1

Condensed matter physics

Condensed matter physics constructs and studies the behavior of simplified models designed to capture the universal physics of material systems. Phenomena of interest include magnetism, phase transitions, superconductivity, frustrated systems, topological phases, and the interplay of thermalization and many-body localization in closed systems. While many seminal models can be studied analytically in certain limits (e.g., the 1D and 2D classical Ising model), a number of seemingly innocuous models have proven exceedingly difficult to solve. This has led to some models, such as the Fermi–Hubbard model, becoming a proving ground for classical numerical methods. While there has been significant progress in recent decades in understanding the physics of these models through numerical simulation, it is still a challenging problem for many models and parameter regimes. As observed by Feynman [391], quantum computers have a natural advantage over their classical counterparts for simulating the simple Hamiltonians studied in condensed matter physics. While Feynman's proposal was more focused on analog simulation, digital quantum simulation of condensed matter systems has evolved into a major research direction. In this chapter, we focus on models whose end-to-end complexities have been well studied in the literature: the Fermi–Hubbard model, spin models, and the Sachdev–Ye–Kitaev (SYK) model.

The authors are grateful to Ashley Montanaro and Nobuyuki Yoshioka for reviewing this chapter.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

1.1 Fermi–Hubbard model

Overview

The Fermi–Hubbard model was originally introduced as a simplified model of electrons in materials [546], closely related to the tight-binding model. It displays a wide range of behaviors including metallic, insulating, and antiferromagnetic phases. The model has more recently found applicability in studying high-temperature superconductivity. The 2D Fermi–Hubbard model has a complex phase diagram that appears to reproduce universal (rather than chemical-specific) features of the phase diagram of cuprate high-temperature superconductors.

General analytic solutions are not known beyond 1D chains or specific parameter regimes—see [59] for a discussion—which has motivated the use of numerical methods to understand the physics of the Fermi–Hubbard model. More recently, there has been increased interest in understanding the nonequilibrium properties of the model such as its behavior following a quench.

Based on the current estimates, quantum simulation of Fermi–Hubbard models requires considerably fewer resources than simulations of molecules or solving optimization problems. This makes the Fermi–Hubbard model a promising candidate for early demonstrations of quantum advantage.

A detailed case study on the Fermi–Hubbard model is presented in [101], including descriptions of the parameters to probe open scientific questions and estimates of the utility of these computational capabilities.

Actual end-to-end problem(s) solved

The Fermi–Hubbard Hamiltonian on $M/2$ lattice sites is given by

$$H = -t \sum_{\sigma\in\{\uparrow,\downarrow\}} \sum_{\langle i,j\rangle} (c^\dagger_{i\sigma} c_{j\sigma} + c^\dagger_{j\sigma} c_{i\sigma}) + U \sum_i n_{i\uparrow} n_{i\downarrow} \,, \tag{1.1}$$

where $c_{i\sigma}$ are fermionic operators and $n_{i\sigma} \equiv c^\dagger_{i\sigma} c_{i\sigma}$ is the number operator, with t denoting the strength of the kinetic term, U the onsite interaction strength, and $\langle i, j\rangle$ a sum over nearest-neighbor lattice sites, given a lattice geometry. It is also possible to consider longer-range hopping terms, the inclusion of site-dependent chemical potentials, or additional "orbitals" per site.

Quantum simulation provides insights into both equilibrium and nonequilibrium physics. With regards to equilibrium physics, the primary computational task is to resolve and probe the properties of the phase diagram of the Fermi–Hubbard model, as a function of lattice geometry, parameter values (t, U), doping (the expected number of fermions divided by the number of sites), and temperature. This is achieved by preparing the thermal state $\rho \propto \mathrm{e}^{-\beta H}$ (with

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

$\beta = 1/k_B T$, where k_B is the Boltzmann constant and T the temperature) or at zero temperature the ground state $|E_0\rangle$ for the Fermi–Hubbard Hamiltonian instantiated by the given parameters, and measuring the expectation values of a set of physical observables to error ϵ. A thorough discussion of this end-to-end problem (at zero temperature) is provided in [1031], where it is shown how to perform the following steps:

- Prepare mean-field states in a given phase, for example, a BCS superconducting ground state.
- Adiabatically evolve from the mean-field Hamiltonian to the final Fermi–Hubbard Hamiltonian. The absence of a phase transition confirms the predicted phase.
- Measure observables, including density correlation functions $(n_{i\uparrow}+n_{i\downarrow})(n_{j\uparrow}+n_{j\downarrow})$, pair correlation functions $c^\dagger_{i\sigma} c^\dagger_{j\sigma'} c_{k\sigma'} c_{l\sigma}$, and dynamical correlation functions $\langle E_0|\mathrm{e}^{\mathrm{i}Ht} A \mathrm{e}^{-\mathrm{i}Ht} B|E_0\rangle$ (for operators A, B and ground state $|E_0\rangle$).

The difficulty of this problem depends on the parameter regime under consideration. The ground state in the weak coupling regime of $U < 4t$ is well understood, but questions remain in the intermediate ($4t \leq U \leq 6t$) and strong ($U > 6t$) regimes [847]. Challenges include precisely determining the phase boundaries and understanding the nature of the superconducting phase [399]. Progress has been made on this latter question in recent years, for example, by showing the absence of a superconducting phase at the physically relevant parameters of $U \sim 8t$ and $1/8$th doping (see [847] for a more detailed discussion). Calculations are made challenging by small energy differences between competing phases, as well as the need to extrapolate from finite simulations to the thermodynamic limit.

The simulation of nonequilibrium quantum dynamics is of interest for modeling materials driven by an external field (e.g., an ultrafast laser pulse or an applied voltage), or following a quench in the Hamiltonian. Classically simulating nonequilibrium quantum dynamics has so far proven challenging and is a less well-studied problem than probing the equilibrium physics of the model. Example applications include modeling ultrafast spintronics (whereby lasers are used to manipulate spin degrees of freedom to control and store information) [1094], understanding photo-induced phase transitions [812], and clarifying the nature of thermalization in isolated quantum systems following a quench [837].

Dominant resource cost/complexity

Mapping the problem to qubits: Simulation of the Fermi–Hubbard model is most naturally performed in the second-quantized representation, as the regime

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

of interest is usually close to half-filling (for comparison, we refer to Chapter 2 on simulating molecules). The Jordan–Wigner mapping between fermions and qubits is typically used. Locality-preserving mappings have also been developed, which map fermionic operators to qubit operators acting on a constant number of qubits [1010, 344]. For an $L \times L$ lattice, we require $M = 2L^2$ qubits to simulate the spinful Fermi–Hubbard model using the Jordan–Wigner mapping.

