3

Nuclear and particle physics

Simulating nuclear and particle physics is an inherently quantum problem. There have been proposals to use quantum computers to accelerate simulations of quantum field theories, nuclear physics, neutrino physics, and quantum gravity [104]. In this chapter, we focus on the simulation of quantum field theories and nuclear physics, as these have received the most attention in the literature to date and are the closest to having end-to-end resource estimates available. While not covered explicitly in this chapter, the building blocks of quantum algorithms for data analysis in high-energy physics [340] can be found in Chapter 20 on variational quantum algorithms and Chapter 9 on machine learning. For existing reviews of quantum computing for nuclear and particle physics, we direct the reader to [104, 844, 92, 404, 632, 105, 110].

The authors are grateful to Zohreh Davoudi and John Preskill for reviewing this chapter.

3.1 Quantum field theories

Overview

We seek the static and dynamic properties of quantum field theories, specifically gauge field theories and scalar field theories. Gauge field theories describe the interactions between matter and/or gauge degrees of freedom and can be classified by their symmetry groups, such as U(1) (describing quantum electrodynamics), SU(2) (the weak interaction), and SU(3) (quantum chromodynamics). Scalar field theories describe interactions between scalar fields, such as the Higgs field or ϕ^4 theory.

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

Interacting quantum field theories are typically not analytically solvable, and techniques such as perturbation theory are only accurate in some parameter regimes. For example, low-energy quantum chromodynamics (QCD), relevant to quark confinement and hadron formation, cannot be treated perturbatively. As such, complex scattering processes at particle accelerators are currently treated with a combination of first-principles calculations and approximate phenomenological methods.

To tackle quantum field theories numerically from first principles, lattice field theory is employed. The Lagrangians arising from lattice field theory can be numerically solved on classical computers using Euclidean Monte Carlo methods, which have proven highly efficient and accurate for a number of static quantities, including hadron masses, static matrix elements of (primarily) time-local operators between hadronic states, and even certain properties of light nuclei [332, 333, 44]. However, these classical Monte Carlo methods become intractable due to a sign problem in two regimes: (i) at high fermion density (of considerable scientific interest for understanding the decomposition of neutron stars and large atomic nuclei) and (ii) in simulations of real-time dynamics (e.g., scattering problems). Hamiltonian formulations of these problems are challenging due to the size of the required Hilbert space. As such, there have been a number of proposals to use quantum computers for calculating the static and dynamic properties of matter described by scalar and gauge field theories. For further background, see [844, 104, 761] and references therein.

Actual end-to-end problem(s) solved

Classical computational methods for lattice field theories have produced a number of insights, including high-precision computations of fundamental quantities (such as the muon's magnetic moment and quark masses), tests of beyond-the-Standard-Model physics (such as charge conjugation and parity (CP) violation and beyond-Higgs theories), and nuclear cross sections with dark matter candidates or neutrinos. For a more complete and detailed list, we refer the reader to [333, Page 6].

We primarily focus on the case of lattice field theories in the Hamiltonian formulation, which explicitly separates temporal and spatial degrees of freedom [640] and discretizes the d-dimensional space using an L^d lattice (which may be noncubic). Matter degrees of freedom (e.g., quarks, scalar fields) are placed on the vertices of the lattice. Gauge degrees of freedom (e.g., the value of the electromagnetic field) are placed on the links between lattice sites. Dynamical simulations proceed by initializing the system in a desired state [82], performing time evolution under the Hamiltonian, and measuring relevant observables. See [588] for an example of simulating scalar field theories. Static

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

simulations aim to prepare a state of interest, such as the ground state of a collection of quarks representing a composite hadron, and then measure observables of interest, including binding energies, as well as structural and reaction properties.

Quantum simulations of lattice field theories may be incorporated as part of a larger (multiscale) computational workflow. For example, when studying scattering processes (such as those that occur at particle accelerators), it is not necessary to simulate the entire scattering process on a quantum (or classical) computer. Instead, the scattering cross section can be separated into short- and long-distance contributions, which can be computed using perturbative and nonperturbative (e.g., quantum or classical simulation of lattice gauge theories) methods, respectively [413, 103].

