

Resource Assessment of Classical and Quantum Hardware for Post-Quench Dynamics

Joseph Vovrosh^{✉,*}, Tiago Mendes-Santos, Hadriel Mamann, Kemal Bidzhiev, Fergus Hayes[✉],
Bruno Ximenez[✉], Lucas Béguin, Constantin Dalyac, and Alexandre Dauphin^{✉,†}
PASQAL SAS, 24 rue Emile Baudot-91120 Palaiseau, Paris, France

 (Received 8 December 2025; revised 5 June 2026; accepted 22 June 2026; published 15 July 2026)

We estimate the run-time and energy consumption of simulating nonequilibrium dynamics on neutral-atom quantum computers in analog mode, directly comparing their performance to classical methods, namely matrix product states and neural quantum states. By collecting both experimental data from a quantum processing unit (QPU) in analog mode and numerical benchmarks, we enable accurate predictions of run-time and energy consumption for large-scale simulations on both QPUs and classical systems through fitting of theoretical scaling laws. Our analysis shows that neutral-atom devices are already operating in a competitive regime, achieving comparable or superior performance to classical approaches while consuming significantly less energy. These results demonstrate the potential of analog neutral-atom quantum computing for energy-efficient simulation and highlight a viable path toward sustainable computational strategies.

DOI: [10.1103/kfy5-q531](https://doi.org/10.1103/kfy5-q531)

I. INTRODUCTION

The development in the late 1980s of new algorithms tailored for quantum computers marked the beginning of a new interdisciplinary field at the intersection of quantum physics, information theory, and computer science. Quantum computing's original promise, grounded in complexity theory, stemmed from the discovery of algorithms offering exponential speedups over classical methods and from evidence that quantum computers may be able to solve problems that are intractable for nondeterministic classical algorithms [1,2]. With recent technological advances demonstrating key building blocks for fault-tolerant quantum computers, more systematic resource estimations have been performed for various applications to evaluate their utility [3–9].

These results, however, describe asymptotic regimes for fault-tolerant quantum computers. In contrast, current quantum processors operate at the scale of hundreds of qubits, with a limited circuit depth, and with noise, far from the asymptotic limit where such separations are expected to appear. It is therefore timely to evaluate today's devices through resource-based metrics, such as run-time efficiency and energetic cost on specific problems. Beyond

computational metrics, energy consumption has become an increasingly important dimension of evaluation [10,11]. This is particularly relevant given that classical methods, such as large-scale AI or high-performance scientific computing, are demanding ever-growing amounts of energy [12–14]. As computing contributes a growing share of global energy consumption, energy-to-solution benchmarks for quantum processors are needed to evaluate whether they can provide not only computational advantages but also more sustainable pathways to scaling [11].

This reframing shifts the focus toward identifying concrete regimes in which existing hardware can already demonstrate practical benefits over classical approaches. This strategy is also adopted in two recent independent works [15,16] proposing definitions of quantum advantage. In particular, the authors use concepts such as quantum utility [15] or empirical quantum advantage [16]. For example, a $100\times$ speedup in run-time with respect to classical computers with the same accuracy means that computations that currently take months on a classical computer could be done in a day on a quantum computer. Given a limited time budget, this speedup would allow a broader scan of the parameter space.

While such benchmarks have been developed for superconducting [17,18], trapped-ion [19,20], and photonic platforms [21,22], neutral-atom quantum processing units (QPUs) remain comparatively not yet much explored [23–25]. This gap is especially relevant for current devices operated in analog mode [26], where precisely controlled Rydberg interactions among trapped atoms provide a natural substrate for quantum simulation [27]. These systems have already enabled advances in probing ground-state

* Contact author: joseph.vovrosh@pasqal.com

† Contact author: alexandre.dauphin@pasqal.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.




FIG. 1. Resource estimation for simulating post-quench dynamics of an Ising Hamiltonian, showing both computational time and energy consumption on classical and quantum hardware. The classical data are measured from actual simulations, while the QPU data are estimated based on performance data. The figure highlights the crossover point where classical simulations become less efficient and quantum hardware is expected to offer an advantage.

physics [28–31], exploring nonequilibrium dynamics [32–36], and tackling graph-based problems [37], including optimization [38,39] and machine learning [40–42] applications.

In this work, we address this gap by analyzing neutral-atom analog QPUs through the lens of resource-based benchmarking. Rather than contrasting classical and analog approaches in general terms, we focus on identifying a concrete problem relevant for quantum materials and a parameter regime where the QPU may achieve advantages over state-of-the-art classical methods, whether in run-time performance or energy efficiency.

For the classical baseline in our comparison, we consider matrix product state (MPS)-based methods, as they represent the most widely used and well-developed numerical approach for simulating quantum systems of medium size, systems consisting of tens of qubits [43–45]. MPS techniques have a long track record of accuracy and efficiency for medium-sized systems, and their performance characteristics are well understood, making them a suitable benchmark for assessing quantum hardware.

As an example of an emergent classical method, we also include neural quantum states (NQS) [46], which have recently demonstrated comparable performance relative to established tensor-network approaches in certain problem settings [47]. Unlike MPS, the practical limits of NQS methods remain less clearly defined, offering an interesting point of comparison as their capabilities continue to be explored.

The computational cost of the QPU remains largely independent of the evolution time (or pulse duration) t , though it scales exponentially with system size N when requiring perfect register preparation. Classical simulation methods considered in this work, MPS and NQS, incur an exponential cost for time evolution due to entanglement growth, despite their underlying polynomial scaling with system size. We find that for sufficiently long evolution times, the QPU is expected to outperform classical methods in both run-time and energy consumption, as quantified by these scalings, including the relevant pre-factors, see Fig. 1.

The remainder of this paper is structured as follows. Section II defines the computational problem addressed in this work. Sections IV–VI compare computation times between employing a neutral-atom analog QPU and classical simulation methods for simulating post-quench dynamics. Sections VII and VIII present an analysis of energy consumption across these methods. Finally, Sec. IX summarizes our findings and outlines directions for future research.

II. PROBLEM CHOICE

We consider an Ising-like Hamiltonian on the square lattice, a problem of relevance for many applications in quantum materials [48]. This problem has been extensively studied both theoretically and experimentally for its ground-state properties in the antiferromagnetic regime [28–31] and for its dynamics that can lead to exotic phenomena such as quantum scars [33], coarsening dynamics [34], confinement and roughening dynamics [49,50], and Kibble-Zurek mechanisms [51]. In particular, Refs. [24,25] explored post-quench dynamics regimes where the distribution of the quantum state in the Z-basis follows a Haar random distribution and showed an advantage in the accuracy of sampling such states on a neutral-atom QPU with respect to their classical counterpart. While the ground-state properties of this stoquastic Hamiltonian can be well captured by classical methods, its dynamics can be very complex for parameter regimes that are nonperturbative [52]. We therefore consider a quench compatible with our Rydberg atom quantum processing unit.

In a typical setup, ^{87}Rb atoms are first loaded into a magneto-optical trap, then transferred to individual optical tweezers generated using spatial light modulators or acousto-optic deflectors and typically separated by a few microns. In the analog mode, the qubit states are

encoded as $|g\rangle = |5S_{1/2}, F=2, m_F=2\rangle \equiv |0\rangle$ and $|r\rangle = |6S_{1/2}, m_J=1/2\rangle \equiv |1\rangle$. These two states form an effective two-level system, allowing us to map the neutral-atom qubits onto spin-1/2 particles. The resulting system is governed by

$$H = \sum_{i<j} V_{ij} n_i n_j + \frac{\Omega}{2} \sum_i \sigma_i^x - \Delta \sum_i n_i, \quad (1)$$

where $\sigma_i^{x/z}$ are the Pauli- x/z operators acting on site i , and $n_i = \frac{1}{2}(1 + \sigma_i^z)$ is the projector onto the Rydberg state of qubit i . Here, Ω is the Rabi frequency of the laser driving the $|0\rangle \leftrightarrow |1\rangle$ transition, and Δ is the detuning of the laser from the Rydberg transition. The interaction term V_{ij} describes the van der Waals interaction between atoms in the Rydberg state, typically scaling as $V_{ij} = C_6/|r_i - r_j|^6$, where C_6 is the interaction coefficient and $|r_i - r_j|$ is the distance between atoms i and j .