Accessing the Hamiltonian: Quantum algorithms for simulating the Fermi–Hubbard model require access to the Hamiltonian. This is typically provided by block-encoding or Hamiltonian simulation.[1] The structure in the Fermi–Hubbard Hamiltonian reduces the costs of these subroutines. For example, performing a block-encoding using the linear combinations of unitaries (LCU) technique requires access to a PREPARE unitary and a SELECT unitary (we refer to Section 10.2 for definitions). The PREPARE unitary requires preparing a quantum state from classical data. Because the Fermi–Hubbard Hamiltonian has a small number of unique coefficients, the cost of this unitary can be reduced. Combining the results of [75, 1069, 225], one can implement an $(M(2t + U/8), O(\log(M)), \epsilon)$-block-encoding (see Eq. (10.1) for definition) of the Fermi–Hubbard Hamiltonian using

$$O\left(M + \log(M/\epsilon)\right)$$

non-Clifford gates.

As another example, the costs of product formula approaches for Hamiltonian simulation can exploit the fact that many terms in the Fermi–Hubbard Hamiltonian commute, due to their locality. We will explicitly discuss these costs below.

State preparation:

- Classical trial states: Approximate eigenstates obtained from a classical calculation can be prepared as quantum trial states; examples include Slater determinant states [1031], linear combinations of D Slater determinants (with complexity $\widetilde{O}(D)$ [394]–$O(MD)$ [998]), and matrix product states with bond dimension χ (with complexity $O(M\chi^2)$ [394]).
- Quantum trial states: Parameterized quantum circuits, in conjunction with variational quantum algorithms, have been proposed for preparing approxi-

[1] Hamiltonian simulation is used to explicitly simulate dynamics but can also be used implicitly to provide access to the Hamiltonian for use in static calculations, for example, in quantum phase estimation.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

mate energy eigenstates (see §NISQ implementations, below). Like classical trial states, the states prepared by these circuits can be used as inputs to other quantum algorithms that further refine the initial state, such as eigenstate filtering, or quantum phase estimation. Initial resource estimates for the Fermi–Hubbard model can be found in [215].

- Eigenstate preparation: There exist quantum algorithms that can prepare energy eigenstates using QSVT-based eigenstate filtering [688], where the cost scales as $1/\gamma$ with γ the overlap of the initial state with the desired eigenstate. Alternatively, adiabatic state preparation can be used, with a cost that depends on the gap between energy levels along the adiabatic path. Adiabatic state preparation was proposed as a method of classifying the phase diagram of the Fermi–Hubbard model [1031]. A discrete version of the adiabatic approach based on qubitization and quantum phase estimation (QPE) was numerically investigated in the context of preparing ground states of the Fermi–Hubbard model [673], and showed promising results for the small system sizes considered (see also [998]).
- Thermal states: A number of algorithms have been developed for preparation of thermal states, also known as Gibbs states. The most promising variants of these "Gibbs sampling" algorithms depend on the mixing time of a Markov chain (similar to classical Monte Carlo approaches for preparing Gibbs states), which is currently undetermined for the Fermi–Hubbard model.

Time evolution:

- As discussed above, Trotter approaches for Hamiltonian simulation can exploit beneficial features of the Fermi–Hubbard Hamiltonian, such as locality, fixed particle number, and commutativity of the terms [280, 300, 963]. For a Fermi–Hubbard model with η fermions on M lattice sites, pth-order Trotter methods can simulate time evolution for time τ up to error ϵ using

$$O\left(\frac{5^p M \eta^{1/p} \tau^{1+1/p}}{\epsilon^{1/p}}\right)$$

gates. Explicit gate counts for Trotterization can be obtained from [630, 300, 225, 909], which have focused on constant prefactors for low-order product formulas, rather than the asymptotic scaling.

Post-Trotter methods, such as [479], using quantum signal processing as a building block, can achieve similar scaling in M and τ. A suboptimal approach (i.e., not using the method of [479]), briefly discussed in [392], has

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

a gate complexity of approximately

$$44M^2(2t + 3U/8)\tau$$

T gates to simulate time evolution for time τ using quantum signal processing, neglecting logarithmic dependence on the error of the simulation.[2]

Measuring observables:

- Energies: QPE can be used to measure the energy eigenvalues of the Fermi–Hubbard Hamiltonian. Given access to (i) an initial state $|\psi\rangle$ that has sufficient overlap $\gamma = |\langle\psi|E_j\rangle|$ with the target eigenstate $|E_j\rangle$ and (ii) a unitary $U = f(H)$ that encodes the eigenspectrum of the Hamiltonian with a known, classically invertible relationship f, we can use QPE to project into the desired eigenstate and provide an estimate of the eigenphase ϕ_i of U—which can then be converted into an estimate of the eigenenergy of H using $\phi_i = f(E_i)$. QPE makes

$$O\left(\gamma^{-2}\|f'(H)\|^{-1}\epsilon^{-1}\log(\theta^{-1})\right)$$

calls to the unitary U encoding the spectrum of the Hamiltonian, where θ is the failure probability, and ϵ is the desired precision in the eigenenergy of H.[3]

A common choice for the unitary encoding the Hamiltonian is $U \approx e^{iHt}$ approximated via quantum algorithms for Hamiltonian simulation, where the approximation error must be balanced against the error from QPE. Using $U \approx e^{iHt}$ implemented via a second-order product formula results in a T gate count of $O(M^{3/2}/\Delta E^{3/2})$ to resolve the energy of the Fermi–Hubbard model to precision ΔE, neglecting the cost of initial state preparation and the dependence on the overlap and failure probability [630, 225]. Another common choice is to perform QPE on a quantum walk operator $W(H)$ which acts like $e^{i\arccos(H/\alpha)}$, where α is the normalization of the block-encoding of H. The operator $W(H)$ can be implemented exactly via qubitization [841, 139]. This results in a T gate scaling of $O(M^2/\Delta E)$, also neglecting the cost of initial state preparation and the dependence on the overlap and failure probability [75].

[2] Note that in [392], M is defined as the number of lattice sites, and so corresponds to $M/2$ here.

[3] It is possible to improve the complexity to $O(\gamma^{-1}\|f'(H)\|^{-1}\epsilon^{-1}\log(\theta^{-1}))$ using amplitude amplification if a sufficiently precise estimate of the eigenvalue is known, or to $O((\gamma^{-2}\Delta^{-1} + \epsilon^{-1})\|f'(H)\|^{-1}\log(\theta^{-1}))$ by exploiting knowledge of the gap Δ between the energy eigenstates to perform rejection sampling [139].

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

- Other observables: There have been few studies considering the costs of measuring observables other than the ground state energy using fault-tolerant quantum algorithms. In general, it is important to minimize the number of calls to the unitary U_ψ that prepares the desired state, as this is typically considered the dominant cost. Reference [1031] discussed methods for measuring density correlation functions $(n_{i\uparrow} + n_{i\downarrow})(n_{j\uparrow} + n_{j\downarrow})$, pair correlation functions $c^\dagger_{i\sigma} c^\dagger_{j\sigma'} c_{k\sigma'} c_{l\sigma}$, and dynamical correlation functions $\langle E_0 | e^{iHt} A e^{-iHt} B | E_0 \rangle$ (for operators A, B and ground state $|E_0\rangle$), including approaches for nondestructively measuring some of these observables. Some of these approaches can now be reframed as performing amplitude estimation [637] on U_O, a unitary block-encoding of the observable O with subnormalization factor α_O [853]. The measurement of similar observables using these modern algorithmic tools was studied in [9].

