Fast, accurate, high-resolution simulation of large-scale Fermi-Hubbard models on a digital quantum processor

Gavin S. Hartnett, Khadijeh Sona Najafi, Aleksei Khindanov, Haoran Liao,
Michael Schutzman, Michael R. Hush, Michael J. Biercuk, and Yuval Baum
Q-CTRL, Los Angeles, CA USA and Sydney, NSW Australia

The Fermi-Hubbard model provides a paradigmatic description of strongly correlated electrons relevant to various problems in materials science. Despite its importance, simulating its non-equilibrium dynamics exactly is extremely challenging beyond modest scales due to exponential expansion of the Hilbert space with increasing system size. Quantum computers have been postulated as useful tools to address these computational bottlenecks, but have been limited by both machine size and the impact of hardware noise and error. Here, we report experimental digital quantum simulation of the one-dimensional Fermi-Hubbard model on a superconducting quantum processor at a scale beyond the reach of exact statevector simulation and challenging for state-of-the-art tensor-network methods. We encode this problem using up to 120 qubits through an efficient mapping that reduces circuit complexity, and we improve accuracy through error suppression to simulate dynamical evolution using up to 90 Trotter steps. From a vacancy defect introduced in the middle of an $L = 31$-site (62-qubit) Néel initial state, we directly observe spin-charge separation to $t = 9$ in natural units using up to 90 Trotter steps, and quantitatively extract velocity ratios v_c/v_s which match classical simulations across a range of model parameters. We then extend experiments to $L = 60$ (120 qubits) and long evolution times to $t = 6$ using 30 Trotter steps; Quantum-processor outputs agree quantitatively with approximate classical simulations performed using a time-dependent variational principle (TDVP) solver; increasing the TDVP bond dimension through $\chi = 4096$ expands the range of evolution times within which agreement has RMSE $\sim 1\%$ before the approaches diverge. Owing to the large scale of the simulation and the use of efficient overhead-free error-suppression techniques, for simulated evolution times at the limit of quantum/classical agreement ($t \gtrsim 5$ in natural hopping units), the wall-clock runtime of the quantum processor is up to 3000× faster than an optimized TDVP simulation using $\chi = 4096$. These results establish contemporary digital quantum processors as a versatile, quantitatively accurate, and competitive platform for the study of fermionic many-body dynamics in regimes where leading classical methods can become prohibitively expensive.

Strongly correlated electron systems carry special significance in condensed matter physics, hosting phenomena ranging from Mott-insulator transitions and unconventional superconductivity to strange-metal transport and quantum magnetism [1–3]. The Fermi-Hubbard model captures the interplay between kinetic energy and onsite Coulomb repulsion with a minimal set of parameters, and provides the canonical theoretical framework for studying these phases of matter [4, 5]. In one dimension, the model's spectrum and thermodynamic properties can be calculated exactly via the Bethe ansatz [6, 7], though closed-form expressions for dynamical quantities and real-time evolution are generally not obtainable and instead require numerical methods. The low-energy physics away from half-filling is described by Tomonaga–Luttinger liquid theory [8], while the half-filled system is a Mott insulator with a charge gap for all $U > 0$. In two dimensions, the model is widely believed to contain the key ingredients of high-temperature superconductivity in cuprates [9–11], yet no exact solution exists and its phase diagram at finite doping remains incompletely understood.

In either setting, numerical simulations are critical to understanding the broad range of real-time dynamics in the Fermi-Hubbard model—how correlations build up after a quantum quench, how energy redistributes between spin and charge sectors, and how dressed excitations such as lattice vacancies propagate at finite energy density. Simulation via exact diagonalization is limited to small system sizes due to the unfavorable growth of the Hilbert space with particle count [12]: $L \lesssim 14$ sites for the 1D Fermi-Hubbard model [13]. As a result, a range of approximate numerical methods have been developed to overcome this challenge, including quantum Monte Carlo [14], dynamical mean-field theory [15], approximate Heisenberg simulation of Pauli operators [16–22], and tensor-network methods [23, 24]. In the latter class, time-dependent variational principle (TDVP) solvers are particularly powerful in one dimension [23, 24]; the `ITensor` package in particular [25] has become a workhorse in the computational condensed-matter research community, with over 1,250 published papers using it since launch in 2015. Despite its popularity, the computational cost of accurate simulation via TDVP grows exponentially with time and system size after a quench. This arises because the required bond dimension grows exponentially with the entanglement entropy, which in turn grows linearly in time before saturating at a volume-law value $\sim L$.

Quantum computers provide a well-established framework for the efficient simulation of local many-body Hamiltonians. Analog quantum simulators based on ultracold atoms have realized Fermi-Hubbard and spin models at large scales, including single-site resolution of spin and charge degrees of freedom [26–41], but generally face limits on tunability imposed by the physical mechanisms used to realize interaction Hamiltonians. In the alternative digital approach, the time-evolution operator is decomposed into a sequence of discrete circuit elements using Trotter techniques, enabling programmable dynamical control over model parameters [35, 42]. It is known that the real-time evolution generated by a local Hamiltonian acting on n degrees of freedom for time t can be approximated to accuracy ϵ using a number of quantum gates that scales polynomially in n, t, and $1/\epsilon$ [43–45], making digital quantum simulation attractive from a computational complexity standpoint. This approach has also gained experimental support with recent digital quantum simulations of the Fermi–Hubbard model executed on superconducting processors approaching 100 qubits [46–49], as well as trapped-ion platforms reaching 56 qubits [50, 51].

Here, we report digital quantum simulation of the 1D Fermi-Hubbard model on a superconducting processor for particle numbers and evolution times that are far beyond the reach of exact-diagonalization methods, and challenging for state-of-the-art simulation methods. Our approach introduces an efficient compilation scheme that encodes the problem using a fermion-to-qubit mapping placing all hopping and interaction gates on adjacent pairs of qubits, plus use of a fermionic SWAP (fSWAP) network. These techniques are combined with a comprehensive set of runtime error-reduction strategies which do not incur execution overhead and push the limits of our simulations up to: (widest) $L = 60$, 120 qubits, $t = 6$ in natural units of the inverse hopping amplitude t_h^{-1}, and 30 Trotter steps; (deepest) $L = 31$, 62 qubits, $t = 9$, and 90 Trotter steps. Our simulations begin by confirming known phenomena using $L = 31$ and up to 90 Trotter steps, demonstrating spin-charge separation by tracking the propagation of charge and spin wavefronts from a single-vacancy Néel initial state. Quantitative extraction of a charge/spin velocity ratio v_c/v_s over a range of interaction parameters yields results in good agreement with TDVP-based classical simulations. We then enter an otherwise unexplored regime, studying relaxation dynamics of initial Fock states for $L = 60$ using 120 qubits. Over all lattice sites we observe agreement between quantum and classical TDVP simulations within root-mean-square error (RMSE) $\lesssim 1\%$, up to $t \approx 5.2$ using bond dimension $\chi = 4096$. Beyond this evolution time, the agreement between quantum and classical simulations diverges, with RMSE reaching $\sim 4\%$ at $t \approx 6$; in this range of evolution times the accuracy of the digital quantum simulation is indeterminate. We include a wall-clock runtime comparison between quantum and classical simulation methods, showing up to 3000× faster execution using the digital quantum simulator compared against an industry-standard implementation of TDVP [25].

The Fermi-Hubbard model, defined over a general graph $G = (V, E)$, is:

$$H = -\sum_{(i,j)\in E} t_{ij} \sum_{\sigma} \Big(c^\dagger_{i\sigma} c_{j\sigma} + c^\dagger_{j\sigma} c_{i\sigma} \Big) + \sum_{i\in V} U_i \, n_{i\uparrow} \, n_{i\downarrow} - \sum_{i\in V} \sum_{\sigma} \mu_{i\sigma} \, n_{i\sigma} \, . \tag{1}$$

Here, $c^\dagger_{i\sigma}$ is the creation operator for lattice site i and spin $\sigma \in \{\uparrow, \downarrow\}$, and $n_{i\sigma} := c^\dagger_{i\sigma} c_{i\sigma}$ is the number operator for site i and spin σ. The first term represents a hopping kinetic interaction, with strength controlled by a (possibly inhomogeneous) coupling t_{ij}. The second term represents the onsite interaction, controlled by the coupling U_i. For $U_i > 0$, doubly occupied orbitals are energetically penalized, and the interaction is repulsive; for $U_i < 0$ the situation is reversed, and the interaction is attractive. Lastly, $\mu_{i\sigma}$ is the chemical potential, which encodes onsite energies due to external potentials or disorder.

Equation (1) represents a broad family of electronic systems. In this work, we focus on the case of a homogeneous one-dimensional chain of fermions to facilitate the direct benchmarking of quantum computer performance against classical alternatives. The model is formulated as

$$H = -t_h \sum_{i=0}^{L-2} \sum_{\sigma} \left(c^\dagger_{i\sigma} c_{i+1,\sigma} + c^\dagger_{i+1,\sigma} c_{i\sigma} \right) + U \sum_{i=0}^{L-1} n_{i\uparrow} n_{i\downarrow} - \mu \sum_{i,\sigma} n_{i\sigma} \, , \tag{2}$$

where t_h is the uniform hopping strength. We will implicitly measure time in units of $1/t_h$ and thus often set $t_h = 1$ in what follows. We emphasize that none of the following claims rely on the existence of translation symmetry or the integrability of the model in this parameter regime.

Our framework for digital quantum simulation is as follows. First, an initial state $|\psi_0\rangle$ is prepared; in this work, we will always take $|\psi_0\rangle$ to be a Fock state characterized by the set of occupation numbers for each of the $2L$ orbitals:

$$|n_{0,\uparrow} n_{0,\downarrow} \cdots n_{L-1,\downarrow}\rangle = \prod_{i=0}^{L-1} \prod_{\sigma\in\{\uparrow,\downarrow\}} \left(c^\dagger_{i\sigma} \right)^{n_{i,\sigma}} |\mathrm{vac}\rangle \, , \tag{3}$$

where $|\mathrm{vac}\rangle$ is the zero-particle vacuum state, and $n_{i,\sigma} \in \{0, 1\}$ are the occupation numbers (Fig. 1).

In our studies we explore relaxation dynamics following a quantum quench from a half-filled Néel state $|\psi_0\rangle = |{\downarrow\uparrow\downarrow\uparrow} \cdots\rangle$ (including instances with an inserted local vacancy), where $|{\downarrow}\rangle \equiv |0, 1\rangle$ and $|{\uparrow}\rangle \equiv |1, 0\rangle$ are shorthands for single-particle occupations. Quench dynamics are observed by approximately time-evolving the initial state under the action of the model Hamiltonian by applying an n_{step} Trotter circuit. Each Trotter step approximates the time evolution for a step size of Δt; the

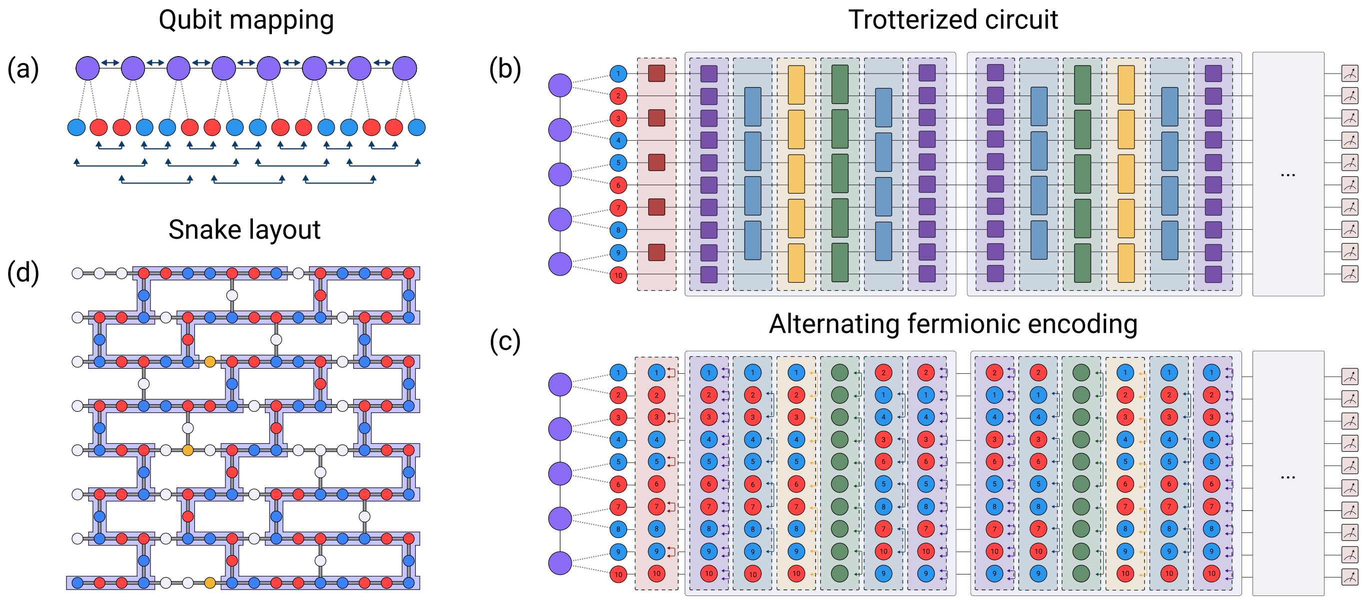

Fig. 1. Application-aware compilation for fermionic simulation. (a) The fermion-to-spin mapping represents a length-L chain of sites with $L = 8$ here (top row, purple circles) as a system of $2L$ qubits, one for each site/spin combination (bottom row, red and blue circles). We use a pair-interleaved ordering $\downarrow\uparrow\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\cdots$. The nearest-neighbor hopping interactions (indicated by the $\leftrightarrow$ arrows in the top row) map to "short-hop" and "long-hop" terms (bottom arrows). (b) The Trotterized time-evolution circuit with layers color-coded for simplicity. Red indicates state preparation. Purple indicates the single-qubit R_Z for onsite and chemical potential. Blue indicates short hopping terms, onsite interactions are yellow, and fSWAP for spin exchange are green. Subsequent Trotter layers are mirrored to improve the Trotter error scaling of the overall circuit. (c) As a companion to (b), this panel tracks the degree of freedom encoded by a given qubit. The fSWAP layer exchanges adjacent spin-up and spin-down qubits. Alternating red and blue shading indicates qubit pairs used to encode a single fermion. (d) Embedding the qubitized spin chain into a heavy-hex coupling graph on an IBM quantum device. Qubits are colored according to which spin they encode—spin-up (red) or spin-down (blue). Qubits in use have blue shading around the chain. The snaking layout is chosen to avoid qubits with low gate fidelity (indicated by orange color coding) relative to the median.

entire circuit therefore evolves the system up to a total time of $t = n_{\text{step}}\,\Delta t$ (up to maximum value $T = 60 \times \Delta t$) resulting in the state

$$|\psi_t\rangle = \prod_{k=1}^{n_{\text{step}}} U_k\,|\psi_0\rangle \ , \tag{4}$$

where U_k is the unitary corresponding to the k-th Trotter layer. Once ψ_t has been prepared, the state is measured in a suitable basis depending on the desired expectation values we wish to estimate. Here, we are primarily interested in observables in the occupation (Fock) basis (equivalently, Z-basis), which allows access to both particle and spin densities and correlations.

Our experimental implementation on quantum hardware relies on *application-aware compilation*. By framing the simulation problem directly in terms of the relevant physical quantities—the Hamiltonian, initial state, simulation time, and target observables—we co-optimize the circuit design and compilation procedure. Specifically, accounting for device topology as part of the Trotterization scheme yields significant performance gains, ultimately contributing to the viability of the large-scale simulations executed in this work.

Application-aware compilation begins with the task of mapping fermionic variables to spin variables, which is accomplished through the Jordan–Wigner transformation [52]:

$$c_J = \frac{1}{2}\left(\prod_{k=0}^{J-1} Z_k\right)(X_J + iY_J)\,. \tag{5}$$

To fully specify the transformation, the two fermion indices, the site index i and the spin index σ, must be mapped to the single index $J = J(i,\sigma)$, resulting in a qubitized Hamiltonian. We seek a mapping that avoids introducing both high-weight and non-local Paulis in H to ensure short-depth Trotter circuits that may be efficiently compiled for IBM Heron devices with heavy-hex topology [53] (Fig. 1(a)). The *pair-interleaved* ordering $\{c_{0\downarrow}, c_{0\uparrow}, c_{1\uparrow}, c_{1\downarrow}, c_{2\downarrow}, c_{2\uparrow}, c_{3\uparrow}, c_{3\downarrow}, \cdots\}$ satisfies this criterion when used in conjunction with an fSWAP network [54, 55]. The qubitized Hamiltonian is then

$$
\begin{aligned}
H = -\frac{t_h}{2} \sum_{j=0}^{L-2} \Big(& X_{2j+1} X_{2j+2} + Y_{2j+1} Y_{2j+2} \\
& + X_{2j} Z_{2j+1} Z_{2j+2} X_{2j+3} + Y_{2j} Z_{2j+1} Z_{2j+2} Y_{2j+3} \Big) \\
+ \frac{U}{4} \sum_{i=0}^{L-1} & \left(\mathbb{1} - Z_{2i} \right) \left(\mathbb{1} - Z_{2i+1} \right) . \qquad (6)
\end{aligned}
$$

Next, we address the implementation of Hamiltonian terms as circuit elements. For our 1D system, the pair-interleaved (↓↑↑↓↓↑↑↓ $\cdots$) ordering introduces an asymmetry in the kinetic term leading to weight-2 "short-hopping" terms of the form $XX + YY$, and weight-4 "long-hopping" terms of the form $XZZX + YZZY$. The short-hopping terms may be readily implemented in a Trotter scheme via two-qubit R_{XX+YY} rotation gates. The long-hopping terms require additional care; we employ a layer of fSWAPs to exchange the spin-up and spin-down fermions, $c_{i\uparrow} \leftrightarrow c_{i\downarrow}$, which transforms short-hops into long-hops, and vice versa. The onsite interaction and chemical potential terms together yield two-qubit R_{ZZ} rotations between the spin-up and spin-down qubits at each site $(2i, 2i+1)$ and single-qubit R_Z rotations on every qubit, with the latter combining the Z contributions from both terms.

The fermion ordering utilized in our experiments leads to an efficient circuit construction, with the complete circuit as implemented illustrated in Fig. 1(b-c). The first step consists of a sequence of X gates on select qubits to prepare a desired Fock state. A single Trotter layer is then applied, beginning with a one-qubit layer that implements the R_Z rotations needed for the onsite and chemical potential terms (purple). Next, the short-hopping terms are implemented by a layer of R_{XX+YY} gates (blue), and a layer of R_{ZZ} gates implements the onsite interactions (yellow). A subsequent layer of fSWAP gates (green) performs the spin-exchange operation which permutes the fermions such that the ordering alternates every Trotter layer from ↓↑↑↓ $\cdots$ to ↑↓↓↑ $\cdots$, and so on. This converts the as-yet unimplemented long-hopping interactions into short-hopping interactions, which are then executed with a second layer of R_{XX+YY} gates (blue), followed by a symmetrized R_Z layer (purple). Subsequent Trotter layers are mirrored to improve the Trotter error scaling of the overall circuit.

In this construction, a single Trotter step is a first-order approximation with error $\mathcal{O}(\Delta t^2)$; however, the mirrored contraction of two adjacent steps enjoys an enhanced symmetry that yields second-order accuracy with $\mathcal{O}(\Delta t^3)$ error. Thus, the total Trotter error after n steps is $\mathcal{O}(n\Delta t^3)$ for n even and $\mathcal{O}(\Delta t^2) + \mathcal{O}(n\Delta t^3)$ for n odd, where the quadratic term arises from the single un-paired step. Due to the local nature of the qubitized Hamiltonian, each step requires only constant depth in the system size L. For each choice of Hamiltonian parameters, we empirically select the Trotter step Δt that balances total evolution time against algorithmic error (see Supplementary Material, Sec. I for detailed circuit compilation and Sec. III A for Trotter error analysis.)

With the Trotter circuits defined, we utilize the Q-CTRL compilation pipeline in `Fire Opal` [56–58] to transpile and schedule the circuits for execution on the IBM quantum processor. First, the circuits are compiled using the native gate set on IBM Heron devices. This includes fractional gates $R_{ZZ}(\theta)$, $R_Z(\theta)$, in addition to the fixed-angle gates CZ, X, and $\sqrt{X}$, but does not include the fSIM(θ, ϕ) gate specifically tailored for fermionic simulations and available on other platforms [55]. Next, we embed the ordered fermions into the device connectivity as a snaking path that is natively compatible with the heavy-hex topology of IBM Heron processors (Fig. 1(d)). A layout selection process [59, 60] avoids poorly performing qubits while maintaining the 1D chain connectivity. The subsequent workflow used here integrates deterministic error suppression [56, 58, 61] and randomized compiling to reduce runtime errors, and a lightweight post-processing step to compensate for both readout errors and incoherent decay (see Supplementary Material, Sec. IV). With dynamical decoupling included, our largest circuits use 120 qubits and consist of 9,057 two-qubit gates organized into 152 layers. All experiments are *executed* on `ibm_boston`, a 156-qubit IBM Heron device, with 20,000 shots per circuit.

Our first study probes spin-charge separation [33], the dynamical signature of fractionalization in 1D interacting systems. We prepare a chain of length $L = 31$ in the Néel state with a localized vacancy at the central site. Time evolution under the Trotterized Hamiltonian generates charge and spin disturbances that spread through the chain (Fig. 2). We observe that the sudden quench of the highly ordered Néel background induces a transient equilibration period before ballistic propagation emerges [49] and remains visible out to the longest evolution times indicating the persistence of coherent dynamics over the deepest circuits in use here. This phenomenology, along with the impact of finite-length effects, is confirmed by auxiliary TDVP simulations (see Supplementary Material, Sec. V).

In the Fermi-Hubbard model, the elementary excitations of the integrable spectrum are spinons (charge 0, spin 1/2) and holons (charge $\pm e$, spin 0), which propagate independently with distinct velocities v_s and v_c [6–8, 62, 63]. Visual examination of the upper (charge) and lower (spin) rows of Fig. 2(a-f) clearly indicates that the observed disturbance propagation occurs at differing rates from the site of the initial localization. Further, we observe that as we enter the strong-coupling limit we see a sharpening of the spin distribution and greater deviation from the charge sector (e.g., Fig. 2(c) and (f)). In these simulations the spatial extent of $L = 31$ is critical to providing a broad window enabling direct exploration of the wavefront dynamics.

