


Benchmarking Digital Quantum Simulations Above Hundreds of Qubits Using Quantum Critical Dynamics

Alexander Miessen^{1,2,†}, Daniel J. Egger¹, Ivano Tavernelli¹, and Guglielmo Mazzola^{2,*}

¹IBM Quantum, IBM Research—Zurich, Rüschlikon 8803, Switzerland

²Institute for Computational Science, University of Zurich, Winterthurerstrasse 190, Zurich 8057, Switzerland

 (Received 17 April 2024; revised 7 July 2024; accepted 13 September 2024; published 6 November 2024)

The real-time simulation of large many-body quantum systems is a formidable task, that may only be achievable with a genuine quantum computational platform. Currently, quantum hardware with a number of qubits sufficient to make classical emulation challenging is available. This condition is necessary for the pursuit of a so-called quantum advantage, but it also makes verifying the results very difficult. In this paper, we flip the perspective and utilize known theoretical results about many-body quantum critical dynamics to qualitatively benchmark digital quantum hardware and various error mitigation techniques. In particular, we benchmark against known universal scaling laws in the Hamiltonian simulation of a time-dependent transverse-field Ising Hamiltonian of up to 133 qubits. Incorporating only basic error mitigation and suppression methods, our study shows reliable control up to a two-qubit gate depth of 28, featuring a maximum of 1396 two-qubit gates, before noise becomes prevalent. These results are transferable to applications such as Hamiltonian simulation, variational algorithms, optimization, or quantum machine learning. We demonstrate this on the example of digitized quantum annealing for optimization and identify an optimal working point in terms of both circuit depth and time step on a 133-site optimization problem.

DOI: [10.1103/PRXQuantum.5.040320](https://doi.org/10.1103/PRXQuantum.5.040320)

I. INTRODUCTION

Quantum critical dynamics occur when a quantum system reaches a critical point (CP), characterized by a non-analytic change in the system's ground-state energy as a function of a Hamiltonian parameter. The system's correlation length and relaxation time are maximal at the CP and diverge in the thermodynamic limit. Crossing it, the system undergoes a quantum phase transition (QPT). As a result, system details no longer affect macroscopic quantities, causing the emergence of universal behavior, a key property of critical phenomena [1–3].

QPTs are experimentally realized by changing the control parameters of the system Hamiltonian over time. However, understanding the real-time dynamics of many-body quantum systems close to the critical point is a formidable task. Many believe that only quantum simulators [4], i.e., controllable quantum systems that can emulate others [5], can tackle this problem at scale [6,7]. Quantum simulators

were first realized with ultracold dilute gas in an optical lattice [8] and have since been implemented on a variety of platforms [6,7,9–11]. Most of these platforms are *analog* quantum simulators, which are subject to calibration errors and decoherence.

A parallel approach is to perform the simulation on *digital* quantum computers using suitable algorithms compiled to the native basis gate set of the hardware [12–17]. However, current digital machines are prone to errors, such as calibration errors, cross-talk, and decoherence. In the future, quantum error correction could potentially enable fault-tolerant simulations of many-body quantum systems on digital machines [18]. Reference [19] gives a perspective on the relative strengths and weaknesses of analog and digital platforms.

Crucially, quantum critical dynamics have implications beyond condensed-matter and statistical physics. For example, QPTs occur ubiquitously in quantum optimization, where a quantum algorithm helps solve a classical optimization problem. Such applications are among the most anticipated and economically impactful use cases for quantum computers [20]. Quantum annealing (QA) is an algorithm to find ground states that can be used to solve combinatorial optimization problems [21]. It evolves an easy-to-prepare ground state of one Hamiltonian to the unknown ground state of another problem Hamiltonian, which corresponds to the classical optimization problem to

*Contact author: guglielmo.mazzola@uzh.ch

†Contact author: lex@zurich.ibm.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

solve. If the evolution is adiabatically slow, the system follows the instantaneous ground state of the time-dependent Hamiltonian and ends in the solution of the optimization problem. Particularly for QA, QPTs create algorithmic bottlenecks [21–23], as they imply an energy gap between the ground and the first excited state that closes in the thermodynamic limit. For finite annealing times t_f , the evolution may therefore not be adiabatic due to a diverging relaxation time at the CP. This mechanism produces defects that carry through to the final state [24].

In a noise-free, closed system, the density of defects is a nonincreasing function of t_f , and several scaling regimes can be identified as t_f increases [24,25]. The quantum Kibble-Zurek mechanism (QKZM) [24,26,27] quantifies the relationship between t_f and the number of defects produced during the annealing run in a regime of sufficiently slow, yet finite, t_f . It predicts a universal scaling, a power-law decay, of the density of defects in the final solution as a function of t_f and the system’s critical exponents [28], i.e., a set of numbers characterizing the system’s behavior near its CP. It has been the subject of a range of largely analog experimental studies [29–33].

We first present an application-oriented benchmarking method that utilizes the predictability of known universal scaling laws, such as the Kibble-Zurek (KZ) scaling. Those attributes make it ideal for an intuitive and easily scalable metric to assess the quality of large-scale quantum simulations. This is increasingly important as digital quantum computing devices and algorithms have left the infantile stage of a few tens of qubits with error rates prohibiting more than a handful of two-qubit gates [34,35]. Today’s quantum devices can exceed 100 qubits at error rates that, combined with error mitigation and suppression techniques [36], allow a coherent simulation of thousands of two-qubit gates [14,37].

Many benchmarking methods employ randomized circuits. This makes them more comparable, objective, and device agnostic while accounting for, e.g., qubit connectivity and basis gate set [38]. Moreover, randomness helps collect many of the error sources into the same benchmarking process. Examples of such methods span an entire field of research, including variants of randomized benchmarking [39–41], cross-entropy benchmarking [42], and quantum volume [43]. These methods are crucial to assess gate fidelities and characterize different kinds of noise. However, they do not appropriately represent the capabilities of current quantum hardware to run a specific application. Indeed, applications for noisy quantum devices have circuits with highly structured repeated layers that are often tailored to the connectivity and the basis gate set of the hardware. It is therefore of little surprise that such simulations achieve much higher gate counts and qubit numbers than one would predict from generic benchmarks. Even though there are very promising new benchmarking methods designed to overcome these issues [44], they still

provide only a fidelity without relating to applications. Pure application benchmarks are the other extreme. Either by testing against a set of well-understood problems and, possibly, corresponding solutions [45,46], or based on specific applications, such as quantum optimization [47] or discrete time crystals [48].

We go a step further and propose a concrete, application-oriented benchmarking scheme to assess how many two-qubit gate layers of a given application circuit can be reliably simulated given a specific device and, importantly, error mitigation and suppression (EMS) method. A schematic of the principle is shown in Fig. 1(a). The result is not an abstract fidelity but simply the number of circuit layers that can be reliably simulated, which can be directly transferred to other applications. Here, we focus on the KZ scaling, though other universal scalings, e.g., at short times, can also be used [25].

This benchmark is transferrable to applications in time evolution, optimization, quantum machine learning, and variational algorithms. To showcase this, we apply the benchmark to combinatorial optimization specifically as the underlying circuits directly implement digitized QA or the quantum approximate optimization algorithm (QAOA). QAOA is an annealing-inspired variational ansatz, with variational layers representing time steps in digitized QA, used in conjunction with a variational quantum eigensolver (VQE) [49]. Recent studies indicate that solving hard optimization problems requires many QAOA layers [50,51]. At the same time, the optimally converged parameter values of QAOA in the large-layer limit reproduce annealing schedules, meaning QAOA implements digitized QA with variationally found annealing schedules. It is therefore important to better understand digitized QA. A relevant question is whether the algorithmic error stemming from a finite time step is detrimental or possibly even beneficial, similar to what happens in simulated QA [52]. Here, we identify an optimal working point with respect to the time step and number of circuit layers to minimize the residual energy given the finite hardware resources and how it depends on the system’s minimum energy gap.

In summary, we explore two directions—benchmarking and optimization—and show how the former can be used to guide the design of the latter. Section II provides the necessary technical background to QA, the QKZM, and Trotterized time evolution. In Sec. III, we introduce and implement our qualitative application-oriented benchmark of quantum simulation against universal behavior. Here, we present the main results of this work—a comparison of different levels of EMS on two quantum processors with up to 133 qubits using our QKZM-based benchmarking. Section IV studies digitized QA, using the same underlying circuits as in the benchmarking experiments, for solving optimization problems and showcases how results from Sec. III are transferable to related applications. We conclude in Sec. V.

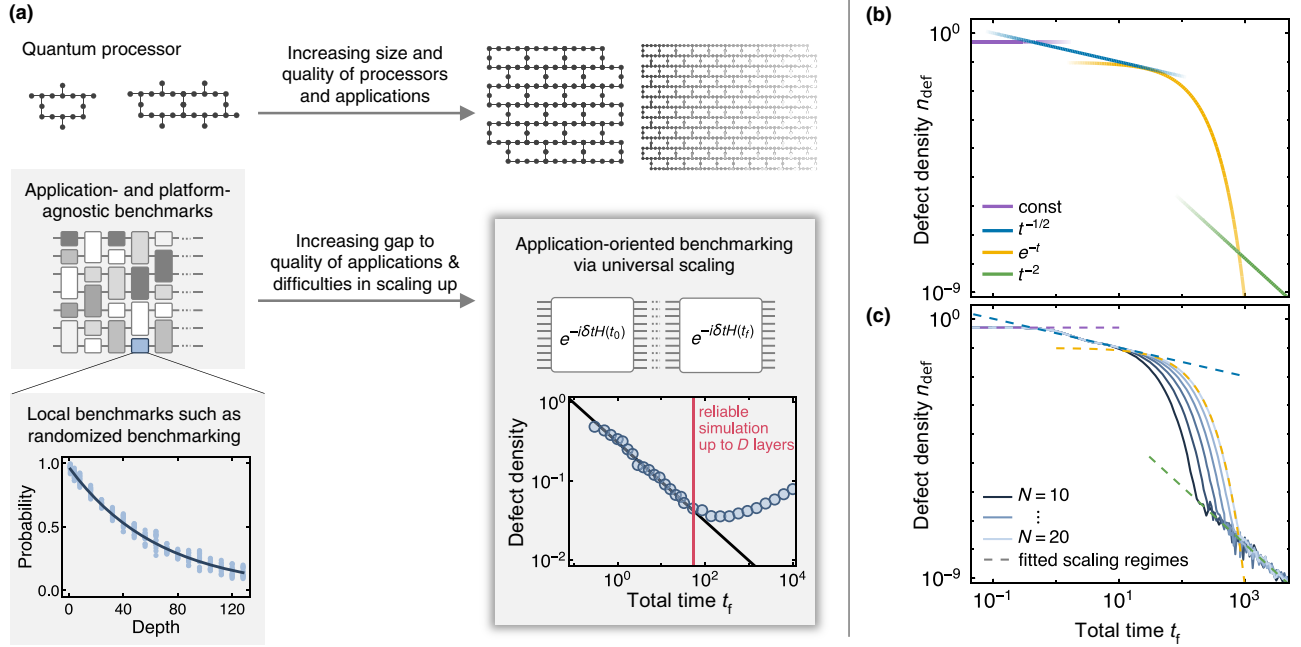


FIG. 1. Application-oriented benchmarking of quantum simulations. (a) Benchmarking experiments give detailed insight into device characteristics such as error rates and decoherence times. As devices scale, however, it is desirable to have a benchmarking method that resembles the applications that are executed. By simulating critical dynamics and measuring an observable that follows a known universal scaling, we track for how many time steps, which correspond directly to the number of circuit layers, we can reproduce the expected universal scaling. Here, we measure the density of defects produced in quantum annealing, predicted by the Kibble-Zurek scaling. This provides a direct and concrete metric predicting how many layers of two-qubit gates can be reliably simulated given, e.g., a certain device, qubit subset, and error mitigation technique. Note that (a) contains only schematic plots and shows no actual data. (b) For finite system sizes away from the thermodynamic limit, the density of defects exhibits two scaling regimes beyond the Kibble-Zurek scaling that are theoretically known (see main text). These can benchmark deep circuits. (c) The departure from Kibble-Zurek scaling is determined by the system size and all three scaling regimes can be identified in state-vector simulations with $N = 10, \dots, 20$ spins (solid lines). Shown here are results from Trotterized simulations with $\Delta t = 0.01$ for $t_f \leq 0.1$ and $\Delta t = 0.1$ otherwise.