  A recent approach [549, 49] based on the quantum gradient estimation algorithm of [430] simultaneously computes the value of K (noncommuting) observables O_j. The algorithm makes $\widetilde{O}(K^{1/2}/\epsilon)$ calls to U_ψ and $U_\psi^\dagger$ (or $R_\psi = I - 2|\psi\rangle\langle\psi|$) and either $\widetilde{O}(K^{3/2}/\epsilon)$ calls to gates of the form e^{ixO_j} [549] or $\widetilde{O}(K/\epsilon)$ calls to a block-encoding of the observables [49]. The algorithm also requires $O(K \log(1/\epsilon))$ additional qubits. This approach has been considered in the context of measuring fermionic reduced density matrices and dynamic correlation functions [549].

Existing resource estimates

There have been a number of logical resource estimates for algorithms targeting both static and dynamic properties of the Fermi–Hubbard model. In Table 1.1, we present approximate resource estimates for simulations of the 2D 10×10 spinful Fermi–Hubbard model. The table presents the number of logical qubits and gates required to run the algorithm; these can be converted into physical resource estimates via methods for fault-tolerant quantum computation.

References [75, 1069] applied qubitization-based QPE to calculate the ground state energy to constant additive error. For a lattice with M spin sites, using the compilation of [75], the number of T gates scales as approximately (neglecting the dependence on the overlap and failure probability) [75, Eq. (61)]

$$\#T \propto \frac{(4t + U)M^2}{\Delta E},$$

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

and the number of logical qubits scales as approximately [75, Eq. (62)]

$$\#\text{Qubits} \sim M + \log\left(\frac{(2t + 0.5U)M^4}{\Delta E}\right).$$

References [630, 225] applied second-order Trotter-based QPE to calculate the ground state energy. In both references, rigorous but potentially loose upper bounds on the Trotter error are computed. For a lattice with M spin orbitals, using the compilation of [225], the number of T gates scales as approximately (neglecting the dependence on the overlap and failure probability) [225, Eqs. (C3), (D6), (D10), (E17), (F10)]

$$\#T \propto t\sqrt{t + U}\left(\frac{M}{\Delta E}\right)^{3/2},$$

and the number of logical qubits scales as approximately [225, Table II]

$$\#\text{Qubits} \sim (1 + \kappa)M,$$

where κ is a free parameter that controls the number of ancilla qubits used for a compilation technique known as Hamming weight phasing (which reduces the cost of applying identical arbitrary angle rotation gates in parallel) [421, 630], set to $\kappa = 0.25$ in [225] and in our Table 1.1.

Problem and method	# T gates	# Logical qubits	Parameters
Ground state energy via qubitized QPE [75, 1069]	$\sim 10^8$	~ 236	$U/t = 4$ and $\Delta E = 0.01t$
Ground state energy via Trotterized QPE [225, 630]	$\sim 5 \times 10^6$	~ 250	$U/t = 8$ and $\Delta E = 0.005E_{\text{tot}}$
Dynamics via fourth-order Trotter [392]	4.6×10^5	200	$T = 10/t$, $U = t$, and $\epsilon \leq 1\%$

Table 1.1 Logical resource estimates for quantum phase estimation (QPE) and dynamics simulation applied to a 2D 10×10 Fermi–Hubbard model. The QPE circuits target an energy error of ΔE. In the second row, E_{tot} denotes the ground state energy. The dynamics simulation runs for time T, and targets an error of less than 1% in a spatially averaged intensive observable, with Trotter errors bounded numerically via extrapolated small-scale simulations. The presented gate counts are for a single run of the circuit. For QPE, the number of required runs depends on the overlap between the initial state and the ground state (inverse polynomial dependence), as well as the desired failure probability of QPE. For dynamics simulations, the number of circuit repetitions depends on the precision to which one wants to estimate a given observable. The parameters for each problem vary between different rows of the table, and so cannot be directly compared (although the different methods for the same problem, e.g., ground state energy estimation, could be compared by adjusting the analyses in the original papers to the desired matching parameter values).

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

The methods described above for encoding the Hamiltonian spectra (qubitization and Trotter) can also be used to simulate the dynamics of the Fermi–Hubbard model. Trotter methods can be applied directly, while qubitization can be combined with quantum signal processing (QSP) to perform Hamiltonian simulation. In [392], a comparison was made between fourth-order Trotterization and qubitization+QSP for simulating time evolution of a 10×10 Fermi–Hubbard model. Trotter was determined to be the more efficient method, although this conclusion hinges on a Trotter decomposition with large steps (justified via numerical simulations). We note that the Trotter decompositions and analyses in [225, 392] are different, which hampers an immediate comparison. It may also be fruitful to compare with Hamiltonian simulation algorithms designed explicitly for simulating local Hamiltonians [479] (see discussion in [75]).

Caveats

In general, preparing the ground state of the Fermi–Hubbard model is known to be a hard problem, even for a quantum computer. This task has been proven QMA-hard for the Fermi–Hubbard model with a site-dependent magnetic field [910] and for the Fermi–Hubbard model with a site-dependent kinetic term strength (i.e., $t \to t_{ij}$ in Eq. (1.1)) [808]. While the complexity class of the canonical Fermi–Hubbard model is not yet known, when preparing the ground state via QPE or eigenstate filtering methods, it is necessary to prepare an initial state with an overlap that decays no worse than polynomially with system size; otherwise, the overall complexity will be superpolynomial. While numerical simulations on small system sizes have shown encouraging results [998, 673], it is still an open question as to whether this property holds for sufficiently large system sizes to enable extrapolation to the thermodynamic limit.

It is also important to note that this extrapolation of measured properties, computed at a range of finite system sizes, to the thermodynamic limit, has been observed to contribute a significant proportion of the uncertainty and errors in classical methods [666], and will also afflict quantum simulations.

