CO, 2022), p. 476.
[48] MultiQopt compilation benchmarks, https://entanglednetworks.com/multiqopt. Accessed: 02-12-2022.
[49] M. Motta, C. Sun, A. T. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. Brandão, and G. K. L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, Nat. Phys. **16**, 205 (2020).
[50] D. Amaro, C. Modica, M. Rosenkranz, M. Fiorentini, M. Benedetti, and M. Lubasch, Filtering variational quantum algorithms for combinatorial optimization, Quantum Sci. Technol. **7**, 015021 (2022).
[51] 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).