We identify the wavefronts of charge and spin disturbances through the use of tracer correlators $C_i^c(t)$

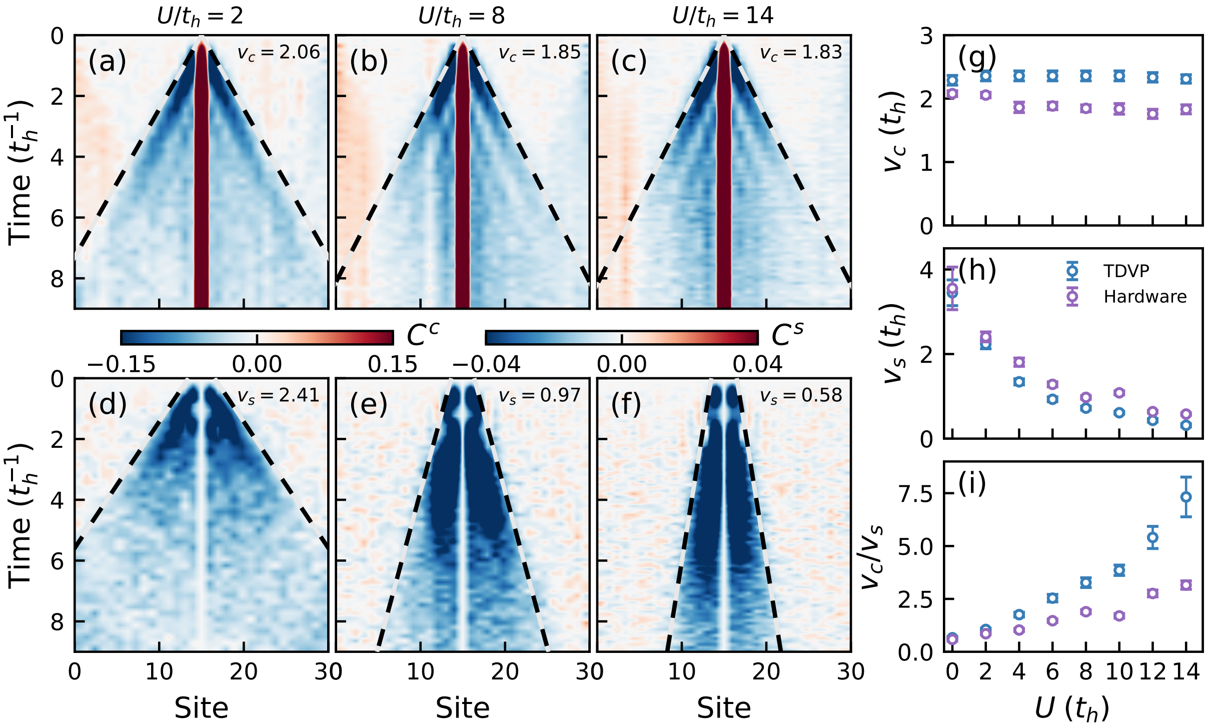

Fig. 2. Spin-charge separation exhibited in the evolution of a central vacancy in a Néel initial state for $L = 31$ over a range of repulsive couplings in each column ($U > 0$). All simulations run to a total time $t = 9\,t_h^{-1}$. Simulations in panels (a,d) use 45 Trotter steps, (b,e) use 60 Trotter steps, and (c,f) use 90 Trotter steps. (a–c) Heatmaps of the charge tracer correlator $C_i^c(t) := \langle n_{i,\uparrow}(t) + n_{i,\downarrow}(t)\rangle - \langle n_{i,\uparrow}(0) + n_{i,\downarrow}(0)\rangle$. Time flows top-to-bottom, and the horizontal axis indicates site location on the one-dimensional chain. The wavefront boundary, as identified by the protocol detailed in Supplementary Material, Sec. V A, is shown as a dashed line, and the velocity of the wavefront is reported in each panel. (d–f) Heatmaps of the spin tracer observable, $C_i^s(t) := 4(\langle S_i^z(t) S_{i_*}^z(t)\rangle - \langle S_i^z(t)\rangle\langle S_{i_*}^z(t)\rangle)$, where i_* is the site of the initial vacancy. Axes oriented as above. The same wavefront extraction has been applied as above. (g–i) Extracted wavefront velocities as a function of $U/t_h \in \{0, 2, 4, 6, 8, 10, 12, 14\}$. All TDVP simulations run on a lattice of size $L = 23$ to a total time $t = 8.9\,t_h^{-1}$, using a time step size of $0.1\,t_h^{-1}$ and max bond dimension of $\chi = 1024$. For the digital quantum simulations we use the following conditions: For $U/t_h \in \{0, 2, 4\}$ Trotter step size $\Delta t = 0.2\,t_h^{-1}$; for $U/t_h \in \{6, 8, 10\}$ $\Delta t = 0.15\,t_h^{-1}$; for $U \in \{12, 14\} t_h^{-1}$ $\Delta t = 0.1\,t_h^{-1}$. The error bars are computed using 2000 bootstrap samples and correspond to one standard deviation.

and $C_i^s(t)$, defined as $C_i^c(t) := \langle n_{i,\uparrow}(t) + n_{i,\downarrow}(t)\rangle - \langle n_{i,\uparrow}(0) + n_{i,\downarrow}(0)\rangle$, the change in the per-site electron density, and $C_i^s(t) := 4(\langle S_i^z(t) S_{i_*}^z(t)\rangle - \langle S_i^z(t)\rangle\langle S_{i_*}^z(t)\rangle)$, where $S_i^z = (n_{i,\uparrow} - n_{i,\downarrow})/2$ is the z-spin at site i, and i_* is the site of the central vacancy. Because the initial Néel state lacks connected correlations, this observable exclusively captures the dynamic magnetic disturbance of the propagating spinon.

We extract wavefront velocities through a simple detection algorithm (see Supplementary Material, Sec. V A) applied within the ballistic regime, permitting direct comparison with numerical simulation. Fig. 2(g–i) shows the extracted velocities and velocity ratios compared against TDVP calculations and exhibits quantitative agreement up to the largest values of U. Small systematic deviations between quantum hardware and TDVP classical simulation results are largely attributable to hardware noise and their potential impact on wavefront extraction. See Supplementary Material Sec. V for additional data sets derived from both digital quantum simulation and TDVP classical simulation.

We next extend digital quantum simulations to scales beyond prior demonstrations using systems with $L = 60$ (120 qubits). We execute simulations that begin in a Fock state and evolve under Trotterized time evolution for a total simulation time up to $t = 6\,t_h^{-1}$, corresponding to simulations employing up to 30 Trotter steps. Fig. 3(a)(e) show the average per-site spin-up occupancy arising from digital quantum simulation as a function of evolution time for initial states excluding and including a localized vacancy. The digital quantum simulations exhibit complex dynamical evolution of the charge density that remains visible by eye out to the maximum evolution time. This includes prominent oscillation of the staggered spin state, edge effects, and (where appropriate) spreading of localized vacancies. Such effects are highlighted in 1D time-evolution slices taken at individual sites (Fig. 3(b-d) and (f-h)) showing the oscillatory behavior for both spin species, $\langle n_{i,\uparrow}\rangle$ and $\langle n_{i,\downarrow}\rangle$.

At this system scale, the only relevant method to de-

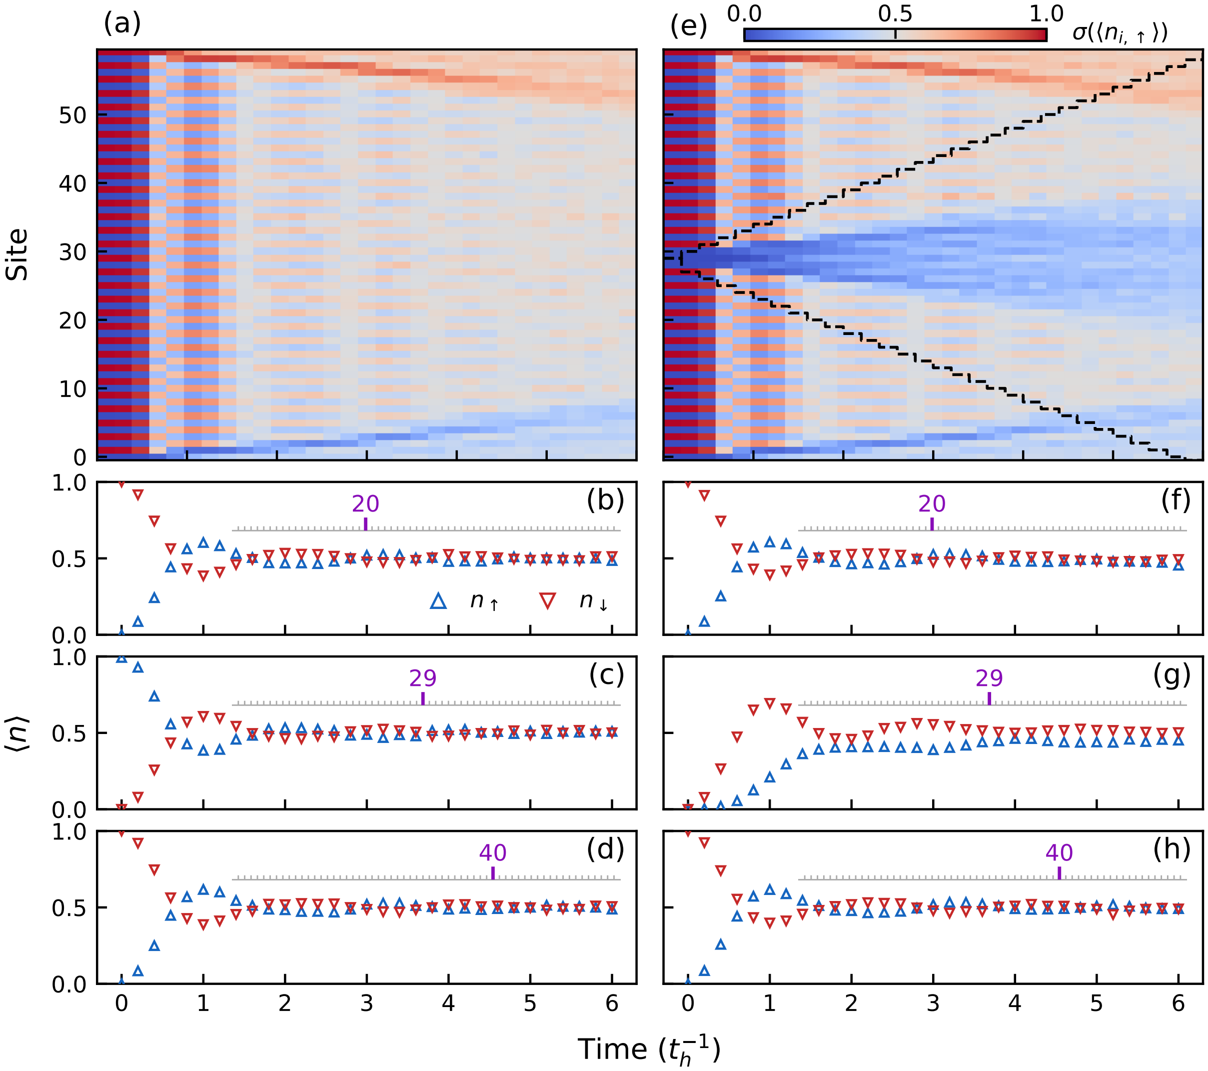

Fig. 3. Digital quantum simulation of the Fermi-Hubbard model for two initial states of an $L = 60$ chain with interaction strength $U/t_h = -2$: a Néel state (*left*) and a Néel state with a central vacancy defect (*right*). (a) Heatmap showing the evolution of the per-site spin-up occupation $\langle n_{i,\uparrow}\rangle$. To enhance the visibility of small late-time occupations, the colormap applies the transformation $\sigma(\langle n_{i,\uparrow}\rangle)$, where $\sigma(x) = [1 + \tanh(k(x - 1/2))/\tanh(k/2)]/2$. (b)–(d) Time evolution of the per-site spin occupations, $\langle n_{i,\uparrow}\rangle$ and $\langle n_{i,\downarrow}\rangle$, for selected sites. The site shown in each panel is indicated by the purple marker in the one-dimensional chain schematic shown in the inset. (e) Heatmap analogous to (a) for an initial state in which the central site ($i = 29$) is a vacancy. The black line indicates the causal light cone. (f)–(h) Time evolution plots analogous to (b)–(d) for the same set of sites. Error bars are smaller than the markers and have been omitted for clarity.

termine the accuracy of the digital quantum simulation is to compare against approximation methods. We rely on TDVP as the primary high-accuracy classical baseline, utilizing maximum bond dimensions up to $\chi = 4096$ to validate the digital quantum simulation results. For further details on the classical-simulation benchmarking approach, see Supplementary Material Sec. II where we also include comparisons against the Pauli path propagation (PPP) technique.

The general performance of TDVP in capturing the dynamics exhibited by the digital quantum simulator is illustrated in Fig. 4(a–b). First, we compare the digital quantum simulation outputs for a single site and spin against TDVP with different bond dimensions up to $\chi = 4096$. The quantum data and TDVP simulations exhibit good agreement to increasing evolution times as χ grows, before eventually diverging. This is consistent with the expectation that TDVP incurs singular value decomposition (SVD) truncation errors that restrict its ability to capture highly entangled states at long simulation times and finite χ, and that the method systematically approaches exactness as $\chi \to \infty$.

We quantitatively evaluate this phenomenology over the entire 60-site lattice by calculating the root-mean-

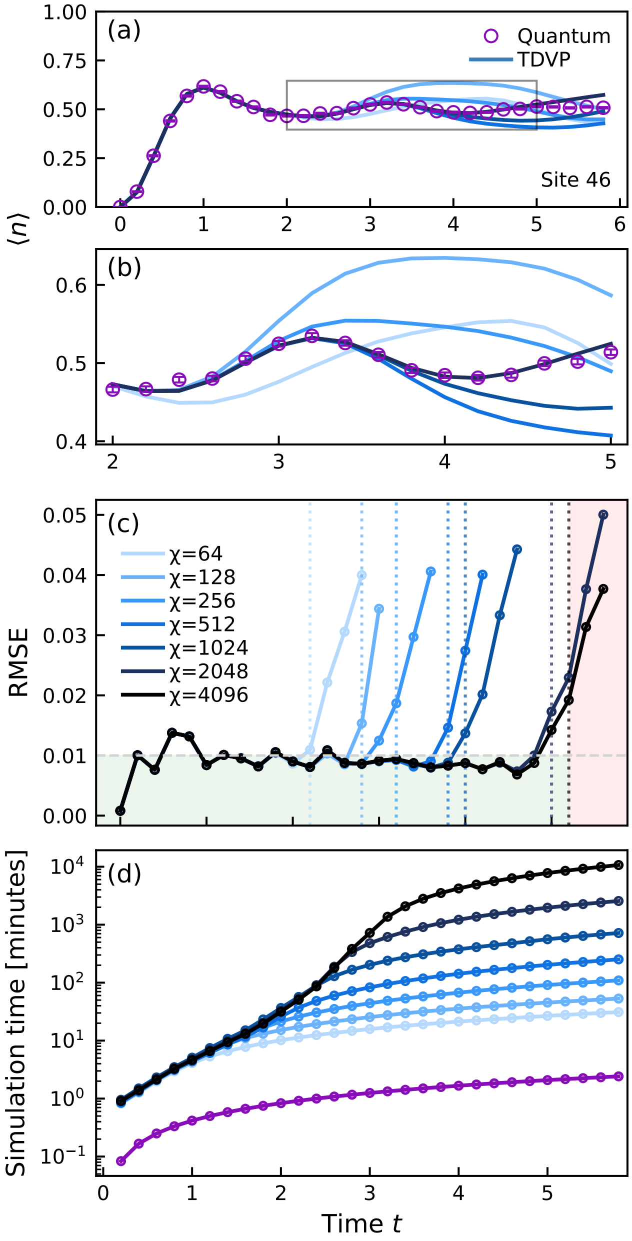

Fig. 4. Quantum simulator outputs benchmarked against TDVP for a range of bond dimensions, $\chi \in \{64, 128, 256, 512, 1024, 2048, 4096\}$. (a) Time evolution of the occupation expectation value of a representative spin-orbital, $\langle n_{46,\uparrow}\rangle(t)$ for site $i = 46$. Quantum hardware results are shown as purple circles; TDVP results are shown as solid lines, with color indicating bond dimension χ. The light rectangular box indicates the region of zoomed-in dynamics treated in panel (b). (b) Zoom of (a), showing that larger-χ TDVP simulations track the quantum data to later times before diverging. (c) RMSE calculated between the quantum and TDVP occupation values across all 120 spin orbitals, shown for all tested values of χ at the discrete times for which quantum measurements exist (markers). The dashed horizontal line and green shading indicate the estimated noise floor, RMSE $= 0.01$, and vertical dotted lines mark the last step where each curve represents $< 2\%$ error. Red shading at right for $t \gtrsim 5.2 t_h^{-1}$ indicates the zone in which correctness of quantum and classical simulations is indeterminate. (d) Wall-clock runtime required to simulate evolution to time t using quantum hardware (purple lines and markers) and TDVP (blue lines and markers, shaded by χ).

square error (RMSE) between the different simulation approaches averaged over all sites and spins: RMSE$(\langle \boldsymbol{n}\rangle_1, \langle \boldsymbol{n}\rangle_2) = \sqrt{\frac{1}{2L}\sum_{i,\sigma}\left(\langle n_{i\sigma}\rangle_1 - \langle n_{i\sigma}\rangle_2\right)^2}$. As shown in Fig. 4(c), the RMSE remains within $\sim 1\%$ to increasing evolution times up to $t \approx 5.2\,(1/t_h)$ as χ increases up to 4096. The level of agreement at shorter times is consistent with the expected computational error arising from shot noise and Trotterization in the digital quantum simulation ($\sim 1\%$), and even at the most extreme evolution time $t = 6\,(1/t_h)$, agreement remains within 4%. For evolution times beyond $t \approx 5.2\,(1/t_h)$, where agreement between quantum and classical simulations diverges at the largest values of χ, it is not possible to know which simulation methodology most accurately reflects the true system dynamics. We note the visual continuity of digital quantum simulation results shown over all lattice sites (Fig. 3(a) and (e)) as evidence that it is not likely that accumulated errors arising from decoherence or gate error have undermined the potential correctness of the digital quantum simulation at the longest evolution times treated.

An additional point of comparison between the simulation methodologies relates to the wall-clock runtime required for execution. The execution time for the digital quantum and TDVP simulations implemented here as a function of evolution time is shown in Fig. 4(d). In these experiments, the digital quantum simulation over all time steps requires approximately 2 minutes and 46 seconds of wall-clock QPU execution for all circuits and repetition over 20,000 shots. The runtime is independent of the number of sites simulated, L, and only grows linearly with the number of simulated Trotter steps. By contrast, TDVP runtime scales as $\mathcal{O}(L\chi^3)$ per step, indicating explicit dependence on both evolution time and system size.

Reflecting the fact that beyond a particular evolution time we cannot assert "correctness" for either simulation methodology, we explicitly compare runtimes at the point of quantum/classical divergence, where RMSE $> 2\%$. This occurs near evolution time $t = 5.2\,(1/t_h)$ for TDVP using $\chi = 4096$. For these specific parameters, the quantum wall-clock runtime is around two minutes and TDVP execution exceeds 100 hours. The longest TDVP execution for $t = 6\,(1/t_h)$ and $\chi = 4096$ exceeds 160 hours.

Our implementation of TDVP is optimized according to discussions with the `ITensor` [25] developers. We employ the `ITensor` Julia implementation on a 32vCPU instance (AWS c7i.8xlarge, 64GB RAM), and use all computational parallelization supported by `ITensorMPS.jl`. We identify computational bottlenecks in the algorithm arising from essential singular value decomposition and Krylov exponentiation steps, which for $\chi = 4096$ inhibit full use of the available computational cores. We also empirically validate negligible improvements doubling the classical compute engine to 64 cores. Accordingly, large-scale parallelization is likely not possible or at best highly constrained for `ITensor`, consistent with previous studies that show rapidly declining benefit from parallelization

TABLE I. A comparison of system sizes, evolution times, and Trotter steps across recent demonstrations of simulating the Fermi-Hubbard model on digital quantum simulators. The last two rows correspond to this work. ††Executed over 8 different days with a total device runtime (for all 8 twirling instances) of 36 hours, using $\sim 4\times$ fewer samples, at 1,280 shots, than this work. **The mirrored construction of our 30 (90)-step Trotter circuits makes them equivalent to a circuit with 15 (45) second-order Trotter steps. See Sec. I B for details.

System	**Lattice sites**	**Evolution time (t_h^{-1})**	**Trotter steps**	**Trotter order**	**Platform / Ref.**
2D 6×6 square	36	2	4	1st & 2nd	Quantinuum [51]
2D 7×4 torus	28	2	4	2nd	Quantinuum††[50]
2D 6×6 square	36	1.2	3	2nd	Google [46]
1D chain	52	5	10	1st & 2nd	IBM [48]
1D chain (this work, deepest)	31	9	90**	1st (mirrored**)	IBM
1D chain (this work, widest)	60	6	30**	1st (mirrored**)	IBM

at the scales of classical CPU cores used here [64]. At present, there is no readily available support for GPU acceleration that incorporates the particle and spin conservation symmetries required to reproduce the runtimes achieved in this work [65]. Further details of the computational configuration and commentary on the limits of parallelization for TDVP are presented in Supplementary Material Sec. II A.

We cannot exclude the possibility that the classical computational runtime could be improved via future GPU acceleration, modifications to the underlying algorithm, or complete replacement with a novel computational method, including the development of narrow-purpose tooling capable of only simulating the specific conditions and parameters treated in this study. Despite the value that such capabilities could bring to the community, we are not aware that any such classical acceleration is readily available to the core user base employing TDVP and related tools in their simulation tasks.