Dominant resource cost/complexity

In this section, we focus predominantly on the simulation of dynamics in lattice gauge theories (LGTs), as the majority of studies to date have considered this application. In the standard formulation, one allocates one qubit per fermion (or antifermion) type per site of an $N = L^d$ lattice. Each gauge degree of freedom (one in U(1), three in SU(2), eight in SU(3)) requires its own register associated with each edge between lattice sites. The quantum numbers associated to the gauge degrees of freedom are encoded in binary, up to a maximum cutoff value Λ, so the corresponding register requires $\log(\Lambda)$ qubits. It was shown in [991] that for time evolution performed with fixed lattice spacing, the cutoff can be set as $\Lambda = \Lambda_0 + \widetilde{O}(T) \cdot \text{polylog}(N/\epsilon)$, where Λ_0 is the maximum initial value of the gauge fields, T is the time evolution duration, and ϵ is the resulting error in the final state. Hence, the overall number of qubits required to store the state of the system scales as

$$O\left(L^d \log\left(\Lambda_0 + T \ \text{polylog}\left(\frac{L^d}{\epsilon}\right)\right)\right).$$

Algorithms for implementing time evolution under LGT Hamiltonians are presented in [991, 927, 593, 852, 292, 334, 873]. It is necessary to (approximately) maintain gauge invariance during the simulation, which can be achieved either by the choice of formulation, by actively protecting symmetries [487, 994], or by detecting and eliminating the gauge-violating states [960, 864]. As an example of the first option, one can calculate the desired Hamiltonian matrix elements on the fly using Clebsch–Gordon coefficients [213], but this is expensive in terms of elementary quantum operations [593]. An alternative approach described in [292] encodes only the

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

physical transitions in the SU(3) gauge theory. This method requires many controlled operations and a large classical precomputing overhead.

Existing resource estimates

The number of T gates required to simulate instances of the lattice Schwinger model (U(1) LGT in $d = 1$ with both matter and gauge degrees of freedom) was studied in [927]. That work considered the resources required to perform Trotterized time evolution and estimate the electron-positron pair density. The most complex simulations analyzed (64 lattice sites, cutoff of $\Lambda = 8$) required 5×10^{13} T gates per shot, and 333 logical qubits. Such a circuit would need to be repeated $O(1/\epsilon^2)$ times to estimate the pair density to accuracy ϵ; this overhead could be improved to $O(1/\epsilon)$ using quantum amplitude estimation at the expense of increased gate depth [927]. These estimates were later improved in the small system-size, long-time, or low-error regime using algorithms based on qubitization with quantum signal processing [892]. Note that a simulation of the 64-site lattice Schwinger model with $\Lambda = 8$ is well within the range of classical simulations [389, 734].

Reference [593] performed similar resource estimates for the simulation of dynamics in U(1), SU(2), and SU(3) LGTs for $d = 3$. These resource estimates were performed for synthesizing the time evolution operator, with choice of simulation parameters inspired by the following tasks: computing transport coefficients relevant to the study of quark-gluon plasmas, simulating heavy ion collisions, and computing the hadronic tensor of the proton—although we note that the costs of initial state preparation and observable measurements are not included in these resource estimates. Logical qubit counts ranged from 10^4 to 10^8, while T gate counts ranged from 10^{17} to 10^{56}. The large constant factors present in these resource estimates stem partly from the use of quantum arithmetic (e.g., constituting 99.998% of the gate count in the most expensive calculations [593]), and partly from the decomposition of the plaquette term, which is exponential in the number of colors. The large gate counts also arise from the value chosen for the error in the time evolution operators $\epsilon = 10^{-8}$, which may be overly conservative when viewed in conjunction with other sources of error.

These gate counts can be reduced significantly using the approach of [334], which investigated alternative ways of representing the simulation, allowing for an improved grouping of terms in the product formula. These ideas were illustrated for SU(2) in 1+1D, but can likely be generalized to more complex, higher-dimensional LGTs.

Nevertheless, we note that any implementation scaling as $\Omega(TL^3)$ (i.e., linearly in spacetime volume) already faces a factor of 10^8 for $T = L = 100$,

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

highlighting the potentially large resource counts of simulating quantum field theories. Recent work presented an algorithm for simulating the time evolution of LGTs that achieves this optimal complexity, up to polylogarithmic factors [873]. This work uses a number of subroutines for Hamiltonian simulation [479, 718, 615] and requires locality-preserving fermion- and boson-to-qubit mappings. Resource estimates were carried out for U(1), SU(2), and SU(3) LGTs. For simulating time evolution in an SU(3) LGT on a lattice with $T = L = 100$, the algorithm required approximately 10^{21} T gates and 6×10^7 logical qubits. Despite the large improvements compared to [593], the significant discrepancy from the expected lower bound of 10^8 discussed above suggests that there is further opportunity for optimization in the implementation of algorithms for simulating LGTs.

Caveats

Additional investigation is required to better quantify the theoretical uncertainties arising from discretization, finite-volume, and Hilbert space truncation effects of quantum computing formulations, as well as the algorithmic errors present in quantum simulation algorithms applied to lattice and scalar field theories.