Finally, it will be necessary to repeat simulations a large number of times. In order to measure a single observable to precision ϵ we require $O(1/\epsilon^2)$ incoherent repetitions of the simulation, or $O(1/\epsilon)$ using methods based on amplitude estimation. To map out and compute properties of the phase diagram or extract the phase following a quench, we may need to measure a large number of observables. In some cases, it may be necessary to re-prepare the initial state for each observable.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

Comparable classical complexity and challenging instance sizes

The Fermi–Hubbard model has been a fertile environment for the development and testing of classical numerical methods for both static and dynamical properties. To the best of our knowledge, the largest exact diagonalization calculations performed to date are on systems with 17 fermions in 22 sites (44 spin sites), requiring 7.1 terabytes of memory [1063]. State-of-the-art approximate methods for computing the phase diagram include quantum Monte Carlo methods (variational QMC, determinantal QMC, diagrammatic MC, auxiliary-field QMC, diffusion MC), density matrix renormalization group (DMRG), coupled cluster methods, and impurity methods (dynamical mean-field theory, density matrix embedding theory), among others. These methods typically have an approximation parameter (e.g., the excitation degree in coupled cluster or the bond dimension in DMRG) which influences the scaling of the algorithm and the accuracy of the simulation. Modern numerical studies of the Fermi–Hubbard model typically cross-validate using a number of simulation methods [666, 900]. For example, [666] benchmarked a range of methods and performed sufficiently large and accurate simulations for extrapolation to the thermodynamic limit. That work concluded that "the ground-state properties of a substantial part of the Hubbard model phase space are now under numerical control," but that some uncertainties still remain for $4t \leq U \leq 8t$ and dopings near half-filling. For a recent review of numerical simulations of the Fermi–Hubbard model, we refer the reader to [847]. We also refer to [1054], which benchmarked a number of variational classical methods on a range of condensed matter systems, including the Fermi–Hubbard model.

The simulation of *dynamics* of the Fermi–Hubbard model appears to be more challenging for classical methods. For example, [326, 392] concluded that simulating the dynamics of a 10×10 lattice would be infeasible for tensor network techniques. Other classical approaches for simulating time evolution of the Fermi–Hubbard model include nonequilibrium extensions of dynamical mean-field theory [43] or Floquet methods [812].

Speedup

The speedup of quantum algorithms for computing static properties of the Fermi–Hubbard model, such as its ground state energy, is difficult to determine. In general, we know that closely related models are QMA-hard (see §Caveats, above) and so should be exponentially difficult for both classical and quantum computers. Assuming an initial state that has overlap with the target eigenstate that decays no faster than polynomially, then QPE can be used to efficiently measure the eigenenergy and project into the desired eigenstate. It does so with cost poly$(M/\Delta E)$, where the precise scaling depends on

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

the quantum algorithm used. Exact classical methods such as exact diagonalization have a cost that scales exponentially with M or $1/\Delta E$. Approximate classical methods scale with an approximation parameter (e.g., bond dimension, number of excitations) which will depend on both M and ΔE. For example, [670, Fig. 4] shows the convergence of a tensor network calculation for the 2D Fermi–Hubbard model as a function of bond dimension and system size. For the small systems studied (up to 16×4 sites) the plots are consistent with the bond dimension scaling polynomially in $1/\Delta E$, with a weak dependence on the system size. If this holds for larger system sizes and across a range of system parameters, this would suggest that quantum algorithms provide only a polynomial speedup for computing the ground state energy.

Simulating the dynamics of the Fermi–Hubbard Hamiltonian requires polynomial resources using quantum algorithms, scaling almost linearly both in M and in the evolution time τ. By using carefully engineered interactions (e.g., deviating significantly from a square lattice) it can be shown that simulating the dynamics of the Fermi–Hubbard model on a planar graph is a BQP-complete problem and so is expected to be hard for classical computers in the worst case [93]. Supporting this observation, all known classical methods appear to scale exponentially in system size and simulation accuracy. For example, [392] used tensor network (matrix product state) approaches for simulating the dynamics of the Fermi–Hubbard model following a quench. When truncating the bond dimension to facilitate efficient classical simulation, they found that errors in the observables grew exponentially with time. While this supports the conclusion of an exponential quantum speedup, we note that classical approaches will likely continue to improve and be applied to increasingly large system sizes.

NISQ implementations

There have been a number of proposals (and experimental demonstrations) for simulating the Fermi–Hubbard model on NISQ hardware. Ground state calculations can be performed using the variational quantum eigensolver (VQE) [576, 871, 872, 218, 215], and experimental demonstrations have been carried out on lattices of size 1×8 and 2×4 using 16 superconducting qubits, yielding qualitative agreement with theoretical expectation [951].

Dynamics can be simulated using Hamiltonian simulation (typically Trotter methods) [300] and have been demonstrated for an 8×1 lattice on 16 superconducting qubits [64].

The simple Hamiltonian of the Fermi–Hubbard model makes it well suited to realization in analog quantum simulators, including ultracold atoms in optical lattices, trapped ions, and neutral atom arrays. It has been argued that some

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

local observables can be robust to errors in the simulation [836, 392], enabling analog simulations to already surpass classical methods for simulating dynamics. Nevertheless, it can be challenging to cool the analog fermionic system to its ground state. We refer the reader to [326, 461] for additional discussion on analog simulation.

Reference [300] considered an approach with flavors of both digital and analog simulation, moving to a cost model based on evolution time, rather than number of gates. This reduces the cost of simulating Trotterized dynamics of the Fermi–Hubbard model, compared to purely gate-based approaches.

Outlook

The Fermi–Hubbard model provides a long-standing and physically relevant computational challenge. The low gate counts and modest number of logical qubits required to compute ground state energies could make quantum algorithms competitive with leading classical approaches in challenging regimes. We note that further research is required to ascertain the costs for initial state preparation for these calculations. For the less well-studied task of simulating the dynamics of the Fermi–Hubbard model, quantum algorithms currently provide an exponential speedup over known classical algorithms. Nevertheless, as the Fermi–Hubbard Hamiltonian is sufficiently simple to be realized in many controlled physical systems, future fault-tolerant quantum computers will also have to compete against analog quantum simulators.

1.2 Spin models

Overview

Classical and quantum spin systems are prototypical models for a wide range of physical phenomena including magnetism, neuron activity, simplified models of materials and molecules, and networks. Studying the properties of spin Hamiltonians can also provide useful insights in quantum information science.

A number of scientific and industrial problems can be mapped onto finding the ground or thermal states of classical or quantum spin models, for example, solving combinatorial optimization problems, training energy-based models in machine learning, and simulating low-energy models of quantum chemistry [983].

Simulating the *dynamics* of quantum spin models is primarily of interest for quantum information science, as well as condensed matter physics or chemistry. For instance, such simulations are relevant for interpreting nuclear mag-

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

netic resonance (NMR) [920, 806] or related spectroscopy experiments [273, 749].

Because of the natural mapping between spin-1/2 systems and qubits, as well as the fact that interactions in spin models are typically local, the resources required to simulate simple spin models using quantum algorithms can be much lower than for problems in areas like quantum chemistry or cryptography.

While our discussion will focus on quantum algorithms designed to be run on fault-tolerant quantum computers, the simple Hamiltonians of spin models are naturally realized in many physical systems. This has led to the use of analog simulators [157, 417], such as arrays of trapped ions or neutral atoms, for simulating the static and dynamic properties of interesting spin models. We will comment briefly on this below.