II. DIGITIZED QUANTUM ANNEALING AND DEFECT PRODUCTION

Given a (mixing) Hamiltonian H_M with an easy-to-prepare ground state and a problem Hamiltonian H_P whose ground state we wish to compute, we construct a time-dependent combination of the two,

$$H(s) = A(s)H_M + B(s)H_P. \quad (1)$$

Here, $s = t/t_f$ is the time $t \in [0, t_f]$ rescaled by the total annealing time t_f . $A(s)$ and $B(s)$ are the annealing schedules such that $H(0) = H_M$ and $H(1) = H_P$. The system is initially prepared in the ground state of H_M and evolved for time t_f . If this evolution is adiabatically slow, i.e., for large enough t_f , the system remains in the ground state of the instantaneous Hamiltonian and ends in the desired ground state of H_P at $t = t_f$ [21].

The prototypical Hamiltonian considered in QA is an N -site transverse-field Ising model (TFIM) with nearest-neighbor interactions,

$$H_M = - \sum_i \sigma_i^x, \quad H_P = - \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z, \quad (2)$$

with Pauli matrices σ_i^x and σ_i^z acting on site i and couplings J_{ij} between neighboring sites indicated by $\langle i,j \rangle$. The ground state of $H(0) = H_M$ is $|+\rangle^N$. We apply linear schedules $A(s) = (1 - s)$ and $B(s) = s$ for simplicity, resulting in

$$H(s) = -A(s) \sum_i \sigma_i^x - B(s) \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z. \quad (3)$$

Unless otherwise stated, we consider an ordered Ising model with uniform couplings $J_{ij} \equiv J$.

The TFIM is the simplest model exhibiting a QPT [2] in the thermodynamic limit $N \rightarrow \infty$. In the simple case of uniform couplings $J > 0$ ($J < 0$), the CP at $J_c = A/B$ separates the paramagnetic phase ($A \gg JB$) from the doubly degenerate (anti)ferromagnetic phase ($A \ll JB$) with all spins (anti)aligned. In the disordered case, the latter is substituted by a glassy phase, i.e., an energy landscape with many deep local minima [23,53,54].

A. The Kibble-Zurek mechanism

The QKZM describes the nonequilibrium dynamics of the systems in the realistic setting of crossing this CP with finite annealing time t_f , leading to defects in the final solution. Reference [24] provides an excellent account of the QKZM and the resulting density of defects scaling in a TFIM, both by employing the original (classical) reasoning of the Kibble-Zurek mechanism (KZM) and in terms of a fully quantum description based on Landau-Zener (LZ) theory. Here, a defect refers to the wrong orientation of a spin and the density of defects to the number of defects averaged over all sites. In the ferromagnetic phase specifically, defects are domain walls between antialigned spins. We now summarize the classical description of the KZM following Ref. [24]. Classically, the CP is characterized by a diverging correlation length ξ and relaxation time τ . Far away from the CP, τ is small enough for the system to instantaneously relax to equilibrium. As the system approaches the CP, τ becomes equal to and, eventually, grows beyond the time scale on which H changes. When this happens, the system can no longer relax to its instantaneous ground state, with reactions to changes slowing down until they halt completely. This is usually referred to as the critical slowdown and subsequent freezing out. Past the CP, as τ decreases again, the system unfreezes and, crucially, continues evolving from approximately the frozen-out state, resulting in defects in the final state. In summary, the growth of the system's relaxation time at the CP, which is finite for finite system sizes, determines the necessary rate of change of the system $ds/dt = 1/t_f$ required to track the instantaneous ground state. This is the essence of the KZM. It predicts the density of defects as a function of a finite annealing time t_f with a power-law decay [28,30]

$$n_{\text{def}} \propto t_f^{-\frac{d\nu}{1+z\nu}}, \quad (4)$$

fully determined by the system dimension d and its critical exponents ν and z . In one dimension (1D), with $d = 1, z = 1, \nu = 1$, this becomes

$$n_{\text{def}} \propto t_f^{-1/2}. \quad (5)$$

In the quantum description, the existence of a QPT and all its related quantities originate from the closing of the system's energy gap at the CP. For finite system sizes $N < \infty$, the minimum gap between the ground state and the first excited state is small but finite and decreases as $\propto 1/N$. It closes in the thermodynamic limit $N \rightarrow \infty$, causing both the correlation length ξ and relaxation time τ to diverge. Assuming annealing times fast enough to produce at least one defect on average, LZ theory yields the same KZ scaling in a quantum setting [24].

For uniform couplings J , defects in the ferromagnetic solution after annealing appear as domain walls between

antialigned spins and can be measured by $\sigma_i^z \sigma_j^z$ correlators over all N_e edges of the spin lattice, i.e., all nearest-neighbor correlators. The density of defects is measured as the density of spin misalignments (domain walls) in 1D,

$$n_{\text{def}} = \frac{1}{2N_e} \sum_{(i,j)}^{N_e} \left(1 - \sigma_i^z \sigma_j^z\right). \quad (6)$$

For example, an open chain of N spins has $N_e = N - 1$ edges and a periodic chain has $N_e = N$ edges. In arbitrary higher-dimensional systems, the scaling of the number of misaligned spins may not exactly match the scaling of the number of defects, i.e., domains of spin alignment. This is particularly true when the system features few but large domains. In that case, one could directly count the number of defects, taking into account the graph topology or, for instance, measure the average number of domains from the number of spin misalignments along different directions [55].

B. Density of defect scaling across various regimes

As mentioned above, the QKZM is not expected to hold for every possible annealing time t_f . Deviations from the KZ scaling occur at both very short and very long t_f [25,56,57]. For finite system sizes and very slow anneals, i.e., large t_f , Ref. [24] describes an exponential drop in defect density following the regime of KZ scaling that is captured by LZ dynamics. In that case, when the anneal is slow enough to never freeze out and to produce less than one defect in the chain on average, the number of defects is proportional to the LZ probability p_{LZ} of exciting the system,

$$n_{\text{def}} \propto p_{\text{LZ}} \approx \exp\left(-b \frac{t_f}{N^2}\right), \quad (7)$$

with b a constant. Following Refs. [58,59], for adiabatically slow anneals, i.e., even larger t_f , the scaling reads

$$n_{\text{def}} \approx p_{\text{LZ}}(t_f) + \frac{1 - 2p_{\text{LZ}}(t_f)}{at_f^2}, \quad (8)$$

where a is a constant that depends on the derivatives of the annealing schedules [58]. The expressions of constants a and b can be found in Refs. [24,59].

The theoretical predictions for the three scaling regimes in Eqs. (5), (7), and (8) are shown in Fig. 1(b). Density of defect scalings obtained from ideal state-vector simulations in Fig. 1(c) with system sizes up to $N = 20$ spins show that the beyond-KZM scaling regimes, i.e., Eqs. (7) and (8) are finite-size effects and that their onset is controlled by the system size.

C. Product formulas for time-dependent Hamiltonians

Ultimately, QA implements the time evolution under a time-dependent Hamiltonian $H(t)$. The unitary time evolution operator obtained by integrating the time-dependent Schrödinger equation reads

$$U(t_f, t_0) = \mathcal{T} \exp \left(-i \int_{t_0}^{t_f} dt H(t) \right). \quad (9)$$

The time-ordering \mathcal{T} accounts for any noncommutativity of the Hamiltonian with itself at different times, $[H(t), H(t')] \neq 0$.

To implement QA on a digital quantum computer, we must decompose Eq. (9) into quantum gates [12]. The simplest digitized implementation of Eq. (9) for a Hamiltonian $H(t) = \sum_{l=1}^L a_l(t) H_l$ is via a first-order product formula (PF), which approximates the integral as a Riemann sum, $\int_0^{t_f} dt H(t) = \lim_{n \rightarrow \infty} \sum_{m=1}^n H(m\Delta t) \Delta t$ with $\Delta t = t_f/n$ [60,61]. Importantly, the time ordering in the Trotterized $U(t_f, t_0)$ is enforced with a “right-to-left ordering,” resulting in

$$U \approx \prod_{m=n}^1 e^{-i\Delta t \sum_l a_l(m\Delta t) H_l}. \quad (10)$$

Therefore, in contrast to the case of a time-independent Hamiltonian, splitting the exponent into discrete time steps introduces a first approximation. Typically, a digital quantum computer cannot natively implement e^{-iH} . Therefore, the Hamiltonian at time step $m\Delta t$ is further decomposed using a first-order PF, or any higher-order PF,

$$U \approx \prod_{m=n}^1 \prod_{l=1}^L e^{-ia_l(m\Delta t) H_l \Delta t}. \quad (11)$$

Each of these exponentials can be directly represented by the hardware native basis gates of the digital machine. Note that, in principle, any other quantum algorithm for quantum dynamics could be used to decompose the time-evolution operator as long as it allows for time-dependent Hamiltonians [12,62].

III. APPLICATION-ORIENTED BENCHMARKING THROUGH UNIVERSAL SCALING

Here, we suggest a direct approach to benchmark close to applications. We simulate the time evolution under a time-dependent Hamiltonian and benchmark the accuracy with which a known universal scaling is replicated. A schematic of our method is shown in Fig. 1(a). We propose to measure, for example, the density of defects after digitized QA as a function of total annealing time t_f . Given a fixed time step, the time t_f directly corresponds

to the number of circuit layers in the Trotterized time-evolution circuit. In the KZ regime, the density of defects will decrease according to the predicted scaling up to a certain threshold t_f , after which hardware noise becomes dominant and leads to a deviation from the KZ scaling. The number of circuit layers for which the expected scaling of n_{def} is observed is the number of circuit layers for which reliable simulation was achieved.

In particular, decoherence of the system will lead to a deviation of the defect density from the predicted KZ scaling. This has been previously confirmed in several numerical and analog experiments that studied the effect of dissipation on the KZ scaling [28,31,63,64]. These studies show that different kinds of noise can cause decoherence of the system. However, the behavior of the density of defects as noise accumulates, and whether it increases or decreases, depends on the noise.

In this regard, it is important to note that observing a $t_f^{-1/2}$ scaling may not be a sufficient condition to claim coherent quantum dynamics. For instance, a purely classical diffusion model may reproduce the same density of defect scaling [31,65]. This is particularly important for analog quantum annealers. In the presence of sufficiently strong thermal noise that destroys coherence, the density of defects could still decrease with t_f , since the machine may behave like a classical thermal annealer (at low temperature) or as a machine featuring incoherent quantum tunneling events [66–68]. These processes could still lower the density of defects. Indeed, this effect has been observed in Ref. [31], necessitating additional analysis to reasonably prove the existence of coherent annealing [69].