For example, discretization of the continuous field theory to the lattice setting introduces a number of nuances (which are also present in classical approaches but must be considered afresh in quantum calculations). As discussed in [970, 598], discretization of the fermion field breaks the Lorentz invariance of the fermion kinetic term, which introduces unphysical additional flavors of fermions (known as the fermion doubling problem). This issue can be mitigated in several established ways, each with their own merits and drawbacks for quantum simulation [747]. It is also necessary to carefully track other errors resulting from discretization and ensure that these vanish when scaling and extrapolating to the continuum limit [588, 231].

As noted in [104, Section 6b] and [292, 105], there are a number of possible representations/basis sets that can be used for the gauge degrees of freedom, and it is currently unclear which choice is optimal for quantum simulation.

Comparable classical complexity and challenging instance sizes

The end-to-end scattering processes typically considered at particle accelerators are too complex to be solved from first principles and are instead tackled using a range of approximate techniques [413]. These computations often include parameters obtained from first-principles LGT calculations on simpler systems, and they typically proceed through a Lagrangian formulation, rather than a Hamiltonian formulation. This leads to Monte Carlo sampling of a path

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

integral in Euclidean spacetime, the application of which to dynamical problems or static problems with high fermion density is limited by the fermionic sign problem [793]. For example, the phase diagram of QCD and the existence of exotic phases at extreme densities, in and out of equilibrium, have eluded classical methods. Nevertheless, classical approaches have been very effective for static problems with low fermion density and for dynamical scattering problems at low energy and low inelasticity; for a review of current state-of-the-art computations and limitations, see [333, 585] and companion whitepapers referenced therein.

Recent work has begun to investigate using tensor network methods to simulate the Hamiltonian formulation of LGTs; see, for example, [389] ($d = 2, L = 16$, U(1) LGT with gauge field cutoff $\Lambda = 1$) and [734] ($d = 3, L = 8$, U(1) LGT with gauge field cutoff $\Lambda = 1$). Like quantum simulations, tensor network approaches are sign-problem free and thus may be of interest in regimes out of reach of conventional Monte Carlo–based approaches. However, tensor network approaches are currently limited to small system sizes, and often need to be verified by comparing to other methods, as they do not come with provable guarantees on the bond dimension required for capturing the entanglement structure of the states present in LGTs. For recent reviews on the use of tensor networks to simulate LGTs, we refer the reader to [91, 735].

Speedup

For simulations with a sign problem, classical Monte Carlo methods are exponentially costly in system size [997]. In addition, it was observed that the bond dimensions required for tensor network approaches increase rapidly with system size [734], suggesting the potential for exponential quantum speedups for dynamical problems. This suggestion is reinforced by the BQP-completeness of the simulation of certain field-theoretic processes [589]. Nevertheless, the constant prefactors for quantum simulations of LGTs are currently high, and we require the (currently underexplored) ability to efficiently prepare initial states of interest.

NISQ implementation

A number of works have investigated the simulation of scalar or gauge field theories on noisy digital, or analog, quantum simulators. A common strategy is to map the lattice field theory Hamiltonian to that of a bosonic system, such as cold atoms or trapped ions; see, for example, [104, 105, 632, 92] and references therein. While these techniques appear promising for simple Hamiltonians, such as the Schwinger model, it may be challenging to engineer the more complicated interactions required in nonabelian gauge field theories. There

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

have also been works applying variational algorithms to lattice field theories, such as [641, 67, 699], as well as digital simulations of the time dynamics of the lattice Schwinger model [631, 387].

Outlook

Investigations into how quantum computers can be used to complement classical methods for simulating lattice field theories are advancing rapidly. While quantum computers can, in principle, efficiently simulate the complex scattering experiments performed in particle accelerators, the resources required to do so would be impractical using currently known techniques. Future work must determine the best targets for quantum simulations and reduce asymptotic scaling as well as constant prefactors. In particular, the qubit encoding (currently scaling as $O(L^d)$ qubits for a lattice in d spatial dimensions with each dimension having L sites) means that a large number of logical qubits will likely be required for computations of interest where, as illustrated by examples above, we may consider $L = 10$–100 to challenge classical approaches.