Actual end-to-end problem(s) solved

The most commonly studied spin models are those with pairwise interactions, referred to as 2-local Hamiltonians. We note that the interactions are not necessarily geometrically local, although this will be present in many models of physical systems. Given a graph $\mathcal{G}$ with N vertices $\{i\}$ and L edges $\{[i, j]\}$, we associate a classical or quantum spin with each vertex, and an interaction between spins with each edge. We can also add one-body interactions acting on individual spins. The Hamiltonian can then be written as

$$H = \sum_i \sum_{\alpha\in\{x,y,z\}} B_i^\alpha \sigma_\alpha^i + \sum_{[i,j]} \sum_{\alpha,\beta\in\{x,y,z\}} J_{ij}^{\alpha\beta} \sigma_\alpha^i \sigma_\beta^j , \tag{1.2}$$

where $\{\sigma_x^i, \sigma_y^i, \sigma_z^i\}$ denote the Pauli operators X_i, Y_i, Z_i acting on site i, and $\{B_i^\alpha\}, \{J_{ij}^{\alpha\beta}\}$ are coefficients. For classical spin Hamiltonians, the sums are restricted to Z operators. The Hamiltonian in Eq. (1.2) encompasses a wide range of spin models, including the following:

- The classical Ising model

$$H = \sum_i B_i Z_i + \sum_{ij} J_{ij} Z_i Z_j , \tag{1.3}$$

  which also describes the Hamiltonians arising from quadratic unconstrained binary optimization (QUBO) problems.
- The (quantum) transverse-field Ising model (TFIM)

$$H = B \sum_i X_i + J \sum_{ij} Z_i Z_j . \tag{1.4}$$

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

- The Heisenberg model with a site-dependent magnetic field, defined in 1D with nearest-neighbor interactions by

$$H = \sum_j \left(B_j Z_j + J^x X_j X_{j+1} + J^y Y_j Y_{j+1} + J^z Z_j Z_{j+1} \right). \tag{1.5}$$

Across the different models, we can vary the dimension, locality of interactions (e.g., nearest-neighbor vs. fully connected vs. power-law), and values of the site-dependent coefficients in comparison to the interaction terms. The models can be extended beyond 2-local by considering couplings of 3 or more spins—see, for example, p-spin models, which are p-local [345]. The above definitions can be extended from spin-1/2 systems to higher-dimensional spin operators by generalizing the Pauli operators with their higher-dimensional counterparts.

For classical spin models, we seek to prepare the ground or thermal states of the model, as these may encode, for example, the solution to a combinatorial optimization problem, or a probability distribution that can be used for generative modeling. For quantum spin models, we similarly seek to compute ground or thermal states. However, because these are not classical states that can be easily extracted, we typically wish to sample observables with respect to these states. Examples include the energy, the magnetization of the system, and correlations between sites. In dynamics simulations of quantum systems, we seek to determine how observables of interest vary as a function of evolution time. Examples include the magnetization (used to infer the Hamiltonian in nuclear magnetic resonance [533] or related [172] experiments), or the growth of correlations between sites to probe thermalization. Digital Hamiltonian simulation can additionally extract certain quantities that may be hard to directly measure in experiments (e.g., time reversal) [402]. Since studies of quench dynamics often require preparation of simple states, such as product states or the ground states of classically solvable Hamiltonians, and the measurement of local observables, propagation under the Hamiltonian typically dominates the simulation cost. For lattice systems with N spins in D spatial dimensions, it is conventional to consider evolution times that scale as $\Omega(N^{1/D})$, as the system must evolve for at least this long in order for information to propagate across the system due to the Lieb–Robinson bound [261].

Reference [101] details a number of applications of simulating quantum spin models, with relevance to systems studied at Los Alamos National Laboratory, including the parameters of end-to-end simulations. Reference [374] presents an end-to-end assessment of using quantum computers to simulate quantum spin Hamiltonians for applications in NMR spectroscopy.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

Dominant resource cost/complexity

For a system of N spin-$\frac{1}{2}$ particles, we require N qubits to represent the state of the system. For N spin-S particles, the problem can be mapped to qubits in different ways, for example, using $N\lceil \log_2(2S+1) \rceil$ qubits [897] or using $2NS$ qubits [749].

Quantum algorithms for preparing the ground or Gibbs states of classical spin systems are discussed in detail in Chapter 4 (combinatorial optimization), and Section 9.2 (energy-based machine learning models), respectively. We will restrict our discussion to the resources required for performing time evolution of quantum spin models. The reason for this is that quantum algorithms for preparing ground or thermal states require similar primitives for Hamiltonian access to algorithms for time evolution (e.g., block-encodings or Hamiltonian simulation itself) and use these in conjunction with either (i) eigenstate filtering approaches [689, 688] based on quantum singular value transformation, (ii) adiabatic state preparation, (iii) quantum phase estimation from a trial state, or (iv) quantum algorithms for thermal state preparation, that is, Gibbs sampling. More detailed discussions of these algorithms and their caveats can be found in the corresponding sections, as well as in the discussion of quantum algorithms for simulating molecules and materials (Section 2.1) or the Fermi–Hubbard model (Section 1.1), where preparing (approximate) eigenstates is the primary topic of interest. All of these algorithms depend on either an overlap between the trial state and the target state, the minimum gap along an adiabatic path, or the mixing time of a Markov chain—all of which are difficult to bound in the general case.

When simulating the time evolution of spin systems via Hamiltonian simulation, the most efficient algorithms exploit the locality of interactions in the Hamiltonian, and the resulting commutation structure. For 2-local spin-1/2 systems on a D-dimensional lattice with nearest-neighbor geometric locality, algorithms with almost optimal gate complexity are known for performing time evolution. Reference [280] showed that the gate complexity of the $(2k)$th-order product formula scales as

$$O\left((Nt)^{1+1/2k}/\epsilon^{1/2k}\right)$$

to simulate time evolution for time t to accuracy ϵ, using a Hamiltonian given in the Pauli access model. Note that this expression suppresses the 5^{2k} constant prefactor present in $(2k)$th-order Trotter (see Eq. (11.1)). Similarly, [479] gave an algorithm with complexity $O(Nt \cdot \mathrm{polylog}(Nt/\epsilon))$ for Hamiltonians given in the sparse access model. In contrast, note that approaches that are asymptotically optimal in the black-box setting, such as quantum signal processing, have a gate complexity of $O(N^2 t + N \log(1/\epsilon))$ using a block-encoding based on lin-

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

ear combinations of unitaries (LCU)—the normalization of the block-encoding is $\alpha = O(N)$ and the gate complexity to implement the block-encoding is also $O(N)$.