These results demonstrate that digital quantum simulation of the 1D Fermi-Hubbard model can now be performed at scales and in regimes that challenge or exceed the reach of the best available classical methods with accuracy and speed meeting end-user expectations. Achieving strong quantitative agreement between the outputs from both simulation methods for $L = 31$ builds confidence in the accuracy of the digital quantum simulation results before pushing to more challenging scales and models. To the best of our knowledge, the measurements we have performed—characterized by the maximum dynamical evolution time ($t = 9$), number of lattice sites calculated ($L = 60$), number of qubits (120), number of Trotter steps employed (90), circuit depth (452 layers), number of two-qubit gates ($> 13,800$), and quantitative agreement with classical benchmarks (RMSE $\lesssim 1\%$)—represent the largest and most accurate digital quantum simulations of the 1D Fermi-Hubbard model to date.

Achieving these results has required significant innovation in the execution of Trotterized time evolution on digital quantum simulators relative to previous experiments (see Table I). For the largest circuits we execute, our specialized compilation reduces the overall circuit depth (measured in two-qubit-gate layers) by $> 80\%$ and the number of two-qubit gates by $> 60\%$ relative to a baseline `Qiskit Nature` Jordan–Wigner transformation and the `Qiskit` transpiler (level-3 optimization) [66]. Relative to special-purpose methods for Trotterized circuits e.g., the *interleaved* ordering used in Ref. [48], and for comparable circumstances (equivalent to $L = 52$ sites and 10 second-order Trotter steps), our approach uses 40.5% fewer two-qubit gates and has 60% lower circuit depth. Further, our demonstrations include some of the largest circuits executed on quantum computers overall, with over 13,800 two-qubit gates ($L = 31$, $t = 9$, 90 Trotter steps), enabled by a customized error-reduction pipeline that incurs no execution overhead. Ensuring there is no exponential growth in the number of circuit executions commonly required for error mitigation [67] has proven essential to preserving potential time-to-solution benefits with digital quantum simulators. Details on circuit characteristics and error suppression, strategies are presented in Supplementary Material Sec. I.

The digital quantum simulation approach we employ enables wide flexibility in calculated observables; measurement at arbitrary U/t_h, and including spatially dependent hopping, interactions, and chemical potential are achievable by simply adjusting the appropriate angles of the R_Z, R_{ZZ}, and R_{XX+YY} gates in the quantum circuit. Accordingly, this platform allows the full interaction dependence of any physical observable to be explored in a single experiment. This tunability—combined with the ability to use arbitrary initial product states, measure any local observable, and scale the chain length in a straightforward and penalty-free manner—is a structural advantage of the use of digital quantum simulators over analytic methods as well as classical and analog-quantum simulation alternatives.

Natural next steps to consider include extensions to address impurity doping and two-dimensional lattices. Digital quantum simulation of two-dimensional Fermi-Hubbard models [55] with efficient two-dimensional fSWAP networks can greatly enhance the investigation of regimes such as finite doping, where conventional classical methods are often severely limited by the fermionic sign problem. Ongoing work must also address residual challenges in digital quantum simulators for them to realize their full potential as new computational tools in condensed matter physics. Measurement of higher-order correlation functions, including the single-particle spectral function and the dynamical structure factor, can provide stringent tests of Luttinger liquid theory, but also requires lower noise floors than current hardware provides.

Nonetheless, we are enthusiastic that these results support the view that useful digital simulation of certain physically relevant condensed matter physics problems with pre-fault-tolerant machines is now possible. We are excited by the possibility that modest expansion of these results may enable access to the regime believed to harbor the mechanism of high-temperature superconductivity in the cuprates [2, 3, 11]—a problem that has resisted exact classical treatment for nearly four decades.

Acknowledgments: We thank Garnet Chan for insightful conversations regarding Pauli path propagation and tensor network simulation methods, and Miles Stoudemire for valuable discussions concerning TDVP and the `ITensor` package. We also thank Alexey Gorshkov, Zhi-Yuan Wei, and Mathi Raja for illuminating discussions on the physics of the Fermi-Hubbard model and for ongoing collaborations on related topics. We acknowledge and thank Michael Brett from AWS for the provision of access to a computational cluster used in the conduct of classical benchmarking simulations. Finally, we are grateful to all our colleagues at Q-CTRL whose technical work has supported the results presented in this paper.

[1] M. Imada, A. Fujimori, and Y. Tokura, Rev. Mod. Phys. **70**, 1039 (1998).

[2] P. A. Lee, N. Nagaosa, and X.-G. Wen, Rev. Mod. Phys. **78**, 17 (2006).

[3] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, Nature **518**, 179 (2015).

[4] J. Hubbard, Proc. R. Soc. London A **276**, 238 (1963).

[5] M. C. Gutzwiller, Phys. Rev. Lett. **10**, 159 (1963).

[6] E. H. Lieb and F. Y. Wu, Phys. Rev. Lett. **20**, 1445 (1968).

[7] F. H. L. Essler, H. Frahm, F. Göhmann, A. Klümper, and V. E. Korepin, *The One-Dimensional Hubbard Model* (Cambridge University Press, 2005).

[8] F. D. M. Haldane, J. Phys. C **14**, 2585 (1981).

[9] P. W. Anderson, Science **235**, 1196 (1987).

[10] F. C. Zhang and T. M. Rice, Phys. Rev. B **37**, 3759 (1988).

[11] D. P. Arovas, E. Berg, S. A. Kivelson, and S. Raghu, Annu. Rev. Condens. Matter Phys. **13**, 239 (2022).

[12] E. Dagotto, Rev. Mod. Phys. **66**, 763 (1994).

[13] M. Innerberger, P. Worm, P. Prauhart, and A. Kauch, The European Physical Journal Plus **135**, 922 (2020).

[14] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).

[15] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996).

[16] T. Begušić, J. Gray, and G. K.-L. Chan, Science Advances **10**, eadk4321 (2024).

[17] T. Begušić and G. K.-L. Chan, PRX quantum **6**, 020302 (2025).

[18] T. Begušić, K. Hejazi, and G. K. Chan, The Journal of Chemical Physics **162** (2025).

[19] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, arXiv preprint arXiv:2505.21606 (2025).

[20] E. Fontana, M. S. Rudolph, R. Duncan, I. Rungger, and C. Cîrstoiu, npj Quantum Information **11**, 84 (2025).

[21] A. Angrisani, A. A. Mele, M. S. Rudolph, M. Cerezo, and Z. Holmes, arXiv preprint arXiv:2501.13101 (2025).

[22] C.-J. Lin, H. Gharibyan, and V. P. Su, Quantum **10**, 2014 (2026).

[23] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Phys. Rev. B **94**, 165116 (2016).

[24] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, Ann. Phys. **411**, 167998 (2019).

[25] M. Fishman, S. R. White, and E. M. Stoudenmire, SciPost Phys. Codebases , 4 (2022).

[26] J. W. Britton, B. C. Sawyer, A. C. Keith, C. C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, Nature **484**, 489 (2012).

[27] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, Nature **462**, 74 (2009).

[28] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, Nature **467**, 68 (2010).

[29] U. Las Heras, L. García-Álvarez, A. Mezzacapo, E. Solano, and L. Lamata, EPJ Quantum Technology **2**, 1 (2015).

[30] J.-M. Reiner, M. Marthaler, J. Braumüller, M. Weides, and G. Schön, Physical Review A **94**, 032338 (2016).

[31] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, Nature **545**, 462 (2017).

[32] J. Vijayan, P. Sompet, G. Salomon, J. Koepsell, S. Hirthe, A. Bohrdt, F. Grusdt, I. Bloch, and C. Gross, Science **367**, 186 (2020).

[33] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, Y.-A. Chen, B. Chiaro, R. Collins, S. J. Cotton, W. Courtney, S. Demura, A. Derk, A. Dunsworth, D. Eppens, T. Eckl, C. Erickson, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, J. A. Gross, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, M. Marthaler, O. Martin, J. M. Martinis, A. Marusczyk, S. McArdle, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, C. Mejuto-Zaera, X. Mi, *et al.*, arXiv preprint arXiv:2010.07965 10.48550/arXiv.2010.07965 (2020).

[34] J. Koepsell, D. Bourgund, P. Sompet, S. Hirthe, A. Bohrdt, Y. Wang, F. Grusdt, E. Demler, G. Salomon,

C. Gross, and I. Bloch, Nature **599**, 211 (2021).
[35] S. Stanišić, J. Haferkamp, D. Hangleiter, J. Eisert, and R. Brierley, Nat. Commun. **13**, 5743 (2022).
[36] L. C. Céleri, D. Huerga, F. Albarrán-Arriagada, E. Solano, M. Garcia de Andoin, and M. Sanz, Physical Review Applied **19**, 064086 (2023).
[37] D. González-Cuadra, D. Bluvstein, M. Kalinowski, R. Kaubruegger, N. Maskara, P. Naldesi, T. V. Zache, A. M. Kaufman, M. D. Lukin, H. Pichler, *et al.*, Proceedings of the National Academy of Sciences **120**, e2304294120 (2023).
[38] D. Srinivasan, A. Beyer, D. Zhu, P. Srikanth, S. Churchill, K. Mehta, S. K. Sridhar, K. Chakrabarti, D. W. Steuerman, N. Chopra, *et al.*, arXiv preprint arXiv:2411.07778 (2024).
[39] U. E. Khodaeva, D. L. Kovrizhin, and J. Knolle, Physical Review Research **6**, 013032 (2024).
[40] A. Michel, L. Henriet, C. Domain, A. Browaeys, and T. Ayral, Physical Review B **109**, 174409 (2024).
[41] J. Y. Z. Jee, N. Makadia, J. H. Pham, G. C. de Miranda, M. J. Biercuk, A. Shankar, and R. N. Wolf, Experimental realisation of topological spin textures in a penning trap (2026), arXiv:2604.13872 [quant-ph].
[42] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, Phys. Rev. B **102**, 235122 (2020).
[43] S. Lloyd, Science **273**, 1073 (1996).
[44] D. S. Abrams and S. Lloyd, Phys. Rev. Lett. **79**, 2586 (1997).
[45] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Phys. Rev. X **11**, 011020 (2021).
[46] F. Alam, J. L. Bosse, I. Čepaitė, *et al.*, arXiv preprint arXiv:2510.26845 10.48550/arXiv.2510.26845 (2025).
[47] S. J. Evered, M. Kalinowski, A. A. Geim, T. Manovitz, D. Bluvstein, S. H. Li, N. Maskara, H. Zhou, S. Ebadi, M. Xu, J. Campo, M. Cain, S. Ostermann, S. F. Yelin, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin, Nature **645**, 341 (2025).
[48] T. A. Chowdhury, V. Korepin, V. R. Pascuzzi, and K. Yu, Appl. Phys. Rev. **13**, 011434 (2026), arXiv:2509.14196.
[49] Y.-T. Lee, B. Pokharel, J. Cohn, A. Schleife, and A. Banerjee, Phys. Rev. Lett. **136**, 050603 (2026).
[50] F. Alam, J. L. Bosse, I. Čepaitė, *et al.*, arXiv preprint arXiv:2510.26300 10.48550/arXiv.2510.26300 (2025).
[51] E. Granet, H. Dreyer, *et al.*, arXiv preprint arXiv:2511.02125 10.48550/arXiv.2511.02125 (2025).
[52] P. Jordan and E. Wigner, Z. Phys. **47**, 631 (1928).
[53] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Physical Review X **10**, 011022 (2020).
[54] Z. Jiang, A. Kalev, W. Mruczkiewicz, and H. Neven, Quantum **4**, 276 (2020).
[55] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Phys. Rev. Lett. **120**, 110501 (2018).
[56] P. S. Mundada, A. Barbosa, S. Maity, Y. Wang, T. Merkh, T. Stace, F. Nielson, A. R. Carvalho, M. Hush, M. J. Biercuk, and Y. Baum, Phys. Rev. Appl. **20**, 024034 (2023).
[57] A. Kakkar, S. Marsh, Y. Wang, P. Mundada, P. Coote, G. Hartnett, M. J. Biercuk, and Y. Baum, arXiv preprint arXiv:2511.21831 (2025).
[58] Q-CTRL, Fire Opal, https://q-ctrl.com/fire-opal (2025), [Online].
[59] G. S. Hartnett, A. Barbosa, P. S. Mundada, M. Hush, M. J. Biercuk, and Y. Baum, Quantum **8**, 1542 (2024).
[60] Y. Wang, E. Ginez, J. Friel, Y. Baum, J.-S. Kim, A. Shih, and O. Green, arXiv:2508.21287 (2026).
[61] P. Coote, R. Dimov, S. Maity, G. S. Hartnett, M. J. Biercuk, and Y. Baum, PRX Quantum **6**, 010332 (2025).
[62] S. Tomonaga, Prog. Theor. Phys. **5**, 544 (1950).
[63] J. M. Luttinger, J. Math. Phys. **4**, 1154 (1963).
[64] P. Secular, N. Gourianov, M. Lubasch, S. Dolgov, S. R. Clark, and D. Jaksch, Phys. Rev. B **101**, 235123 (2020).
[65] ITensor Developers, Running on GPUs—ITensors.jl documentation, `https://itensor.github.io/ITensors.jl/dev/RunningOnGPUs.html` (2025), accessed: 2026-05-02.
[66] The Qiskit Nature developers and contributors., Qiskit nature 0.6.0 (2023).
[67] E. van den Berg, Z. K. Minev, A. Kandala, and K. Temme, Nature Physics **19**, 1116–1121 (2023).

Supplementary Material for "Fast, accurate, high-resolution simulation of large-scale Fermi-Hubbard models on a digital quantum processor"

Gavin S. Hartnett, Khadijeh Sona Najafi, Aleksei Khindanov, Haoran Liao, Michael Schutzman, Michael R. Hush, Michael J. Biercuk, Yuval Baum

Q-CTRL, Los Angeles, CA USA and Sydney, NSW Australia

SUPPLEMENTARY CONTENTS

I. APPLICATION-AWARE COMPILATION AND ERROR SUPPRESSION

In this section, we describe the end-to-end procedure used to perform the digital quantum simulation of the one-dimensional Fermi-Hubbard model. This workflow includes mapping the fermionic degrees of freedom to qubits, constructing Trotterized time-evolution circuits, and compiling those circuits for execution on IBM heavy-hex devices. We refer to this broader workflow as *application-aware compilation*, distinguishing it from standard compilation pipelines, which typically treat the quantum circuit as a "black box" input, employing general-purpose optimization techniques to produce a hardware-native output. In contrast, our approach leverages the specific structure of the Fermi-Hubbard model to optimize the circuit construction and compilation.

To make this supplement self-contained and establish notation, we begin by restating the model Hamiltonian:

$$H_{\mathrm{FH}} = -t_h \sum_{i=0}^{L-2} \sum_{\sigma} \left(c^\dagger_{i\sigma} c_{i+1,\sigma} + \mathrm{h.c.}\right) + U \sum_{i=0}^{L-1} n_{i\uparrow} n_{i\downarrow} - \mu \sum_{i,\sigma} n_{i\sigma} \equiv H_{\mathrm{K}} + H_U + H_\mu \, . \tag{S.1}$$

Here, L is the length of the chain, with sites indexed as $i = 0, 1, \ldots, L-1$. Each site supports spin-up and spin-down orbitals. The $c^\dagger_{i\sigma}$ ($c_{i\sigma}$) are creation (annihilation) operators at site i and for spin $\sigma \in \{\uparrow, \downarrow\}$, and $n_{i\sigma} = c^\dagger_{i\sigma} c_{i\sigma}$ is the corresponding site/spin number operator. The first term, H_{K}, represents the kinetic energy due to nearest-neighbor hopping, the second term, H_U, is an onsite interaction (with the Coulomb interaction strength U) that only contributes if both orbitals at a given site are occupied, and the final term, H_μ, is the chemical potential term (with the chemical potential μ). In the following, we adopt units where the hopping amplitude $t_h = 1$, effectively measuring energy in units of t_h and time in units of t_h^{-1}.

A. Jordan–Wigner transformation

To simulate fermionic dynamics on a qubit-based quantum processor, the fermionic creation and annihilation operators must be mapped to qubit Pauli operators. To that end, we employ the standard Jordan–Wigner (JW) transformation, which is particularly efficient for one-dimensional models as it maintains the locality of nearest-neighbor terms [1]. For an ordered set of fermionic modes indexed by $J \in \{0, 1, \ldots, 2L-1\}$, the JW transformation is defined as

$$c_J = \left(\prod_{k=0}^{J-1} Z_k\right) \sigma^-_J \, , \qquad c^\dagger_J = \left(\prod_{k=0}^{J-1} Z_k\right) \sigma^+_J \, , \tag{S.2}$$

where $\sigma^\pm_J = \frac{1}{2}(X_J \mp i Y_J)$. Implementing this mapping requires an explicit ordering $J(i, \sigma)$, which assigns a unique qubit index J to each pair of physical site $i \in \{0, \ldots, L-1\}$ and spin $\sigma \in \{\uparrow, \downarrow\}$ indices. In this work, we introduce a *pair-interleaved* ordering, specifically designed to minimize routing overhead on heavy-hex device topologies. We define this ordering as the sequence

$$\{c_{0\downarrow}, c_{0\uparrow}, c_{1\uparrow}, c_{1\downarrow}, c_{2\downarrow}, c_{2\uparrow}, c_{3\uparrow}, c_{3\downarrow}, \ldots\} \to \{c_0, c_1, \ldots c_{2L-1}\} \, , \tag{S.3}$$

and it is related to an equivalent ordering by a global spin-flip (↑↔↓). For simplicity, we focus on the convention defined above, as the results are invariant under this choice. A detailed comparison between this ordering and standard JW orderings—focusing on their respective routing costs and circuit depths—is presented in Sec. I E.

Using the pair-interleaved ordering, we perform the JW transformation (S.2) of the Fermi-Hubbard Hamiltonian Eq. (S.1), upon which the kinetic part of the Hamiltonian H_{K} splits into “short-hopping” (H_{S}) and “long-hopping” (H_{L}) components. As such, we write $H_{\mathrm{K}} = H_{\mathrm{S}} + H_{\mathrm{L}}$, where

$$H_{\mathrm{S}} = -\frac{1}{2} \sum_{J=0}^{L-2} \left(X_{2J+1} X_{2J+2} + Y_{2J+1} Y_{2J+2}\right) \, , \tag{S.4a}$$

$$H_{\mathrm{L}} = -\frac{1}{2} \sum_{J=0}^{L-2} \left(X_{2J} Z_{2J+1} Z_{2J+2} X_{2J+3} + Y_{2J} Z_{2J+1} Z_{2J+2} Y_{2J+3}\right) \, . \tag{S.4b}$$

In Eq. (S.4b), the $Z_{2J+1} Z_{2J+2}$ operators are JW “Z-strings”, which account for the fermionic anti-commutation relations between non-adjacent modes in our ordering.

Following the JW transformation, the interaction and chemical potential terms of the Hamiltonian (S.1) can be grouped into one- and two-qubit components: $H_U + H_\mu = H_{1\mathrm{Q}} + H_{U,2\mathrm{Q}}$. After discarding irrelevant additive constants, these components can be written as

$$H_{1\mathrm{Q}} = \frac{1}{2}\left(\mu - \frac{U}{2}\right) \sum_{J=0}^{2L-1} Z_J \, , \qquad H_{U,2\mathrm{Q}} = \frac{U}{4} \sum_{J=0}^{L-1} Z_{2J} Z_{2J+1} \, , \tag{S.5}$$

and the full Fermi-Hubbard Hamiltonian (S.1) is thus transformed into

$$H_{\mathrm{FH}} = H_{\mathrm{L}} + H_{\mathrm{S}} + H_{1\mathrm{Q}} + H_{U,2\mathrm{Q}} \, . \tag{S.6}$$

B. Trotterization

To simulate the dynamics governed by the Hamiltonian in Eq. (S.6), we employ a standard Trotterization technique, discretizing the full evolution operator $e^{-iH_{\mathrm{FH}}t}$ into n_{step} Trotter steps of duration Δt, such that $t = n_{\mathrm{step}}\Delta t$. For each individual Trotter step, its evolution operator is decomposed using a first-order Trotterization scheme

$$\begin{aligned} e^{-iH_{\mathrm{FH}}\Delta t} &= e^{-i\frac{\Delta t}{2}H_{1\mathrm{Q}}} e^{-i\Delta t H_{\mathrm{L}}} e^{-i\Delta t H_{U,2\mathrm{Q}}} e^{-i\Delta t H_{\mathrm{S}}} e^{-i\frac{\Delta t}{2}H_{1\mathrm{Q}}} \\ &\quad + \frac{\Delta t^2}{2}\left([H_{\mathrm{S}} - H_{\mathrm{L}}, H_{U,2\mathrm{Q}}] - [H_{\mathrm{L}}, H_{\mathrm{S}}]\right) + \mathcal{O}(\Delta t^3)\,, \end{aligned} \tag{S.7}$$

leading to an $\mathcal{O}(\Delta t^2)$ scaling of the Trotterization error. Since each exponentiated component in Eq. (S.7) consists of mutually commuting terms, it can be further decomposed into product sequences of individual qubit rotations.

Specifically, the $H_{1\mathrm{Q}}$ and $H_{U,2\mathrm{Q}}$ terms are implemented via R_Z and R_{ZZ} rotations with angles $(\mu - U/2)\Delta t/2$ and $U\Delta t/2$, respectively. The kinetic terms, H_{S} and H_{L}, are mapped to pairs of R_{XX}/R_{YY} and R_{XZZX}/R_{YZZY} rotations with angles Δt. Notably, the $H_{1\mathrm{Q}}$ contribution is symmetrized in the style of a second-order Trotter expansion. This targeted symmetrization reduces the error contribution from the $H_{1\mathrm{Q}}$ terms without increasing the physical gate depth, as R_Z rotations are implemented as "virtual" phase shifts in superconducting hardware.