This Hamiltonian captures the essential physics of Rydberg-mediated interactions in neutral-atom arrays, providing a versatile platform for exploring strongly correlated quantum dynamics. Throughout this work, we consider the time evolution (post-quench) of the initial $|00\cdots 0\rangle$ state under the time-independent Hamiltonian with parameters $\Omega/2\pi = 2$ MHz, $\Delta = \frac{1}{2} \sum_j V_{0j}$ in which 0 is the index of a central atom in the square lattice with nearest neighbor atomic separation $R = (C_6 h_x / (2 \Omega))^{1/6}$ such that $h_x = 2.5$. These are chosen in accordance with Ref. [52]. We will exclusively consider measurements of the local observables $\langle Z_i \rangle$ and $\langle Z_i Z_j \rangle$ routinely available with neutral-atom hardware.

It is important to note that QPUs are noisy, and this limits their coherence time. While the precise characterization of the noise is beyond the scope of this paper, we would like to emphasize key points. Neutral atoms are noisy, but they are very well-isolated systems, allowing for a deep understanding of dominant noise sources. By controlling the noise, it has been demonstrated that these QPUs can reach regimes where classical methods are not accurate anymore for global quantities such as the fidelity [25]. In this work, we are considering local observables, and in that case, the errors due to the noise remain bounded with the increase in the system size [53–56]. We therefore expect that the combination of experimental control of the noise, the increase in the system size, and some potential error mitigation can allow the measurement of local observables with bounded errors.

III. RESOURCE CONSIDERATIONS FOR CLASSICAL AND QUANTUM METHODS

To enable a meaningful comparison between classical simulations and QPU, it is essential to clearly define the relevant computational resources in each case. In this

work, the comparison is based on the same physical task: the time evolution of the Hamiltonian in Eq. (1) and the estimation of the local observables $\langle Z_i \rangle$ and $\langle Z_i Z_j \rangle$. Both classical and quantum approaches prepare the full many-body wave function, but the cost of extracting information from it differs between the two platforms.

For the QPU, the dominant run-time contributions are the duration of the analog pulse sequence and the number of measurement shots required to estimate expectation values with the desired statistical precision. In the measurement stage, a single experimental shot yields simultaneous site-resolved readout of all atoms in the array. Consequently, all single-site observables $\langle \sigma_i^z \rangle$ and all two-point correlators $\langle \sigma_i^z \sigma_j^z \rangle$ accessible within the measurement basis are obtained in parallel. The measurement overhead, therefore, does not scale with the system size N beyond the statistical scaling associated with finite sampling.

Current neutral-atom devices are subject to several experimental imperfections that ultimately limit the usable evolution time. Typical parameters include finite atomic temperature $T \approx 20$ μ K, dephasing with a coherence time $T_{\text{deph}} \sim 20$ μ s, and shot-to-shot fluctuations in the control parameters, with relative variations of order 1% in the Rabi frequency Ω and detuning noise $\sigma_\delta \approx 0.3$, rad/ μ s. State-preparation and measurement errors are routinely characterized, while calibration-based mitigation procedures are commonly used to partially suppress their effects in neutral-atom analog simulations with negligible additional overhead and therefore do not affect the run-time considerations discussed here [57,58]. In the following, we therefore restrict our analysis to pulse durations that are comparable to those routinely implemented in current neutral-atom experiments and that are known to produce physically meaningful many-body dynamics [35,59,60]. Accordingly, all subsequent resource estimates are understood to hold in a regime of pulse durations and control parameters that remains within experimentally validated operating conditions of current neutral-atom devices, where systematic errors do not qualitatively alter the observed many-body dynamics.

For simulations based on classical methods, the computational cost is governed by the resources required to accurately reproduce the observables of interest to a desired precision. In principle, classical algorithms can always be executed with arbitrarily limited resources at polynomial cost, for example, a tensor network simulation with low bond dimension or a neural quantum state simulation with a minimal number of variational parameters. However, such calculations would generally yield poor approximations to the true dynamics. In order to meaningfully assess the scaling of classical methods, it is therefore necessary to require convergence of the relevant observables. In the following, we impose convergence criteria directly on the local quantities considered in this

work, rather than on more demanding global metrics such as the full many-body fidelity. Explicitly accounting for this convergence ensures that the reported classical computational cost reflects the resources required to obtain physically meaningful results.

This resource framework ensures that both approaches are assessed using consistent physical criteria. The comparison is based on identical Hamiltonian dynamics and identical local observables, while the respective resource requirements reflect the natural scaling limitations of each platform. In the following sections, we describe in detail the classical simulation methods and the experimental implementation of the QPU protocol within this resource model.

IV. RESOURCE ESTIMATION FOR QUANTUM SIMULATION

In this section, we estimate the total time required to simulate the quantum dynamics described in Sec. II using a neutral-atom QPU. Extracting meaningful physical information, such as the expectation value of a local observable, typically requires many shots to build sufficient statistical accuracy.

Atom loading into the tweezer array is probabilistic, with each trap typically filled with a probability of about 50%. To reliably prepare a defect-free register of N atoms, experiments therefore begin with a larger reservoir array containing roughly $2N$ traps. After imaging the occupancy pattern, occupied traps are dynamically rearranged using movable optical tweezers to assemble the desired register. The number of required atom transport operations grows with the size of the target register, so the preparation time increases approximately with the number of atoms. Consequently, the overall repetition rate of the device (shots per second) decreases as the register size increases.

A further complication is that not all shots executed on current quantum hardware yield usable data. In practice, imperfections during initialization or rearrangement can lead to defects, the mechanisms of which are sketched in Fig. 2(a). To address this issue, we model an effective zero-defect rate and derive phenomenologically the scaling behavior of the total QPU simulation time, incorporating contributions from both the required number of defect-free shots and imperfections in the hardware. This allows for a more realistic estimate of the time cost associated with executing the full protocol on a quantum device. The parameters used in the model correspond to experimentally measurable quantities that are routinely characterized on a daily basis. These parameters are continuously monitored over time in order to evaluate and track the rearrangement performance of the QPU.

We begin by estimating the probability that a given shot has no defects. The initial loading of atoms into optical tweezers is intrinsically probabilistic. Because of light-assisted collisions, if two atoms are captured in the same




FIG. 2. (a) Illustration of the different processes required during the rearrangement step. (b) The scaling of repetition rate as a function of system size. (c) Comparison between the mathematical model (dashed line) describing the probability to have zero defects in the register and experimental data (dots), N being the number of atoms in the register. Here, the error bars are given by the standard deviation of the defect-free distribution over time.

trap, both are lost, leading to a maximum filling probability of 50% per site. As a consequence, the initial array of traps (the layout) should be larger than the register we aim to prepare (at least 2 times larger): it contains numerous vacancies, and a rearrangement step is required in order to prepare the desired configuration (the register) in which every site is occupied by a single atom [61,62]. This rearrangement process consists of moving atoms from filled sites to empty ones until the register is fully populated and dumping any additional atoms. The success of this procedure can be quantified by the probability of obtaining a defect-free register, i.e., an array with no missing atoms after rearrangement and removal of unwanted atoms. Several processes contribute to this probability; we now describe them in turn.

The first contribution arises from the transfer of the atoms: each atom that needs to be relocated has a probability p_{transf} of being successfully picked up, transported, and released at its target site. For N_{transf} such operations, the contribution is $(p_{\text{transf}})^{N_{\text{transf}}}$.

The second contribution comes from the imperfect dumping of atoms lying outside the register. Removing such atoms requires successfully picking them up, which occurs with probability $p_{\text{pick-up}}$. For N_{dump} dumps, the success factor is $(p_{\text{pick-up}})^{N_{\text{dump}}}$.

In addition, traps not involved in either transfers or dumps may still suffer from accidental loading of atoms

with probability p_{acci} . Hence, the probability that no accidental events occur in the $N_{\text{traps}} - N_{\text{transf}} - N_{\text{dump}}$ idle traps is $(1 - p_{\text{acci}})^{N_{\text{traps}} - N_{\text{transf}} - N_{\text{dump}}}$.

