

Fermilab

Quantum-centric supercomputing for materials science: A perspective on challenges and future directions

FERMILAB-PUB-24-0001-SQMS

arXiv:2312.09733

DOI: 10.1016/j.future.2024.04.060

Accepted Manuscript

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Quantum-centric Supercomputing for Materials Science: A Perspective on Challenges and Future Directions

Yuri Alexeev,¹ Maximilian Amsler,² Marco Antonio Barroca,^{3,4} Sanzio Bassini,⁵ Torey Battelle,⁶ Daan Camps,⁷ David Casanova,^{8,9} Young Jai Choi,¹⁰ Frederic T. Chong,¹¹ Charles Chung,¹² Christopher Codella,¹² Antonio D. Córcoles,¹² James Cruise,¹³ Alberto Di Meglio,¹⁴ Ivan Duran,¹² Thomas Eckl,² Sophia Economou,¹⁵ Stephan Eidenbenz,¹⁶ Bruce Elmegreen,¹² Clyde Fare,¹² Ismael Faro,¹² Cristina Sanz Fernández,¹² Rodrigo Neumann Barros Ferreira,³ Keisuke Fuji,¹⁷ Bryce Fuller,¹² Laura Gagliardi,^{18,1} Giulia Galli,^{11,1} Jennifer R. Glick,¹² Isacco Gobbi,¹⁹ Pranav Gokhale,²⁰ Salvador de la Puente Gonzalez,¹² Johannes Greiner,¹² Bill Gropp,²¹ Michele Grossi,¹⁴ Emanuel Gull,²² Burns Healy,²³ Matthew R. Hermes,¹⁸ Benchen Huang,¹¹ Travis S. Humble,²⁴ Nobuyasu Ito,²⁵ Artur F. Izmaylov,^{26,27} Ali Javadi-Abhari,¹² Douglas Jennewein,⁶ Shantenu Jha,^{28,29} Liang Jiang,¹¹ Barbara Jones,³⁰ Wibe Albert de Jong,³¹ Petar Jurcevic,¹² William Kirby,³² Stefan Kister,¹² Masahiro Kitagawa,¹⁷ Joel Klassen,³³ Katherine Klymko,³⁴ Kwangwon Koh,³⁵ Masaaki Kondo,^{36,25} Doğa Murat Kürkçüoğlu,^{37,38} Krzysztof Kurowski,³⁹ Teodoro Laino,⁴⁰ Ryan Landfield,²⁴ Matt Leininger,⁴¹ Vicente Leyton-Ortega,²⁴ Ang Li,^{42,43} Meifeng Lin,²⁸ Junyu Liu,¹¹ Nicolas Lorente,^{44,8} Andre Luckow,⁴⁵ Simon Martiel,¹² Francisco Martin-Fernandez,¹² Margaret Martonosi,⁴⁶ Claire Marvinney,²⁴ Arcesio Castaneda Medina,¹⁹ Dirk Merten,¹⁹ Antonio Mezzacapo,^{12,*} Kristel Michielsens,⁴⁷ Abhishek Mitra,¹⁸ Tushar Mittal,¹² Kyungsun Moon,¹⁰ Joel Moore,^{48,49} Sarah Mostame,¹² Mario Motta,¹² Young-Hye Na,¹² Yunseong Nam,¹² Prineha Narang,^{28,50} Yu-ya Ohnishi,⁵¹ Daniele Ottaviani,⁵ Matthew Otten,⁵² Scott Pakin,¹⁶ Vincent R. Pascuzzi,¹² Edwin Pednault,¹² Tomasz Piontek,³⁹ Jed Pitera,⁵³ Patrick Rall,³² Gokul Subramanian Ravi,²² Niall Robertson,⁵⁴ Matteo A. C. Rossi,⁵⁵ Piotr Rydlichowski,³⁹ Hoon Ryu,⁵⁶ Georgy Samsonidze,⁵⁷ Mitsuhsa Sato,²⁵ Nishant Saurabh,⁵⁸ Vidushi Sharma,⁵³ Kunal Sharma,¹² Soyoung Shin,³⁰ George Slessman,⁵⁹ Mathias Steiner,³ Iskandar Sitdikov,¹² In-Saeng Suh,²⁴ Eric D. Switzer,^{8,60} Wei Tang,⁴⁶ Joel Thompson,⁶¹ Synge Todo,⁶² Minh C. Tran,³² Dimitar Tenev,⁶³ Christian Trott,⁶⁴ Huan-Hsin Tseng,²⁸ Norm M. Tubman,⁶⁵ Esin Tureci,⁴⁶ David García Valiñas,¹² Sofia Vallecorsa,¹⁴ Christopher Wever,⁶⁶ Konrad Wojciechowski,³⁹ Xiaodi Wu,⁶⁷ Shinjae Yoo,²⁸ Nobuyuki Yoshioka,⁶² Victor Wen-zhe Yu,¹ Seiji Yunoki,^{25,68} Sergiy Zhuk,⁵⁴ and Dmitry Zubarev¹²