Spin Hamiltonians with power-law interactions were studied in [993, 286], that is, where the interaction strength between spins i and j depends inversely on a power of the distance between the spins, denoted by $\|i-j\|_2$. (This implies that the model is translationally invariant.) For a D-dimensional lattice with 2-local interactions with interaction strengths scaling as $1/\|i-j\|_2^{\alpha}$, $(2k)$th-order Trotter gives a gate complexity scaling as (as above, suppressing the 5^{2k} constant prefactor present in $(2k)$th-order Trotter) [286]

$$\widetilde{O}\left(\begin{array}{ll} N^{3-\frac{\alpha}{D}(1+1/2k)+1/k} t^{1+1/2k} \epsilon^{-1/2k} & \text{if } 0 \leq \alpha < D, \\ N^{2+1/2k} t^{1+1/2k} \epsilon^{-1/2k} & \text{if } \alpha \geq D \end{array} \right).$$

Further improvements are possible in cases where the Hamiltonian coefficients are efficiently computable by an oracle, or if certain truncations can be performed [721]. Focusing on the $D = 1$ case, if one were to directly apply quantum signal processing based on a block-encoding via the linear combination of unitaries approach, the scaling of the gate complexity would be

$$\widetilde{O}\left(N^2 t + N \log(1/\epsilon)\right).$$

These asymptotic complexities are complemented by the constant prefactor analyses discussed in the following section.

For estimating expectation values of observables to precision ϵ, one can either consider directly sampling and then re-preparing the state of interest (scaling as $O(1/\epsilon^2)$), or coherent approaches based on amplitude estimation (scaling as $O(1/\epsilon)$, but requiring a longer coherent circuit depth). Measurements of simple observables, such as the magnetization, can be obtained through the computational basis measurements on single qubits. For more complicated observables, one can consider the approaches in [853, 549, 49], discussed in more detail in Chapter 2 (quantum chemistry).

Existing resource estimates

A number of logical resource estimates for simulating the dynamics of spin systems and for finding their ground states via quantum phase estimation have been reported in the literature. In such calculations, it is necessary to optimize the constant prefactor contributions from implementing the algorithmic primitives used. A detailed comparative study on simulating the dynamics of a 1D nearest-neighbor Heisenberg model (Eq. (1.5)) was reported in [283], comparing the logical qubit and T gate counts of product formulas, Taylor series, and

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

quantum signal processing. The two most efficient approaches are shown in the first two rows of Table 1.2.

On a fault-tolerant quantum computer, arbitrary angle rotation gates must be synthesized using a sequence of T and Clifford gates [625]. The number of T gates to synthesize a group of parallel rotation gates can be reduced if they share the same angle [421, 630, 225], a method known as Hamming weight phasing. This technique can be exploited in fault-tolerant compilations of algorithms simulating physical spin systems, which often exhibit features such as translational invariance.

In addition to the entries given in Table 1.2, fault-tolerant approaches to simulating NMR [806, 374] and muon spectroscopy [749] experiments, which are effectively spin model simulations, have been considered.

Caveats

When formulated as a decision problem, determining the ground state energy for 2-local classical and 2-local quantum spin models is NP-complete [94, 726] and QMA-complete [607], respectively. As such, we do not expect quantum algorithms to provide efficient solutions to these problems in the general case. Nevertheless, given the success of classical heuristics for these problems, one may hope to observe a similar benefit from quantum heuristic algorithms, such as Monte Carlo–style Gibbs sampling algorithms.

In contrast, simulating the dynamics of spin models is a BQP-complete problem [705]; it is likely one of the most simple beyond-classical calculations that could be performed on a future fault-tolerant quantum computer. While such a computation would be of great scientific interest, providing new insights in quantum information and many-body physics, it is currently unclear whether dynamics simulations of large systems will have a direct impact on industrially relevant applications.

Comparable classical complexity and challenging instance sizes

Exact classical simulations of quantum spin models are exponentially costly in system size. Exact simulations that consider a time evolution long enough for information to propagate across the system (as per the Lieb–Robinson bound) are limited to around 50 spins using the largest classical supercomputers [492, 283].

Approximate classical algorithms for studying quantum spin systems include tensor network approaches and quantum Monte Carlo (QMC) methods. These methods provide empirically accurate results for computing the ground states of physically motivated spin systems, in particular those with local interactions, in low dimensions. For example, the ground states of local, gapped 1D

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

Problem	Method	# Spins	# T gates	# Logical qubits	Parameters
1D Heisenberg dyn.	QSP	50	2.4×10^9	67	$B_j \in [-1, 1], J^x = J^y = J^z = 1, t = N, \epsilon = 10^{-3}$ [283]
1D Heisenberg dyn.	Trotter (sixth order)	50	1.8×10^8	50	$B_j \in [-1, 1], J^x = J^y = J^z = 1, t = N, \epsilon = 10^{-3}$ [283]
2D NN TFIM[a] dyn.	Trotter (fourth order)	100	1.7×10^5	100	$t = 10/J, B = J, \epsilon = 10^{-2}$ [392, 147]
2D $1/r^2$ TFIM dyn.	Trotter (fourth order)	100	1.5×10^7	100	$t = 10/J, B = J, \epsilon = 10^{-2}$ [392]
2D Heisenberg ground state with nearest- and next-nearest-neighbor interactions	Qubitized QPE	100	10^8	112	$\epsilon = 10^{-2}, J_1 = 1, J_2 = 0.5, B_j = 0$ [1069]

Table 1.2 Logical resource estimates for quantum phase estimation (QPE) and dynamics simulation (dyn.) applied to different spin models. The presented gate counts are for a single run of the circuit. The results presented in rows 1 and 2 can be compared to each other, and both target an error of $\epsilon = 10^{-3}$ in the operator norm distance between the ideal and implemented time evolution unitary. While [283] presents both analytic and empirical Trotter error bounds, the gate count presented in the table is that resulting from the empirical bound, though we remark that more recent analytic bounds are close to matching the empirical bounds [286]. The results presented in rows 3 and 4 can be compared to each other, and determine the number of Trotter steps used empirically by targeting an error of $\epsilon = 10^{-2}$ in a particular spatially averaged local observable, and then extrapolating this behavior to larger system sizes.

[a] 2D nearest-neighbor transverse-field Ising model.

Hamiltonians have area law entanglement [504]; thus, they can be efficiently represented by matrix product states, a type of tensor network. The situation in higher dimensions is less clear, as an area law for gapped 2D local Hamiltonians has not been proven, and such a relation does not imply efficient representation via classical tensor network approaches such as projected entangled pair states (PEPS) [503]. Indeed, the 2D local Hamiltonian problem with the constraint that the ground state obeys area laws is QMA-complete [545]. Other approximate classical methods, including QMC approaches, can also be effective for preparing low energy states of spin systems. For example, [1054] provides benchmarks for both tensor network and QMC-based variational methods applied to spin systems in a range of lattice geometries.

In contrast, these methods are less accurate when performing simulations of quantum spin dynamics [912, 907]. In many of these systems, the entanglement entropy grows linearly with time [219], resulting in a cost that grows exponentially with time for tensor network approaches targeting fixed accuracy (see, e.g., [921, 501] for counterexamples). For example, it was claimed in [392] that simulations of the dynamics of the 2D TFIM for $N = 100$ spins would be far beyond the current capabilities of tensor network methods [392].