Finally, atoms within the register may be lost during this process due to collisions with background gas or imperfect detection with probability p_{loss} . For the $N_{\text{register}} - N_{\text{transf}}$ atoms that were not moved, the survival probability is $(1 - p_{\text{loss}})^{N_{\text{register}} - N_{\text{transf}}}$.

Assuming these processes are statistically independent, the total probability of preparing a defect-free register is then given by the product of these contributions

$$P_{\text{defect-free}} = (p_{\text{transf}})^{N_{\text{transf}}} (p_{\text{pick-up}})^{N_{\text{dump}}} \times (1 - p_{\text{acci}})^{N_{\text{traps}} - N_{\text{transf}} - N_{\text{dump}}} \times (1 - p_{\text{loss}})^{N_{\text{register}} - N_{\text{transf}}}. \quad (2)$$

These quantities can be extracted experimentally by comparing pre- and post-rearrangement images while taking into account the moves computed by the rearrangement process predetermined by the Hungarian algorithm [63], which determines the repetition rate of the QPU shown in Fig. 2(b), see Supplemental Material [64] for details.

In order to verify the validity of the phenomenological model introduced above, we perform measurements on our QPU using a layout consisting of 200 traps and varying the register size between 30 and 100 atoms. For each configuration, we extract the relevant observables by comparing the pre- and post-rearrangement images. Averaging over the different register sizes, we find [65]

$$p_{\text{transf}} = (98.9 \pm 0.1)\%, \quad p_{\text{pick-up}} = (99.8 \pm 0.1)\%, \\ p_{\text{acci}} = (0.09 \pm 0.02)\%, \quad p_{\text{loss}} = (0.9 \pm 0.1)\%. \quad (3)$$

Figure 2(c) shows the comparison between the prediction of Eq. (2) and direct measurements of the defect-free probability, obtained by counting the fraction of experimental realizations in which all register sites were occupied, and all traps outside the register were successfully emptied. The excellent agreement between the experimental data and the analytical model demonstrates that Eq. (2) captures the essential processes governing the preparation of defect-free registers in our system.

Given Eq. (2), we are now in a position to estimate the total simulation time required to accurately measure a local observable of a prepared quantum state. Specifically, consider a scenario in which we prepare a quantum state $|\psi(t)\rangle$, where t is the pulse duration and aim to estimate expectation values of the form $\langle \sigma_i^z \rangle$ and $\langle \sigma_i^z \sigma_j^z \rangle$, where i and j label specific qubits in the system. Since both σ_i^z and $\sigma_i^z \sigma_j^z$ have eigenvalues ± 1 , the shot noise in these measurements can be modeled using a binomial distribution. Let p denote the probability of measuring the outcome $+1$. Then, after n measurement shots, the estimated expectation value of an

observable $\hat{O} \in \{\sigma_i^z, \sigma_i^z \sigma_j^z\}$ is given by $\langle \hat{O} \rangle = 2p - 1$, with an associated standard deviation $\sigma_O = 2\sqrt{p(1-p)/n}$.

To achieve a precision of α with 95% confidence, we require that the standard deviation satisfy $\sigma_O \leq \alpha/2$. Solving for the number of required shots under this constraint yields

$$n = \frac{16p(1-p)}{\alpha^2}. \quad (4)$$

In the worst-case scenario, where $p = 0.5$, and for a tolerance $\alpha = 0.05$, we find that $n = 1600$ shots are needed. However, this estimate assumes that all shots are usable, which is generally not the case. As discussed above, imperfections in the filling of the atomic register can lead to defects, such as missing atoms, that render some shots unusable. We now extend our estimate to account for this zero-defect probability.

Assuming the probability of obtaining a defect-free register is given by Eq. (2), and aiming to collect m usable (defect-free) shots, the probability of obtaining at least m such shots out of a total of n attempts follows a binomial distribution

$$\sum_{k=m}^n \binom{n}{k} P_{\text{defect-free}}^k (1 - P_{\text{defect-free}})^{n-k}. \quad (5)$$

With this, given a system size N , we are able to calculate the number of shots required to accurately estimate a local observable and, with the known repetition rate, the corresponding run-time. In the conclusions and outlook, we discuss potential technological improvements to reduce this run-time.

V. RESOURCE ESTIMATION FOR CLASSICAL SIMULATIONS

Several computational approaches have been developed to study the time evolution and ground-state properties of many-body quantum systems. In this work, we consider matrix product state (MPS) based methods, as they are the most widely adopted ansatz for simulating quantum quenches in moderately sized systems [66]. We also consider a time-dependent variational Monte Carlo (tVMC) approach based on neural networks (e.g., restricted Boltzmann machines or Convolutional Neural Networks) to represent many-body wave functions, commonly known as the neural quantum states (NQS) approach [46,47].

We begin by deriving the dominant contributions of computational complexity and run-time scaling for both MPS and NQS simulations. This analytical groundwork provides insight into the expected computational cost as a function of key problem parameters. To validate and calibrate these theoretical expressions, we fit them to empirical data obtained from simulations performed on a

single NVIDIA A100 40 Gb SXM GPU for MPS due to the difficulty of parallelization (see Supplemental Material [64] for a comparison of CPU and GPU run times). For NQS simulations, we perform simulations with multiple of these GPUs in certain cases (we comment more on this below). This approach enables us to generate accurate and practically relevant estimates of the performance of such simulations, which are valid in different regimes of post-quench dynamics.

A. MPS simulations

We analyze the key operations involved in the time evolution of a quantum state in MPS form to estimate the computational complexity. The total evolution is partitioned into n discrete time steps, each executed using the two-site time-dependent variational principle (TDVP) algorithm [67,68], which enables an accurate treatment of the long-range interactions described by Eq. (1). Since each time step involves a similar sequence of operations, the total computational cost can be estimated by analyzing the cost associated with a single representative step and subsequently scaling by n . To quantify the per-step cost, we examine the dominant computational operations required to propagate the state. Two contributions dominate the run-time: the construction of the effective environments [69,70] (also referred to as the left and right “baths”) and the application of the effective two-site Hamiltonian using the Lanczos algorithm [71].

The construction of the environment requires contracting the MPS tensors with the matrix product operator (MPO) representation of the Hamiltonian in Eq. (1). In a two-dimensional square lattice, each bulk qubit is most strongly coupled to its four nearest neighbors—two along the horizontal and two along the vertical directions—while additional weaker couplings to more distant sites arise from the long-range nature of the interactions. When mapped to a 1D MPO representation of the Hamiltonian, the two vertical couplings correspond to strong quasi-long-range interactions. As a result, for a square lattice of N qubits, the MPO bond dimension scales with system size as $h = \mathcal{O}(\sqrt{N})$, see Supplemental Material [64] for further details.

It follows that bath construction or update involves tensor contractions between the bath and an MPS tensor, contributing terms that scale as $\mathcal{O}(\sqrt{N}\chi^3)$, and contractions between the resulting intermediate tensor and an MPO tensor, which scale as $\mathcal{O}(N\chi^2)$. Moreover, since the left and right baths must each be computed approximately N times, the overall expected complexity is $\mathcal{O}(N^{3/2}\chi^3 + N^2\chi^2)$.

Following bath construction, the application of the effective two-site Hamiltonian involves performing a Krylov subspace evolution of the two-site tensor using the Lanczos algorithm. This process is typically iterated k times per step and includes additional tensor contractions and a singular value decomposition to restore the canonical

MPS form. These operations produce contributions with similar scaling, such that the total cost of one TDVP step can be modeled as the sum of both dominant terms. Aggregating all contributions and absorbing constant prefactors and repeated Krylov steps into fitted coefficients, the expected run-time per TDVP step can be expressed as

$$\langle \Delta t \rangle \approx \alpha N^{3/2} \chi^3 + \beta N^2 \chi^2, \quad (6)$$

where α and β are constants determined by hardware characteristics and algorithmic implementation details. The first term dominates in regimes where the entanglement, and hence the bond dimension χ , is large, while the second term becomes significant in systems with modest entanglement but large system size N . This model provides a predictive framework for assessing the computational demands of MPS simulations and for informing optimization strategies, such as bond dimension truncation or interaction locality adjustments, to improve run-time efficiency.