3.2 Nuclear physics

Overview

Nuclear physics describes the behavior of individual nuclei, as well as that of dense nucleonic matter, such as neutron stars. The structure of nuclei can be approximately described using the shell model (see [337] for an overview), a phenomenological model with parameters fitted to experimental observations. However, high-accuracy descriptions of nuclear structure, exotic nuclei, accurate scattering cross sections, and nonequilibrium phenomena require a first-principles treatment. Describing the properties of nuclei from first principles (e.g., lattice quantum chromodynamics simulations) is beyond the reach of analytic and current computational capabilities for all but the simplest nuclei [369, 108, 332]. Nevertheless, one can often integrate out the short-range physics to obtain effective field theories (EFTs) that describe the interactions of nucleons. The prototypical example is chiral effective field theory, which describes the interactions of nucleons and pions (pions, which have mass less than 6× smaller than that of the proton, are mediators of the residual strong interaction between pairs of nucleons). The parameters of the EFT can be inferred from experiments or directly from lattice quantum chromodynamics (QCD) calculations, resulting in a many-body Hamiltonian that describes the formation, structure, and potential decay of nuclei.

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

Actual end-to-end problem(s) solved

An EFT provides a many-body Hamiltonian describing how nucleons interact. This quantum many-body problem can be tackled using a range of classical methods, which employ different mathematical approaches to approximately solve the nuclear many-body Schrödinger equation including quantum Monte Carlo (QMC) methods [232], the no-core shell model (NCSM) [100], the coupled cluster (CC) method [485], the self-consistent Green's function (SCGF) method [230], the in-medium similarity renormalization group (IM-SRG) method [958], and nuclear lattice methods [658] (a related review article discussing inputs to these calculations is given by [987]).

A common problem is to prepare the ground state of a collection of nucleons, in order to compute nuclear binding energies and determine if a given nucleus is stable (e.g., determining the long lifetime of ^{14}C [740, 485]). Simulations can also be used for computing scattering cross sections, in order to analyze experiments on nucleus-neutrino scattering [877], beta decay [472], and nuclear reactions. Reactions, such as nuclear fission and nuclear fusion, can also be studied using explicitly time-dependent approaches [115], although these have higher computational costs than static computations and are often based on semiclassical, mean-field, or other phenomenological models. Simulating both fusion and fission reactions has a number of use cases, such as an improved understanding of nuclear astrophysics, where reactions commonly occur at energies too high or too low to be replicated in experiments [796].

Dominant resource cost/complexity

The prototypical EFT for nuclear interactions is chiral effective field theory. At very low energies, the theory can be expressed as a convergent perturbative expansion (chiral perturbation theory). However, for the larger energies relevant to multi-nucleon systems, the theory becomes nonperturbative. Despite several notable successes, developing EFTs that converge across a wider range of scenarios remains an active area of research [987, 489].

At low energies (below the rest mass of the pion), it can be appropriate to apply pionless EFTs, which integrate out the pions, leading to implicit interactions between nucleons, including a 3-nucleon contact interaction required at leading-order by renormalization. At higher energies, it is necessary to explicitly account for the effect of pions. Pionfull EFTs are typically studied numerically, as it can be difficult to obtain analytic predictions [987]. In a formulation known as the one-pion-exchange Hamiltonian, pions are integrated out, leading to a long-range two-body interaction between nucleons, which decays exponentially with distance. In an alternative formulation, known as dynam-

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

ical pion EFT, the (relativistic) pions and their interaction with nucleons are explicitly simulated.

A common formulation of EFT simulations is to project the problem onto a lattice in position or momentum space [667]. For quantum simulations formulated on a lattice, a typical second-quantized mapping uses 4 qubits per lattice site for nucleons (two isospin degrees of freedom and two spin degrees of freedom), although additional qubits may be required if using a fermion-to-qubit mapping that maintains the locality of the fermionic Hamiltonian (e.g., 6 qubits per site [1028]). If simulated explicitly (e.g., dynamical pion EFT), the value of the spin-0 pion field at each lattice site can be stored using a number of qubits scaling logarithmically with the pion energy cutoff by storing its quantum number in binary.

An alternative approach is to project the EFT Hamiltonian onto a single-particle basis, commonly harmonic oscillator eigenfunctions [153]. In second quantization, a qubit is required per single-particle (iso)spin mode. However, this mapping can lead to long-range interactions between modes, and in the most general case, up to $O(N^6)$ distinct terms for an N-mode system [104, 957].

Quantum algorithms that prepare energy eigenstates (or good approximations thereof) scale either as $1/\gamma$ (where γ is the overlap of the initial state with the desired eigenstate) [688], or with the minimum gap size along an adiabatic or thermalizing path. If we are only interested in measuring the energy of the state, this can be obtained using the quantum phase estimation algorithm, which also projects the system into the corresponding energy eigenstate. The cost of this approach scales as $O(1/\gamma^2)$ in terms of the original overlap γ, which can be improved to $O(1/\gamma)$ using amplitude amplification at the expense of increased circuit depths. Once the desired state has been prepared, observables can be measured to precision ϵ with complexity $O(1/\epsilon^2)$ (direct sampling) or $O(1/\epsilon)$ (using amplitude estimation, also requiring coherent state preparation, e.g., via amplitude amplification).