In our digital quantum setting, the situation is different: the presence of hardware noise never yields a behavior that resembles a classical thermal limit. To prove this, in Appendix A, we study the impact of simulated hardware noise with varying strength. We observe that introducing hardware noise qualitatively changes n_{def} as a function of t_f . Even for large t_f and accumulating noise, we do not recover a decreasing $n_{\text{def}}(t_f)$, contrary to the noisy analog quantum annealing case [31]. Therefore, the density of defects following the expected scaling constitutes evidence of a noise-free digital simulation up to a threshold t_f . This argument is further corroborated by the fact that the noise model is in very good agreement with hardware results (Appendix A).

The proposed benchmark has several advantages. First and foremost, it yields a concrete, intuitive, and scalable metric; the number of simulable circuit layers as they can be found in countless applications. Even more so since Hamiltonian simulation is a prime application for quantum computers and a building block for many other applications, for example, quantum phase estimation [70] and sampling algorithms [71]. Second, our method can benchmark hardware and EMS algorithms separately or

in combination. Third, our method is scalable since no classical verification is required. Fourth, since the KZ scaling is defined in the thermodynamic limit $N \rightarrow \infty$, it is particularly well suited to benchmark large digital quantum computers without scaling issues. Finally, since finite-size scaling regimes are also well understood [cf. Sec. II B and Figs. 1(b) and 1(c)], the method is applicable in the setting of scaling to large circuit depths at qubit counts far below the thermodynamic limit.

A. Experimental setup

We now discuss benchmarking different levels of EMS on two quantum processors through digitized QA and measuring defect densities. An N -spin Ising model with uniform couplings $J = 1$ and linear schedules (cf. Sec. II) is mapped to N qubits. The qubit register is initialized as $|\psi(t=0)\rangle = |+\rangle^N$, the ground state of $H(t=0) = H_M$, and time evolved under the time-dependent Hamiltonian in Eq. (3). The time evolution is implemented using a first-order PF as in Eq. (11) with a time step of Δt . Throughout this section, we use a time step of $\Delta t = 0.5$. We justify the choice of this time step in Appendix B, showing results of noiseless state-vector simulations with different time steps and comparing them to continuum results. Although smaller time steps would yield higher accuracy, more Trotter steps, i.e., deeper circuits, would be required to explore the relevant KZ regime. Our results show that $\Delta t = 0.5$ is a reasonable compromise to avoid significant algorithmic errors and allow for sufficiently large annealing times. Hardware improvements will eventually allow for smaller time steps and thereby increase the overall accuracy of the results, particularly at short time scales. After annealing for a total time t_f , we measure the defect density via Eq. (6) to compare to the expected scaling in Eq. (5). Expectation values of observables are estimated from measurements with QISKIT’s Estimator primitive [72].

We employ the IBM Heron processor `ibm_torino` and the IBM Eagle processor `ibm_sherbrooke` for our simulations, with 133 and 127 qubits, respectively, and heavy-hexagonal qubit connectivity [73]. Hardware characteristics such as error rates are given in Appendix D. Eagle and Heron processors differ in their native two-qubit gates and how they are physically realized. `ibm_sherbrooke`’s native two-qubit gate is an echoed cross-resonance (ECR) gate implemented via a dispersive coupling mediated by a fixed-frequency resonator [74,75]. By contrast, `ibm_torino`’s native two-qubit gate is a CZ gate realized through tunable-frequency couplers. In a tunable coupler architecture, the coupling element is frequency tunable, and driving it can create different two-qubit gates [76,77]. Most importantly, this results in `ibm_torino` having a roughly $6\times$ shorter two-qubit gate time compared to `ibm_sherbrooke` and a higher two-qubit gate quality [44].

On both devices, we compare simulations with no EMS to readout error mitigation (REM) [78,79] as well as to REM combined with dynamical decoupling (DD) [80]. In addition, `ibm_sherbrooke` allows for pulse-efficient transpilation of our circuits, another method of error suppression [81]. Pulse-efficient transpilation scales hardware native cross-resonance pulses and their echoes (combined, this makes up an ECR gate on `ibm_sherbrooke`) and is thus not compatible with `ibm_torino`’s native CZ gates. Appendix C gives a brief technical summary of each method.

B. Results: 1D chain

The Trotter circuit for the time evolution of a spin chain requires at least two layers of R_{ZZ} gates per Trotter layer, or time step. Since each R_{ZZ} gate is transpiled to two hardware native two-qubit gates (ECR or CZ), the final circuit for a 1D Ising model has a two-qubit gate depth of $4 \times N_t$ where N_t is the number of time steps. Figure 2 shows the density of defect scaling obtained from digitized QA on subsets of 100 qubits connected through an open line (see insets), which was chosen to minimize cumulative two-qubit gate errors along the line (cf. Appendix D). Owing to noise, the experimental results differ from the noiseless state-vector results in Fig. 1(c), which show a monotonic decrease in defect density with increasing t_f . Hardware errors only make it possible to follow the KZ scaling and subsequent scaling regimes up to a certain threshold time. Beyond this point (indicated by gray shading), the defect density deviates from the $t^{-1/2}$ scaling and eventually increases, as the noise accumulates with increasing circuit depth as predicted by the noisy emulation of the algorithm.

This threshold time depends on the hardware and the EMS employed. The top panel of Fig. 2 shows the results for `ibm_torino` comparing no EMS with only REM, and with REM combined with DD. While we can simulate two Trotter layers with a two-qubit gate depth of 8 and a total of 396 CZ gates without any level of EMS, adding REM already significantly improves the results. We can reliably simulate up to five Trotter layers with a two-qubit gate depth of 20 and 990 CZ gates in total. Assuming `ibm_torino`’s median CZ gate time of 84 ns for all two-qubit gates, this corresponds to a circuit execution time of 1.7 μ s. This is remarkable given that REM is the simplest and easiest-to-implement form of error mitigation (EM). We attribute the negligible impact of the DD sequence on the `ibm_torino` results to the reduced cross-talk of the device compared to `ibm_sherbrooke`.

The picture is different on `ibm_sherbrooke`, shown in the bottom panel of Fig. 2. We compare the same EMS methods as on `ibm_torino` in addition to pulse-efficiently transpiled circuits. The results without any EMS and only REM are slightly worse than on `ibm_torino`. However, adding DD to counter static ZZ cross-talk results

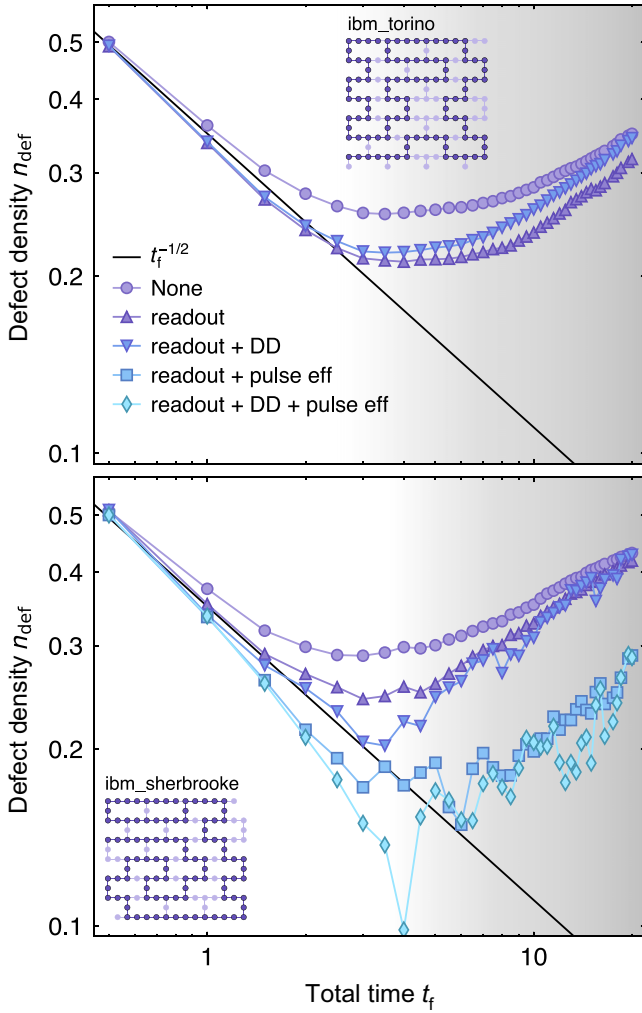


FIG. 2. Density of defects scaling in 1D from 100-qubit circuits. Hardware results comparing the density of defects scaling on subsets of 100 qubits on *ibm_torino* (top) and *ibm_sherbrooke* (bottom) employing different levels of EMS. Each point corresponds to one additional Trotter layer or time step of $\Delta t = 0.5$. Raw simulation results with no EMS (None) are compared to results using combinations of REM, DD error suppression, and pulse-efficient transpilation. Insets show the qubit layout of the respective processor with the chosen qubit subset highlighted. The gray shading indicates decoherence and deviation from the expected KZ scaling shown by the black solid line.

in substantial gains on this device. According to our metric, we achieve a reliable simulation of seven circuit layers with a two-qubit gate depth of 28 and 1386 ECR gates in total. Assuming *ibm_sherbrooke*'s median ECR gate time of 533 ns [73] for all two-qubit gates, this equates to a circuit execution time of roughly 15 μ s. Moreover, we show in Appendix F that the correlations between defects show the characteristic nonmonotonic fingerprint of a genuine quantum KZM [69,82]. Pulse-efficient transpilation fairs similarly, though we observe that the curve

dips below the expected scaling. We attribute this to rotation errors in the scaled two-qubit pulses at small angles [81], see Appendix C. As the two-qubit gate angles $\theta = -2JB(m\Delta t/t_f)\Delta t = -2Jm\Delta t^2/t_f$, with $m\Delta t = t$, are proportional to the annealing schedule, increasing t_f means smaller and smaller angles at the beginning of an anneal. This causes a growing disparity between the target operation of the circuit before transpilation and the final sequence of gates with growing t_f . Fixing these rotation errors requires custom calibrations [83], which are difficult at scale through a cloud-based quantum computing service. Since our goal is to compare out-of-the-box EMS methods, we did not conduct any custom calibrations.

C. Results: Heavy-hexagonal lattice

Our approach is not limited to one-dimensional spin systems. In fact, such critical behavior and the resulting universal scaling manifests in lattice geometries beyond 1D [56] and can hence benchmark any qubit connectivity, provided that critical exponents and scaling are known.

The processors we use have a 2D heavy-hexagonal qubit connectivity. In this case, theoretical comparisons are less straightforward than in the 1D case. Figure 3(a) shows noiseless state-vector simulations comparing the density of defects scaling of a 1D periodic chain, a heavy-hex lattice, and a 2D square lattice with $N = 21$ sites each. Here, the density of defects in 1D and heavy-hex lattices coincide on the $t^{-1/2}$ scaling before the exponential drop-off at larger t_f , attributed to finite-size effects, while the 2D square lattice shows a steeper scaling. However, this could be an artifact of the small system sizes since a heavy-hex lattice of 21 sites might not be large enough to show genuine 2D scaling. We consider the $t^{-1/2}$ scaling line as a reference line also for experiments on the heavy-hex lattice, even though deviations from the 1D scaling are possible for this geometry. Given also that this 2D system is only sparsely connected, with an edges-over-sites ratio of 1.2, meaning correlations will spread slowly across the lattice and may not acquire 2D character within a few time steps, we still adopt Eq. (6) to measure the density of defects.

The Trotter circuit for a Hamiltonian Eq. (3) on a heavy-hex lattice is made depth-optimal by grouping the couplings according to the edge-coloring shown in Fig. 3(b). Each Trotter layer in the circuit therefore consists of three layers of R_{ZZ} gates. Again, each R_{ZZ} gate is transpiled into two hardware-native two-qubit gates, resulting in a total two-qubit gate depth of $6 \times N_t$.