B. NQS simulations

We now discuss the expected system-size scaling of the complexity (and the corresponding run-time) of the NQS approach. More details about the method are provided in the Supplemental Material, Sec. III [64] and in Refs. [47,72,73]. Here, instead, we summarize the key aspects of the tVMC algorithm that determine its computational complexity.

Similarly to tensor network approaches, the NQS aims to represent the amplitudes of quantum states, parametrized by a set of variational parameters, $\vec{\theta} = (\theta_1, \theta_2, \dots, \theta_{N_v})$, taking the form of a neural network, with N_v denoting the total number of variational parameters. The dynamics of the system is described by updating the set $\{\theta_k\}$ at each time step by solving the first-order differential equation $S_{k,k'} \dot{\theta}_{k'} = -iF_k$, where the matrix S (called the Fisher matrix) has dimension $N_v \times N_v$, and F (the force vector) has dimension N_v .

The complexity of NQS simulations is thus related to the computational cost of evaluation (i) S and F , and (ii) the required inversion of the S matrix to solve the differential equation.

On the one hand, the gradients and the local energy defining S and F is evaluated by Monte Carlo sampling. For this, one must generate N_{MC} samples using the Metropolis Algorithm. For each MC sample, a forward-pass evaluation of the neural network is required, whose complexity $\mathcal{C}_s$ is characterized by the number of floating-point operations (FLOPs). To leading order, the total complexity related to this part is then given by $\mathcal{O}(N_{\text{MC}}N) \times \mathcal{C}_s$, where we assume that for each MC sample, one typically consider N local updates. Other aspects can also affect the complexity. For instance, when the Hamiltonian contains a number of




FIG. 3. Run-time of a single 1-ns time step ($dtJ \sim 0.01$, in which J is the energy scale of the system) for a quantum quench to the Ising model, measured on an NVIDIA A100 GPU, as a function of system size. Results are shown for both MPS with varying bond dimensions in teal and NQS with different types of networks (containing varying numbers of parameters) in purple. Theoretical scaling fits are overlaid and show good agreement with the data; see Supplemental Material at [64] for details.

nondiagonal terms that scale as $\mathcal{O}(N^\alpha)$ (with $\alpha > 1$), additional overhead is introduced to access the local energy.

On the other hand, the complexity of inverting the matrix S scales as $\mathcal{O}(N_v^3)$.

In the simulations considered here, the complexity of part (i) is typically the dominant contribution, since we work in the regime $N_{MC} \gg N_v$. To leading order, the number of FLOPs required to perform a single time step with the tVMC approach thus scales as $\mathcal{O}(N_{MC}N) \times \mathcal{C}_s$. Assuming that the number of variational parameters scales as $N_v = \alpha \mathcal{O}(N)$, where α is a parameter that depends on the particular architecture of the NQS, and taking into account that for shallow convolutional neural network (CNN) or restricted Boltzman machine (RBM) $\mathcal{C}_s = \mathcal{O}(N^2)$, we obtain that the complexity of the NQS algorithm follows the scaling: $\mathcal{O}(\alpha N_{MC}N^3)$.

It is worth mentioning that the actual run-time of the simulations depends on additional factors that are not directly captured by the complexity analysis above. Examples include the batch size of samples used during the neural network forward pass and the number of GPUs available to perform independent Monte Carlo sampling in parallel; in the latter case, one often observes an approximately linear speedup with the number of GPUs [72]. As we discuss in the following section, empirical results for the run-time of tVMC simulations with shallow-CNN or RBM architectures are well described by a polynomial function of N of degree 3.

C. Results

In Fig. 3, we present simulation results of the quench protocol described in Sec. II, exploring how the computational complexity scales for both MPS and NQS methods. Specifically, we report the average run-time required to compute a single time step, with a resolution of approximately $dt = 1$ ns, for each method across varying system parameters. Note, this is such that $dtJ \sim 0.01$ where $J = C_4/4r^6$ is the energy scale of the system, in which r is the nearest neighbor separation in the square lattice. Superimposed on these data are the corresponding theoretical fits derived previously, which show remarkably good agreement with the observed run times, lending strong support to the validity of our scaling analysis and allowing for predictions beyond what is possible with currently available hardware.

However, before we can directly compare classical and QPU simulation times, it is important to note that classical methods do not always yield exact results, their accuracy depending on the complexity of the quantum state being simulated. Therefore, in the following section, we introduce convergence criteria to determine the minimum classical resources required to achieve reliable results.

VI. SCALING OF CONVERGED CLASSICAL SIMULATIONS

To assess the validity of both MPS and NQS simulations, we employ independent convergence criteria tailored to each method. For MPS, we rely on physically motivated indicators: convergence is assumed when the total energy of the system is conserved and the underlying symmetries of the lattice, such as translational invariance or spin conservation, are preserved throughout the simulation. These criteria ensure that the MPS faithfully captures the essential features of the quantum state. For NQS, we consider estimates of the TDVP error.

In this section, we shift our focus from the simulation time required for a single small time step to the total time needed to simulate a quench protocol over a given pulse duration. That is, rather than considering the cost of advancing the wave function by $dt \approx 1$ ns, we evaluate the cumulative computational resources required to prepare the full wave function and measure an observable of interest at varying times t after the quench. This approach allows us to capture the scaling of the total simulation effort as a function of the evolution time, providing a more direct comparison with the duration of quantum experiments.

A. Benchmarking MPS against QPU run-time

For the problem considered in Sec. II, we employ physically motivated criteria to assess the reliability of the classical simulations. Because the Hamiltonian is spatially uniform, the dynamics must preserve the symmetries

of the underlying square lattice, namely the D_8 symmetry group. Following Ref. [52], we therefore use the preservation of these geometric symmetries as the primary empirical benchmark for MPS convergence and simulation reliability in the dynamical regimes considered here. In contrast to fixed truncation-error cutoffs, whose relation to observable accuracy can be highly problem- and observable-dependent, symmetry-based diagnostics directly probe physically relevant deviations in the simulated dynamics. As demonstrated in Ref. [52], monitoring relative symmetry violations provides a practical and empirically robust indicator for identifying the onset of MPS breakdown in the regimes relevant to the present work.

To quantify symmetry preservation, we define the symmetry error of an MPS simulation as the maximum difference between the observable matrix and any of its symmetry-equivalent counterparts under the D_8 dihedral group. That is,

$$\epsilon_{D_8} = \max_{g \in D_8} \left\| \langle \hat{O} \rangle - g \cdot \langle \hat{O} \rangle \right\|, \quad (7)$$

where g denotes an element of the D_8 symmetry group acting on the observable matrix $\langle \hat{O} \rangle$. We then impose the convergence conditions on MPS simulations that the D_8 symmetry error must be less than 40%. Given any results that obey these criteria, we consider the results reliable according to this diagnostic, see Ref. [52] for details.

In Fig. 4(a), we present the minimum required simulation times for MPS simulations of the quench protocol under consideration, as a function of pulse duration. Note that, assuming a shot frequency on the order of 1 Hz and pulse durations on the order of 100 ns, any variation in the pulse duration has a negligible impact on the results. For each pulse duration, we perform a series of simulations with progressively increasing bond dimensions, in increments of 100, until convergence was achieved according to the criteria described above. The reported simulation time corresponds to the minimum bond dimension that satisfies the convergence requirements, thereby representing the lower bound on classical simulation time for that parameter regime. The black region indicates where simulations could not be converged due to memory limitations on a 40 Gb RAM A100 GPU, highlighting the growing computational cost in more entangled regimes.

The results in Fig. 4(a) enable a direct comparison between MPS and QPU simulation times across varying system sizes, with the corresponding data summarized in Fig. 4(b). For short post-quench dynamics, in which the total entanglement in the system remains low and tensor network methods are applicable, we observe that classical MPS simulations consistently outperform the QPU, even for the largest system sizes considered and are expected to continue doing so beyond this range due to their more favorable scaling behavior. However, for medium-duration




FIG. 4. (a) Minimum simulation time required for MPS to achieve convergence as a function of pulse duration t and system size N . The black region highlights the pulse durations in which 40 Gb of RAM was not enough to achieve convergence. (b) Comparison of total simulation time between classical MPS and estimates for QPU execution across varying system sizes for increasing pulse durations.

post-quench dynamics (≥ 500 ns), the performance of MPS deteriorates rapidly with system size, while the QPU maintains a relatively constant execution time, leading to a clear advantage in favor of the QPU. Considering that the coherence time of a neutral-atom QPU operating in analog mode can reach up to $6 \mu\text{s}$, these results highlight a substantial regime in which quantum hardware can outperform classical MPS methods.