The four-qubit rotations R_{XZZX} and R_{YZZY} associated with H_{L} cannot be straightforwardly implemented on hardware and must be further decomposed into deeper sequences of two-qubit gates. To perform this decomposition efficiently, we employ fermionic swap networks [2]. In this framework, the fSWAP operator

$$\mathrm{fSWAP} = \mathrm{SWAP}\cdot \mathrm{C}Z = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} = \mathrm{fSWAP}^\dagger = \mathrm{fSWAP}^{-1} \tag{S.8}$$

is utilized to exchange fermionic modes. Specifically, given our pair-interleaved ordering Eq. (S.3), a layer of fSWAP gates,

$$\mathcal{F} = \bigotimes_{J=0}^{L-1} \mathrm{fSWAP}_{2J,2J+1} = \mathcal{F}^\dagger = \mathcal{F}^{-1}\,, \tag{S.9}$$

effectively swaps the spin-↑ and spin-↓ modes on every site ($\mathcal{F}\, c_{i,\uparrow(\downarrow)}\, \mathcal{F}^\dagger = c_{i,\downarrow(\uparrow)}$). Crucially, this transformation maps the short- and long-hopping terms onto one another:

$$\mathcal{F}\, H_{S(L)}\, \mathcal{F}^\dagger = H_{L(S)}\,. \tag{S.10}$$

This identity can be verified by observing how fSWAP conjugates the constituent Pauli strings in Eq. (S.4),

$$\mathrm{fSWAP}(XZ)\mathrm{fSWAP}^\dagger = IX\,, \quad \mathrm{fSWAP}(YZ)\mathrm{fSWAP}^\dagger = IY\,, \quad \mathrm{fSWAP}(IZ)\mathrm{fSWAP}^\dagger = ZI\,. \tag{S.11}$$

By inserting Eq. (S.10) into the Trotter step Eq. (S.7), we can replace the hardware-intensive "long-hops" with "short-hops", yielding

$$\begin{aligned} e^{-iH_{\mathrm{FH}}\Delta t} &= \mathcal{F} e^{-i\frac{\Delta t}{2}H_{1\mathrm{Q}}} e^{-i\Delta t H_{\mathrm{S}}} \mathcal{F} e^{-i\Delta t H_{U,2\mathrm{Q}}} e^{-i\Delta t H_{\mathrm{S}}} e^{-i\frac{\Delta t}{2}H_{1\mathrm{Q}}} \\ &\quad + \frac{\Delta t^2}{2}\left([H_{\mathrm{S}} - H_{\mathrm{L}}, H_{U,2\mathrm{Q}}] - [H_{\mathrm{L}}, H_{\mathrm{S}}]\right) + \mathcal{O}(\Delta t^3)\,, \end{aligned} \tag{S.12}$$

where we have used the fact that $\mathcal{F}$ commutes with $H_{1\mathrm{Q}}$ (as can seen from the last identity of Eq. (S.11)). Consequently, through the strategic use of two fSWAP layers, we reduce the complex four-qubit rotations to a sequence of simpler two-qubit R_{XX} and R_{YY} rotations, significantly decreasing the total circuit depth.

An equivalent, *mirrored* Trotter decomposition is obtained by interchanging H_{S} and H_{L} in Eq. (S.7),

$$\begin{aligned} e^{-iH_{\mathrm{FH}}\Delta t} &= e^{-i\frac{\Delta t}{2}H_{1\mathrm{Q}}} e^{-i\Delta t H_{\mathrm{S}}} e^{-i\Delta t H_{U,2\mathrm{Q}}} \mathcal{F} e^{-i\Delta t H_{\mathrm{S}}} e^{-i\frac{\Delta t}{2}H_{1\mathrm{Q}}} \mathcal{F} \\ &\quad + \frac{\Delta t^2}{2}\left([H_{\mathrm{L}} - H_{\mathrm{S}}, H_{U,2\mathrm{Q}}] - [H_{\mathrm{S}}, H_{\mathrm{L}}]\right) + \mathcal{O}(\Delta t^3)\,, \end{aligned} \tag{S.13}$$

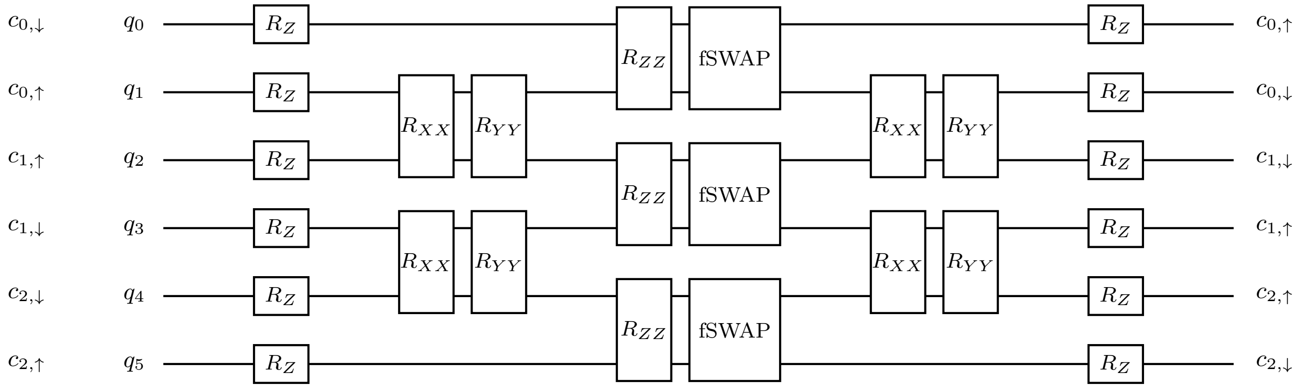

Fig. S1. Schematic quantum circuit implementing a single Trotter step U_{step} of Eq. (S.15), together with the operation $\mathcal{P}$ relabeling fermionic modes at the end, shown for $L = 3$ sites. The two layers of R_Z rotations realize the one-qubit term $H_{1\text{Q}}$ of Eq. (S.5), the two layers of R_{XX} and R_{YY} rotations realize the "short hops" in the Hamiltonian H_{S} of Eq. (S.4a), and the layer of R_{ZZ} rotations implements the two-qubit onsite interaction term $H_{U,2\text{Q}}$ of Eq. (S.5). The "long hops" H_{L} of Eq. (S.4b) are converted into the "short hops" using the layer of fSWAPs, $\mathcal{F}$, and the final virtual permutation of the fermionic modes $\mathcal{P}$, denoted by the mismatch between the input and the output qubit wire labels.

where H_{L} has been mapped to H_{S} via the $\mathcal{F}$ layers. To minimize circuit depth, we alternate between Eq. (S.12) for odd steps and Eq. (S.13) for even steps. Consequently, an odd-even pair of Trotter steps decomposes as

$$\begin{aligned} e^{-2iH_{\text{FH}}\Delta t} &= e^{-i\frac{\Delta t}{2}H_{1\text{Q}}} e^{-i\Delta t H_{\text{S}}} \mathcal{F} e^{-i\Delta t H_{U,2\text{Q}}} e^{-i\Delta t H_{\text{S}}} e^{-i\frac{\Delta t}{2}H_{1\text{Q}}} \\ &\times e^{-i\frac{\Delta t}{2}H_{1\text{Q}}} e^{-i\Delta t H_{\text{S}}} \mathcal{F} e^{-i\Delta t H_{U,2\text{Q}}} e^{-i\Delta t H_{\text{S}}} e^{-i\frac{\Delta t}{2}H_{1\text{Q}}} + \mathcal{O}(\Delta t^3) \\ &\equiv U_{\text{step}}^2 + \mathcal{O}(\Delta t^3)\,, \end{aligned} \tag{S.14}$$

where we have utilized the fact that $\mathcal{F}$ commutes with $H_{U,2\text{Q}}$ and defined an operator approximating a single Trotter step,

$$U_{\text{step}} = e^{-i\frac{\Delta t}{2}H_{1\text{Q}}} e^{-i\Delta t H_{\text{S}}} \mathcal{F} e^{-i\Delta t H_{U,2\text{Q}}} e^{-i\Delta t H_{\text{S}}} e^{-i\frac{\Delta t}{2}H_{1\text{Q}}}\,. \tag{S.15}$$

Notably, the leading-order commutators from Eq. (S.12) and Eq. (S.13) cancel out when steps are paired. This results in an $\mathcal{O}(\Delta t^3)$ local Trotter error scaling, consistent with a second-order Trotterization scheme achieved through mirroring. Therefore, for an even number of steps, the full evolution operator is

$$e^{-iH_{\text{FH}}t} = U_{\text{step}}^{n_{\text{step}}} + \mathcal{O}(n_{\text{step}}\Delta t^3)\,, \quad n_{\text{step}} = 2k\,. \tag{S.16}$$

For an odd number of steps, the final step (which is odd) is appended to the even-step evolution of Eq. (S.16) using Eq. (S.12). The concluding layer of fSWAPs in this case can be implemented virtually by a classical relabeling of the fermionic modes (measurement outcomes) at the end of the evolution; we denote this relabeling as $\mathcal{P}$. The full evolution operator for an odd number of steps thus yields

$$e^{-iH_{\text{FH}}t} = \mathcal{P}U_{\text{step}}^{n_{\text{step}}} + \mathcal{O}(\Delta t^2) + \mathcal{O}[(n_{\text{step}} - 1)\Delta t^3]\,, \quad n_{\text{step}} = 2k + 1\,, \tag{S.17}$$

where the quadratic Trotter error $\mathcal{O}(\Delta t^2)$ stems from the error in Eq. (S.12). While the quadratic error dominates at short times $[(n_{\text{step}} - 1)\Delta t \approx t \ll 1]$, the cubic error becomes dominant as the number of steps increases for long-time simulations ($t \gg 1$). Thus, for the timescales of interest, we effectively achieve second-order accuracy.

The quantum circuit implementing the Trotter step U_{step} and the virtual permutation $\mathcal{P}$ at the end is schematically depicted in Fig. S1. Crucially, the pair-interleaved ordering allows for the compilation of circuits onto a one-dimensional topology with minimal fSWAP overhead. By extension, this efficiency translates directly to heavy-hex architectures, which natively embed the required one-dimensional connectivity.

We emphasize that the error analysis presented in this subsection is performed for the evolution operators and thus provides a worst-case error bound for physical observables. In practice, the error in an observable depends on the initial state and the form of that observable and may be significantly lower. The appropriate value for the Trotter step should thus be determined and validated empirically; such an analysis is presented in Sec. III A.

C. Circuit compilation: Gate decompositions and layout selection

To execute the Trotterized evolution on quantum hardware, the abstract circuit, such as the one depicted in Fig. S1, must be transpiled into the device's native gate set and mapped onto its topology. We utilize the `ibm_boston` processor, which features a heavy-hex topology. The pair-interleaved JW ordering, introduced in Sec. I A, allows for an optimal embedding of the fermionic chain into a 1D qubit line topology that is natively supported by the heavy-hex lattice. Examples of this embedding for 120 qubits are illustrated in Fig. S2, where different colors distinguish the spin-↑ (marked red) and spin-↓ (marked blue) species. Crucially, this mapping enables the implementation of the required connectivity without incurring any routing overhead beyond the fSWAPs already integrated into the Trotter steps (Sec. I B).

The native instruction set on `ibm_boston` includes fractional $R_X(\theta)$ and $R_{ZZ}(\theta)$ gates, alongside fixed-angle X, $\sqrt{X}$, and CZ gates. During compilation, we utilize these fractional gates to numerically decompose circuit blocks into the most efficient native representations, thereby minimizing the total gate count and circuit depth. All circuit identities below are defined up to an irrelevant global phase.

1. *Kinetic term decompositions*

Each $R_{XX}(\theta)R_{YY}(\theta)$ block in the Trotter circuit (see Fig. S1) is decomposed as follows,

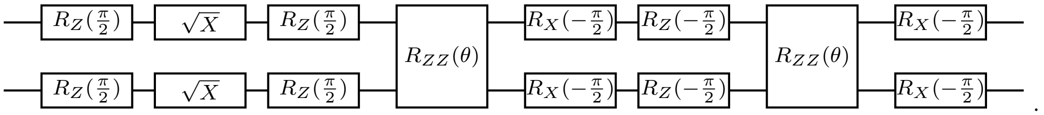

Furthermore, the adjacent blocks of R_{XX}, R_{YY} and R_Z gates at the boundary between the k-th and $(k+1)$-th Trotter steps,

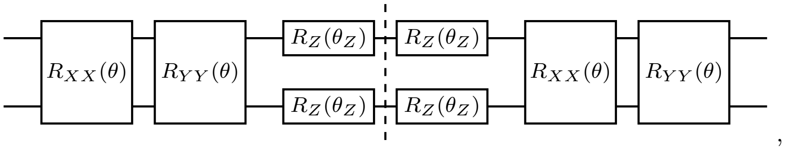

are combined and resynthesized into a consolidated block,

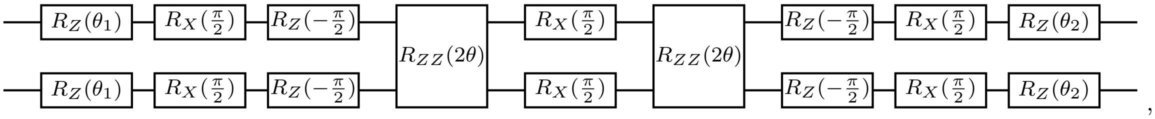

where the angles θ_1 and θ_2 are determined numerically from the input angles θ and θ_Z.

2. *Interaction term and* fSWAP *decompositions*

Each combined $R_{ZZ}(\theta) \cdot$ fSWAP block is compiled as

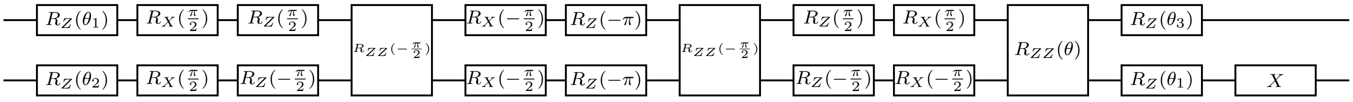

where angles $\theta_{1,2,3}$ are determined numerically based on the input angle θ.

TABLE S1. Two-qubit circuit depths (D_{2Q}), total gate counts (N_{2Q}), and aggregated QPU execution times (T_{QPU}) for various experimental configurations. Em-dashes (—) indicate configurations not explicitly reported in this study. The reported T_{QPU} represents the cumulative wall-clock execution time for the main circuits, while $T_{\mathrm{QPU}}^{\mathrm{REM+DR}}$ includes the additional noise characterization overhead for post-processing (readout error mitigation and decay recovery, see Sec. IV). Both T_{QPU} and $T_{\mathrm{QPU}}^{\mathrm{REM+DR}}$ report actual QPU usage only, and do include time needed for compilation, data retrieval from IBM Cloud, and post-processing.

	$n_{\mathrm{step}} = 30$				$n_{\mathrm{step}} = 60$				$n_{\mathrm{step}} = 90$			
System size L	D_{2Q}	N_{2Q}	T_{QPU}	$T_{\mathrm{QPU}}^{\mathrm{REM+DR}}$	D_{2Q}	N_{2Q}	T_{QPU}	$T_{\mathrm{QPU}}^{\mathrm{REM+DR}}$	D_{2Q}	N_{2Q}	T_{QPU}	$T_{\mathrm{QPU}}^{\mathrm{REM+DR}}$
10 sites (20 qubits)	152	1,457	2m46s	4m25s	302	2,897	5m36s	8m42s	452	4,337	8m46s	13m26s
31 sites (62 qubits)	152	4,649	2m46s	4m25s	302	9,239	5m36s	8m42s	452	13,829	8m46s	13m26s
60 sites (120 qubits)	152	9,057	2m46s	4m25s	—	—	—	—	—	—	—	—

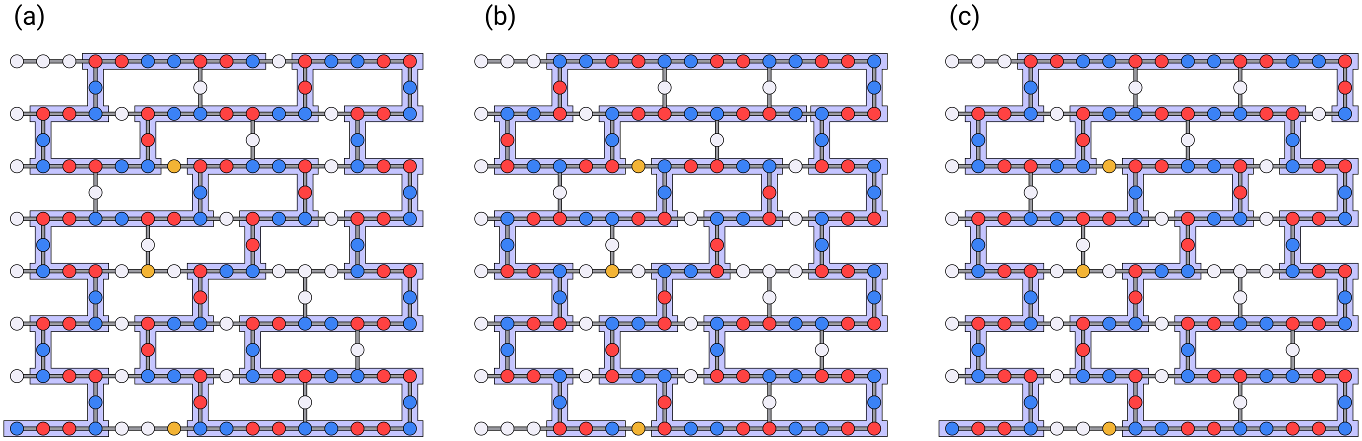

Fig. S2. Device layout and qubit mapping on `ibm_boston`. Representative layouts for $L = 60$ sites (120 qubits) utilized in our experiments. Blue and red circles denote qubits encoding spin-up and spin-down fermionic modes, respectively. These alternate every Trotter step due to the spin-exchange fSWAP layer. White circles indicate unused qubits, while orange circles highlight qubits excluded from the layout due to gate and readout fidelities significantly below the device median. The chosen layouts avoid these qubits.

3. Total depth and gate complexity

Following these decomposition rules, the total two-qubit circuit depth (D_{2Q}) and gate count (N_{2Q}) can be straightforwardly evaluated. For n_{step} Trotter steps, the two-qubit depth is

$$D_{2Q} = 5n_{\mathrm{step}} + 2. \tag{S.18}$$

Crucially, D_{2Q} is independent of the system size L, allowing us to scale the simulation without increasing its execution time. The total number of two-qubit gates scales linearly with L,

$$N_{2Q} = (5L - 2)n_{\mathrm{step}} + 2(L - 1) - 1, \tag{S.19}$$

where the final -1 accounts for the cancellation of a constituent CX gate within the fSWAP sequence. This compiler optimization occurs because the control qubit of that gate is initialized in a known computational basis state ($|0\rangle$ or $|1\rangle$). Table S1 shows D_{2Q}, N_{2Q}, and the aggregated QPU execution times for various experimental configurations $\{L, n_{\mathrm{step}}\}$. We distinguish between the time required for the primary circuits alone (T_{QPU}) and the total time including noise characterization overhead ($T_{\mathrm{QPU}}^{\mathrm{REM+DR}}$). For our largest-scale simulations by system size ($L = 60$, $n_{\mathrm{step}} = 30$), the circuits utilize $N_{2Q} = 9{,}057$ two-qubit gates at a depth of $D_{2Q} = 152$. Meanwhile, our deepest simulations ($L = 31$, $n_{\mathrm{step}} = 90$) reach a two-qubit depth of $D_{2Q} = 452$ with $N_{2Q} = 13{,}829$ two-qubit gates.

4. Layout selection

Once the circuit is transpiled into hardware-native gates and its routing is established, it must be mapped onto a specific physical qubit register. This task is known as layout selection, and it can be framed as a subgraph isomorphism problem on the device coupling graph. Typically, this is addressed using the VF2++ algorithm [3, 4], which identifies viable isomorphic subgraphs based on the circuit's connectivity. These candidates are then ranked using a scoring function that incorporates the device's backend calibration data—retrieved from the most recent calibration of `ibm_boston`—including gate error rates, coherence times (T_1, T_2), and readout fidelities.

For simulations up to $2L < 60$ qubits, we find that VF2++ efficiently identifies high-fidelity layouts. Specifically, the algorithm consistently avoids qubits with fidelities significantly below the device median. We observed that the `ibm_boston` processor contains three such persistently high-noise qubits (marked in orange in Fig. S2) whose locations remain fixed.

However, for circuits with $2L \geq 60$ qubits (such as the 62- and 120-qubit experiments we perform), the increased complexity of the search space often causes VF2++ to converge on suboptimal solutions that include high-noise qubits. Due to the nature of error propagation in Trotterized circuits, the presence of even a single high-noise qubit can significantly degrade the simulation's overall fidelity. To avoid this, we employ a variation of the recently introduced Δ-Motif algorithm [5], which offers a highly parallelizable approach to the subgraph isomorphism problem. For 120-qubit circuits, this approach allows us to exhaustively identify and score all $108,988$ possible 1D-chain layouts on the `ibm_boston` heavy-hex lattice.

At the time of our experiments, the `ibm_boston` device contained three notably high-noise qubits. An exhaustive search revealed that only 21 of the 108,988 possible layouts successfully excluded these faulty components. From this filtered subset of high-fidelity candidates, we selected the layout with the highest aggregate fidelity score for runtime execution The exact chosen layout depends on the device's backend calibration data at the time of the circuit execution. Examples of the configurations we used in our experiments are shown in Fig. S2.

D. Error suppression and Pauli twirling

State-of-the-art quantum hardware remains limited by decoherence and residual error sources, even when employing sophisticated application-aware compilation to minimize circuit depth and bypass high-noise qubits. These noise processes accumulate during execution, reducing fidelity and effectively limiting the circuit depth over which a meaningful signal can be resolved.

To counteract these effects, we utilize circuit-level error suppression and noise tailoring via Q-CTRL's `Fire Opal` pipeline [6–9]. Unlike the post-processing mitigation techniques discussed in Sec. IV, these methods—including hardware-optimized dynamical decoupling [7]—require no additional sampling overhead or the execution of auxiliary noise characterization circuits. Due to the large gate counts employed in circuits executed here and limitations on the IBM Platform's support for direct analog-level gate-waveform optimization [6, 10, 11], we further layer Pauli twirling[12, 13] in the compilation stage, in addition to our standard pipeline.