Utilizing all 133 and 127 qubits of *ibm_torino* and *ibm_sherbrooke*, respectively, we see an even starker difference between both chips in defect density [Fig. 3(c)]. Taking the 1D scaling as a reference line as argued above, we can reliably simulate up to three Trotter layers on *ibm_torino* employing only REM. With 150 edges, this equals a two-qubit gate depth of 18 and 900 CZ

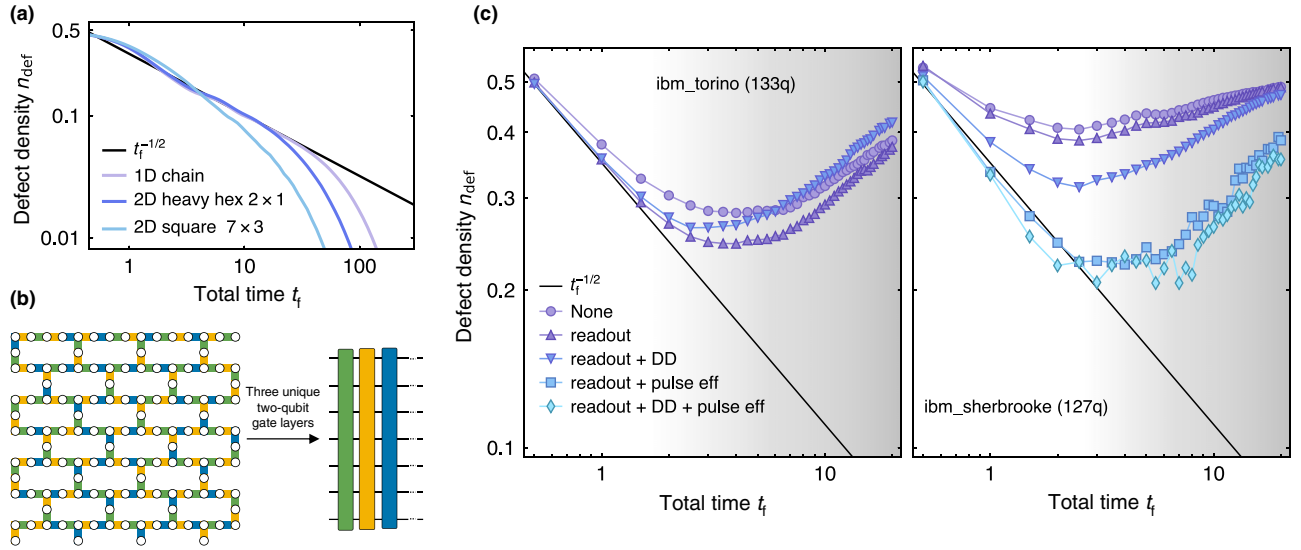


FIG. 3. Density of defect scaling on a heavy-hexagonal lattice from 133- and 127-qubit circuits. (a) State-vector results comparing the density of defect scaling in a periodic chain, a heavy-hexagonal lattice consisting of two heavy-hex cells (2×1), and a 7×3 square lattice, all with $N = 21$ spins. The heavy-hexagonal lattice exhibits the same $t^{-1/2}$ scaling as predicted for 1D systems, whereas defects in a true 2D lattice show a steeper scaling. (b) The couplings on the heavy-hexagonal lattice of *ibm_torino* and *ibm_sherbrooke* are grouped such that each Trotter layer consists of three layers of R_{ZZ} gates. (c) Hardware results comparing the density of defect scaling on *ibm_torino* (left) and *ibm_sherbrooke* (right) employing different levels of EMS. Each point corresponds to one additional Trotter layer or time step of $\Delta t = 0.5$. Raw simulation results with no EMS (None) are compared to results using combinations of REM, DD error suppression, and pulse-efficient transpilation. The gray shading indicates decoherence and deviation from the expected KZ scaling shown by the black solid line.

gates in total. Again, assuming *ibm_torino*'s median CZ gate time of 84 ns for all two-qubit gates, this implies a circuit execution time of $1.5 \mu\text{s}$ using merely REM. *ibm_sherbrooke*, on the other hand, requires more involved EMS, which becomes even more apparent when using all qubits of the device. Although DD again provides sizable improvements, we do not reproduce the expected KZ scaling. However, pulse-efficient transpilation combined with REM extends the circuit depth up to five Trotter layers before noise becomes prevalent. This is equivalent to a Trotter circuit with a CNOT gate depth of 30 and 1440 CNOT gates in total. In this case, the circuit duration reduces to roughly $6 \mu\text{s}$ due to the pulse-efficient transpilation.

IV. DIGITIZED QUANTUM ANNEALING FOR OPTIMIZATION

The proposed benchmark provides a reasonable prediction of hardware and error mitigation capabilities for a set of problems broader than the nonequilibrium quantum dynamics of a ferromagnetic Ising model. Estimating how many circuit layers can be reliably executed before noise prevails is useful in many applications. For instance, variational algorithms [46,84] for many-body spin systems feature circuits very similar to the ones considered in the previous section. Circuits with the same structure

may also be used in quantum machine learning [85–87]. Finally, the entangling blocks used previously can be adapted to QAOA or digitized QA of disordered spin models.

We now turn to combinatorial optimization, where the couplings J_{ij} in Eq. (2) can be arbitrary in magnitude and sign. Optimization problems defined by this type of Hamiltonian may fall within the NP-complete complexity class, depending on the connectivity of the underlying graph, and are frequently employed as benchmarks for quantum optimization [88,89]. It is therefore of interest to determine whether the results obtained from benchmarking the QKZM and identifying an optimal working point are transferable to an optimization context. This is not straightforward since the QKZM is connected with second-order phase transitions, while the bottlenecks in quantum annealing of hard optimization problems are due to first-order transitions in the glassy phase [23,53,54].

QA has been little explored on digital quantum computers [90]. Instead, the focus has been mostly on QAOA, an annealing-inspired variational ansatz and the most popular near-term algorithm for quantum optimization [49, 80,91,92]. However, optimizing the variational parameters in QAOA is costly in itself [93] and, once found, optimal variational parameters reduce QAOA to QA in the large-layer limit [50]. This does not mean that the optimal QAOA parameters can be directly derived from

annealing schedules. Rather, the variational optimization results in optimized annealing schedules. Nonetheless, precisely because of this relationship, annealing schedules can be used to initialize QAOA parameters, thereby reducing the training cost of the variational ansatz [93,94]. It is therefore sensible to study digitized QA in more detail. Previous experimental realizations of digitized QA featured only up to nine qubits [95].

Since digitized QA requires discretizing the time evolution, a central question is what effect does the choice of the time step have on the success of the annealing. In principle, a large time-step error could, counterintuitively, even improve the performance of digitized QA. This effect has been observed in classical path-integral simulations of QA, where an unconverged, finite Trotter step is beneficial for tunneling between configurations [52], reconciling earlier expectations of quantum speedup [96] with observations [88]. While investigating this aspect we identify an optimal time step for digitized QA given a fixed number of circuit layers or time steps N_t . Fixing N_t is meaningful since, in reality, one can only access a limited amount of resources such as the expendable number of gates, determining both the computational cost and, more importantly before achieving fault tolerance, the amount of error introduced by hardware noise. The success of QA relies on as large as possible annealing times t_f , which is achieved by increasing the time step Δt for a fixed number of time steps N_t . On the other hand, the errors of time-evolution algorithms typically scale with Δt [12]. PFs are derived on the assumption of small time steps $\Delta t \ll 1$ and algorithmic errors scale with $\mathcal{O}(\Delta t^2)$ [61]. This naturally introduces a trade-off between QA performance and the algorithmic error due to Δt [94]. Note that, despite this trade-off, having a converged time step is no requirement here. Instead, the sole objective is to obtain the best possible solution, i.e., the lowest energy. This differs from Sec. III where we chose the largest possible time step that is still reasonably close to the continuum while allowing us to reach as large as possible t_f .

Since the solution to a classical optimization problem is a single bitstring, we are usually interested in individual measurement samples rather than expectation values. Therefore, we use QISKIT's Sampler primitive [72] throughout this section to obtain individual measurement samples, i.e., bitstrings of measurement outcomes. Many EM techniques apply only to expectation values. By contrast, error suppression methods such as DD and pulse-efficient transpilation apply to sampling. However, as discussed in Sec. III, DD does not yield substantial improvements on `ibm_torino`. Furthermore, IBM Heron processors, such as `ibm_torino`, are based on tunable couplers and thus do not allow for pulse-efficient transpilation, which was designed for cross-resonance-based hardware [81]. To make results comparable across different devices, we therefore only use REM in this

section. Furthermore, we will compute the residual energy,

$$E_{\text{res}} = E(t_f) - E_0, \quad (12)$$

where $E(t_f) = \langle \psi(t_f) | H(t_f) | \psi(t_f) \rangle$ and E_0 is the exact ground state obtained through exact diagonalization for small system sizes and using CPLEX [97] for large systems.

A. Dependence of the residual energy on the time step and spectral gap

We first study the dependence of the digitized QA result on the time step on a small system of a periodic 12-qubit chain. Three instances of disorder with couplings randomly sampled from $J_{ij} \in [-1, 1]$ are chosen with varying minimum energy gaps between the ground and the first parity-preserving excited state since a smaller gap entails longer annealing times and makes finding a solution generally more difficult. Note that the Hamiltonian Eq. (3) and therefore its spectrum is a function of s . The minimum energy gap takes on its smallest value at the critical point. Furthermore, the first excited state of opposite parity becomes degenerate with the ground state as $s \rightarrow 1$. The spectra of the three different Hamiltonians as a function of s are shown in Fig. 4(a) with minimum gaps ranging across one order of magnitude. Figure 4(b) displays the residual energy after annealing the respective system from (a) on `ibm_auckland` [98] with fixed time step $\Delta t \in \{0.1, \dots, 2.0\}$ as a function of the number of time steps (depth). Here, we observe that the smallest residual energies are reached after approximately 5–10 Trotter layers, owing to decoherence. The observation of a minimum is compatible with both numerical predictions [64] and experimental results from analog simulation [31,99]. Furthermore, it is analogous to what was observed in the benchmarking experiment, Fig. 3(c), where we also located a minimum defect density after roughly five–ten Trotter layers.

For each Δt , the minimum residual energy achieved over all circuit depths is plotted in Fig. 4(c). For all three systems, the lowest residual energy is obtained with a time step of $\Delta t > 1$, specifically $1.0 < \Delta t < 1.5$. Noiseless state-vector simulations confirm that this is not merely a consequence of hardware noise, as seen in Appendix E, yielding an optimal time step of $1.2 < \Delta t < 1.4$. However, in state-vector simulation, the existence of a finite optimal time step is likely due to the finite range of depths chosen. Therefore, choosing an infinitesimal time step seems to be both (i) inefficient in noise-free simulations and (ii) impractical in real experiments. On the other hand, a too large time step implies algorithmic errors.