B. Comparison between NQS and QPU run-time

Time-dependent VMC simulations provide a promising alternative to tensor network approaches such as MPS, given their ability to efficiently represent certain highly entangled states and the parallelizable implementation of tVMC codes that typically allows for an almost linear reduction in run-time with the number of GPUs [72]. It remains unclear, however, whether tVMC simulations can reach convergence for the particular type of quench considered in this work and time scales of order ≈ 4000 ns expected for neutral-atom QPUs. Recent studies suggest that the observed limitations in capturing the dynamics in post-quench dynamics are not due to an intrinsic lack of representability of NQS but rather arise from numerical instabilities in solving the TDVP [47] equation or from finite Monte Carlo sampling [74].

To assess the convergence of NQS simulations, we consider the following quantity:

$$r^2(t) = \frac{\mathcal{D}(|\psi_{\vec{\theta}(t+\delta t)}\rangle, e^{-i\hat{H}\delta t}|\psi_{\vec{\theta}(t)}\rangle)^2}{\mathcal{D}(|\psi_{\vec{\theta}(t+\delta t)}\rangle, |\psi_{\vec{\theta}(t)}\rangle)^2}, \quad (8)$$

where $\mathcal{D}$ is the Fubini-Study distance between two quantum states [46,47]. The numerator of the expression mentioned above represents the optimization objective used to define the equation of motion for the variational parameters $\vec{\theta}$, and the denominator is introduced as a natural scale. The r^2 can be obtained using Monte Carlo




FIG. 5. (a) Minimum simulation time required for NQS to achieve convergence as a function of quench duration and system size. The black region highlights the quench durations in which simulations do not converge. (b) Comparison of total simulation time between classical NQS and estimates for QPU execution across varying system sizes for increasing pulse durations. Only simulations with R^2 below a certain threshold (see text) are considered.

sampling by considering a second-order expansion in the parameter δt [47]. The integrated error up to a time t , i.e., $R^2(t) = \int_0^t r^2(t') dt'$, can be used as a figure of merit for the convergence of the NQS results.

For a given system size N , the simulations are performed with a fixed NQS architecture. As a reference of NQS architecture, we consider a two-layer CNN containing a number of variational parameters that scale as $O(N)$ [47]. In the Supplemental Material [64], we present the colormap of R^2 as a function of system size and evolution time. The results indicate that tVMC simulations with CNN employed face difficulties in reaching convergence for intermediate post-quench durations, specifically for $t > 400$ ns (i.e., $tJ \approx 1.0$). This behavior is reflected in the rapid growth of R^2 with time, see Supplemental Material [64] for more information.

In Fig. 5(a), we highlight only the regime of parameters that can be considered reliable (i.e., cases for which $R^2 < 0.05$). A more precise discussion of the convergence of NQS results for different observables and architectures is presented in Ref. [52].

Figure 5(b) reports a direct comparison between NQS simulations and QPU run-time estimates, with NQS simulations executed on 4 NVIDIA A100 40 GB SXM GPUs. Although NQS approaches preserve polynomial run-time scaling, analogous to that of the QPU, our results indicate that, already at intermediate post-quench durations, the QPU exhibits a clear run-time advantage for the estimation of local observables.

VII. ENERGY CONSUMPTION

In this section, we analyze the power requirements of both MPS and NQS classical simulations and quantum simulation executed on a neutral-atom QPU in analog mode. By comparing their respective energy footprints,

we aim to assess which approach is more efficient as system size and simulation complexity scale up.

While simulation time is a key performance metric, it is not the only factor to consider when evaluating different simulation approaches. Power consumption is also a critical consideration, particularly in the context of large-scale simulations where energy efficiency can significantly impact both operational cost and sustainability.

For both the MPS and NQS methods, we measure the average power consumption of an A100 GPU during simulation runs across a range of ansatz complexities, Figs. 6(a) and 6(b). The estimation is performed by monitoring power consumption with the NVIDIA command *nvidia-smi dmon*, which logs GPU utilization, memory usage, and power draw throughout the simulation. As expected, the average power usage increases with the numerical complexity of the task, reflecting the higher computational demands of more entangled or less compressible quantum states. Importantly, this measurement captures only the power drawn by the GPU as reported by *nvidia-smi* monitoring tool and does not include the energy consumed by the host CPU system, memory, cooling infrastructure, or network interconnects required to run the simulation.

To estimate the power consumption of the QPU, we use the data presented in Fig. 6(c), which are based on experimental measurements of the electrical power draw of the QPU system collected over a two-month period. These measurements monitor the total power consumption of the hardware racks supplying the system, including the optical cores, electronic control racks, laser systems, and laser control hardware. The time-averaged power over this measurement period is taken as the representative power draw of the device during typical operation.

Unlike classical accelerator hardware such as GPUs, the power consumption of a neutral-atom QPU is largely independent of computational workload. The dominant power contributions arise from continuous-wave laser systems, optical hardware, and supporting electronics that operate continuously in order to maintain trapping, cooling, and calibration of the atomic array. These subsystems therefore run at approximately constant power regardless of whether user circuits are actively executed. Although some radio-frequency control electronics are pulsed during each experimental cycle, the system typically operates at a fixed repetition rate (≈ 1 Hz) even when idle in order to maintain calibration and thermal stability. Consequently, the overall system power draw remains highly stable, varying by only about $\pm 10\%$ during operation. The long-term averaged measurement therefore provides a representative estimate of the QPU power consumption during typical operation.

Consequently, our comparison is slightly biased in favor of the GPU-based simulations, as it does not account for their full system-level power footprint, whereas the QPU




FIG. 6. (a),(b) Average power consumption of the NVIDIA A100 GPU during MPS and NQS simulations with increasing system size for various bond dimensions. (c) Total power consumption of a neutral-atom QPU operating in analog mode.

estimate reflects the complete device power draw during typical operation. Nevertheless, this distinction will not affect the main conclusion of the following section: due to the QPU’s favorable run-time scaling, it still consumes less total energy than classical simulation methods for the same post-quench dynamics task for medium-duration times.

VIII. SYSTEM-SIZE EXTRAPOLATION OF RUN TIME AND ENERGY CONSUMPTION

In this section, we discuss the extrapolation of resource requirements for classical simulations and contrast them with QPU estimates.

This analysis is possible for MPS simulations since their errors can be systematically decreased to an arbitrarily low level by enlarging the bond dimension, albeit at the

expense of higher computational cost and memory requirements. Given that the same level of control is not currently understood for the NQS approach, in the following, we focus exclusively on extrapolations of MPS simulations.

By combining the fitted run-time data from Fig. 4 with the known memory scaling behavior of MPS methods (see Supplemental Material [64]), we can extrapolate both simulation times and memory requirements for large-scale simulations. Additionally, assuming that the GPU operates at its maximum rated power consumption of 0.4 kW under heavy computational load [75], we can estimate the total energy usage for these simulations. This assumption is supported by the observed increase in GPU power consumption with larger N and bond dimension, as illustrated in Fig. 6(a). In Table I, we present projected simulation times, memory demands, and power consumption for a 4 μ s pulse duration (~ 10 J) applied to large square lattices, evaluated across a range of bond dimensions. Even for moderately sized systems of $O(100)$ atoms and post-quench durations on the order of a few microseconds, our results clearly show that the QPU outperforms classical methods by several orders of magnitude in both run-time and energy efficiency.

It is known that the entanglement generated after the quench grows approximately linearly with time across the system, leading to an exponential increase in the MPS bond dimension needed to accurately represent the state. In 2D lattices, this growth is further amplified by the width of the lattice so that longer pulse durations or larger system sizes can require substantially higher bond dimensions than those considered here. Consequently, the projected simulation times, memory requirements, and energy consumption presented in Table I likely underestimate the true computational cost for very long pulse durations with MPS methods.