Pauli twirling is a noise tailoring strategy that applies random single-qubit Pauli gates before and after each two-qubit gate. These transformations are chosen such that each resulting gate (and thus the entire circuit) remain unitarily equivalent to the original. The randomly sampled Pauli gates are then absorbed into the existing single-qubit rotations of the circuit. This ensures that each circuit randomization—or "twirl"—maintains a gate count and execution time virtually identical to those of the original circuit. This process effectively converts coherent gate errors—which can constructively interfere and lead to significant systematic biases—into stochastic Pauli errors, and works well in circumstances where the effective orientation of noise processes in the Pauli basis is not well characterized. Beyond making the noise more amenable to subsequent mitigation, this transformation suppresses the accumulation of residual coherent errors not otherwise suppressed, typically reducing the overall error magnitude by ensuring that imperfections add incoherently rather than constructively. For the parameterized $R_{ZZ}(\theta)$ gates utilized in this work, we employ pseudo-twirling (or partial twirling), restricting to the subset of the Pauli group that commutes with the ZZ interaction.

In our implementation, a target circuit requiring N_{shot} is compiled into N_{twirl} randomized variations, each executed with $N_{\text{shot}}/N_{\text{twirl}}$ shots. After execution, the results from these variations are aggregated into a single distribution. Consequently, Pauli twirling introduces no additional shot overhead. For all experiments, we used $N_{\text{twirl}} = 10$, which provides an optimal balance between coherent error reduction and compilation time. Using $N_{\text{shot}} = 20{,}000$ total shots, we executed 10 unique circuit randomizations with 2,000 shots each.

Figure S3 quantifies the impact of Pauli twirling on the measured fermionic occupations $\langle n_{i\sigma} \rangle$. In Fig. S11(a), we evaluate the RMSE relative to classical TDVP benchmarks as a function of evolution time. We observe that Pauli twirling consistently suppresses the RMSE across the entire time domain by an average of $\sim 1\%$. The performance

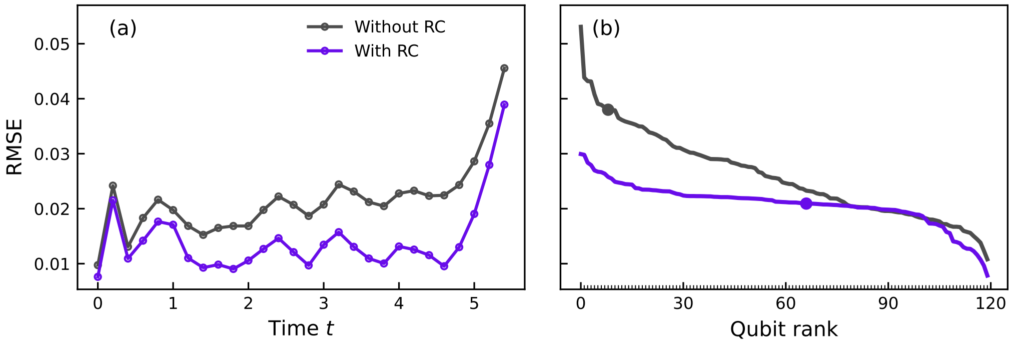

Fig. S3. Impact of Pauli twirling (AKA Randomized compiling, RC). Data is shown for the $U = -2$, $L = 60$ system prepared in a Néel initial state. The RMSE of the fermionic occupations $\langle n_{i\sigma} \rangle$ is calculated relative to classical TDVP benchmarks (maximum bond dimension $\chi = 2048$) and presented as a function of evolution time (a) and qubit rank (b). All experiments were performed with dynamical decoupling enabled. No post-processing or error mitigation techniques were applied. (a) The RMSE is computed across all qubits for each time step. (b) The RMSE is computed across all time steps for each qubit, with both datasets (with and without RC) separately sorted by magnitude. To emphasize that the sorting is independent for each dataset, a single representative qubit (index 30) is marked with a distinct marker on both lines. Error bars are smaller than the markers and have been omitted for clarity.

gain from Pauli twirling generally increases with the number of Trotter steps until $t \approx 5$, at which point the RMSE begins to drastically rise due to either the effects of hardware decoherence, the breakdown of the TDVP approximation, or both (as detailed in the main text). This behavior is consistent with suppression of the compounding of coherent errors with increasing circuit depth.

Figure S11(b) presents the RMSE calculated for each individual qubit across the full evolution period. To illustrate the global improvement in fidelity, both datasets with and without Pauli twirling are independently sorted by magnitude. This visualization demonstrates that the benefits of noise tailoring are distributed across the entire register, rather than being confined to a subset of qubits.

Quantitatively, the application of Pauli twirling yields a 21.7% reduction in the total RMSE across all time steps; within the $t \in [0, 5.2]$ interval—the regime where the simulation remains most reliable—the reduction in RMSE is 34.5%.

E. Comparative analysis of fermion-to-qubit orderings and routing efficiency

In this subsection, we compare our proposed *pair-interleaved* fermion-to-qubit ordering in the JW transformation [Eq. (S.3)] with commonly-used alternative orderings for digital quantum simulations of the Fermi-Hubbard model. In one dimension, the pair-interleaved ordering yields the lowest routing overhead on heavy-hex device topologies.

A common choice in the literature is the *block* ordering [14–16], which sequentially groups all spin-up modes followed by all spin-down modes:

$$\{c_{0\uparrow}, c_{1\uparrow}, \ldots, c_{(L-1)\uparrow}, c_{0\downarrow}, c_{1\downarrow}, \ldots, c_{(L-1)\downarrow}\} \to \{c_0, c_1, \ldots, c_{L-1}, c_L, c_{L+1}, \ldots, c_{2L-1}\} . \tag{S.20}$$

Applying the JW transformation to the Fermi-Hubbard Hamiltonian under this ordering yields the following expression for the kinetic term

$$H_{\mathrm{K}} = -\frac{1}{2} \sum_{\substack{J=0 \\ J \neq L-1}}^{2L-2} \left(X_J X_{J+1} + Y_J Y_{J+1} \right) . \tag{S.21}$$

Unlike the pair-interleaved approach [see Eq. (S.4)], this kinetic term is free of "Z-strings". Consequently, the Trotterized evolution of H_{K} can be executed on a (linear) JW chain using strictly local two-qubit gates, circumventing the need for fSWAP gates.

However, this simplification for the kinetic term comes at a cost for the interaction term. While the one-qubit component $H_{1\mathrm{Q}}$ remains identical to that in Eq. (S.5), the two-qubit interaction term becomes

$$H_{U,2\mathrm{Q}} = \frac{U}{4} \sum_{J=0}^{L-1} Z_J Z_{J+L} . \tag{S.22}$$

This term requires executing R_{ZZ} rotations between qubits separated by a distance of L along the JW chain. On hardware with a square-grid topology, a JW chain can be folded such that the spin-up and spin-down sectors run parallel, allowing indices J and $J+L$ to remain physically adjacent [14, 16]. On a heavy-hex architecture, however, embedding such a perfectly folded structure is topologically prohibited. As a result, implementing these long-range inter-sector interactions necessitates extensive fSWAP routing, rendering the block ordering highly inefficient for heavy-hex devices.

Additionally, our pair-interleaved ordering naturally accommodates spin-flip terms of the form $c^\dagger_{i,\uparrow}c_{j,\downarrow} + c^\dagger_{i,\downarrow}c_{j,\uparrow}$, should such interactions be present in the fermionic Hamiltonian, for example due to magnetic impurities. By contrast, in a block ordering where the spin-up and spin-down sectors are represented as two *disconnected* JW chains [16], such spin-flip interaction terms cannot be incorporated directly.

Another common choice in the literature is the *interleaved* ordering [17], defined by the sequence

$$\{c_{0\uparrow}, c_{0\downarrow}, c_{1\uparrow}, c_{1\downarrow}, \ldots, c_{(L-1)\uparrow}, c_{(L-1)\downarrow}\} \to \{c_0, c_1, c_2, c_3, \ldots, c_{2L-2}, c_{2L-1}\}\,. \tag{S.23}$$

Under this mapping, the one-qubit component $H_{1\mathrm{Q}}$ remains identical to Eq. (S.5). However, unlike the block ordering, this interleaved sequence enables the direct, nearest-neighbor implementation of the two-qubit interaction term on a heavy-hex topology

$$H_{U,2\mathrm{Q}} = \frac{U}{4}\sum_{J=0}^{L-1} Z_{2J}Z_{2J+1}\,. \tag{S.24}$$

Conversely, the kinetic hopping term transforms into weight-three Pauli operators, as the "Z-strings" reduce to a single intermediate Pauli Z operator

$$H_{\mathrm{K}} = -\frac{1}{2}\sum_{J=0}^{2L-3}\left(X_J Z_{J+1} X_{J+2} + Y_J Z_{J+1} Y_{J+2}\right)\,. \tag{S.25}$$

Implementing these next-nearest-neighbor hoppings on a heavy-hex lattice requires routing through the intermediate qubits via fSWAP gates. Specifically, the interleaved ordering necessitates two layers of fSWAP gates per Trotter step. In contrast, our proposed pair-interleaved ordering reduces this routing overhead to just a single fSWAP layer per Trotter step, rendering it significantly more hardware-efficient.

To quantify this advantage, we compare our pair-interleaved ordering against the interleaved ordering utilized in Ref. [17] for 1D Fermi-Hubbard dynamical simulations on IBM Heron devices. Reference [17] reports a two-qubit depth of $D_{2\mathrm{Q}} \approx 26$ per second-order Trotter step. In comparison, to achieve the equivalent algorithmic evolution using our symmetrized Trotterization (two mirrored first-order steps), our compilation requires a two-qubit depth of only $D_{2\mathrm{Q}} = 10$.

This efficiency scales highly favorably for deeper circuits. For the largest simulation reported in Ref. [17] ($L = 52$ sites, 10 second-order Trotter steps), the interleaved ordering in that study yields $D_{2\mathrm{Q}} = 263$ and $N_{2\mathrm{Q}} = 8{,}844$, as reported in Tab. III of that study. For the same system size and equivalent Trotter sequence, our pair-interleaved ordering requires only $D_{2\mathrm{Q}} = 102$ and $N_{2\mathrm{Q}} = 5{,}261$. This represents a substantial 61.2% reduction in two-qubit circuit depth and a 40.5% reduction in total two-qubit gate count.

II. CLASSICAL BASELINES

In this section, we detail the classical simulation algorithms used to benchmark our quantum hardware results. We focus on two methods capable of scaling to the large system sizes explored in our experiments (up to $L = 60$): the time-dependent variational principle (TDVP) and the Pauli path propagation (PPP) algorithm. For both approaches, we rely on publicly available software implementations executed on commodity hardware. In addition to TDVP and PPP, for the purpose of validation and testing, we also considered several classical simulation methods that do not scale to large system sizes. These include exact dynamics (ED) and simulation of Trotter circuits. These are separately discussed in Sec. III.

A. Time-dependent variational principle

TDVP [18, 19] is a widely-used algorithm for classically simulating the time evolution of quantum systems.

Before discussing the details of the algorithm, we first quantitatively probe agreement of the results of TDVP simulations against exact diagonalization in Fig. S4, for a small system of $L = 4$ sites. The top panels Fig. S4(a-d) depict the number operator expectation values $\langle n_{i,\sigma} \rangle$ for each site i and spin σ. The bottom panels Fig. S4(e-h) show the RMSE, which remains small throughout the entire evolution. The failure of the two methods to exactly agree is primarily attributable to Trotter error.

With this validation in hand we now discuss the algorithmic structure, its scaling, and computational bottlenecks. An MPS is evolved by projecting the Schrödinger equation onto the tangent space of the variational manifold and splitting the tangent-space projector into local terms that can each be integrated exactly. Two variants exist, distinguished by whether the projector is built from single-site or two-site tangent vectors; we use the two-site version (2TDVP), which adapts the bond dimension on the fly at the cost of an SVD truncation. These local flows are composed in a DMRG-like sweep: at each bond, the two-site tensor is forward-evolved under an effective Hamiltonian, SVD-factorized with truncation, and the resulting single-site tensor is back-evolved before the sweep proceeds; a reverse sweep then yields a symmetric second-order integrator. The method applies to any Hamiltonian with an efficient MPO representation, incurs no projection error for nearest-neighbor Hamiltonians, and mirrors a DMRG sweep with the local eigensolver replaced by a local matrix exponential.

Two-site TDVP for a 1D MPS evolves the state by sweeping through the chain, applying a Krylov exponential of the effective two-site Hamiltonian at each bond and truncating the resulting SVD. Five parameters control the cost. The chain length L and the on-site physical dimension d are fixed by the problem ($d = 4$ for Fermi-Hubbard). The MPO bond dimension w, i.e. the size of the auxiliary index in the matrix-product representation of H, is fixed by the Hamiltonian—in our case, $w = 6$. The MPS bond dimension χ, which sets how much entanglement the state can represent, is the only knob we actually turn: it is controlled indirectly via `maxdim` (a hard cap) and `cutoff` (the truncation tolerance on discarded singular values), and the value entering the cost is the actual $\chi(t)$ the state has reached at time t, not the cap. Finally, the Krylov subspace dimension k is the size of the Lanczos subspace built at each local matrix-exponential solve; it is adaptive, governed by a maximum cap with default value 30. The dominant per-bond cost is the matvec (matrix-vector multiplication) inside the Krylov solver, which scales as $\mathcal{O}(kd^2w\chi^3)$, plus an $\mathcal{O}(d^3\chi^3)$ truncation SVD; environment updates are subleading at $\mathcal{O}(dw\chi^3)$. Of these parameters, L, d, w, k are effectively constants of a given run. A full sweep visits $L-1$ bonds, so a symmetric step (one forward + one backward sweep, second-order in Δt_{TDVP}) costs $\mathcal{O}(L\chi^3)$ in time and $\mathcal{O}(L\chi^2)$ in memory, with prefactors set by k, d, w. The binding constraint is χ. Since under a quench entanglement grows linearly ($S(t) \sim v_E t$), forcing $\chi(t) \sim e^{v_E t}$ at fixed truncation error and pushing the total runtime to $\mathcal{O}(LTe^{3v_E T}/\Delta t_{\mathrm{TDVP}})$.

Next, we detail the implementation. We use the `ITensorMPS.jl` Julia implementation within the `ITensor` package [20]. TDVP involves a time step Δt_{TDVP} which is *a priori* independent of the Trotter step $\Delta t_{\mathrm{Trotter}}$. To enable pointwise comparison between the TDVP and quantum-hardware simulations, we require the two to be commensurate: $\Delta t_{\mathrm{Trotter}} = \ell\, \Delta t_{\mathrm{TDVP}}$ for some positive integer ℓ, so that every ℓ-th TDVP time point coincides with a Trotter time point. Holding $\Delta t_{\mathrm{Trotter}}$ fixed, increasing ℓ amounts to refining the TDVP time step. In our simulations, we take $\ell = 1$, so the two sets of time points coincide. For the SVD truncation, we used a cutoff value of $\varepsilon = 10^{-8}$, which was applied at each step. All simulations were executed on a single 32vCPU instance (AWS `c7i.8xlarge`, 64 GB RAM). We save the full MPS after each step and extract expectation values as needed to compare against the output of the quantum computer, for instance as shown in Figs. 3 and 4 in the main text.

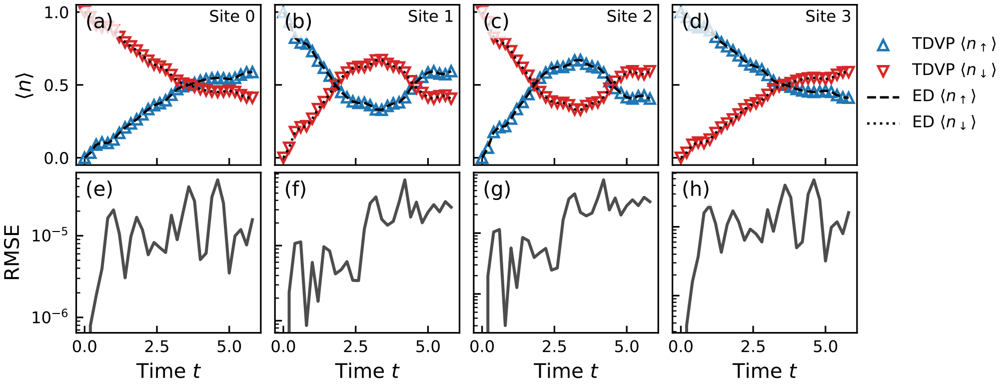

Fig. S4. Validation of TDVP for small system size ($L = 4$). (a) The occupation number expectation values $\langle n_{i,\uparrow} \rangle$, $\langle n_{i,\downarrow} \rangle$, as computed by exact dynamics (ED) (triangular markers) and TDVP with $\chi = 256$ (dashed and dotted line) for each site and spin, as a function of time. (b) The per-site RMSE between the TDVP and ED simulated values.

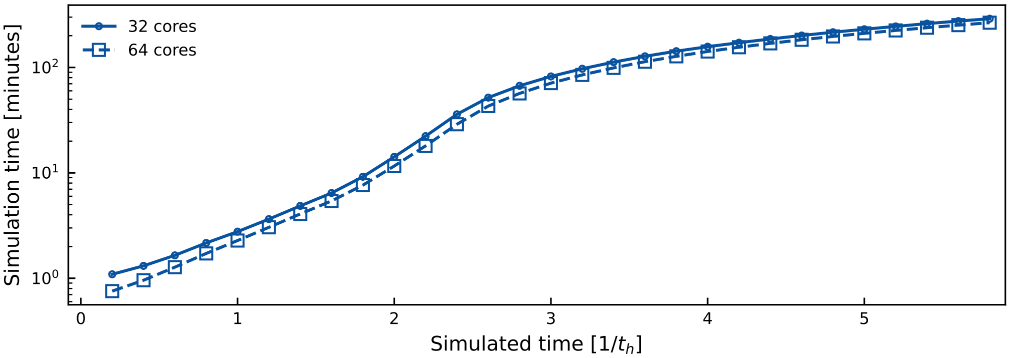

Fig. S5. Impact of CPU paralellization on overall TDVP runtime for $L = 30$ and $\chi = 1024$. There are small differences in how MPS is written to RAM vs Disk in these plots, but this has negligible impact on overall runtime or scaling.

Parallelization across multiple CPU cores provides the most straightforward potential route to acceleration, but in practice yields only modest gains [21] due to the structure of TDVP. We validate this directly by expanding the computational cluster used throughout this study to employ 64 cores, demonstrating negligible impact on overall execution time in Fig. S5, with diminishing enhancement at the longest simulation times of interest here. As described above, TDVP execution requires a loop within which one must perform two matrix exponentiations, SVD, and a large tensor contraction; these steps may be individually optimized but must then be serialized, and then serialized again over multiple sweeps. We explore the impact of parallelization on these individual tasks, focusing on SVD. Table S2 directly reveals the diminishing impact of paralellization on the SVD step up to 32 cores, indicating little opportunity for further speedups at larger cluster sizes. Because the SVD implementation in TDVP use a linear algebra library (BLAS within LAPACK) that scales as $\sim n_{\min}/32$, with $n_{\min}$ the minimum matrix dimension, for the largest matrices treated we only employ 26 of 32 available cores in the SVD step.

GPU-accelerated tensor network libraries such as cuTensorNet [22] report substantial speedups for *dense* tensor contractions at very large bond dimension, but do not extend to symmetry-adapted TDVP simulations as used here. In our case, enforcing the $U(1) \times U(1)$ symmetry corresponding to particle-number preservation for both spin-up and spin-down particles decomposes tensors into many small sectors, so the computation is dominated by numerous small matrix operations and decompositions, leading to poor GPU utilization and limited practical speedup. This limitation is consistent with the current capabilities of tensor network software, where efficient GPU support for block-sparse symmetric tensors and associated decompositions remains under development [23].

B. Pauli path propagation

In this section, we review Pauli path propagation (PPP), describe our implementation, and present benchmarking results. PPP evaluates expectation values by evolving operators in the Heisenberg picture. For an observable O and time-evolution operator $U(t)$: $\langle O(t)\rangle = \langle U^\dagger(t) O U(t)\rangle$. Expanding O in the Pauli basis, $O = \sum_{P\in\mathcal{P}} c_P P$ where $\mathcal{P}$ is

TABLE S2. Impact of CPU parallelization on the SVD step of TDVP for $\chi = 4096$, corresponding to manipulation of a 1483×837 element matrix

Threads	Time (s)	Speedup	Efficiency
1	0.8877	1.00x	100.0%
2	0.6129	1.45x	72.4%
4	0.4481	1.98x	49.5%
8	0.3439	2.58x	32.3%
12	0.3468	2.56x	21.3%
16	0.3349	2.65x	16.6%
20	0.3274	2.71x	13.6%
24	0.3565	2.49x	10.4%
28	0.3448	2.57x	9.2%
32	0.3336	2.66x	8.3%

the set of n-qubit Pauli operators and c_P the coefficient of Pauli P, yields

$$\langle O(t)\rangle = \sum_{P\in\mathcal{P}} c_P \langle P(t)\rangle \,, \tag{S.26}$$

If $U(t)$ is a Pauli rotation, $U(t) = e^{-it\sigma/2}$, i.e., $U(t) = e^{-it\sigma/2}$ for a particular Pauli σ. Then the evolution of a single Pauli in the expansion of O follows

$$U(t)^\dagger P U(t) = \begin{cases} P & [P,\sigma] = 0, \\ \cos(t)P + i\sin(t)\,\sigma P & \{P,\sigma\} = 0, \end{cases} \tag{S.27}$$

where $[\cdot,\cdot]$ and $\{\cdot,\cdot\}$ denote the commutator and anticommutator, respectively. Since a general unitary can be decomposed into a sequence of such rotations, repeated application of Eq. (S.27) suffices to evolve an arbitrary operator.

The central challenge is that, for generic dynamics, operator support grows rapidly: each anticommuting rotation induces a branching of terms, leading to exponential growth in the number of Pauli strings. Truncation schemes reduce the number of retained terms but do not eliminate this underlying growth, and therefore introduce uncontrolled approximation error that accumulates over time.

The special case of Clifford evolution avoids this issue entirely—because the Clifford group is the normalizer of the Pauli group, each Clifford conjugation simply maps a single Pauli to another Pauli, and no binary branching of terms is encountered (for the anti-commutation branch, either $\cos(t) = 0$ or $\sin(t) = 0$).