Many physical systems are subject to strong interactions with their environment, which limits their coherence times. In these cases, the behavior of the system can often be reproduced by simulating a smaller number of spins (e.g., $N \leq 30$) and accounting for the interactions with the environment through physically motivated heuristics [1045]. Such simulations (accessible via open-source software libraries) are used to analyze NMR [533] and muon spectroscopy experiments [172].

Speedup

The speedup for computing the ground states of quantum spin Hamiltonians over classical approximate methods (such as tensor networks or QMC) is currently an open research question; it depends on the complexity of being able to prepare good approximations of the ground state using quantum algorithms for cases where classical trial states are unable to efficiently and accurately represent the ground state [670].

The simulation of quantum spin dynamics appears to be exponentially costly using all known classical methods. As such, quantum algorithms for Hamiltonian simulation would provide an exponential speedup for this task. This would likely provide insights in quantum information and many-body physics. As an example, such systems could study the competition and interplay between thermalization and many-body localization in quantum systems.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

NISQ implementations

Quantum spin models are commonly used as benchmark systems for NISQ algorithms—for example, finding ground states [595], simulating dynamics [880], and probing thermalization [762]. For instance, variational algorithms have been applied to spin systems with up to 24 qubits, showing rapid convergence to the ground state as the number of layers of the variational circuit was increased [174, 604].

The Hamiltonians of spin models are also naturally realized in a wide range of physical systems, including trapped ions or neutral atoms [157, 417]. For example, recent experiments in neutral atom systems have studied the dynamics of on the order of 200 spins, which went beyond the capabilities of classical simulation via matrix product state approaches [371, 906]. Analog simulators are already an important tool providing new scientific insights, and they set a high bar for the future performance of fault-tolerant approaches to simulating spin systems. Nevertheless, analog simulators are restricted in the complexity of the models that they can simulate (e.g., it may be more challenging to study site-dependent impurity models or composite spin–fermion systems), and they are susceptible to errors from miscalibration and interactions with the environment.

Outlook

Simulating the behavior of spin systems is arguably one of the most natural tasks for quantum computers, while being exponentially costly using all known classical methods. Such simulations can provide important insights into questions in quantum information science and many-body physics. Spin system simulations are also relevant to condensed matter physics and chemistry, since spin systems can act as models for more complex systems in those fields.

Logical resource estimates for quantum algorithms simulating spin systems are among the lowest known for beyond-classical tasks. Nevertheless, analog quantum simulators are already able to natively simulate the dynamics of hundreds of spins. In order to surpass these capabilities, digital approaches may need to consider more complex observables or target better accuracies that are only achievable using devices capable of quantum error correction.

In addition, for many systems of scientific interest in related fields, such as chemistry or condensed matter physics, decoherence-inducing interactions with the environment often limit the required simulation sizes. Identifying applications where accurate dynamics simulation of large spin models is required would increase the impact and applicability of quantum algorithms in this area.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

1.3 SYK model

Overview

The Sachdev–Ye–Kitaev (SYK) model [889, 622] is a simplified model of a quantum black hole that is strongly coupled and "maximally chaotic," but still analytically tractable. This remarkable and, to date, unique combination of properties has led to great activity surrounding SYK. It has applications in high-energy physics through its connections to black holes and quantum gravity, and it has applications in condensed matter physics as a model of quantum chaos and scrambling, which sheds light on phases of matter in strongly coupled metals [882, 945]. While many interesting properties of the SYK model can be computed analytically in certain limits, not all properties qualify, and questions remain about the behavior of the model outside of these limits—these questions can potentially be addressed numerically by a quantum computer.

Actual end-to-end problem(s) solved

The SYK model has many variants; a common version to consider is the four-body ($q = 4$) N-site Majorana fermion Hamiltonian with Gaussian coefficients

$$H_{\mathrm{SYK}} = \frac{1}{4 \cdot 4!} \sum_{i,j,k,\ell=1}^{N} g_{ijk\ell}\, \chi_i \chi_j \chi_k \chi_\ell \,, \tag{1.6}$$

where χ_i denote Majorana fermion mode operators obeying the anticommutation relation $\chi_i \chi_j + \chi_j \chi_i = 2\delta_{ij} I$ (where I denotes the identity operator, and δ_{ij} the Kronecker delta symbol), and $g_{ijk\ell}$ are coefficients drawn independently at random from a Gaussian distribution with zero mean and variance $\sigma^2 = 3! g^2 / N^3$ (with g the tunable coupling strength). The normalization of Eq. (1.6) matches the convention of [78] and ensures that the ground state energy of H_{SYK} is extensive (i.e., growing linearly in N).

In the limit of a large number of local degrees of freedom $N \to \infty$ and at strong coupling $\beta g \gg 1$ (where β is the inverse of the temperature), analytic predictions can be computed for the asymptotic behavior of some properties. While these arguments are not mathematically rigorous, in practice they provide a consistent picture for the SYK model and provide insights into quantum gravity and quantum chaos. However, questions remain about the wealth of properties out of reach by taking limits or the nonasymptotic regime of parameters. For example, it has been challenging to rigorously calculate the density of states at arbitrary energies and the ground state energy in the large-N limit [316, 78, 511]. These problems can potentially be probed numerically on a quantum computer.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

Generally speaking, this often reduces to performing the following task on the quantum computer: given as input an instance of H_{SYK} (generated by choosing the couplings $g_{ijk\ell}$ at random) and an observable O, estimate the expectation value $\mathrm{tr}(\rho O)$, where ρ could be, for instance, (i) the ground state of H_{SYK}, (ii) the thermal state $\rho \propto \mathrm{e}^{-\beta H_{\mathrm{SYK}}}$, or (iii) a time-evolved state $\rho = \mathrm{e}^{\mathrm{i}H_{\mathrm{SYK}}t}|0\rangle\langle 0|\mathrm{e}^{-\mathrm{i}H_{\mathrm{SYK}}t}$ from an easy-to-prepare initial state $|0\rangle$, among other possibilities. The observable O could be a local operator or even H_{SYK} itself. Another case is for O to be composed of t-dependent time-evolution unitaries $\mathrm{e}^{\mathrm{i}H_{\mathrm{SYK}}t}$.

For example, computing the ground state energy corresponds to taking ρ to be the ground state of H_{SYK} and O to be H_{SYK}, and computing a 4-point out-of-time-ordered correlation function corresponds to taking ρ to be the thermal state at inverse temperature β and O to be $A\mathrm{e}^{\mathrm{i}H_{\mathrm{SYK}}t}B\mathrm{e}^{-\mathrm{i}H_{\mathrm{SYK}}t}A\mathrm{e}^{\mathrm{i}H_{\mathrm{SYK}}t}B\mathrm{e}^{-\mathrm{i}H_{\mathrm{SYK}}t}$, where A and B are few-body operators [555]. In another example, [204, 799] give a detailed proposal to "simulate quantum gravity in the lab" via computing expectation values of observables and states formed via simulation of the SYK model.