To summarize, we observe a trade-off between realizing the largest possible annealing times while keeping the algorithmic error under control. Choosing an unconverged,

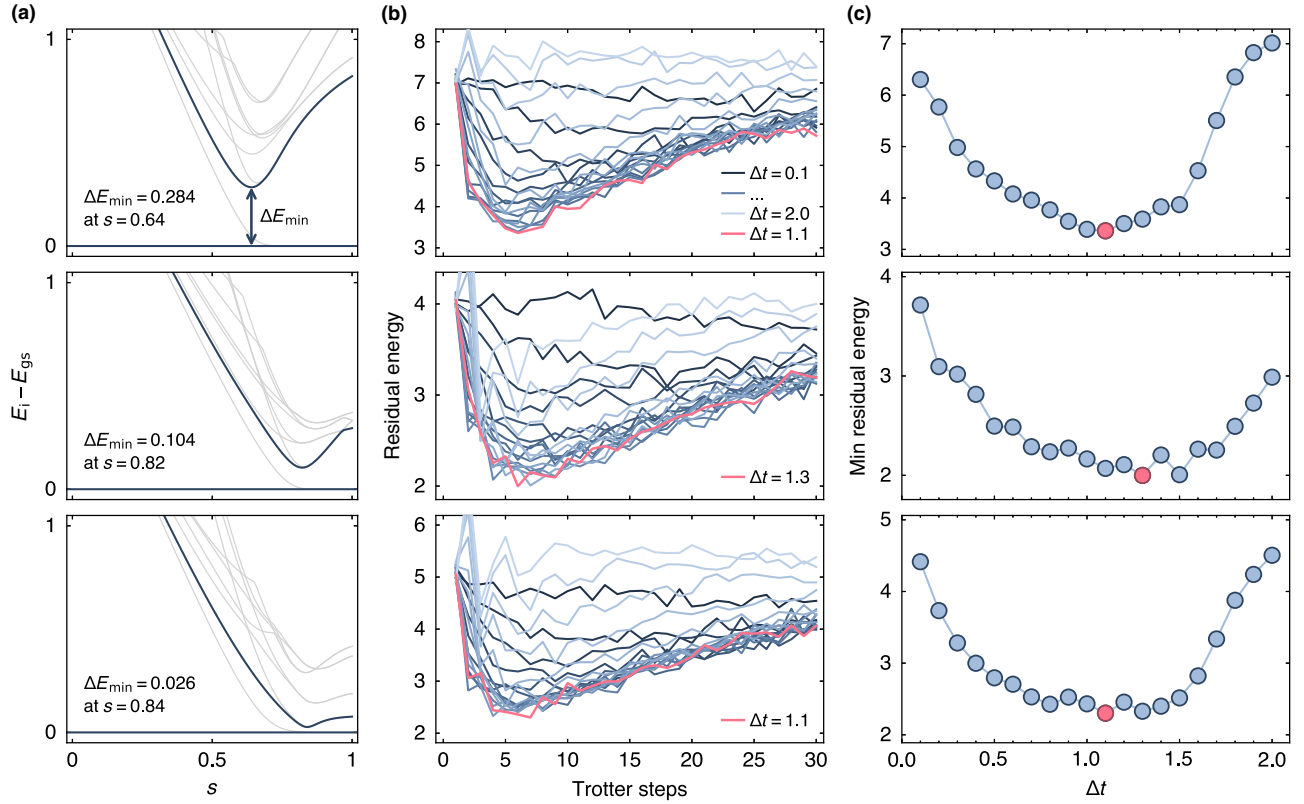


FIG. 4. Hardware results of the residual energy dependence on time step and spectral gap. (a) Spectra relative to the ground state of three different instances of disorder in a 12-qubit periodic chain with coupling coefficients uniformly sampled from $J_{ij} \in [-1, 1]$. The instances were specifically chosen with three different minimum energy gaps, differing by an order of magnitude between the largest (top) and the smallest gap (bottom). (b) Residual energy averaged over 400 samples obtained from QA of the system in (a) on `ibm_auckland` using only REM with fixed time steps $\Delta t \in \{0.1, \dots, 2.0\}$ as a function of circuit depth (number of time steps). (c) Minimum residual energy from (b) as a function of the time step, i.e., each point corresponds to the minimum over one curve in (b).

large time step is indeed advantageous when using digitized QA for optimization. This is of further relevance when initializing QAOA variational parameters mentioned previously. In the context of directly studying QAOA performance, Ref. [94] finds an optimal time step of 0.75 for initializing variational parameters. Increasing Δt even further, however, does not provide any benefits as shown by the state-vector results. This observation is consequential for practical noisy settings and allows us to optimize the number of Trotter steps needed to reach a target annealing time.

B. Optimization of disordered heavy-hexagonal graph

We now study a heavy-hexagonal spin lattice of 133 qubits, i.e., `ibm_torino`'s full chip. We consider just one realization of disorder with uniformly sampled couplings $J_{ij} \in [-1, 1]$, and compute E_{res} as a function of circuit depth and time step as before. The ground state is again computed using CPLEX. This time, the system is too large to compute its instantaneous higher excited states, which is why we have no information about the size of its minimum

energy gap. However, the comparison of different systems in the previous section indicates that the dependence of the optimal time step on the spectral gap is small. Figure 5(a) shows the results for time steps $\Delta t \in \{0.5, \dots, 1.5\}$. Again, we observe that, for each time step, the minimum residual energy is obtained at around five–ten Trotter layers, or time steps, after which decoherence begins to dominate. Figure 5(b) reports the minimum E_{res} for each value of the time step. Also here, we observe a continuous decrease in E_{res} with increasing Δt up to a time step of 1.0 to 1.1, after which E_{res} sharply increases. In conclusion, this suggests an unconverged, large time step of $\Delta t \gtrsim 1.0$ to be optimal for digital QA regardless of the system instance and graph connectivity. Further studies will be required to explore the generality of this empirical finding.

C. Consistency with the KZ benchmark

Here, we briefly summarize the relation of our large-scale optimization results in Sec. IV B to the results of Sec. III. In Sec. III, the 1D ferromagnetic system (i.e., uniform

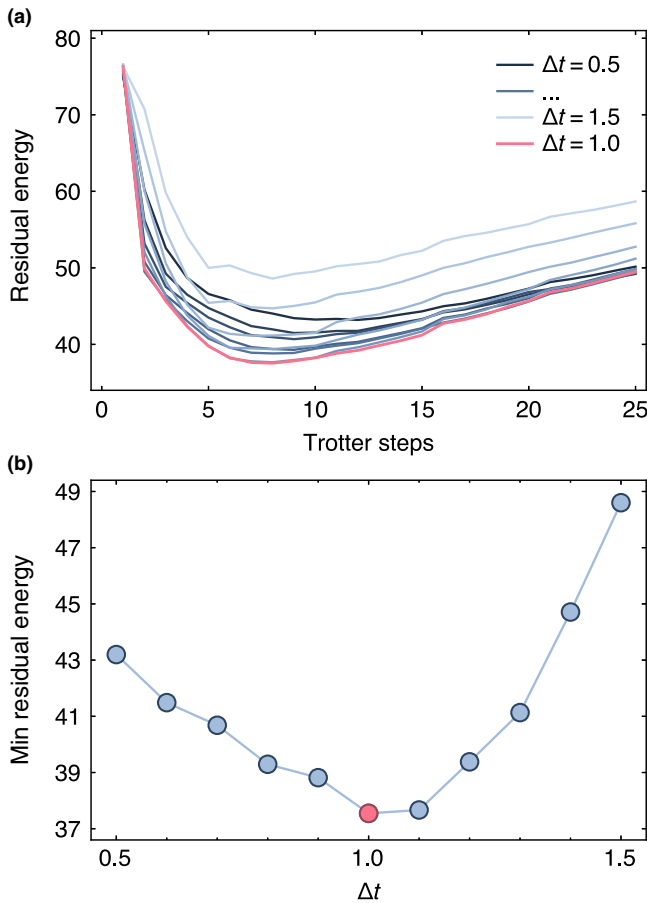


FIG. 5. Hardware results of the residual energy dependence on the time step. (a) Residual energy of a heavy-hexagonal 133-qubit graph with coupling coefficients uniformly sampled from $J_{ij} \in [-1, 1]$, averaged over 10^5 samples obtained from QA on `ibm_torino` with fixed time steps $\Delta t \in \{0.5, \dots, 1.5\}$ as a function of circuit depth (number of time steps). (b) Minimum mean residual energy from (a) as a function of the time step, i.e., each point corresponds to the minimum over one curve in (a).

couplings) simulation on `ibm_torino`, see Fig. 2, indicates that about five Trotter steps can be executed reliably before the density of defects substantially deviates from the expected scaling (referring to the REM curve for consistency with the optimization experiments). Subsequently, at about eight Trotter steps, the density of defects reaches a minimum before it increases due to the accumulated noise. Similar values of three and eight Trotter steps, respectively, are obtained on the heavy-hex lattice, see Fig. 3.

The optimization experiment shares the same connectivity as these benchmarks but has random couplings and different time steps. Depending on the time step used, the minimum energy we measure (which is the disordered analog to the defect density) lies again in the range of five–ten Trotter steps, see Fig. 5(a). This finding is consistent with our results from Sec. III. When the goal is solely to minimize an expectation value, as in applications such

as optimization or variational methods, the location of the minimum in $n_{\text{def}}(t_f)$ can be relevant in addition to the point where it starts deviating from the expected scaling.

This demonstrates the usefulness of our benchmarking routine in an application context. Suppose, for example, that one is interested in using QAOA rather than digitized QA. Since QAOA features the same circuit structure in a closed-loop classical parameter optimization, it is expensive to discover the optimal circuit depth empirically. Running the proposed benchmark beforehand allows one to directly select the optimal QAOA depth according to the hardware capabilities.

V. DISCUSSION AND OUTLOOK

This work presents an application-oriented method to jointly benchmark the hardware and the algorithms. Developing such a method to predict simulation quality, and, importantly, more closely resemble device and algorithm capabilities, is critical. Particularly so in times of rapid algorithmic and hardware advancements. The simple figure of merit that we introduce is the number of Trotter steps, i.e., the threshold depth, before which errors induce deviations from a universal scaling law. While we propose to use digitized QA and the universal KZ scaling in our work, we emphasize the central idea: benchmark against the prediction of a known universal behavior. Other scaling laws, model Hamiltonians, or lattice geometries could be employed as well.

We observe that different machines, equipped with different EMS, provide different threshold depths. Importantly, the simple benchmark we propose can measure the continuous improvements of hardware and algorithms in the near future. Moreover, it is tailored to Hamiltonian simulation, which is among the most anticipated applications with provable quantum advantage in physics and a building block in many other applications.

We test our method on hardware native geometries and the resulting circuits contain only dense layers of R_{ZZ} gates. However, this does not limit its predictive power since all two-qubit gates are eventually transpiled to the same hardware-native two-qubit gate. One can therefore directly transfer results in terms of two-qubit-gate depth to the simulation of more complicated models. Nonetheless, future research could devise more varied models exhibiting universal behavior that could be used analogously to the KZ scaling. For instance, circuits with interleaved two-qubit gates of different kinds that do not commute with each other such that they cannot be arranged in a dense layer of gates, as for example encountered in fermionic models. Furthermore, our method could be used to select best-qubit subsets. By measuring all nearest-neighbor $\sigma_i^z \sigma_j^z$ correlators on a given quantum processor, the defect densities of different qubit subsets can be reconstructed in postprocessing. For example, computing the

density of defects for all qubit subsets of size N in this way would allow selecting the subset that best matches the KZ scaling. Moreover, the known finite-size scaling regimes for large annealing times could one day be used to benchmark very deep quantum circuits once they become accessible with advanced EMS techniques or, eventually, quantum error correction (QEC).

Our benchmark extends beyond quantum many-body simulations to quantum optimization and variational algorithms. We observe consistency between the threshold number of Trotter steps before the noise becomes prevalent, identified using the KZ scaling, and the optimal depth of digitized QA for combinatorial optimization. We demonstrate that, counterintuitively, the time step should indeed be chosen at a value $\gtrsim 1.0$, resulting in significantly improved residual energy values after annealing. This means in practice that digitized QA could be a competitive quantum optimization framework as it (i) avoids a classical optimization loop and (ii) greatly reduces the runtime of QAOA [100], avoiding the costly iterative optimization, affected by hardware and shot noise [93,101]. More generally, our results suggest that current hardware with simple EMS may support variational ansätze (for a similar class of Hamiltonians) comprising about ten Trotter-circuit layers. Our results seem to be robust against several instances of disorder with varying minimum spectral gaps as well as on systems of different sizes and connectivities. Nonetheless, future work should aim at exploring other settings such as solving dense graphs [91].