The above algorithms for preparing states (and related algorithms for performing time evolution in dynamics simulations) require access to the Hamiltonian, typically implemented via block-encoding or Hamiltonian simulation.

Existing resource estimates

The gate costs of Hamiltonian simulation using product formulas for neutrino-nucleus scattering in pionless EFT were studied in [877]. The T-gate costs of Hamiltonian simulation using product formulas, as well as quantum phase estimation, were estimated in [1028] for pionless EFT, one-pion-exchange EFT, and dynamical pion EFT (all formulated on a lattice). As an example, resources to simulate the dynamics of 40 nucleon system ranged from 4×10^{12} T gates

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

and 6000 logical qubits for pionless EFT, to 2.0×10^{24} T gates and 6000 logical qubits for one-pion-exchange EFT, to 5.2×10^{49} T gates and 168,000 logical qubits for dynamical pion EFT (all to an error $\epsilon = 0.1$). The resources required for the latter EFT are considerably higher, due to the qubit overhead and worse error dependence that stems from explicitly simulating the pions. Resource estimates for other parameter regimes can also be found in [1028], where it is noted that higher-order EFTs are required for the accurate simulation of medium- and large-mass nuclei, which may further increase simulation costs.

Caveats

Developing accurate nuclear EFT Hamiltonians describing heavier nuclei is an active research direction. Most studies are currently based on phenomenological models, which are limited in their predictive capabilities. In some systems, some higher-order EFTs provide sufficient accuracy, but only after fitting the EFT coefficients to experimental data. We refer the interested reader to [987] for a more detailed discussion.

Comparable classical complexity and challenging instance sizes

Classical approaches use similar techniques to those developed for the electronic structure problem, such as perturbation theory, Monte Carlo methods, or coupled cluster approaches. References [232, 100, 485, 230, 958, 658, 523] provide excellent overviews of state-of-the-art approaches. Classical methods can provide outstanding agreement with experiments for the binding energies of small nuclei with 20–50 nucleons [523]. As a further example, recent high-accuracy simulations of the ^{100}Sn nucleus have improved the agreement between theory and experiment for observed β-decay rates [472]. Time-dependent simulations of dynamics or nonequilibrium phenomena are more challenging and are an active area of research [115, 796]. We refer readers to [233] for a detailed analysis of the capabilities and requirements of exascale supercomputers in nuclear physics simulations.

Speedup

The majority of classical approaches for the nuclear structure problem are designed to run in polynomial time with respect to the system size but introduce errors due to the use of approximations (e.g., truncating the expansion in coupled cluster methods) [523]. For quantum computers to achieve exponential speedups, one needs to identify systems where (i) classical methods require an exponential increase in resources to obtain accurate results and (ii) it is efficient to prepare an initial state for the quantum computation with at least in-

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press

verse polynomially large (in terms of the system size) overlap with the desired state. Recent initial investigations have explored whether these requirements coexist in chemical systems [670]. We are not aware of similar work in nuclear physics, although it has been noted that the states obtained from classical methods could be used as initial states for quantum algorithms [110].

For simulating the dynamics of nuclear systems, classical methods typically proceed via mean-field methods, and exact simulations are limited to small system sizes due their exponential scaling. In contrast, quantum algorithms are able to simulate the dynamics of nuclear systems with a cost scaling polynomially with the system size and desired accuracy.

NISQ implementation

Almost all of the work to date on demonstrating near-term quantum computing approaches for the nuclear structure problem has focused on variational algorithms, such as [366, 724, 955, 929], or small-scale simulations of dynamics [877, 1000]. There is currently no evidence that near-term quantum devices will be able to implement sufficiently deep circuits to achieve advantage over their classical counterparts with these methods.

Outlook

Nuclear physics presents a classically challenging quantum many-body problem that appears well suited to simulation on quantum computers. While there are similarities to the electronic structure problem in quantum chemistry, which has led to a transfer of several ideas, nuclear Hamiltonians are typically more complex, involving (depending on the formulation) long-range interactions, 3-body interactions, multiple species of nucleons and pions, and interactions with nontrivial spin and isospin dependence. The simulation of nuclear reaction dynamics appears a particularly interesting target, and future work should determine the resources required for end-to-end simulations, including state preparation and measurement of observables.

https://doi.org/10.1017/9781009639651.005 Published online by Cambridge University Press