Depending on the ultimate end-to-end goal, one may need to repeat this calculation for many different O or for many instances of H_{SYK}, for example, to compute an ensemble average.

Dominant resource cost/complexity

Mapping the problem to qubits: To simulate the SYK model on a quantum computer, the Majorana operators are represented by strings of Pauli operators according to the Jordan–Wigner representation (e.g., [407]). As a result, the Hamiltonian H_{SYK} on N Majoranas becomes a linear combination of multi-qubit Pauli operators over $N/2$ qubits. Methods for Hamiltonian simulation given a classical description of the Hamiltonian as a linear combination of Pauli strings (the Pauli access model) typically introduces into the complexity a dependency on the number of terms, N^4, and on the 1-norm of Pauli coefficients, denoted by λ, which for typical SYK instances is seen to be $\lambda = O(gN^{5/2})$ (see [78, Eq. (16)]).

State preparation: To solve the problem of estimating $\mathrm{tr}(\rho O)$, one must be able to prepare the $(N/2)$-qubit state ρ. In some cases, ρ could simply be a product state, which is trivial to prepare. If ρ is the thermal state at inverse temperature β, then algorithms for Gibbs sampling would be used to prepare the state. Due to the chaotic properties of SYK and the fact that the system is expected to thermalize quickly in nature, one expects that Monte Carlo–style Gibbs samplers (e.g., [984, 256, 939, 856, 260]) have a favorable poly(N) gate

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

complexity, but the exact performance is unknown. If ρ is the ground state of H_{SYK}, there are several methods for preparing ρ, including projection onto ρ by measuring (and postselecting) an ansatz state ϕ in the energy eigenbasis using quantum phase estimation (QPE), or by adiabatic state preparation. The cost of either of these methods depends on details such as which ansatz state is used (in particular, its overlap with ρ), the adiabatic path, and the spectrum of H_{SYK}—in both cases, in the absence of evidence to the contrary, the scaling can be exponential in N. In [511], a poly(N)-time quantum algorithm for preparing states ρ achieving a constant-factor approximation to the ground state energy of H_{SYK} was given, which could be used as ρ to probe low-energy properties of the system.

Time evolution: The calculation also requires simulating time evolution by H_{SYK}. This can be because O is a time-evolved operator, because the state ρ corresponds to a time-evolved state, or simply as a subroutine for QPE or Gibbs sampling, mentioned above. Reference [407] proposed a scheme for simulating time evolution using a first-order product formula approach to Hamiltonian simulation. That is, it implements the unitary $e^{iH_{\mathrm{SYK}}t}$ to precision ϵ, with gate complexity $O(N^{10}g^2t^2/\epsilon)$. However, this steep scaling with N suggests that accessing large system sizes will be difficult with this method. Reference [78] later gave a method with better N dependence, achieving gate complexity $O(N^{7/2}gt + N^{5/2}gt\,\mathrm{polylog}(N/\epsilon))$, leveraging qubitization with quantum signal processing. This gate complexity grows more slowly than the number of terms in H_{SYK} (i.e., $O(N^4)$), a feat that is only possible because the simulation method generates the SYK coupling coefficients pseudorandomly: to construct the block-encoding of H_{SYK}, they perform the PREPARE step in the linear combination of unitaries using a shallow quantum circuit composed of polylog(N) random two-qubit gates, producing a state for which the N^4 amplitudes are distributed approximately as independent Gaussians. Further reduction in the gate count would be bottlenecked by the 1-norm λ of the coefficients of H_{SYK}; however, note that [1059] suggests gravitational features may remain even if the Hamiltonian is substantially sparsified, which could reduce the number of terms and the value of λ.

Measuring observables: Finally, given the ability to prepare a purification of ρ and supposing O is unitary (if it is not, it could be decomposed into a sum of unitaries and each constituent computed separately), estimating the expectation value $\mathrm{tr}(\rho O)$ to precision ϵ can be done by overlap estimation, costing $O(1/\epsilon)$ calls to the routine that prepares ρ and to the routine that applies O. If the purification of ρ cannot be prepared, the cost is $O(1/\epsilon^2)$.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

Existing resource estimates

Reference [78] compiled the dominant contributions in their approach to Hamiltonian simulation into Clifford + T gates, and they found that at $N = 100$, implementing e^{iHt} requires fewer than $10^7 gt$ T gates, and at $N = 200$, it requires fewer than $10^8 gt$ T gates. Both of these figures are for a single circuit; that is, they do not include the cost of averaging over many SYK instances. The T-count can be turned into an estimate of the running time and number of physical qubits; see the discussion in Part III on fault-tolerant quantum computation.

Caveats

Existing resource estimates only focus on simulating the dynamics of SYK models, but the proposed classically challenging problems involve static properties such as density of states and properties of thermal states. Probing these static properties in an end-to-end fashion would likely require preparing thermal states, ground states, or other kinds of low-energy states, in addition to being able to implement e^{iHt}. The cost of preparing these states is unknown and difficult to assess analytically. Another caveat is that the gate counts quoted above do not take into account the $O(1/\epsilon)$ scaling of reading out an observable to precision ϵ, or any repetitions for different instances of H_{SYK} required for making inferences about the physics of SYK.

Comparable classical complexity and challenging instance sizes

As mentioned above, one of the reasons that the SYK model is appealing is that many properties can be computed analytically in certain limits. Other properties that would be of interest to numerically compute on a quantum computer require poorly scaling classical methods. Exact diagonalization of systems consisting of more than roughly 50 Majoranas would be challenging due to the exponential growth of the Hilbert space, which has dimension $2^{N/2}$. For example, [316] and [408] gave a variety of numerical results based on exact diagonalization up to $N = 34$ and $N = 36$, respectively.

Speedup

Hamiltonian simulation has poly(N) runtime, an exponential speedup over exact diagonalization, which is the go-to method for classical simulation of SYK-related problems. However, Hamiltonian simulation does not alone solve the same end-to-end problem as exact diagonalization; the persistence of the exponential speedup requires identifying specific interesting properties where the relevant initial states can also be prepared in poly(N) time, which is currently less clear.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press

NISQ implementations

Experimental realizations of the SYK model have been proposed on several different experimental platforms [401, 851, 728]. However, even if these demonstrations can be realized, we do not expect this approach to scale in the absence of quantum error correction.

Outlook

Simulating time evolution of the SYK model on a quantum computer has relatively mild gate cost, due to the model's straightforward mapping to a qubit Hamiltonian. At the same time, it is difficult to simulate the SYK model on a classical computer, owing to its chaotic and strongly coupled nature. However, further work is needed to understand the entire end-to-end pipeline. It has not yet been identified which properties would be most valuable to compute on a quantum computer and how costly they will be. Computing these properties will likely involve far more than a single run of time evolution on a single instance of the SYK model, so the overall cost is likely to be much larger than what initial gate counts in the literature suggest.

https://doi.org/10.1017/9781009639651.003 Published online by Cambridge University Press