In introducing this benchmark, we report some of the largest-scale digital quantum simulations as of today [14, 102–104]. Compared to Ref. [14], we move closer to a physical system by incorporating Trotter error and time-dependent, general two-qubit gate angles, requiring a full gate decomposition of the R_{ZZ} gate into hardware native two-qubit gates. However, we do not introduce sophisticated EMS such as probabilistic error amplification, which could be a goal of future work.

Finally, we emphasize that the purpose of this paper is not to claim direct quantum advantage with the featured experiments, including the optimization demonstration. For one, the sparse-graph optimization problem studied here is not hard. Other quantum or classical platforms may demonstrate better performance than we do [105]. Moreover, we do not claim that the digitized quantum annealing method is the best-performing digital algorithm. Indeed, warm-start methods [106,107] and digitized counterdiabatic approaches [108] all improve the performance of quantum approximate optimization. Instead, the value of our optimization experiments lies in understanding its correlation with the proposed benchmark. Furthermore, the benchmarks introduced here will be instrumental to verify future quantum simulations on sufficiently large hardware, beyond the reach of approximate classical methods.

The numerical data that support the findings of this study are available from the corresponding authors upon reasonable request.

The relevant scripts of this study are available from the corresponding authors upon reasonable request.

ACKNOWLEDGMENTS

We thank Julien Gacon, Elena Peña Tapia, and Almudena Carrera Vazquez for help with the code, and Stefan Wörner, Laurin Fischer, Samuele Piccinelli, and Francesco Tacchino for insightful and inspiring discussions. We also thank Adolfo Del Campo, Jacek Dziarmaga, Andrew King, and Hidetoshi Nishimori for providing useful feedback. G.M. acknowledges financial support from the Swiss National Science Foundation (Grant No. PCEFP2_203455). This research was supported by the NCCR MARVEL, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (Grant No. 205602). We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team.

APPENDIX A: INFLUENCE OF HARDWARE NOISE

The effects of noise such as a finite temperature or bath couplings on QA have been previously studied both numerically [28,64] and experimentally [31] in analog settings. The result is a deviation from the KZ scaling and a subsequent increase in defect density. As discussed in Sec. III, the effect of the noise encountered in a digital simulation is expected to be fundamentally different from that in an analog simulation. Specifically, digital hardware noise includes an accumulation of individual basis gate errors, mainly two-qubit gate errors, throughout the circuit and will not yield a classical thermal limit.

Here, we perform Trotter simulations of digitized QA including a tunable noise model resembling the digital noise we encounter on our devices. Since noise generally leads to decoherence of the system, this further validates our claim that observing a density of defects that follows the KZ scaling implies coherent evolution. Scaling the average device errors across several orders of magnitude allows for a detailed understanding of their influence on the density of defects. In other words, the noisier the system, the sooner we expect decoherence and, in turn, a deviation from the expected scaling.

The noise model is constructed from error rates of `ibm_sherbrooke` obtained through standard calibrations [73]. Specifically, we used QISKIT AER’s method `NoiseModel.from_backend_properties()` [72] that generates a simplified noise model based on calibration data of a specific device. The resulting noise model

includes readout errors on measurements as well as individual basis gate errors consisting of depolarizing noise (tuned through gate errors obtained from calibrations) followed by thermal relaxation noise (tuned through decoherence times T_1, T_2 and gate execution times). To simplify the scaling of the noise, we use error rates obtained from the average of individual gate and qubit properties. Concretely, our model applies the following noise to all qubits,

$$\begin{aligned} T_1 &= T_1^{\text{ave}}/\eta, & T_2 &= T_2^{\text{ave}}/\eta \\ e_{1q} &= \eta e_{1q}^{\text{ave}}, & e_{2q} &= \eta e_{2q}^{\text{ave}}, & e_{ro} &= \eta e_{ro}^{\text{ave}}. \end{aligned} \quad (\text{A1})$$

Here, T_1, T_2 are relaxation time and dephasing time, e_{1q}, e_{2q}, e_{ro} are single-qubit gate, two-qubit gate, and readout errors, respectively, and η is a scaling factor. The average values are $T_1^{\text{ave}} = 266.37 \mu\text{s}$, $T_2^{\text{ave}} = 178.71 \mu\text{s}$, $e_{1q}^{\text{ave}} = 1.25 \times 10^{-3}$, $e_{2q}^{\text{ave}} = 1.10 \times 10^{-2}$, and $e_{ro}^{\text{ave}} = 2.41 \times 10^{-2}$. All gate and measurement times are left unchanged and correspond to those of `ibm_sherbrooke`.

Figure 6(a) shows the density of defects obtained from QA of a periodic 12-qubit system with simulated noise scaled by $\eta \in \{0.01, \dots, 10\}$. The noise-free line is a state-vector simulation of the same system. In addition, we plot real hardware results of simulating the same 12-qubit system on `ibm_sherbrooke`, using only REM. The results of this experiment match the noisy simulations with unscaled noise $\eta = 1$, thereby validating our choice of noise model. Both the noisy and the noise-free simulations as well as the hardware experiment were done with the same time step of $\Delta t = 0.5$ as in Sec. III.

Since the system size that we can simulate exactly with noise is small, the behavior of the density of defects over time appears less smooth compared to larger sizes. These finite-size effects manifest as a drop in the noise-free density of defects after roughly $t_f \approx 10$. The noise-free simulations indicates a monotonically decreasing density of defects with t_f . In the noisy cases instead, after a threshold time t_f^* , the density of defects starts increasing with time. As the noise strength decreases, t_f^* increases. This is further underlined by Fig. 6(b), showing the relative error of the density of defects with respect to the noise-free curve, $|(n_{\text{def}}(\eta) - \tilde{n}_{\text{def}})/\tilde{n}_{\text{def}}|$, with \tilde{n}_{def} the noise-free defect density. While this behavior is compatible with earlier numerical studies on dissipative quantum annealing and analog hardware experiments [28,31,64], the main difference here is that the density of defects always keeps increasing with time after t_f^* . In contrast, dissipative quantum dynamics or annealing experiments [31,32], indicated that the density of defects may decrease again, possibly with the same $-1/2$ exponent, even beyond the coherence limit.




FIG. 6. Density of defect scaling in the presence of simulated hardware noise. (a) All curves show simulated Trotterized time evolution of a 12-qubit chain with $\Delta t = 0.5$. The noise-free reference curve (pink line) represents a state-vector simulation. The hardware curve (purple diamonds) corresponds to results obtained from `ibm_sherbrooke` using only REM. All other curves show results of sampling 10^5 measurements from a simulated circuit and include simulated hardware noise that corresponds to the full noise model of `ibm_sherbrooke`. (b) Error of the density of defects relative to the noise-free results.

APPENDIX B: TROTTER ERROR OF BENCHMARKING EXPERIMENTS

Discretization of the time evolution requires setting a time step, which in turn controls the discretization error, in our case the Trotter error. Ideally, the time step should be sufficiently small to reduce the Trotter error and accurately describe the continuous-time limit of the dynamics. However, smaller time steps require more Trotter circuit layers to reach the same final time. In Sec. III, we adopt a practical approach and utilize a small, yet nonvanishing, time step of $\Delta t = 0.5$.

Figure 7(a) shows state-vector simulations of a periodic 12-qubit chain with time steps varying between $\Delta t = 0.01, \dots, 1.0$ and compares the resulting density of defects to the exact, continuous-time result. The figure shows that our choice of $\Delta t = 0.5$, despite minor deviations from the continuum solution, provides a stable evolution over the entire time range we consider. This is not the case for larger time steps as can be seen by the increased fluctuations. Crucially, this time step allows us to experimentally probe relevant regimes of t_f . More concretely, the KZ scaling




FIG. 7. Density of defects scaling for different discretizations Δt . State-vector simulations of a periodic 12-qubit chain comparing the exact, i.e., nondiscrete time evolution to Trotterized simulations with different time steps Δt . The time step $\Delta t = 0.5$ is highlighted in pink and corresponds to the time step chosen for the hardware experiments in Sec. III.

we seek to benchmark against is not present for small annealing times t_f , and observing the KZM requires probing sufficiently large t_f . However, this is not a limitation of the proposed method. Rather, it is a limitation of the current hardware but as the hardware improves and longer circuits, i.e., more Trotter steps, can be reliably executed, smaller time steps can be adopted to improve the accuracy at smaller t_f .

APPENDIX C: ERROR MITIGATION AND SUPPRESSION

Error suppression refers to methods that actively reduce errors at the circuit level, while error mitigation refers to methods that invert errors in postprocessing. We employ some of the most established techniques and provide brief introductions to each one of them in this Appendix. More detailed descriptions can be found in the respective references.

1. Readout error mitigation

We employ two methods to mitigate readout errors. When expectation values are evaluated directly, i.e., through QISKIT’s Estimator primitive without evaluating single-measurement bitstrings, twirled readout error mitigation (TRES) is used [78]. TRES twirls the readout with X gates and averages a measured expectation value over these twirls. It can therefore only mitigate readout errors in expectation values, and not on individual bitstrings. That

is why, for optimization problems, where individual bitstrings are sampled and processed (with QISKIT’s Sampler primitive), we use matrix-free measurement mitigation (M3) instead of TRES [79].

2. Dynamical decoupling

DD is an error suppression technique that inserts quantum gate sequences during qubit idle times. DD sequences are designed to remove certain system-environment interactions. There exist a plethora of possible DD sequences canceling different errors or canceling errors to different orders [109]. In all experiments with DD, we used a staggered XY4 sequence [80], $Y_0\tau Y_1\tau X_0\tau X_1\tau Y_0\tau Y_1\tau X_0\tau X_1\tau$, with Y_i acting on qubit i and τ being a delay duration. The XY4 = $Y\tau X\tau Y\tau X\tau$ sequence universally cancels all single-qubit interactions up to first order. By staggering them, i.e., introducing a relative shift between the sequences on neighboring qubits, we also cancel static ZZ cross-talk [80].

3. Pulse-efficient transpilation

By default, parameterized two-qubit gates such as $R_{ZZ}(\theta)$ are transpiled to a parameterized single-qubit $R_Z(\theta)$ sandwiched between two maximally entangling, hardware-native two-qubit gates. However, if the backend exposes to the user ECR gates as underlying or being its native two-qubit gate and allows for pulse-level access, such gates can be transpiled to significantly shorter final pulse sequences than obtained with standard transpilation. This leads to an overall reduction of the circuit duration. Details can be found in Ref. [81] but, in essence, this is achieved by scaling the area of the ECR pulses as a function of the gate parameter θ . In our application, having time-dependent annealing schedules, two-qubit gate parameters start off being small at the beginning of our Trotter circuit and grow towards the end. For our circuits, pulse-efficient transpilation yields a reduction in pulse duration of approximately 40% on average.

When the gate angle and the corresponding pulse area become so small that not only the pulse width but also its amplitude are adjusted, a fine amplitude calibration becomes necessary [83]. This is due to nonlinearities in the cross-resonance gate with respect to pulse amplitude.