A family of related methods has been developed to manage this growth, including PPP, Pauli path simulation (PPS), and sparse Pauli dynamics (SPD), along with truncation-based variants such as xSPD. Early work based on near-Clifford perturbation theory organized the expansion in powers of non-Clifford angles after Clifford recompilation [24]. Later work moved to coefficient-threshold truncation, now commonly called SPD, and used it successfully both for the kicked Ising model at 127 qubits and for genuine real-time dynamics in systems including the tilted-field Ising model, XX ladders, and 2D and 3D transverse-field Ising model (TFIM), reaching up to 1331 qubits [25, 26]. Parallel developments include software frameworks such as `PauliPropagation.jl`, as well as analyses of truncation error and extensions to variational settings [27–30].

Given these advances, we evaluate PPP as a classical baseline for comparison with both our digital quantum simulation results as well as TDVP. We use the `PauliPropagation.jl` library (v0.4.1) [27]. The Fermi–Hubbard Hamiltonian with open boundary conditions is evolved using first-order Trotterization with time step $\Delta t = 0.2$ over 30 steps ($t \in [0, 6]$), matching the discretization used throughout this work. Fermionic operators are mapped to qubits via the JW transformation with interleaved ordering.

At each Trotter step, we truncate the Pauli expansion by discarding operators whose weight exceeds a maximum-weight threshold mw or whose coefficient satisfies $|c_P| < \delta_{\min}$, with $\delta_{\min} = 10^{-5}$. To further reduce the number of terms retained at each step, we also discarded terms whose XY weight mw_{xy} (defined as the number of single-qubit Pauli terms in an n-qubit Pauli operator that are either X or Y) exceed half the maximum weight, in other words, only terms satisfying $\mathrm{mw}_{xy} \leq \mathrm{mw}/2$ were retained at each step. The rationale for the asymmetric treatment of X/Y terms relative to Z terms stems from the final measurement being performed in the Z-basis. Any Pauli operator that retains a non-zero XY-weight at the conclusion of the Heisenberg evolution yields a zero expectation value.

While off-diagonal terms can, in principle, rotate back into diagonal Z-type operators under further evolution, we make a statistical quasi-classical assumption that this becomes improbable as the number of X and Y operators in the string grows. This imposes a strict threshold on mw_{xy} compared to the total weight, which prunes trajectories and concentrates computational resources on statistically relevant Pauli paths. At the same time, this statistical assumption introduces an additional source of approximation error, as discarded terms are not systematically controlled.

Simulations were performed for system sizes $L \in \{10, 20, 30, 40, 50, 60\}$, with initial state given by the half-filled Néel state $|\uparrow\downarrow\uparrow\downarrow\ldots\rangle$.

Fig. S6(a-b) shows the time trace results for a PPP simulation of the central site occupation numbers, $\langle n_{L/2,\uparrow}\rangle$ and $\langle n_{L/2,\downarrow}\rangle$, for a chain with $L = 60$ sites and a range of maximum-weight values mw. The result of a TDVP simulation for $\chi = 2048$ is shown for reference. Larger values of the mw cut-off lead to PPP simulations that track the TDVP reference for longer. However, the errors do not grow monotonically—they exhibit oscillations as shown in Fig. S6(c-d), a behavior observed in previous work [25, 26].

Next, we directly compare PPP against the hardware results. Figure S7(a) shows the root-mean-square error (RMSE) of PPP relative to a hardware simulation for an $L = 30$ site system. Figure S7(b) also reports the wall-clock times required to simulate the evolution of a single site (corresponding to the two expectation values $\langle n_{L/2,\uparrow}\rangle$ and $\langle n_{L/2,\downarrow}\rangle$). Although the wall-clock time grows rapidly with the maximum-weight cut-off, the accuracy remains merely comparable in scale to the TDVP baseline. To obtain the full set of $2L$ occupation numbers, the simulation would

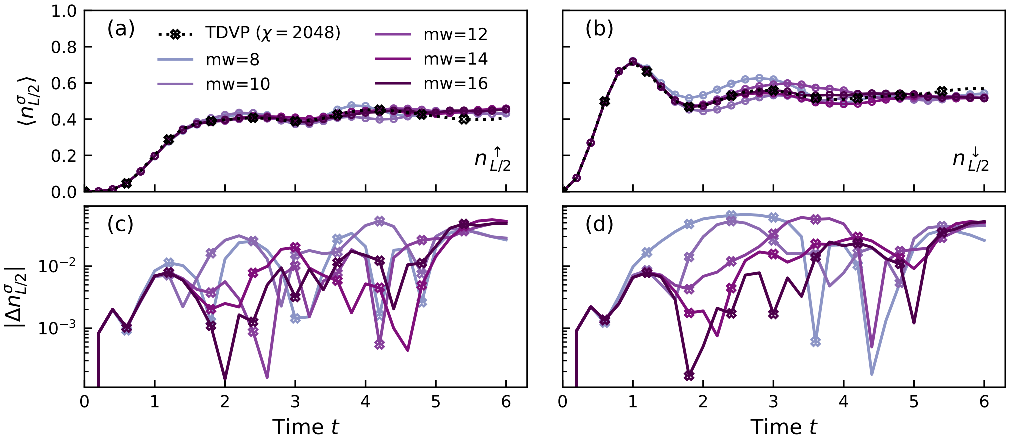

Fig. S6. Pauli Path Propagation (PPP) simulation results and error. PPP was used to simulate the occupation expectation values for the central site in an $L = 60$ chain. The starting state was the Néel state with a vacancy (hole) at the central site. Panels (a) and (b) show the time evolution of $\langle n^{\uparrow}_{L/2} \rangle$ and $\langle n^{\downarrow}_{L/2} \rangle$ for maximum weight thresholds $\mathrm{mw} \in \{8, 10, 12, 14, 16\}$ (light to dark) for the full Pauli weight, and half that for the xy-weight, as described in the main text. The result of a TDVP simulation with $\chi = 2048$ is shown for reference as a black dashed line. Panels (c) and (d) show the corresponding pointwise absolute errors $|\Delta n^{\sigma}_{L/2}| = |n^{\sigma}_{L/2,\mathrm{PPP}} - n^{\sigma}_{L/2,\mathrm{TDVP}}|$, using the $\chi = 2048$ TDVP result as reference.

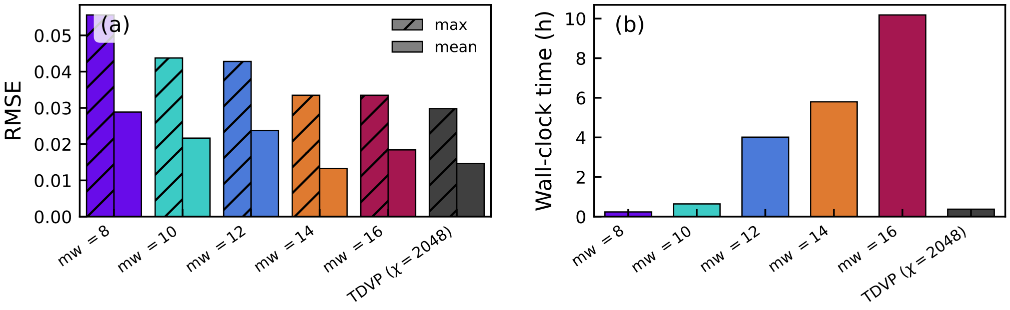

Fig. S7. Central-site RMSE and wall-clock time for $L = 30$ Fermi-Hubbard dynamics compared to a hardware Trotter reference. (a) Time-aggregated RMSE at the central site ($i = \lfloor L/2 \rfloor = 15$, zero-indexed). For each time point $t_k = k\,\Delta t$ ($\Delta t = 0.2$), the per-step RMSE is computed as $\mathrm{RMSE}(t_k) = \sqrt{[(\tilde{n}_{\uparrow} - n^{\mathrm{hw}}_{\uparrow})^2 + (\tilde{n}_{\downarrow} - n^{\mathrm{hw}}_{\downarrow})^2]/2}$, where $\tilde{n}_{\uparrow/\downarrow}$ is the simulated occupation and $n^{\mathrm{hw}}_{\uparrow/\downarrow}$ is the hardware reference occupation. The RMSE is then aggregated across all time points satisfying $t_k \leq 4.5\, t_h^{-1}$ (23 snapshots, $k = 0, \ldots, 22$, including the initial state $t = 0$) using the maximum (hatched, left) and the mean (solid, right) aggregation functions. Methods shown are PPP with maximum Pauli weight $m_w \in \{8, 10, 12, 14, 16\}$ and TDVP with bond dimension $\chi = 2048$. The time is cut-off at $t = 4.5$ to restrict to the region of time for which the TDVP simulation begins to exhibit strong deviations, see Fig. 4 in the main text. (b) Per-site wall-clock time (in hours) required to simulate $t \leq 4.5\, t_h^{-1}$ for each method. For PPP, the cost is the total time recorded at the central site across the same 23 snapshots. For TDVP, the total runtime across all $L = 30$ sites is divided by L to give a per-site cost comparable to PPP.

need to be repeated L times, and an entirely new simulation would be required to obtain higher-point observables, for instance $\langle n_{i,\uparrow} n_{j,\uparrow} \rangle$.

Taken together, these results highlight two key limitations: the accumulation of uncontrolled truncation error due to operator growth, and the rapid increase in computational cost with truncation thresholds. As a result, we find PPP to be significantly less competitive than TDVP for the present problem as a classical simulation benchmark.

III. ERROR ANALYSIS AND VALIDATION

This section provides a characterization of the primary error sources inherent in our digital quantum simulations, specifically algorithmic Trotter error and temporal device drift. Before detailing these contributions, we briefly establish the treatment of statistical shot noise.

Consider a Z-basis Pauli operator P_Z. After N_{shot} shots, the sample proportions of the eigenvalue outcomes (± 1) yield the estimated probabilities $\hat{p}_+ = N_+/N_{\text{shot}}$ and $\hat{p}_- = 1 - \hat{p}_+$, where the hat denotes a statistical estimate. The binomial variance of the estimated expectation value $\widehat{\langle P_Z \rangle}$ is thus given by

$$\mathrm{Var}(\widehat{\langle P_Z \rangle}) = \frac{4\hat{p}_+(1-\hat{p}_+)}{N_{\text{shot}}} = \frac{1 - \widehat{\langle P_Z \rangle}^2}{N_{\text{shot}}} \,. \tag{S.28}$$

These errors are then propagated to other reported quantities, such as the orbital occupations and the spin-spin correlation function.

A. Trotter error and step-size selection

The two-qubit depth of the Trotterized time evolution circuit scales with the number of Trotter steps n_{step} as $D_{2\text{Q}} = 7 + 5(n_{\text{step}} - 1)$. Consequently, the two-qubit gate fidelities of contemporary hardware limit our simulations to a maximum of roughly 30 steps, or a two-qubit depth of 152, to ensure the physical signal remains distinct from hardware noise. With this limitation on the number of steps, the next consideration is the size of the time increment, Δt. This choice involves a fundamental trade-off: larger values of Δt allow us to simulate longer total evolution times $T = n_{\text{step}}\,\Delta t$, but at the cost of increased Trotter error. We choose the time step size Δt to achieve a favorable balance: the circuit gives a total evolution time that is long enough to capture the relevant dynamical features, while keeping the Trotter error, which scales as $\mathcal{O}(\Delta t^2)$ for odd steps and $\mathcal{O}(\Delta t^3)$ for even steps, reasonably small at each time step. The appropriate Δt value depends on the strength of the onsite coupling U—stronger couplings correspond to larger commutator error terms in the Trotter approximation, which require smaller Δt to compensate.

To determine the appropriate Δt for each interaction strength U, we evaluate the Trotter error on small systems. Specifically, we classically simulate the Trotterized circuit at three step sizes—our selected Δt alongside $\Delta t \pm 0.05$—and compare the resulting site occupations against the continuous-time dynamics obtained via exact diagonalization of the Hamiltonian (Fig. S8). We find that using $\Delta t + 0.05$ increases the Trotter error by about one order of magnitude, while using $\Delta t - 0.05$ significantly shortens the total evolution time. Although this analysis is performed for a small system size, we expect the same choice of Δt to provide a similarly favorable balance at larger system sizes.

Furthermore, in Fig. S9, we compare the overall hardware error with the Trotter error for a 10-site example. The hardware error, measured against the exact dynamics, contains contributions from Trotterization, sampling error, and hardware noise. Here, the sampling error—which is much smaller with 20,000 shots—is already included in the Trotter simulation, and therefore forms only a small part of the Trotter error that can be neglected. From Fig. S9, by inspecting the ratio between the RMSE of noiseless simulation and that of hardware, we can see that the Trotter error can account for up to about 60% of the overall hardware error in this example. At this small system size where exact dynamics is possible, we observe that the hardware error is comparable to the Trotter error; however, we expect hardware errors to be dominant at much larger system sizes.

B. Assessment of hardware stability and temporal drift

Superconducting quantum processors are susceptible to temporal drift and fluctuations, necessitating frequent recalibration by dedicated hardware teams [31–33]. Our experiments were conducted via the standard IBM Quantum API [4] on a publicly available device, without performing custom calibrations immediately prior to execution. Accessing the hardware at various intervals relative to the manufacturer's calibration cycle introduces temporal variability that might be absent in a dedicated hardware demonstration that includes full device calibration. To quantify this variability and its impact on our results, we performed a temporal drift study by collecting multiple datasets for the same model parameters using the `ibm_boston` device over several weeks. Crucially, our data acquisition was independent of the hardware's calibration schedule. For each execution, we optimized the qubit layout based on the quality of the device at the time of data collection, following the protocol detailed in Sec. I C.

Figure S10 shows the root-mean-square error (RMSE) of the fermionic occupations $\langle n_{i\sigma} \rangle$ as a function of time for three distinct parameter sets, with each dataset collected at three separate time points over a six-week period. The RMSE is calculated across all sites and spins relative to classical TDVP benchmarks (max bond dimension $\chi = 2048$). To maintain consistency, this data has undergone standard post-processing techniques—the measurement mitigation and decay recovery—we have used throughout our study. The stability of the RMSE values across these disparate time points, characterized by fluctuations of $\sim 2.7\%$ for $L = 30$ [Fig. S10(a)] and $\sim 0.8\%$ for $L = 60$ [Fig. S10(b)], demonstrates that hardware fluctuations do not qualitatively alter the observed dynamics. This consistency confirms the robustness of our experimental findings against temporal device drift.

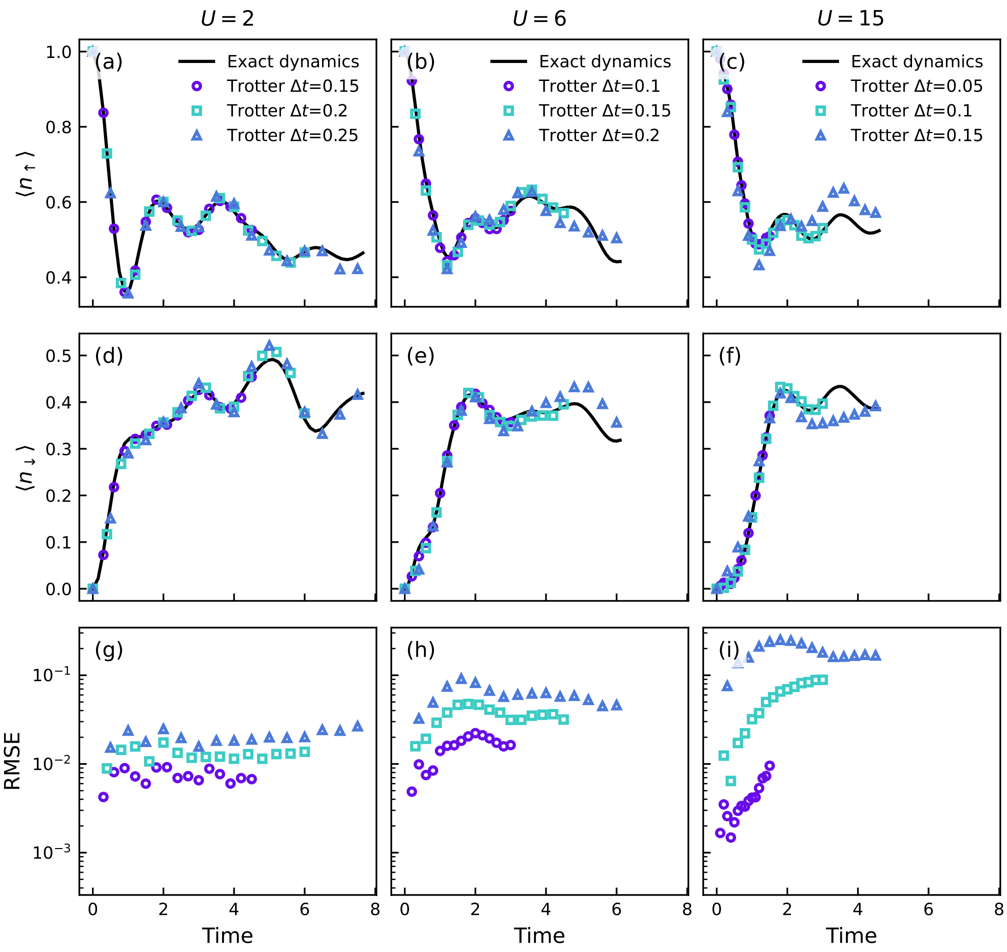

Fig. S8. First-order Trotter error for the 1D Fermi–Hubbard model on $L = 10$ sites with open boundaries, hopping $t = 1$ and chemical potential $\mu = 0$. Each column corresponds to a different on-site interaction strength: $U = 2$ (left), $U = 6$ (middle), and $U = 15$ (right). The initial state is Néel-like $|\uparrow\downarrow\uparrow\downarrow\cdot\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$ with a holon in the middle site $(\cdot)$. All Trotter simulations use 30 steps at varying step sizes Δt, giving different total evolution times $T = 30\,\Delta t$. The top and center rows show the spin-up ($\langle n_\uparrow\rangle$) and spin-down ($\langle n_\downarrow\rangle$) occupation at the middle site (site 5), comparing the exact dynamics (solid black line) with Trotter results (open markers) for each Δt. The bottom row shows the root-mean-square error (RMSE) $\sqrt{\sum_L \mathrm{SE}(t)/L}$, where $\mathrm{SE}(t) = \left(\langle n_\uparrow\rangle_{\mathrm{exact}} - \langle n_\uparrow\rangle_{\mathrm{Trotter}}\right)^2 + \left(\langle n_\downarrow\rangle_{\mathrm{exact}} - \langle n_\downarrow\rangle_{\mathrm{Trotter}}\right)^2$.

IV. POST-PROCESSING

Quantum devices suffer from systematic control imperfections, readout infidelities, and unavoidable environment-induced decoherence, and deep quantum circuits naturally dampen toward a mixed state at the end of hardware execution. In addition to the hardware- and circuit-level error suppression and noise tailoring techniques (see Section I D), some other classes of error reduction techniques have been devised to mitigate this computational bottleneck—these include error mitigation techniques [34–37] that mitigate errors on the expectation values by statistically averaging out some of the noise effect (e.g., readout error mitigation [6, 38–40], zero-noise extrapolation [41], probabilistic error cancellation [35], and damping reversal [14, 42, 43]), and error detection techniques that post-select measurement shots based on violated symmetries [44, 45]; We employed some of the above-mentioned techniques or their variants through post-processing in our experiments, which we detail below. Altogether, they help extend the effective depth and coherence of the hardware execution.

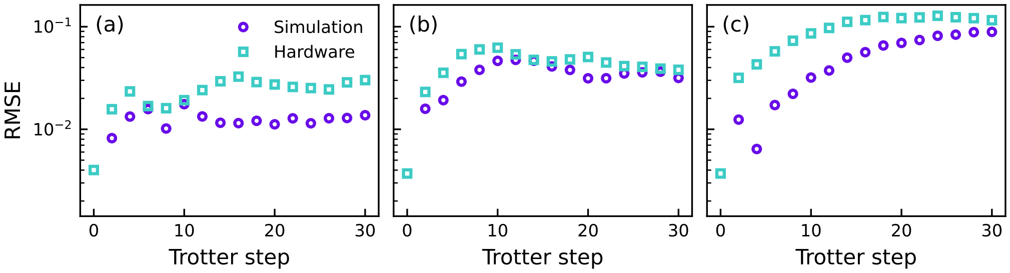

Fig. S9. Same parameter setting and initial state as Fig. S8. Each panel corresponds to a different on-site interaction strength and its corresponding Trotter step size: (a) $U = 2$, $\Delta t = 0.2$, (b) $U = 6$, $\Delta t = 0.15$, and (c) $U = 15$, $\Delta t = 0.1$. We show the root-mean-square error (RMSE) $\sqrt{\sum_L \mathrm{SE}(t)/L}$, where $\mathrm{SE}(t) = \left(\langle n_\uparrow\rangle_{\mathrm{exact}} - \langle n_\uparrow\rangle_{\mathrm{target}}\right)^2 + \left(\langle n_\downarrow\rangle_{\mathrm{exact}} - \langle n_\downarrow\rangle_{\mathrm{target}}\right)^2$, of two different targets: noiseless simulation and hardware execution. Error bars are smaller than the markers and have been omitted for clarity.

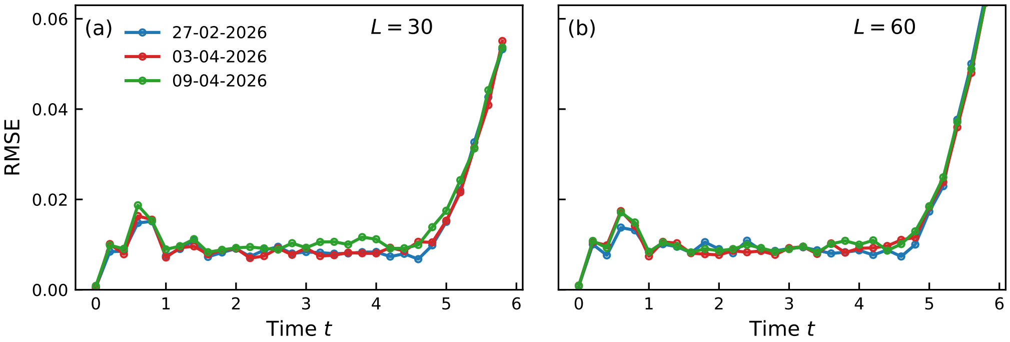

Fig. S10. The root-mean-square error (RMSE) of the fermionic occupations $\langle n_{i\sigma}\rangle$ as a function of the evolution time is shown for the Néel initial state at $U = -2$ for (a) $L = 30$ and (b) $L = 60$ sites. For each case, the RMSE is computed across all sites and spins relative to classical TDVP benchmarks (max bond dimension $\chi = 2048$). Each panel displays three independent datasets (colored lines) collected over a six-week period, illustrating the consistency of the results over time. All data was post-processed using the measurement mitigation and decay recovery protocols described in the text. Error bars are smaller than the markers and have been omitted for clarity.