IX. CONCLUSIONS AND OUTLOOK

In this work, we have numerically demonstrated that the run-time scaling of both MPS and NQS methods for simulating post-quench dynamics of the long-range Ising model on square lattices is consistent with theoretical predictions. Leveraging experimental data from a neutral-atom QPU in analog mode, we estimated the time required to perform equivalent simulations on quantum hardware.

TABLE I. Resource estimation for simulations of large lattice sizes with both MPS methods and a neutral-atom QPU in analog mode for a pulse duration of 4 μ s, or ~ 10 J. (Mem = memory, Energy = energy). Units: y = years, d = days, h = hours.

Problem size	15 $\times$ 15			20 $\times$ 20			25 $\times$ 25		
	Mem	Time	Energy (kWh)	Mem	Time	Energy (kWh)	Mem	Time	Energy (kWh)
Classical MPS ($\chi = 1000$)	150 GB	30 d	0.29 k	600 GB	140 d	1.4 k	1.4 TB	1.0 y	3.6 k
Classical MPS ($\chi = 2000$)	360 GB	93 d	0.88 k	1.4 TB	1.1 y	4.0 k	3.2 TB	2.9 y	10 k
Classical MPS ($\chi = 3000$)	700 GB	220 d	2.1 k	2.8 TB	2.7 y	9.3 k	6.3 TB	6.6 y	23 k
QPU		5.8 h	19		54 h	171		36 d	2.8 k

Our analysis revealed that MPS and NQS classical simulations exhibit unfavorable scaling in both simulation time and energy consumption when convergence is strictly enforced. In contrast, the neutral-atom QPU maintains favorable performance as system size and pulse duration increase for protocols in which device noise remains bounded. Beyond such protocols, error-erasure techniques, which convert dominant physical errors into detectable erasures, will likely allow the QPU to maintain favorable scaling relative to classical methods even in the presence of higher noise levels [76,77].

By combining performance data from both classical GPU-based simulations and QPU experiments, we projected the memory requirements, run times, and energy costs associated with simulating large-scale quantum registers. These projections indicate that analog quantum devices offer a competitive, and in many regimes, superior, alternative to classical simulation methods. Notably, they can achieve comparable accuracy with significantly reduced computational overhead and power consumption, making them a promising platform for future large-scale quantum simulation tasks [78]. It remains an important direction for future work to benchmark neutral-atom QPUs against other classical methods such as Pauli propagation [79–82], tree tensor networks [83–86] and two-dimensional tensor network approaches [87–91] (2DTN). In particular, Ref. [52] shows that 2DTN methods can reliably capture the post-quench time dynamics of the nearest-neighbor transverse-field Ising model, achieving timescales comparable to those reached by the methods used in this work for the largest system sizes considered. Moreover, these methods are expected to scale more favorably with system size.

While our estimates provide a meaningful benchmark for the current capabilities of neutral-atom QPUs, it is important to emphasize that they likely represent an upper bound on the true resource requirements. In practice, several experimental advances could further reduce overhead and improve efficiency, including the use of multiple rearrangement stages to enhance loading efficiency, averaging over symmetrically equivalent qubits to suppress statistical fluctuations, and the development of more sophisticated rearrangement algorithms to minimize atom loss and idle time [92–95]. Another key advance is the implementation of continuous atom reloading, which allows rapid recovery of lost atoms between experimental cycles without disrupting the rest of the register [96,97]. By maintaining high filling fractions with minimal dead time, continuous reloading can substantially increase the achievable shot rate, thereby reducing the overall wall-clock time required for post-quench dynamics experiments. Moreover, if one considers simulations that are intrinsically robust to a small fraction of defects, the scaling of QPU run-time can be significantly improved. In particular, for robustness levels on the order of $O(1\%)$

of the register, the run-time overhead can even become constant with system size. As these and other improvements are implemented, the performance gap between analog quantum devices and classical simulations is expected to widen, reinforcing the long-term promise of neutral-atom platforms for scalable quantum simulation.

ACKNOWLEDGMENTS

We thank Christophe Jurczak, Antoine Browaeys, Thierry Lahaye, Olivier Ezratty, Alexia Auffèves, and Victor Drouin Touchette for carefully reading the manuscript and providing insightful suggestions. We acknowledge funding from the European Union the projects PASQuanS2.1 (HORIZON-CL4-2022-QUANTUM02-SGA, Grant Agreement 101113690).

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

-
- [1] E. Bernstein and U. Vazirani, *Quantum complexity theory*, in *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing* (1993).
 - [2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2010).
 - [3] M. Fellous-Asiani, J. H. Chai, R. S. Whitney, A. Auffèves, and H. K. Ng, *Limitations in quantum computing from resource constraints*, *PRX Quantum* **2**, 040335 (2021).
 - [4] W. van Dam, M. Mykhailova, and M. Soeken, *Using azure quantum resource estimator for assessing performance of fault tolerant quantum computation*, [arXiv:2311.05801](https://arxiv.org/abs/2311.05801).
 - [5] C. Gidney, *How to factor 2048 bit RSA integers with less than a million noisy qubits*, [arXiv:2505.15917](https://arxiv.org/abs/2505.15917).
 - [6] M. P. Harrigan, T. Khattar, C. Yuan, A. Peduri, N. Yosri, F. D. Malone, R. Babbush, and N. C. Rubin, *Expressing and analyzing quantum algorithms with qualtran*, [arXiv:2409.04643](https://arxiv.org/abs/2409.04643).
 - [7] N. Nguyen, T. W. Watts, B. Link, K. S. Williams, Y. R. Sanders, S. J. Elman, M. Kieferova, M. J. Bremner, K. J. Morrell, J. Elenewski *et al.*, *Quantum computing for corrosion-resistant materials and anti-corrosive coatings design*, [arXiv:2406.18759](https://arxiv.org/abs/2406.18759).
 - [8] A. A. Agrawal, J. Job, T. L. Wilson, S. N. Saadatmand, M. J. Hodson, J. Y. Mutus, A. Caesura, P. D. Johnson, J. E. Elenewski, K. J. Morrell, and A. F. Kemper, *Quantifying fault tolerant simulation of strongly correlated systems using the fermi-hubbard model*, [arXiv:2406.06511](https://arxiv.org/abs/2406.06511).
 - [9] N. Bellonzi, A. Kunitsa, J. T. Cantin, J. A. Campos-Gonzalez-Angulo, M. D. Radin, Y. Zhou, P. D. Johnson, L. A. Martínez-Martínez, M. R. Jangrouei, A. S. Brahmachari *et al.*, *Feasibility of accelerating homogeneous catalyst discovery with fault-tolerant quantum computers*, [arXiv:2406.06335](https://arxiv.org/abs/2406.06335).