APPENDIX D: HARDWARE PROPERTIES AND QUBIT SELECTION

For completeness, we report the decoherence times T_1 and T_2 , the two-qubit gate error, and the readout error as reported for the quantum processors on which we execute the circuits, see Fig. 8. These properties were accessed on February 13, 2024 for `ibm_torino`, on February 22, 2024 for `ibm_sherbrooke`, and on November 6, 2023 for `ibm_auckland`. They are indicative of the device’s




FIG. 8. Backend properties. Cumulative distribution functions of the decoherence times T_1 and T_2 , two-qubit gate and readout errors for the three devices used in this work, *ibm_torino*, *ibm_sherbrooke*, and *ibm_auckland*.

performance of when the corresponding experiments in the main text were executed, even though the data reported in the main text were gathered on several different days. The 100-qubit line in Fig. 2 was chosen by computing the cumulative two-qubit gate error along all 100-qubit lines on the respective processor and choosing the one with the smallest error.

APPENDIX E: QUANTUM ANNEALING WITH VERY LARGE TIME STEPS

Here, we confirm that the optimal time step for digitized QA for optimization is $\Delta t > 1$ through ideal state-vector simulations. Figure 9 shows the results for the same systems as in Fig. 4 from the main text, that is, instances of disordered couplings $J_{ij} \in [-1, 1]$. When considering a fixed number of time steps or circuit layers, the minimum residual energy obtained from up to 30 time steps decreases with growing time-step size up to an optimal time step $1.2 < \Delta t < 1.4$ (depending on the system), before sharply increasing, see Fig. 9. Increasing the time




FIG. 9. State-vector simulations of the residual energy dependence on time step and spectral gap. (a) Residual energy obtained from state-vector simulations of QA with fixed time steps $\Delta t \in \{0.1, \dots, 5.0\}$ as a function of circuit depth, i.e., number of time steps. Each row corresponds to the respective spectrum in Fig. 4. (b) Minimum residual energy from (a) as a function of the time step with fixed depth, i.e., each point corresponds to the minimum over one curve in (a).

step even further does not yield any benefit whatsoever and, since it induces significant algorithmic errors, results in randomly fluctuating residual energies with increasing circuit depth.

APPENDIX F: KINK-KINK CORRELATOR

We investigate the correlation between different defects in spin chains after digitized quantum annealing, also termed kink-kink correlation in the literature [31, 110]. Given uniform couplings J , the solution after annealing is ferromagnetic. In this setting, defects are misalignments of spins on edges i between lattice sites, and the correlator between defects i and $i + r$ is defined as

$$C_r^{\text{KK}} = \frac{1}{N_{e,r}} \sum_{i=1}^{N_{e,r}} \frac{\langle K_i K_{i+r} \rangle - n_{\text{def}}^2}{n_{\text{def}}^2}, \quad (\text{F1})$$

where $K_i = 1 - \sigma_i^z \sigma_{i+1}^z$ measures whether there is a defect, i.e., a spin-flip, between sites i and $i + 1$, and $N_{e,r}$ is the number of edges on the graph between edges i and $i + r$.

Figure 10 shows the density of defects (left) of annealing a periodic 12-qubit chain on *ibm_auckland* (a) using

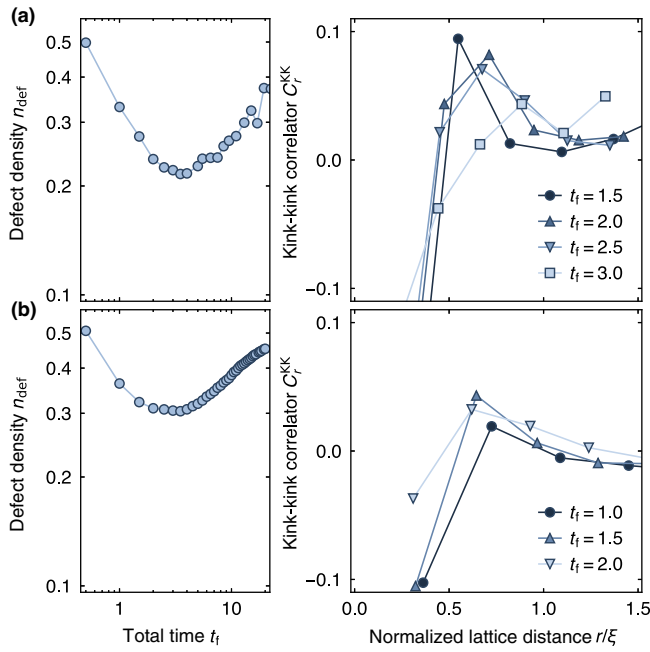


FIG. 10. Kink-kink correlation functions for different final annealing times t_f . (a) Density of defects (left) for a periodic 12-qubit chain on `ibm_auckland` using REM and pulse-efficient transpilation and corresponding kink-kink correlators (right). (b) Density of defects (left) of an open 100-qubit chain on `ibm_sherbrooke` using no EMS and corresponding kink-kink correlators (right).

REM and pulse-efficient transpilation and of a 100-qubit chain on `ibm_sherbrooke` (b) using no EMS. The corresponding kink-kink correlators for different t_f are shown in the respective right panel as a function of the normalized lattice distance r/ξ with $1/\xi = n_{\text{def}}$. The existence of a positive peak, which we observe between $r/\xi = 0.5$ and 1.0 , is expected from theoretical results [69,82,111].

[1] H. E. Stanley, *Phase Transitions and Critical Phenomena* (Clarendon Press, Oxford, 1971), Vol. 7.
 [2] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, Cambridge, 2011), 2nd ed.
 [3] S. L. Sondhi, S. M. Girvin, J. P. Carini, and D. Shahar, Continuous quantum phase transitions, *Rev. Mod. Phys.* **69**, 315 (1997).
 [4] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
 [5] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153 (2014).
 [6] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, K.-M. C. Fu, *et al.*, Quantum simulators: Architectures and opportunities, *PRX Quantum* **2**, 017003 (2021).
 [7] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M.

Linke, G. Pagano, *et al.*, Programmable quantum simulations of spin systems with trapped ions, *Rev. Mod. Phys.* **93**, 025001 (2021).
 [8] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, Cold bosonic atoms in optical lattices, *Phys. Rev. Lett.* **81**, 3108 (1998).
 [9] M. Lewenstein, A. Sanpera, and V. Ahufinger, *Ultracold Atoms in Optical Lattices: Simulating Quantum Many-Body Systems* (OUP Oxford, Oxford, 2012).
 [10] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
 [11] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, *et al.*, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, *Nature* **595**, 233 (2021).
 [12] A. Miessen, P. J. Ollitrault, F. Tacchino, and I. Tavernelli, Quantum algorithms for quantum dynamics, *Nat. Comput. Sci.* **3**, 25 (2023).
 [13] D. S. Abrams and S. Lloyd, Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors, *Phys. Rev. Lett.* **83**, 5162 (1999).
 [14] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
 [15] X. Mi, M. Ippoliti, C. Quintana, A. Greene, Z. Chen, J. Gross, F. Arute, K. Arya, J. Atalaya, R. Babbush, *et al.*, Time-crystalline eigenstate order on a quantum processor, *Nature* **601**, 531 (2022).
 [16] N. Keenan, N. F. Robertson, T. Murphy, S. Zhuk, and J. Goold, Evidence of Kardar-Parisi-Zhang scaling on a digital quantum simulator, *npj Quantum Inf.* **9**, 72 (2023).
 [17] A. Miessen, P. J. Ollitrault, and I. Tavernelli, Quantum algorithms for quantum dynamics: A performance study on the spin-boson model, *Phys. Rev. Res.* **3**, 043212 (2021).
 [18] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
 [19] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, *Nature* **607**, 667 (2022).
 [20] A. Abbas, *et al.*, Quantum optimization: Potential, challenges, and the path forward, *ArXiv:2312.02279*.
 [21] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
 [22] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
 [23] S. Knysh, Zero-temperature quantum annealing bottlenecks in the spin-glass phase, *Nat. Commun.* **7**, 12370 (2016).
 [24] W. H. Zurek, U. Dorner, and P. Zoller, Dynamics of a quantum phase transition, *Phys. Rev. Lett.* **95**, 105701 (2005).
 [25] H.-B. Zeng, C.-Y. Xia, and A. del Campo, Universal breakdown of Kibble-Zurek scaling in fast quenches

- across a phase transition, *Phys. Rev. Lett.* **130**, 060402 (2023).
- [26] T. W. Kibble, Some implications of a cosmological phase transition, *Phys. Rep.* **67**, 183 (1980).
- [27] W. H. Zurek, Cosmological experiments in superfluid helium?, *Nature* **317**, 505 (1985).
- [28] Y. Bando, Y. Susa, H. Oshiyama, N. Shibata, M. Ohzeki, F. J. Gómez-Ruiz, D. A. Lidar, S. Suzuki, A. del Campo, and H. Nishimori, Probing the universality of topological defect formation in a quantum annealer: Kibble-Zurek mechanism and beyond, *Phys. Rev. Res.* **2**, 033369 (2020).
- [29] B.-W. Li, Y.-K. Wu, Q.-X. Mei, R. Yao, W.-Q. Lian, M.-L. Cai, Y. Wang, B.-X. Qi, L. Yao, L. He, *et al.*, Probing critical behavior of long-range transverse-field Ising model through quantum Kibble-Zurek mechanism, *PRX Quantum* **4**, 010302 (2023).
- [30] A. Keesling, A. Omran, H. Levine, H. Bernien, H. Pichler, S. Choi, R. Samajdar, S. Schwartz, P. Silvi, S. Sachdev, P. Zoller, M. Endres, M. Greiner, V. Vuletic, and M. D. Lukin, Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator, *Nature* **568**, 207 (2019).
- [31] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, F. Altomare, A. J. Berkley, S. Ejtemaee, E. Hoskinson, S. Huang, *et al.*, Coherent quantum annealing in a programmable 2000 qubit Ising chain, *Nat. Phys.* **18**, 1324 (2022).
- [32] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, *et al.*, Quantum critical dynamics in a 5000-qubit programmable spin glass, *Nature* **617**, 61 (2023).
- [33] 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* **595**, 227 (2021).
- [34] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits, *J. Appl. Phys.* **132**, 160902 (2022).
- [35] A. C. Vazquez, C. Tornow, D. Riste, S. Woerner, M. Takita, and D. J. Egger, Scaling quantum computing with dynamic circuits, *ArXiv:2402.17833*.
- [36] 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).
- [37] O. Shtanko, D. S. Wang, H. Zhang, N. Harle, A. Seif, R. Movassagh, and Z. Mineev, Uncovering local integrability in quantum many-body dynamics, *ArXiv:2307.07552*.
- [38] M. Amico, H. Zhang, P. Jurcevic, L. S. Bishop, P. Nation, A. Wack, and D. C. McKay, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Bellevue, WA, USA, 2023), p. 692.
- [39] E. Magesan, J. M. Gambetta, and J. Emerson, Characterizing quantum gates via randomized benchmarking, *Phys. Rev. A* **85**, 042311 (2012).
- [40] T. Proctor, S. Seritan, K. Rudinger, E. Nielsen, R. Blume-Kohout, and K. Young, Scalable randomized benchmarking of quantum computers using mirror circuits, *Phys. Rev. Lett.* **129**, 150502 (2022).
- [41] J. Hines, D. Hothem, R. Blume-Kohout, B. Whaley, and T. Proctor, Fully scalable randomized benchmarking without motion reversal, *PRX Quantum* **5**, 030334 (2024).
- [42] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [43] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, Validating quantum computers using randomized model circuits, *Phys. Rev. A* **100**, 032328 (2019).
- [44] D. C. McKay, I. Hincks, E. J. Pritchett, M. Carroll, L. C. G. Govia, and S. T. Merkel, Benchmarking quantum processor performance at scale, *ArXiv:2311.05933*.
- [45] T. Lubinski, S. Johri, P. Varosy, J. Coleman, L. Zhao, J. Necaie, C. H. Baldwin, K. Mayer, and T. Proctor, Application-oriented performance benchmarks for quantum computing, *IEEE Trans. Quantum Eng.* **4**, 1 (2023).
- [46] D. Wu, R. Rossi, F. Vicentini, N. Astrakhantsev, F. Becca, X. Cao, J. Carrasquilla, F. Ferrari, A. Georges, M. Hibat-Allah, *et al.*, Variational benchmarks for quantum many-body problems, *ArXiv:2302.04919*.
- [47] G. C. Santra, F. Jendrzejewski, P. Hauke, and D. J. Egger, Squeezing and quantum approximate optimization, *Phys. Rev. A* **109**, 012413 (2024).
- [48] V. Zhang and P. D. Nation, Characterizing quantum processors using discrete time crystals, *ArXiv:2301.07625*.
- [49] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *ArXiv:1411.4028*.
- [50] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices, *Phys. Rev. X* **10**, 021067 (2020).
- [51] E. Farhi, D. Gamarnik, and S. Gutmann, The quantum approximate optimization algorithm needs to see the whole graph: A typical case, *ArXiv:2004.09002*.
- [52] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, Quantum versus classical annealing of Ising spin glasses, *Science* **348**, 215 (2015).
- [53] A. P. Young, S. Knysh, and V. N. Smelyanskiy, First-order phase transition in the quantum adiabatic algorithm, *Phys. Rev. Lett.* **104**, 020502 (2010).
- [54] T. Jörg, F. Krzakala, G. Semerjian, and F. Zamponi, First-order transitions and the performance of quantum algorithms in random optimization problems, *Phys. Rev. Lett.* **104**, 207206 (2010).
- [55] K. Du, X. Fang, C. Won, C. De, F.-T. Huang, W. Xu, H. You, F. J. Gómez-Ruiz, A. del Campo, and S.-W. Cheong, Kibble-Zurek mechanism of Ising domains, *Nat. Phys.* **19**, 1495 (2023).
- [56] M. Schmitt, M. M. Rams, J. Dziarmaga, M. Heyl, and W. H. Zurek, Quantum phase transition dynamics in the two-dimensional transverse-field Ising model, *Sci. Adv.* **8**, 6850 (2022).
- [57] A. Chandran, A. Erez, S. S. Gubser, and S. L. Sondhi, Kibble-Zurek problem: Universality and the scaling limit, *Phys. Rev. B* **86**, 064304 (2012).