A. Readout error mitigation

To mitigate the impact of readout noise, we employ a simple readout error mitigation (REM) protocol [6, 38–40]. Immediately following the execution of the Trotter time evolution circuits, we execute characterization circuits to quantify the readout fidelity of each qubit. These circuits consist of preparing every qubit in either the $|0\rangle$ or $|1\rangle$ state, followed by measurement in the computational basis. For each experiment, we collect 32,768 readout characterization shots (16,384 for each $|0\rangle$ and $|1\rangle$ initial state), which allow a set of independent single-qubit confusion (assignment) matrices to be constructed:

$$C_i = \begin{bmatrix} p(0|0)_i & p(0|1)_i \\ p(1|0)_i & p(1|1)_i \end{bmatrix}, \quad i \in \{0, \ldots, 2L-1\}. \tag{S.29}$$

Here, $p(\beta|\alpha)_i$ denotes the probability of measuring the outcome $\beta \in \{0,1\}$ given the preparation of state $|\alpha\rangle \in \{0,1\}$ on qubit i, such that $p(0|0)_i + p(0|1)_i = p(1|0)_i + p(1|1)_i = 1$. We assume the full confusion matrix to be given by a tensor product of the single-qubit matrices, $C = \otimes_{i=0}^{2L-1} C_i$, thus omitting inter-qubit readout correlations, which we find to be negligible on the `ibm_boston` processor. In our experiments, typical error probabilities $p(1|0)_i$ and $p(0|1)_i$ range from approximately 0.06% to 9.7%, with a median value of 0.4%. By strategically selecting the circuit layout on the device (see Sec. I C 4), we avoid qubits with exceptionally high readout errors, which in our characterization experiments reached values as high as $\sim 38\%$.

The raw data from the main circuit execution consists of a measured bitstring distribution. Applying a global REM correction to this distribution is computationally prohibitive as the Hilbert space (and thus the number of produced bitstrings after correction) scales exponentially, i.e. the REM procedure explores the full Hilbert space. However, since our goal is to estimate low-weight Pauli-Z observables, $\langle P_Z\rangle$, we circumvent this exponential cost via marginalization.

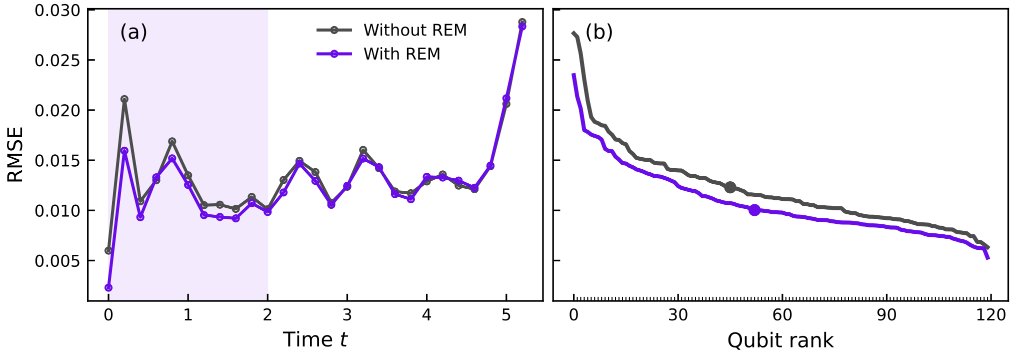

Fig. S11. Impact of readout error mitigation (REM). Data is shown for the $U = -2$, $L = 60$ system prepared in a Néel initial state with a central vacancy. The RMSE of the fermionic occupations $\langle n_{i\sigma} \rangle$ is calculated relative to classical TDVP benchmarks (maximum bond dimension $\chi = 2048$) and presented as a function of evolution time (a) and qubit rank (b). All experiments were performed with deterministic error suppression and Pauli twirling. No additional post-processing or error mitigation techniques were applied. (a) The RMSE is computed across all qubits for each time step. (b) The RMSE is computed over the $t \in [0, 2]$ interval (corresponding to the purple shaded region in (a)) for each qubit, with both datasets (raw and REM-corrected) separately sorted by magnitude. To emphasize that the sorting is independent for each dataset, a single representative qubit (index 30) is marked with a distinct marker on both the raw and REM-corrected lines. Error bars are smaller than the markers and have been omitted for clarity.

For a given observable supported on a subset of qubits, we first marginalize the full output distribution down to that specific support. We then apply the inverse confusion matrices to this reduced distribution to obtain a mitigated marginal distribution, from which the corrected expectation value $\langle \tilde{P}_Z \rangle$ is computed.

Figure S11 quantifies the impact of the REM protocol on the observed fermionic occupations $\langle n_{i\sigma} \rangle$. In Figure S11(a), we plot the RMSE relative to classical TDVP benchmarks as a function of the evolution time. While REM consistently reduces the RMSE, the mitigation gain is most significant during the early stages of the evolution. At later time steps, gate errors become the dominant error source, eventually overshadowing the improvements gained from correcting readout fidelity.

To further resolve the effect during the earlier time steps, Fig. S11(b) displays the RMSE computed for each qubit over the initial time interval $t \in [0, 2]$. Both the raw and REM-corrected data are independently sorted by magnitude to illustrate the global improvement in fidelity across the entire register. Quantitatively, the REM protocol yields a 4.6% reduction in the total RMSE across all time steps, and a more substantial 12.1% reduction within the $t \in [0, 2]$ window. This relatively modest overall improvement reflects the high baseline readout fidelity of the `ibm_boston` processor, as previously indicated by the reported values of the error probabilities within the single-qubit confusion matrices.

B. Decay recovery

To further mitigate noisy expectation values obtained from hardware, we apply a damping-reversal, post-processing pass which we dub *decay recovery* and detail below. Our primary goal with this technique is to empirically correct for the general damping of expectation values caused by hardware noise, as opposed to performing a full inversion of the underlying noise channel. Much like standard mitigation techniques for a depolarizing noise channel, decay recovery relies on extracting damping factors to compute a rescaling factor for the expectation values [8, 42]. However, because actual hardware noise is rarely a simple depolarizing channel, the efficacy of these standard techniques is limited. To address this, our decay recovery protocol makes no assumptions about the specific hardware noise channel. We note the idea of extracting noise-damping factors to rescale observables has been explored previously for the one-dimensional Fermi-Hubbard model [14]. In that work, the authors observed that the damping is largely independent of the interaction strength; by comparing hardware outputs to classical simulations in the tractable, weakly interacting regime, they find the time-dependent noise damping factors and apply them to the strongly interacting regime. While we also compare hardware results against tractable references for calibration, we specifically utilize unitary inversion to isolate the effect of noise from the underlying physics, and our rescaling factors will generally depend on the interaction strength U.

We construct a family of *echo circuits* $\{\mathcal{C}_n\}$ with an even number of Trotter steps $n = 2k$, $k \in \mathbb{Z}$, and where $n = 0$ represents the bare initial state preparation $|\psi_0\rangle$. The n-th echo circuit then applies $n/2$ forward Trotter steps followed

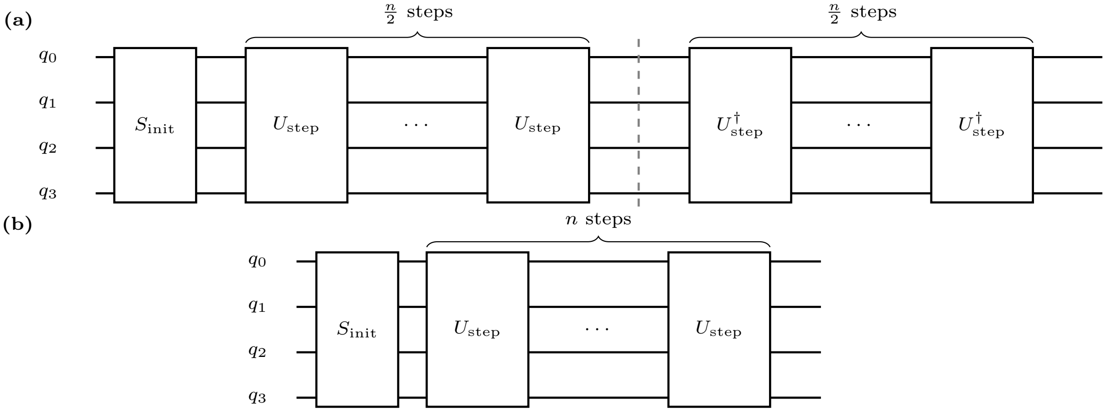

Fig. S12. Decay recovery circuits. (a) The n-th echo circuit used to measure the empirical decay rates. S_{init} prepares an initial Fock state, then $n/2$ forward Trotter steps are applied, followed by $n/2$ inverse Trotter steps. Here, n is assumed to be even. (b) The corresponding time evolution Trotter circuit.

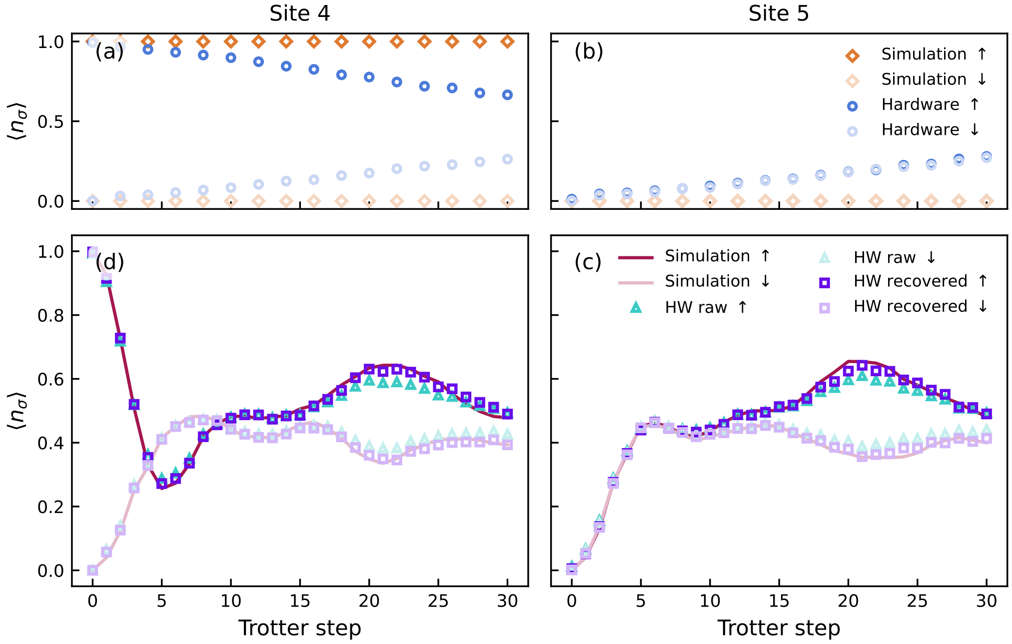

Fig. S13. Decay recovery for a small system with $L = 10$ sites, coupling $U/t_h = 2$, and Trotter step $\Delta t = 0.2$, with the initial state taken to be the Néel state with a central vacancy. Arrows indicate the spin orbitals σ. (a, b) The echo circuit expectation values for sites 4 and 5 (initially spin-up and holon, respectively). (c, d) The site occupation number expectation value for each spin. The raw and mitigated hardware (HW) results are shown, along with the noiseless simulation values. Error bars are smaller than the markers and have been omitted for clarity.

$n/2$ inverse steps:

$$\mathcal{C}_n = \left(U^\dagger_{\text{step}}\right)^{\frac{n}{2}} \left(U_{\text{step}}\right)^{\frac{n}{2}} S_{\text{init}} \,. \tag{S.30}$$

This is depicted in Fig. S12. In the absence of noise, $\mathcal{C}_n$ implements the identity on $|\psi_0\rangle$ for all n. On noisy hardware, the deviation between measured and ideal expectation values grows with circuit depth and varies across physical qubits. By construction, the echo circuit $\mathcal{C}_n$ has the same total depth as a Trotter circuit of n steps.

Next, we detail the decay recovery procedure for a general non-identity Pauli observable P. We assume $\langle\psi_0|P|\psi_0\rangle_n^{\text{echo,ideal}}$ is known and non-vanishing. We note that any observable can be expanded in the Pauli basis, and thus, the decay recovery is applied to each non-identity Pauli term with non-vanishing expectation value individually. Let $\langle P\rangle_n^{\text{hw}}$ denote the measured expectation value for the normal Trotter evolution at step n on hardware, and let $\langle P\rangle_n^{\text{echo,hw}}$ denote the corresponding noisy expectation value measured from the echo circuit on hardware. We note that $\langle P\rangle_n^{\text{echo,hw}}$ can be obtained from hardware only at even Trotter steps $n = 2k$. Therefore, we interpolation between the adjacent even-step echo expectation values to estimate the odd-step echo expectation values: $\langle P\rangle_{2k+1}^{\text{echo, hw}} := \left(\langle P\rangle_{2k+2}^{\text{echo, hw}} + \langle P\rangle_{2k}^{\text{echo, hw}}\right)/2$.

Subsequently, we extract the damping factor from the echo circuit at every Trotter step n as $d_n = \langle P\rangle_n^{\text{echo,hw}}/\langle P\rangle_n^{\text{echo,ideal}}$, and normally d_n is positive. This suggests the simple rescaling $\langle P\rangle_n^{\text{mit}} := \langle P\rangle_n^{\text{hw}}/d_n$, where $\langle P\rangle_n^{\text{mit}}$ is the mitigated Trotter evolution expectation values on hardware. In practice, this bare rescaling often over-compensates for the decay, producing mitigated values that exceed their ideal values. To regularize this behavior, we modify the rescaling procedure as follows:

$$\langle P\rangle_n^{\text{mit}} := \frac{\langle P\rangle_n^{\text{hw}}}{c\, d_n + (1-c)}\,, \quad d_n = \frac{\langle P\rangle_n^{\text{echo,hw}}}{\langle P\rangle_n^{\text{echo,ideal}}}\,, \tag{S.31}$$

where $c \in [0,1]$ is a confidence parameter. By computing a convex combination of the absolute echo factor and 1, we prevent the echo circuits from producing an overly aggressive decay recovery, interpolating between full mitigation ($c = 1$) and raw hardware results ($c = 0$). The confidence parameter c in Eq. (S.31) acts as a single, global scalar applied equally across all sites. To prevent overfitting our error mitigation to the target data, we calibrate this single parameter c strictly out-of-sample. Specifically, we use an independent, small-scale experiment ($L = 10$) where exact noiseless simulation remains tractable. This calibration yielded an optimal value of $c \approx 0.5$, which we then fixed for all subsequent decay recovery at larger system sizes.

Figure S13 illustrates this method for the site occupation $\langle n_{i,\sigma}\rangle = (\langle Z_{i,\sigma}\rangle + 1)/2$, where i and σ index the site and spin. Applying the correction reduces the root-mean-square error (RMSE) (averaged over time) between the hardware and ideal simulation from 2.2×10^{-2} to 1.3×10^{-2}, a 41% improvement. Unless otherwise specified, we apply decay recovery with $c = 0.5$ across all larger-scale experiments to mitigate the noisy expectation values. We note that, in practice, generalizing this procedure requires executing dedicated echo circuits to extract the specific decay factors for every unique initial state and measurement basis.

C. Symmetry post-selection

The Fermi-Hubbard dynamics separately preserve the total number of spin-up and spin-down particles, $[H, N_\sigma] = 0$, where $N_\sigma = \sum_{i=0}^{L-1} n_{i,\sigma}$ is the total number of particles for each spin species σ. Equivalently, total particle number $N_\uparrow + N_\downarrow$ and the z-component of the spin $S^z = (N_\uparrow - N_\downarrow)/2$ are conserved. This is true for both exact and Trotterized time evolution. As our hardware simulations always begin in a state with a definite number of particles, in particular, Fock states of the form

$$|n_{0,\uparrow}\ldots n_{L-1,\downarrow}\rangle = \prod_{i=0}^{L-1}\prod_{\sigma\in\{\uparrow,\downarrow\}} \left(c_{i,\sigma}^\dagger\right)^{n_{i,\sigma}} |\text{vac}\rangle\,, \tag{S.32}$$

the simulated time evolution should be restricted to $\mathcal{H}_{N_\uparrow,N_\downarrow}$, the subset of the full 4^L-dimensional Hilbert space with exactly $N_\sigma = \sum_{i=0}^{L-1} n_{i,\sigma}$ particles for spin species σ. The number of such states is

$$\dim \mathcal{H}_{N_\uparrow,N_\downarrow} = \binom{L}{N_\uparrow} \times \binom{L}{N_\downarrow}. \tag{S.33}$$

Near half-filling, $N_\uparrow = N_\downarrow \approx L/2$, this subset is exponentially large but also exponentially reduced relative to the full Hilbert space. Of course, device noise and imperfections in quantum control cannot be expected to respect the symmetries of the model whose dynamics we wish to digitally simulate. Each gate or idling period will, with some probability, cause the state to exit the $\mathcal{H}_{N_\uparrow,N_\downarrow}$ subspace. We therefore expect that the degree to which the total particle count of each species is not preserved should grow with circuit depth, at least until a saturation point is reached.

This broken symmetry may be put to use for a practical purpose. First, the measured violation rate, $r_{\text{violation}}$, defined as the fraction of measured shots whose Fock states lie outside the particle-number sector of the initial state, provides an indirect diagnostic of device noise and its growth with Trotter step, or equivalently, circuit depth. Second, for Fock-basis observables, one can post-select on the correct particle-number sector by discarding violating shots. In this section, we examine both uses: as a diagnostic for hardware noise and as a possible error-mitigation strategy for Fock-basis expectation values. For the latter, we find that post-selection does not improve the results and often

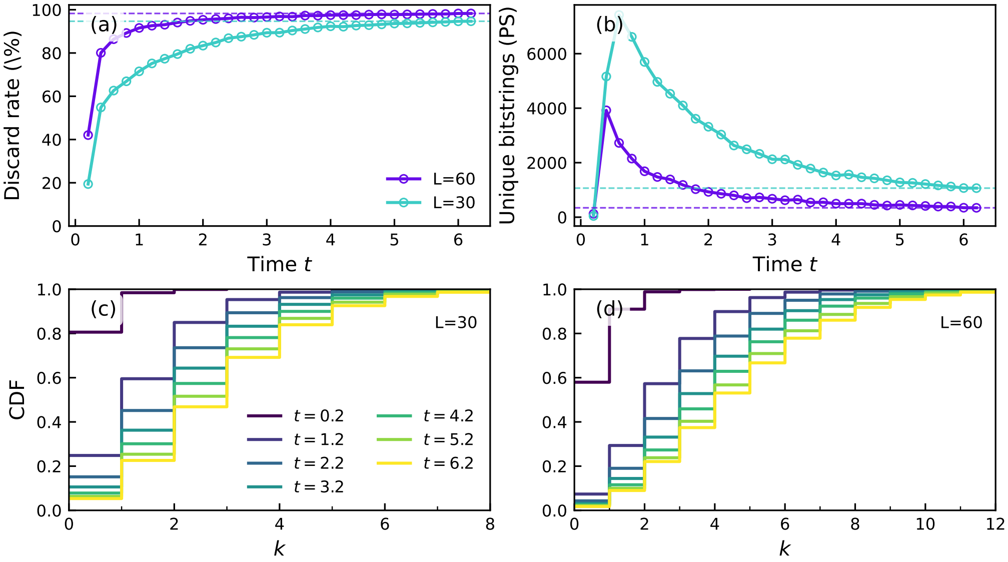

Fig. S14. Hardware noise characterization via particle-number violation. (a) Post-selection discard rate—the fraction of measurements violating total particle-number conservation—as a function of simulation time, for $L = 30$ and $L = 60$ Néel-state simulations with repulsive onsite coupling $U = 4$ and 20,000 shots per step. (b) Number of unique post-selected bitstrings (threshold $K = 0$); the initial count of 1 reflects the Fock-state initial condition. (c,d) Cumulative distribution function (CDF) of the particle-number violation k [Eq. (S.34)] at successive time snapshots for $L = 30$ and $L = 60$ respectively. Early in the simulation most measurements preserve particle number exactly ($k = 0$), with violations confined to small fluctuations ($k \lesssim 4$). As the simulation progresses, the discard rate grows (cf. panels (a,b)) and the k-distribution broadens, reflecting an accumulation of hardware errors that drive the system far from particle-number conservation.

makes them worse; consequently, we do not use this method in the results reported elsewhere in this work. Lastly, we note that violations of a gauge symmetry were recently used to accomplish these tasks in recent work by Ref. [45].

Fig. S14(a) shows the rate of particle-preservation violation for two system sizes, $L = 30$ and $L = 60$. The discard rate quickly grows with simulation time, and eventually saturates at an L-dependent value of 94.7% ($L = 30$) and 98.3% ($L = 60$). The discard rates are higher for the larger system size, consistent with the larger discrepancy between the dimension of the half-filling Hilbert space and the full 2^{2L}-dimensional space. Fig. S14(b) illustrates the number of unique bitstrings after post-selection. The simulation begins with a single Fock state, then initially grows as the Trotterized time evolution generates a superposition over Fock states. After the first few time steps, however, leakage from the symmetry-protected subspace $\mathcal{H}_{N_\uparrow,N_\downarrow}$ becomes significant, and the number of unique bitstrings decreases with the simulation time (or equivalently, with circuit depth).

There is additional structure in the symmetry-violating bitstrings beyond the overall violation rate. To quantify this structure, we define

$$k(s) := \left|N_\uparrow(s) - N_\uparrow^{(0)}\right| + \left|N_\downarrow(s) - N_\downarrow^{(0)}\right|, \tag{S.34}$$

where $N_\sigma(s)$ is the number of σ-spin particles in shot s, and $N_\sigma^{(0)}$ is the corresponding particle number in the initial state $|\psi_0\rangle$. When the shot s is viewed as a measured bitstring, $k(s)$ measures the total particle-number excess or deficit across the two spin sectors. Figures S14(c,d) show the cumulative distribution of k for the same data shown in panels (a,b). At early times, most symmetry-violating shots have small values of k, indicating relatively mild violations of the particle-number constraint. At later times, the distribution shifts toward larger values of k, and very few violating shots correspond to only a single particle-number error, $k = 1$.