- [10] A. Auffèves, *Quantum technologies need a quantum energy initiative*, [PRX Quantum](#) **3**, 020101 (2022).
- [11] D. Jaschke and S. Montangero, *Is quantum computing green? An estimate for an energy-efficiency quantum advantage*, [Quantum Sci. Technol.](#) **8**, 025001 (2023).
- [12] R. Verdecchia, J. Sallou, and L. Cruz, *A systematic review of green AI*, [arXiv:2301.11047](#).
- [13] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, *Green AI*, [arXiv:1907.10597](#).
- [14] E. Suarez, J. Amaya, M. Frank, O. Freyermuth, M. Girone, B. Kostrzewa, and S. Pfalzner, *Energy efficiency trends in HPC: What high-energy and astrophysicists need to know*, [Front Phys.](#) **13**, 1542474 (2025).
- [15] O. Lanes, M. Beji, A. D. Corcoles, C. Dalyac, J. M. Gambetta, L. Henriët, A. Javadi-Abhari, A. Kandala, A. Mezzacapo, C. Porter *et al.*, *A framework for quantum advantage*, [arXiv:2506.20658](#).
- [16] H.-Y. Huang, S. Choi, J. R. McClean, and J. Preskill, *The vast world of quantum advantage*, [arXiv:2508.05720](#).
- [17] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandão, D. A. Buell *et al.*, *Quantum supremacy using a programmable superconducting processor*, [Nature \(London\)](#) **574**, 505 (2019).
- [18] A. Morvan, B. Villalonga, X. Mi, S. Mandr, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, I. K. Drozdov, J. Chau, G. Laun *et al.*, *Phase transitions in random circuit sampling*, [Nature \(London\)](#) **634**, 328 (2024).
- [19] R. Haghsheenas, E. Chertkov, M. Mills, W. Kadow, S.-H. Lin, Y.-H. Chen, C. Cade, I. Niesen, T. Begušić, M. S. Rudolph *et al.*, *Digital quantum magnetism at the frontier of classical simulations*, [Nature \(London\)](#) **653**, 56 (2026).
- [20] F. Góis, M. Pezzutto, and Y. Omar, *Towards energetic quantum advantage in trapped-ion quantum computation*, [arXiv:2404.11572](#).
- [21] H.-L. Liu, H. Su, S.-Q. Gong, Y.-C. Gu, H.-Y. Tang, M.-H. Jia, Q. Wei, Y. Song, D. Wang, M. Zheng *et al.*, *Robust quantum computational advantage with programmable 3050-photon gaussian boson sampling*, [arXiv:2508.09092](#).
- [22] V. Cimini, S. Gherardini, M. Barbieri, I. Gianani, M. Sbroscia, L. Buffoni, M. Paternostro, and F. Caruso, *Experimental characterization of the energetics of quantum logic gates*, [npj Quantum Inf.](#) **6**, 96 (2020).
- [23] H. Zhou, C. Duckering, C. Zhao, D. Bluvstein, M. Cain, A. Kubica, S.-T. Wang, and M. D. Lukin, *Resource analysis of low-overhead transversal architectures for reconfigurable atom arrays*, in *ISCA 25: Proceedings of the 52nd Annual International Symposium on Computer Architecture* (2025), pp. 1432–1448, [10.1145/3695053.3731039](#).
- [24] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, R. Finkelstein, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale *et al.*, *Preparing random states and benchmarking with many-body quantum chaos*, [Nature \(London\)](#) **613**, 468 (2023).
- [25] A. L. Shaw, Z. Chen, J. Choi, D. K. Mark, P. Scholl, R. Finkelstein, A. Elben, S. Choi, and M. Endres, *Benchmarking highly entangled states on a 60-atom analogue quantum simulator*, [Nature \(London\)](#) **628**, 71 (2024).
- [26] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, *Quantum computing with neutral atoms*, [Quantum](#) **4**, 327 (2020).
- [27] A. Browaeys and T. Lahaye, *Many-body physics with individually controlled Rydberg atoms*, [Nat. Phys.](#) **16**, 132 (2020).
- [28] 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 anti-ferromagnets with hundreds of Rydberg atoms*, [Nature \(London\)](#) **595**, 233 (2021).
- [29] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar *et al.*, *Probing topological spin liquids on a programmable quantum simulator*, [Science](#) **374**, 1242 (2021).
- [30] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho *et al.*, *Quantum phases of matter on a 256-atom programmable quantum simulator*, [Nature \(London\)](#) **595**, 227 (2021).
- [31] J. Zhang, S. H. Cantú, F. Liu, A. Bylinskii, B. Braverman, F. Huber, J. Amato-Grill, A. Lukin, N. Gemelke, A. Keesling *et al.*, *Probing quantum floating phases in Rydberg atom arrays*, [arXiv:2401.08087](#).
- [32] 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 \(London\)](#) **551**, 579 (2017).
- [33] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho *et al.*, *Controlling quantum many-body dynamics in driven Rydberg atom arrays*, [Science](#) **371**, 1355 (2021).
- [34] T. Manovitz, S. H. Li, S. Ebadi, R. Samajdar, A. A. Geim, S. J. Evered, D. Bluvstein, H. Zhou, N. U. Koyluoglu, J. Feldmeier, P. E. Dolgirev, N. Maskara, M. Kalinowski, S. Sachdev, D. A. Huse, M. Greiner, V. Vuletić, and M. D. Lukin, *Quantum coarsening and collective dynamics on a programmable simulator*, [Nature \(London\)](#) **638**, 86 (2025).
- [35] J. Vovrosh, J. de Hond, S. Julià-Farré, J. Knolle, and A. Dauphin, *Meson spectroscopy of exotic symmetries of ising criticality in Rydberg atom arrays*, [arXiv:2506.21299](#).
- [36] D. González-Cuadra, M. Hamdan, T. V. Zache, B. Braverman, M. Kornjača, A. Lukin, S. H. Cantú, F. Liu, S.-T. Wang, A. Keesling *et al.*, *Observation of string breaking on a (2 + 1)D Rydberg quantum simulator*, [Nature \(London\)](#) **642**, 321 (2025).
- [37] C. Dalyac, L. Leclerc, L. Vignoli, M. Djellabi, W. da Silva Coelho, B. Ximenez, A. Dureau, D. Dreon, V. E. Elfving, A. Signoles *et al.*, *Graph algorithms with neutral atom quantum processors*, [arXiv:2403.11931](#).
- [38] 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).
- [39] P. Cazals, A. François, L. Henriët, L. Leclerc, M. Marin, Y. Naghmouchi, W. da Silva Coelho, F. Sikora, V. Vitale, R. Watrigant *et al.*, *Identifying hard native instances for the*

- maximum independent set problem on neutral atoms quantum processors, [arXiv:2502.04291](#).
- [40] L.-P. Henry, S. Thabet, C. Dalyac, and L. Henriet, *Quantum evolution kernel: Machine learning on graphs with programmable arrays of qubits*, *Phys. Rev. A* **104**, 032416 (2021).
- [41] I. Cong, S. Choi, and M. D. Lukin, *Quantum convolutional neural networks*, *Nat. Phys.* **15**, 1273 (2019).
- [42] M. Kornjača, H.-Y. Hu, C. Zhao, J. Wurtz, P. Weinberg, M. Hamdan, A. Zhdanov, S. H. Cantu, H. Zhou, R. A. Bravo *et al.*, *Large-scale quantum reservoir learning with an analog quantum computer*, [arXiv:2407.02553](#).
- [43] J. Hauschild and F. Pollmann, *Efficient numerical simulations with tensor networks: Tensor network python (tenpy)*, *SciPost Phys. Lect. Notes* **005** (2018).
- [44] M. Fishman, S. White, and E. M. Stoudenmire, *The itensor software library for tensor network calculations*, *SciPost Phys. Codebases* **004** (2022).
- [45] K. Bidzhiev, S. Grava, P. I. Henaff, M. Mendizabal, E. Merhej, and A. Quelle, *Efficient emulation of neutral atom quantum hardware*, [arXiv:2510.09813](#).
- [46] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, *Science* **355**, 602 (2017).
- [47] M. Schmitt and M. Heyl, *Simulating dynamics of correlated matter with neural quantum states*, [arXiv:2506.03124](#).
- [48] S. Sachdev, *Quantum Phases of Matter* (Cambridge University Press, Cambridge, England, 2023).
- [49] L. Pavešić, D. Jaschke, and S. Montangero, *Constrained dynamics and confinement in the two-dimensional quantum ising model*, *Phys. Rev. B* **111**, L140305 (2025).
- [50] W. Krinitsin, N. Tausendpfund, M. Rizzi, M. Heyl, and M. Schmitt, *Roughening dynamics of interfaces in the two-dimensional quantum ising model*, *Phys. Rev. Lett.* **134**, 240402 (2025).
- [51] 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**, eabl6850 (2022).
- [52] J. Vovrosh, S. Julià-Farré, W. Krinitsin, M. Kaicher, F. Hayes, E. Gottlob, A. Kshetrinmayum, K. Bidzhiev, S. B. Jäger, M. Schmitt *et al.*, *Simulating dynamics of the two-dimensional transverse-field ising model: A comparative study of large-scale classical numerics*, *Phys. Rev. Research* **8**, 023311 (2026).
- [53] R. Trivedi, A. F. Rubio, and J. I. Cirac, *Quantum advantage and stability to errors in analogue quantum simulators*, [arXiv:2212.04924](#).
- [54] 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).
- [55] Y. Cai, Y. Tong, and J. Preskill, *Stochastic error cancellation in analog quantum simulation*, [arXiv:2311.14818](#).
- [56] B. F. Schiffer, A. F. Rubio, R. Trivedi, and J. I. Cirac, *The quantum adiabatic algorithm suppresses the proliferation of errors*, [arXiv:2404.15397](#).
- [57] S. de Léséleuc, D. Barredo, V. Lienhard, A. Browaeys, and T. Lahaye, *Analysis of imperfections in the coherent optical excitation of single atoms to Rydberg states*, *Phys. Rev. A* **97**, 053803 (2018).
- [58] C. Shen and L.-M. Duan, *Correcting detection errors in quantum state engineering through data processing*, *New J. Phys.* **14**, 053053 (2012).
- [59] S. Julià-Farré, A. Michel, C. Domain, J. Mikael, J.-C. Lafoucriere, J. Vovrosh, A. Chahlaoui, D. Claveau, G. Villaret, J. de Hond, L. Henriet, A. Browaeys, T. Ayrat, and A. Dauphin, *Hybrid quantum-classical analog simulation of two-dimensional fermi-hubbard models with neutral atoms*, [arXiv:2510.05897](#).
- [60] L. Leclerc *et al.*, *One-to-one quantum simulation of the low-dimensional frustrated quantum magnet tmmgao_4 with 256 qubits*, [arXiv:2603.20372](#).
- [61] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, *An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays*, *Science* **354**, 1021 (2016).
- [62] K.-N. Schymik, V. Lienhard, D. Barredo, P. Scholl, H. Williams, A. Browaeys, and T. Lahaye, *Enhanced atom-by-atom assembly of arbitrary tweezer arrays*, *Phys. Rev. A* **102**, 063107 (2020).
- [63] W. Lee, H. Kim, and J. Ahn, *Defect-free atomic array formation using the Hungarian matching algorithm*, *Phys. Rev. A* **95**, 053424 (2017).
- [64] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/kfy5-q531> for further details are provided for both QPU estimates, as well as for the implementation and scaling estimates of the classical methods.
- [65] These values are quoted without correcting for vacuum and detection losses.
- [66] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, *Time-evolution methods for matrix-product states*, *Ann. Phys. (Amsterdam)* **411**, 167998 (2019).
- [67] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, *Time-dependent variational principle for quantum lattices*, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [68] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, *Unifying time evolution and optimization with matrix product states*, *Phys. Rev. B* **94**, 165116 (2016).
- [69] G. M. Crosswhite, A. C. Doherty, and G. Vidal, *Applying matrix product operators to model systems with long-range interactions*, *Phys. Rev. B* **78**, 035116 (2008).
- [70] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, *Time-evolution methods for matrix-product states*, *Ann. Phys. (Amsterdam)* **411**, 167998 (2019).
- [71] M. Hochbruck and C. Lubich, *On Krylov subspace approximations to the matrix exponential operator*, *SIAM J. Numer. Anal.* **34**, 1911 (1997).
- [72] M. Schmitt and M. Reh, *jVMC: Versatile and performant variational Monte Carlo leveraging automated differentiation and GPU acceleration*, *SciPost Phys. Codebases* **2** (2022).
- [73] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón, N. Astrakhantsev, and G. Carleo, *NetKet 3: Machine learning toolbox for many-body quantum systems*, *SciPost Phys. Codebases* **7** (2022).