- [58] S. Morita and H. Nishimori, Mathematical foundation of quantum annealing, *J. Math. Phys.* **49**, 125210 (2008).
- [59] T. Caneva, R. Fazio, and G. E. Santoro, Adiabatic quantum dynamics of the Lipkin-Meshkov-Glick model, *Phys. Rev. B* **78**, 104426 (2008).
- [60] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, Quantum simulation of time-dependent hamiltonians and the convenient illusion of Hilbert space, *Phys. Rev. Lett.* **106**, 170501 (2011).
- [61] 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).
- [62] A. Rajput, A. Roggero, and N. Wiebe, Hybridized methods for quantum simulation in the interaction picture, *Quantum* **6**, 780 (2022).
- [63] A. Dutta, A. Rahmani, and A. del Campo, Anti-Kibble-Zurek behavior in crossing the quantum critical point of a thermally isolated system driven by a noisy control field, *Phys. Rev. Lett.* **117**, 080402 (2016).
- [64] L. Arceci, S. Barbarino, D. Rossini, and G. E. Santoro, Optimal working point in dissipative quantum annealing, *Phys. Rev. B* **98**, 064307 (2018).
- [65] J. J. Mayo, Z. Fan, G.-W. Chern, and A. del Campo, Distribution of kinks in an Ising ferromagnet after annealing and the generalized Kibble-Zurek mechanism, *Phys. Rev. Res.* **3**, 033150 (2021).
- [66] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, Evidence for quantum annealing with more than one hundred qubits, *Nat. Phys.* **10**, 218 (2014).
- [67] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babush, V. Smelyanskiy, J. Martinis, and H. Neven, What is the computational value of finite-range tunneling?, *Phys. Rev. X* **6**, 031015 (2016).
- [68] S. V. Isakov, G. Mazzola, V. N. Smelyanskiy, Z. Jiang, S. Boixo, H. Neven, and M. Troyer, Understanding quantum tunneling through quantum Monte Carlo simulations, *Phys. Rev. Lett.* **117**, 180402 (2016).
- [69] J. Dziarmaga and M. M. Rams, Kink correlations, domain-size distribution, and emptiness formation probability after a Kibble-Zurek quench in the quantum Ising chain, *Phys. Rev. B* **106**, 014309 (2022).
- [70] M. Motta and J. E. Rice, Emerging quantum computing algorithms for quantum chemistry, *WIREs Comput. Mol. Sci.* **12**, e1580 (2022).
- [71] D. Layden, G. Mazzola, R. V. Mishmash, M. Motta, P. Wocjan, J.-S. Kim, and S. Sheldon, Quantum-enhanced Markov chain Monte Carlo, *Nature* **619**, 282 (2023).
- [72] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, *et al.*, Quantum computing with Qiskit, [ArXiv:2405.08810](https://arxiv.org/abs/2405.08810).
- [73] IBM Quantum, <https://quantum-computing.ibm.com/>, 2021.
- [74] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, *Phys. Rev. B* **81**, 134507 (2010).
- [75] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, *Phys. Rev. A* **93**, 060302 (2016).
- [76] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, Universal gate for fixed-frequency qubits via a tunable bus, *Phys. Rev. Appl.* **6**, 064007 (2016).
- [77] M. Ganzhorn, G. Salis, D. J. Egger, A. Fuhrer, M. Mergenthaler, C. Müller, P. Müller, S. Paredes, M. Pechal, M. Werninghaus, and S. Filipp, Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform, *Phys. Rev. Res.* **2**, 033447 (2020).
- [78] 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).
- [79] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, Scalable mitigation of measurement errors on quantum computers, *PRX Quantum* **2**, 040326 (2021).
- [80] S. V. Barron, D. J. Egger, E. Pelofske, A. Bärttschi, S. Eidenbenz, M. Lehmkuehler, and S. Woerner, Provable bounds for noise-free expectation values computed from noisy samples, [ArXiv:2312.00733](https://arxiv.org/abs/2312.00733).
- [81] N. Earnest, C. Tornow, and D. J. Egger, Pulse-efficient circuit transpilation for quantum applications on cross-resonance-based hardware, *Phys. Rev. Res.* **3**, 043088 (2021).
- [82] R. J. Nowak and J. Dziarmaga, Quantum Kibble-Zurek mechanism: Kink correlations after a quench in the quantum Ising chain, *Phys. Rev. B* **104**, 075448 (2021).
- [83] A. C. Vazquez, D. J. Egger, D. Ochsner, and S. Woerner, Well-conditioned multi-product formulas for hardware-friendly hamiltonian simulation, *Quantum* **7**, 1067 (2023).
- [84] N. Astrakhantsev, G. Mazzola, I. Tavernelli, and G. Carleo, Phenomenological theory of variational quantum ground-state preparation, *Phys. Rev. Res.* **5**, 033225 (2023).
- [85] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [86] A. Melo, N. Earnest-Noble, and F. Tacchino, Pulse-efficient quantum machine learning, *Quantum* **7**, 1130 (2023).
- [87] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* **1**, 403 (2021).
- [88] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Defining and detecting quantum speedup, *Science* **345**, 420 (2014).
- [89] T. Albash and D. A. Lidar, Demonstration of a scaling advantage for a quantum annealer over simulated annealing, *Phys. Rev. X* **8**, 031016 (2018).
- [90] C. Mc Keever and M. Lubasch, Towards adiabatic quantum computing using compressed quantum circuits, *PRX Quantum* **5**, 020362 (2024).
- [91] S. H. Sack and D. J. Egger, Large-scale quantum approximate optimization on nonplanar graphs with machine learning noise mitigation, *Phys. Rev. Res.* **6**, 013223 (2024).

- [92] S. Hadfield, Z. Wang, B. O’Gorman, E. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator ansatz, *Algorithms* **12**, 34 (2019).
- [93] G. Scrivera, N. Astrakhantsev, S. Pilati, and G. Mazzola, Challenges of variational quantum optimization with measurement shot noise, *Phys. Rev. A* **109**, 032408 (2024).
- [94] S. H. Sack and M. Serbyn, Quantum annealing initialization of the quantum approximate optimization algorithm, *Quantum* **5**, 491 (2021).
- [95] R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. L. Heras, R. Babbush, A. G. Fowler, B. Campbell, Y. Chen, *et al.*, Digitized adiabatic quantum computing with a superconducting circuit, *Nature* **534**, 222 (2016).
- [96] G. E. Santoro, R. Martonák, E. Tosatti, and R. Car, Theory of quantum annealing of an Ising spin glass, *Science* **295**, 2427 (2002).
- [97] IBM ILOG CPLEX Optimizer.
- [98] `ibm_auckland` is one of IBM’s by now retired 27-qubit Falcon chips.
- [99] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, *et al.*, Quantum optimization of maximum independent set using Rydberg atom arrays, *Science* **376**, 1209 (2022).
- [100] J. Weidenfeller, L. C. Valor, J. Gacon, C. Tornow, L. Bello, S. Woerner, and D. J. Egger, Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware, *Quantum* **6**, 870 (2022).
- [101] G. Mazzola, Quantum computing for chemistry and physics applications from a Monte Carlo perspective, *J. Chem. Phys.* **160**, 010901 (2024).
- [102] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Scalable circuits for preparing ground states on digital quantum computers: The Schwinger model vacuum on 100 qubits, *PRX Quantum* **5**, 020315 (2024).
- [103] E. Pelofske, A. Bäertschi, L. Cincio, J. Golden, and S. Eidenbenz, Scaling whole-chip QAOA for higher-order Ising spin glass models on heavy-hex graphs, [ArXiv:2312.00997](https://arxiv.org/abs/2312.00997).
- [104] T. A. Chowdhury, K. Yu, M. A. Shamim, M. L. Kabir, and R. S. Sufian, Enhancing quantum utility: Simulating large-scale quantum spin chains on superconducting quantum computers, *Phys. Rev. Res.* **6**, 033107 (2024).
- [105] D. Rehfeldt, T. Koch, and Y. Shinano, Faster exact solution of sparse maxcut and QUBO problems, *Math. Program. Comput.* **15**, 445 (2023).
- [106] D. J. Egger, J. Mareček, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
- [107] R. Tate, M. Farhadi, C. Herold, G. Mohler, and S. Gupta, Bridging classical and quantum with SDP initialized warm-starts for QAOA, *ACM Trans. Quantum Comput.* **4**, 1 (2023).
- [108] P. Chandarana, N. N. Hegade, K. Paul, F. Albarrán-Arriagada, E. Solano, A. del Campo, and X. Chen, Digitized-counterdiabatic quantum approximate optimization algorithm, *Phys. Rev. Res.* **4**, 013141 (2022).
- [109] N. Ezzell, B. Pokharel, L. Tewala, G. Quiroz, and D. A. Lidar, Dynamical decoupling for superconducting qubits: A performance survey, *Phys. Rev. Appl.* **20**, 064027 (2023).
- [110] L. Cincio, J. Dziarmaga, M. M. Rams, and W. H. Zurek, Entropy of entanglement and correlations induced by a quench: Dynamics of a quantum phase transition in the quantum Ising model, *Phys. Rev. A* **75**, 052321 (2007).
- [111] K. Roychowdhury, R. Moessner, and A. Das, Dynamics and correlations at a quantum phase transition beyond Kibble-Zurek, *Phys. Rev. B* **104**, 014406 (2021).