Symmetry post-selection amounts to discarding all shots with $k(s) \neq 0$. This is well motivated, since any such bitstring necessarily signals a detected error. However, when the discard rate is $\gtrsim 90\%$, strict post-selection also dramatically reduces the effective sample size. This tradeoff may be worthwhile only if the retained data improve enough to compensate for the increased statistical uncertainty. In our investigations, we found that symmetry post-selection does not improve the results, and in many cases can significantly reduce the accuracy. This is shown in Fig. S15, which uses the same dataset as Fig. 4 in the main text. Panel (a) shows the time evolution of the central-site occupation in an $L = 60$ chain for both spin species, comparing the digital quantum simulation with TDVP. Panel (b) zooms in on the same data and includes the symmetry-post-selected results. Post-selection substantially increases the

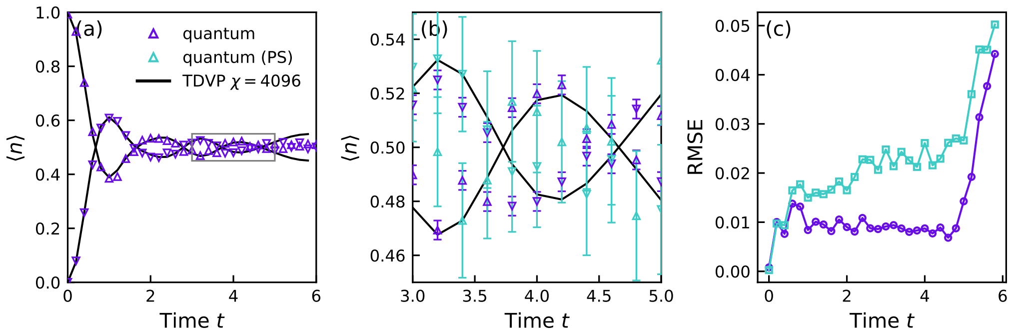

Fig. S15. Impact of symmetry post-selection. All panels show data for a system with $L = 60$ sites, Néel initial state, and coupling $U/t_h = -2$. (a) The spin-up and spin-down occupation numbers for the central site (site 29). The direction of the spin is indicated by the orientation of the triangular markers. Hardware results with measurement mitigation and decay recovery (quantum) are shown in purple circles, TDVP with $\chi = 2048$ is shown in a solid line. (b) Zoomed-in version of (b), with post-selected data shown as well (quantum (PS)). (c) The RMSE between the hardware and TDVP results for both the full hardware data (purple) and symmetry post-selected data (teal).

error bars, and the agreement with TDVP does not appear to improve. A more systematic comparison is shown in panel (c), which reports the RMSE between the quantum results and TDVP, with and without post-selection. Post-selection has little effect at early times, but at intermediate times it significantly worsens the agreement, increasing the RMSE by roughly a factor of two to three.

V. SPIN-CHARGE SEPARATION

In this section, we provide more details for our spin-charge separation experiments and analysis. We first detail our experimental protocol used to extract velocities from hardware and classical simulation results. Lastly, we discuss the comparison between theory and experiment and conduct finite-size scaling analysis.

A. Wavefront detection

In this section, we detail the method used to identify the charge and spin wavefronts for the dynamical quench experiments where the initial state is taken to be the Néel state with a central vacancy defect, that is

$$|\psi_0\rangle = |\uparrow, \downarrow, \ldots, \uparrow, \downarrow, \circ, \downarrow, \uparrow, \ldots, \downarrow, \uparrow\rangle \ . \tag{S.35}$$

Under time evolution, the defect creates a spreading wavefront which travels outwards to the boundary, where it eventually reflects. Due to the charge-spin separation phenomena, this disturbance is a composite structure with contributions from distinct wavefronts for both charge and spin degrees of freedom.

The distinct wavefronts may be isolated by examining observables sensitive to spin and charge separately. For this purpose, we introduce tracer correlators. The tracer correlator for charge is

$$C_i^{\mathrm{c}}(t) = \langle n_{i,\uparrow}(t) + n_{i,\downarrow}(t)\rangle - \langle n_{i,\uparrow}(0) + n_{i,\downarrow}(0)\rangle \,, \tag{S.36}$$

which is simply the difference between the per-site electron density at time t from the value in the initial state. The tracer correlator for spin is the connected spin-spin correlation

$$C_i^{\mathrm{s}}(t) = 4(\langle S_i^z(t) S_{i_*}^z(t)\rangle - \langle S_i^z(t)\rangle\langle S_{i_*}^z(t)\rangle) \,, \tag{S.37}$$

where i_* is the site index of the central defect. Because the initial state is a Fock state, all connected correlators vanish at $t = 0$, so $C_i^s(0) = 0$ without any explicit subtraction — in contrast to the charge tracer, where the subtraction is needed to remove the nonzero background density. The values for these correlators across all sites and across all sampled time points is obtained through either TDVP classical simulations or the results of quantum hardware executions.

We use a wavefront detection algorithm to identify the leading edge of the spreading charge or spin disturbance from the respective tracer correlator. At each time step, the spatial profile of the tracer correlator is smoothed to

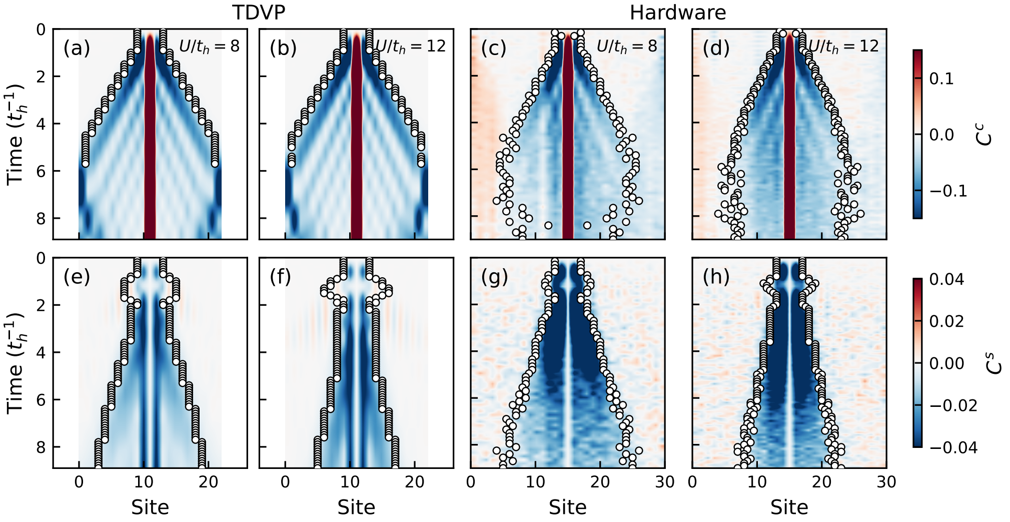

Fig. S16. Wavefront detection algorithm applied to the dynamical quench of a Néel state initialized with a central vacancy defect. White markers indicate the identified left and right boundaries of the propagating wavefronts. Heatmaps depict the charge tracer correlator $C_i^c(t)$ for (a, b) classical TDVP and (c, d) quantum hardware simulations. The corresponding spin tracer correlator $C_i^s(t)$ is shown for (e, f) TDVP and (g, h) hardware simulations. Data is presented for interaction strengths $U \in \{8, 12\} t_h^{-1}$ up to a total evolution time of $T = 9\, t_h^{-1}$. All TDVP simulations use a Trotter step size of $\Delta t_{\text{Trotter}} = 0.1\, t_h^{-1}$. Hardware simulations use $\Delta t_{\text{Trotter}} = 0.15\, t_h^{-1}$ for $U = 8\, t_h^{-1}$ and $\Delta t_{\text{Trotter}} = 0.1\, t_h^{-1}$ for $U = 12\, t_h^{-1}$.

suppress local variations using Gaussian smoothing (we use a smoothing width $\sigma = 1$), and the wavefront position is identified as the innermost site where the signal falls below a chosen fraction p of its instantaneous peak value (we use $p = 30\%$). The velocity is then extracted from the resulting position-time data using a Theil–Sen estimator, a robust non-parametric linear regression that is insensitive to outliers. Finally, for the charge correlator, we enforce the linear fit to have a y-intercept of 0. The full procedure is detailed in Alg. 1.

Algorithm 1 Wavefront Velocity Extraction

```
Inputs: tracer correlator $C_i(t)$, vacancy site $i_*$, smoothing width $\sigma$, threshold fraction $p$
Output: wavefront velocity $v$

for each time step $t$ do
    Spatially smooth $|C_i(t)|$ to suppress local noise
    Walking outward from $i_*$ on both sides and pass the peak $\max_i |C_i(t)|$, use linear interpolation to find the fractional site position where the signal drops below $p \times \max_i |C_i(t)|$
    Average the left and right distances to obtain the wavefront position $x_{\text{wf}}(t)$
end for
Fit $x_{\text{wf}}(t)$ vs. $t$ using Theil–Sen regression to extract $v$
```

In Fig. S16 we illustrate this method on both hardware and TDVP data for different coupling values. The white circles correspond to the detected wavefronts at each time step. From these, a velocity is extracted through the Theil-Sen estimator, which produces the values reported elsewhere and the dashed line boundaries shown in Fig. 2 in the main text.

B. Additional results

In this section, we present extended data supporting our experimental observations of spin-charge separation. Figure S17 provides additional hardware results across a wider range of repulsive couplings than those featured in the main text (Fig. 2). To validate these findings, we performed classical TDVP simulations of identical quench dynamics across the corresponding interaction strengths, with the results detailed in Fig. S18. Velocities are extracted using the wavefront detection algorithm described above (Alg. 1). The extracted velocities are reported in Fig. 2(g-i) in the main text.

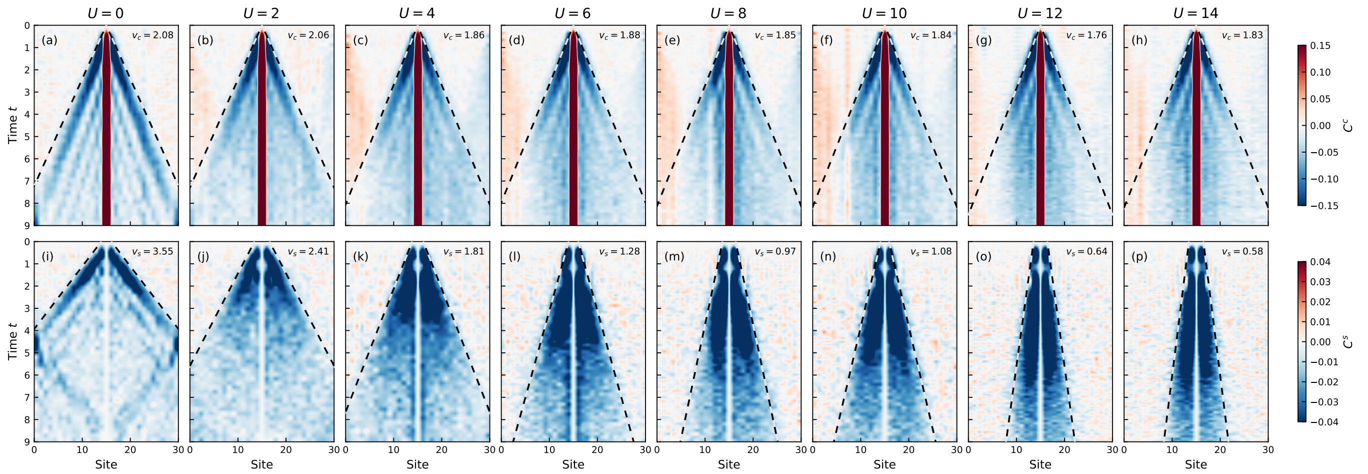

Fig. S17. Quantum hardware simulation of spin-charge separation across varying interaction strengths. The time evolution of a $L = 31$ Néel state initialized with a central vacancy over a range of repulsive couplings U/t_h is simulated. All simulations are evolved to a total time $T = 9\,t_h^{-1}$. For $U/t_h \in \{0, 2, 4\}$ Trotter step size $\Delta t = 0.2\,t_h^{-1}$; for $U/t_h \in \{6, 8, 10\}$ $\Delta t = 0.15\,t_h^{-1}$; for $U \in \{12, 14\}t_h^{-1}$ $\Delta t = 0.1\,t_h^{-1}$. (a–h) Space-time heatmaps of the charge tracer correlator $C_i^c(t)$. The wavefront boundary is indicated by the dashed line. The extracted holon velocity v_h is reported in each panel. (i–p) Space-time heatmaps of the spin tracer observable $C_i^s(t)$. The same wavefront extraction method is applied to determine the spinon velocity v_s. For both sets of panels, a mild Gaussian smoothing has been applied strictly for visualization purposes.

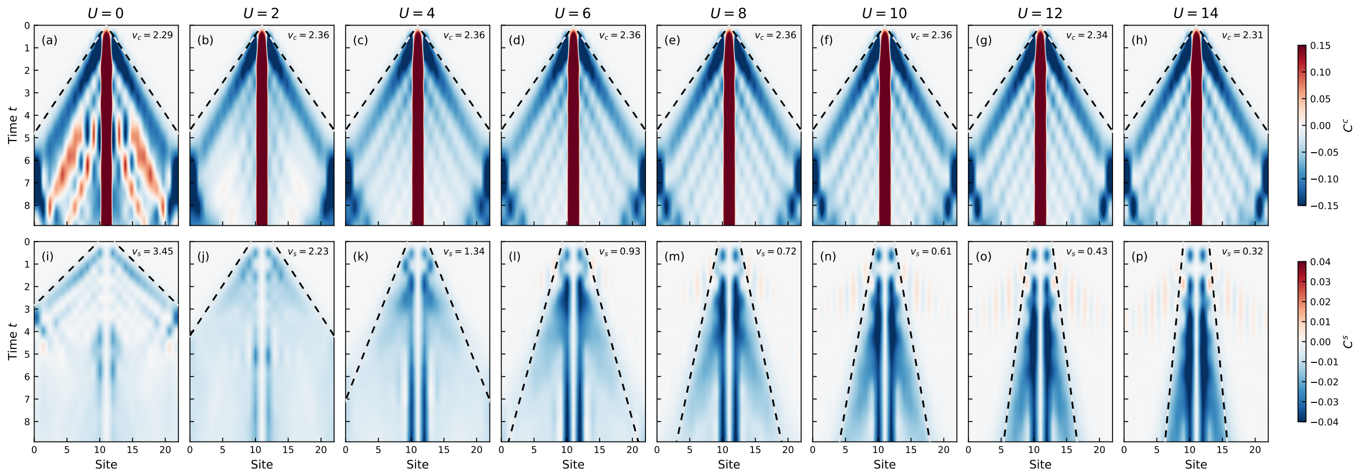

Fig. S18. TDVP simulation of spin-charge separation across varying interaction strengths. All TDVP simulations used a maximum bond dimension of $\chi = 1024$ and $L = 23$. As in Fig. S17, the time evolution of a Néel state initialized with a central vacancy over a range of repulsive couplings U/t_h is simulated. All simulations are evolved to a total time of $T = 8.9\,t_h^{-1}$, with $\Delta t = 0.1$. (a–h) Space-time heatmaps of the charge tracer correlator (the change in hole density, $C_i^c(t)$. The wavefront boundary is indicated by the dashed line. The extracted holon velocity v_h is reported in each panel. (i–p) Space-time heatmaps of the spin tracer observable $C_i^s(t)$. The same wavefront extraction method is applied to determine the spinon velocity v_s. For both sets of panels, a mild Gaussian smoothing has been applied strictly for visualization purposes.

SUPPLEMENTARY REFERENCES

---

[1] P. Jordan and E. Wigner, Z. Phys. **47**, 631 (1928).
[2] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Phys. Rev. Lett. **120**, 110501 (2018).

[3] A. Jüttner and P. Madarasi, Discrete Applied Mathematics **242**, 69 (2018), computational Advances in Combinatorial Optimization.

[4] Qiskit contributors, Qiskit: An open-source framework for quantum computing (2023).

[5] Y. Wang, E. Ginez, J. Friel, Y. Baum, J.-S. Kim, A. Shih, and O. Green, arXiv:2508.21287 (2026).

[6] P. S. Mundada, A. Barbosa, S. Maity, Y. Wang, T. Merkh, T. Stace, F. Nielson, A. R. Carvalho, M. Hush, M. J. Biercuk, and Y. Baum, Phys. Rev. Appl. **20**, 024034 (2023).

[7] P. Coote, R. Dimov, S. Maity, G. S. Hartnett, M. J. Biercuk, and Y. Baum, PRX Quantum **6**, 010332 (2025).

[8] A. Seif, H. Liao, V. Tripathi, K. Krsulich, M. Malekakhlagh, M. Amico, P. Jurcevic, and A. Javadi-Abhari, in *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)* (2024).

[9] G. S. Hartnett, A. Barbosa, P. S. Mundada, M. Hush, M. J. Biercuk, and Y. Baum, Quantum **8**, 1542 (2024).

[10] A. R. R. Carvalho, H. Ball, M. J. Biercuk, M. R. Hush, and F. Thomsen, Phys. Rev. Appl. **15**, 064054 (2021).

[11] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. Carvalho, and M. J. Biercuk, PRX Quantum **2**, 040324 (2021).

[12] J. J. Wallman and J. Emerson, Phys. Rev. A **94**, 052325 (2016).

[13] A. Winick, J. J. Wallman, D. Dahlen, I. Hincks, E. Ospadov, and J. Emerson, Concepts and conditions for error suppression through randomized compiling (2022), arXiv:2212.07500 [quant-ph].

[14] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, Y.-A. Chen, B. Chiaro, R. Collins, S. J. Cotton, W. Courtney, S. Demura, A. Derk, A. Dunsworth, D. Eppens, T. Eckl, C. Erickson, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, J. A. Gross, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, M. Marthaler, O. Martin, J. M. Martinis, A. Marusczyk, S. McArdle, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, C. Mejuto-Zaera, X. Mi, *et al.*, arXiv preprint arXiv:2010.07965 10.48550/arXiv.2010.07965 (2020).

[15] S. Stanišić, J. Haferkamp, D. Hangleiter, J. Eisert, and R. Brierley, Nat. Commun. **13**, 5743 (2022).

[16] F. Alam, J. L. Bosse, I. Čepaitė, *et al.*, arXiv preprint arXiv:2510.26845 10.48550/arXiv.2510.26845 (2025).

[17] T. A. Chowdhury, V. Korepin, V. R. Pascuzzi, and K. Yu, Appl. Phys. Rev. **13**, 011434 (2026), arXiv:2509.14196.

[18] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Physical review letters **107**, 070601 (2011).

[19] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Physical Review B **94**, 165116 (2016).

[20] M. Fishman, S. R. White, and E. M. Stoudenmire, SciPost Phys. Codebases , 4 (2022).

[21] P. Secular, N. Gourianov, M. Lubasch, S. Dolgov, S. R. Clark, and D. Jaksch, Phys. Rev. B **101**, 235123 (2020).

[22] Y. Gao and T. Lubowe, `https://developer.nvidia.com/blog/enabling-matrix-product-state-based-quantum-circuit-simulati` (2023), NVIDIA Technical Blog, Accessed: 2026-05-02.

[23] ITensor Developers, Running on GPUs—ITensors.jl documentation, `https://itensor.github.io/ITensors.jl/dev/RunningOnGPUs.html` (2025), accessed: 2026-05-02.

[24] T. Begušić, K. Hejazi, and G. K. Chan, The Journal of Chemical Physics **162** (2025).

[25] T. Begušić, J. Gray, and G. K.-L. Chan, Science Advances **10**, eadk4321 (2024).

[26] T. Begušić and G. K.-L. Chan, PRX quantum **6**, 020302 (2025).

[27] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, arXiv preprint arXiv:2505.21606 (2025).

[28] E. Fontana, M. S. Rudolph, R. Duncan, I. Rungger, and C. Cîrstoiu, npj Quantum Information **11**, 84 (2025).

[29] A. Angrisani, A. A. Mele, M. S. Rudolph, M. Cerezo, and Z. Holmes, arXiv preprint arXiv:2501.13101 (2025).

[30] C.-J. Lin, H. Gharibyan, and V. P. Su, Quantum **10**, 2014 (2026).

[31] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and J. M. Martinis, Phys. Rev. Lett. **121**, 090502 (2018).

[32] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, npj Quantum Information **5**, 54 (2019).

[33] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, npj Quantum Information **8**, 132 (2022).

[34] W. J. Huggins, S. McArdle, T. E. O'Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, Physical Review X **11** (2021).

[35] E. van den Berg, Z. K. Minev, A. Kandala, and K. Temme, Nature Physics **19**, 1116–1121 (2023).

[36] Y. Kim *et al.*, Nature **618**, 500 (2023).

[37] H. Liao, D. S. Wang, I. Sitdikov, C. Salcedo, A. Seif, and Z. K. Minev, Nature Machine Intelligence **6**, 1478 (2024).

[38] F. B. Maciejewski, Z. Zimborás, and M. Oszmaniec, Quantum **4**, 257 (2020).

[39] S. Bravyi, S. Sheldon, A. Kandala, D. C. Mckay, and J. M. Gambetta, Phys. Rev. A **103**, 042605 (2021).

[40] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, PRX Quantum **2**, 040326 (2021).

[41] K. Temme, S. Bravyi, and J. M. Gambetta, Phys. Rev. Lett. **119**, 180509 (2017).

[42] F. Alam, M. Crichigno, E. Crosson, S. T. Flammia, F. M. Gambetta, M. H. Gordon, M. Kreshchuk, A. Montanaro, A. Nocera, and R. A. Santos, arXiv:2603.12236 (2026).

[43] Google Quantum AI and Collaborators, Nature **646**, 825 (2025).

[44] H. Liao, G. S. Hartnett, A. Kakkar, A. Tan, M. Hush, P. S. Mundada, M. J. Biercuk, and Y. Baum, PRX Quantum **6**, 020331 (2025).

[45] G. S. Hartnett, H. Liao, and E. Rinaldi, arXiv:2604.14094 (2026).