- [74] A. Sinibaldi, C. Giuliani, G. Carleo, and F. Vicentini, *Unbiasing time-dependent variational monte carlo by projected quantum evolution*, [Quantum 7, 1131 \(2023\)](#).
- [75] N. Corporation, *Nvidia A100 tensor core GPU datasheet* (2021), accessed: 2025-09-16.
- [76] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, *Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays*, [Nat. Commun. 13, 4657 \(2022\)](#).
- [77] C. M. Holland, Y. Lu, S. J. Li, C. L. Welsh, and L. W. Cheuk, *Demonstration of erasure conversion in a molecular tweezer array*, [arXiv:2406.02391](#).
- [78] J. Gibson, V. Drouin-Touchette, and S. Kourtis, *Quantum counting in the Rydberg blockade*, [arXiv:2506.19298](#).
- [79] L. Broers and L. Mathey, *Exclusive-or encoded algebraic structure for efficient quantum dynamics*, [arXiv:2404.09312](#).
- [80] L. Broers, R.-Y. Sun, and S. Yunoki, *Scalable simulation of quantum many-body dynamics with or-represented quantum algebra*, [arXiv:2506.13241](#).
- [81] T. Begušić and G. K.-L. Chan, *Real-time operator evolution in two and three dimensions via sparse pauli dynamics*, [PRX Quantum 6, 020302 \(2025\)](#).
- [82] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, *Pauli propagation: A computational framework for simulating quantum systems*, [arXiv:2505.21606](#).
- [83] W. Krinitsin, N. Tausendpfund, M. Heyl, M. Rizzi, and M. Schmitt, *Time evolution of the quantum ising model in two dimensions using tree tensor networks*, [Phys. Rev. B 112, 134310 \(2025\)](#).
- [84] L. Pavešić, D. Jaschke, and S. Montangero, *Constrained dynamics and confinement in the two-dimensional quantum ising model*, [Phys. Rev. B 111, L140305 \(2025\)](#).
- [85] W. Krinitsin, N. Tausendpfund, M. Rizzi, M. Heyl, and M. Schmitt, *Roughening dynamics of interfaces in the two-dimensional quantum ising model*, [Phys. Rev. Lett. 134, 240402 \(2025\)](#).
- [86] L. Pavešić, M. D. Liberto, and S. Montangero, *Scattering and induced false vacuum decay in the two-dimensional quantum ising model*, [arXiv:2509.02702](#).
- [87] J. Tindall, M. Fishman, E. M. Stoudenmire, and D. Sels, *Efficient tensor network simulation of IBM's eagle kicked ising experiment*, [PRX Quantum 5, 010308 \(2024\)](#).
- [88] J. Tindall, A. Mello, M. Fishman, M. Stoudenmire, and D. Sels, *Dynamics of disordered quantum systems with two- and three-dimensional tensor networks*, [Science 392, 868 \(2026\)](#).
- [89] M. S. Rudolph and J. Tindall, *Simulating and sampling from quantum circuits with 2d tensor networks*, [arXiv:2507.11424](#).
- [90] T. Begušić, J. Gray, and G. K.-L. Chan, *Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance*, [Sci. Adv. 10, eadk4321 \(2024\)](#).
- [91] G. Park, J. Gray, and G. K.-L. Chan, *Simulating quantum dynamics in two-dimensional lattices with tensor network influence functional belief propagation*, [Phys. Rev. B 112, 174310 \(2025\)](#).
- [92] B. Cimring, R. El Sabeh, M. Bacvanski, S. Maaz, I. El Hajj, N. Nishimura, A. E. Mouawad, and A. Cooper, *Efficient algorithms to solve atom reconfiguration problems. I. Redistribution-reconfiguration algorithm*, [Phys. Rev. A 108, 023107 \(2023\)](#).
- [93] R. El Sabeh, J. Bohm, Z. Ding, S. Maaz, N. Nishimura, I. El Hajj, A. E. Mouawad, and A. Cooper, *Efficient algorithms to solve atom reconfiguration problems. II. Assignment-rerouting-ordering algorithm*, [Phys. Rev. A 108, 023108 \(2023\)](#).
- [94] F. Afiouni, R. El Sabeh, N. Nishimura, I. El Hajj, A. E. Mouawad, and A. Cooper, *Efficient algorithms to solve atom reconfiguration problems. III. The bird and batching algorithms and other parallel implementations on GPUs*, [Phys. Rev. A 112, 023109 \(2025\)](#).
- [95] R. Lin, H.-S. Zhong, Y. Li, Z.-R. Zhao, L.-T. Zheng, T.-R. Hu, H.-M. Wu, Z. Wu, W.-J. Ma, Y. Gao *et al.*, *AI-enabled rapid assembly of thousands of defect-free neutral atom arrays with constant-time-overhead*, [Phys. Rev. Lett. 135, 060602 \(2025\)](#).
- [96] Y. Li, Y. Bao, M. Peper, C. Li, and J. D. Thompson, *Fast, continuous and coherent atom replacement in a neutral atom qubit array*, [arXiv:2506.15633](#).
- [97] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered *et al.*, *Continuous operation of a coherent 3,000-qubit system*, [Nature \(London\) Vol 646, 1075 \(2025\)](#).