¹Argonne National Laboratory, Lemont, IL 60439, USA

²Corporate Sector Research and Advance Engineering, Robert Bosch GmbH, Robert-Bosch-Campus 1, D-71272 Renningen, Germany

³IBM Research, Rio de Janeiro, 20031-170, RJ, Brazil

⁴Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, 22290-180, RJ, Brazil

⁵CINECA, via Magnanelli 6/3, 40033 Casalecchio di Reno, BO, Italy

⁶Arizona State University, Tempe, AZ, USA

⁷National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory Berkeley, CA, USA

⁸Donostia International Physics Center (DIPC), 20018 Donostia-San Sebastián, Euskadi, Spain

⁹Ikerbasque, Basque Foundation for Science, 48009 Bilbao, Spain

¹⁰Department of Physics, Yonsei University, Seoul 03722, Korea

¹¹University of Chicago, Chicago, Illinois, USA

¹²IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA

¹³Cambridge Consultants part of Capgemini Invent, Cambridge, UK

¹⁴European Organization for Nuclear Research (CERN), Geneva 1211, Switzerland

¹⁵Virginia Tech, Blacksburg, VA 24061, USA

¹⁶Los Alamos National Laboratory, Los Alamos, NM 87545, USA

¹⁷Osaka University, Osaka 560-8531, Japan

¹⁸Department of Chemistry, Chicago Center for Theoretical Chemistry, University of Chicago, Chicago, Illinois, USA

¹⁹Fraunhofer ITWM, Kaiserslautern, Rheinland-Pfalz 67663, DE

²⁰Inflection, Chicago, IL 60622, USA

²¹University of Illinois Urbana-Champaign

²²University of Michigan, Ann Arbor, Michigan 48109, USA

²³Dell Technologies, Research Office

²⁴Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge, Tennessee 37831, USA

²⁵RIKEN Center for Computational Science (R-CCS), Kobe, Hyogo 650-0047, Japan

²⁶Chemical Physics Theory Group, Department of Chemistry, University of Toronto, Toronto, Ontario M5S 3H6, Canada

- ²⁷ Department of Physical and Environmental Sciences,
University of Toronto Scarborough, Toronto, Ontario M1C 1A4, Canada
- ²⁸ Brookhaven National Laboratory, Upton, NY, USA
- ²⁹ Rutgers University, New Brunswick, NJ, USA
- ³⁰ IBM Quantum, Almaden Research Center, San Jose, CA 95120, USA
- ³¹ Applied Mathematics and Computational Research Division,
Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
- ³² IBM Quantum, IBM Research Cambridge, Cambridge, MA 02142, USA
- ³³ Phasecraft Ltd., London, UK
- ³⁴ National Energy Research Scientific Computing Center Lawrence Berkeley National Laboratory Berkeley, CA, USA
- ³⁵ Electronics and Telecommunications Research Institute (ETRI), Daejeon, Republic of Korea
- ³⁶ Keio University, Yokohama, Kanagawa 223-8522, Japan
- ³⁷ Fermi National Accelerator Laboratory, Batavia, IL 60510, USA
- ³⁸ Superconducting Quantum Materials and Systems Center (SQMS),
Fermi National Accelerator Laboratory, Batavia, IL 60510, USA
- ³⁹ Poznań Supercomputing and Networking Center, IBCH PAS Poznań, Poland
- ⁴⁰ IBM Research Europe, Säumerstrasse 4, 8803 Rüschlikon, Switzerland
- ⁴¹ Lawrence Livermore National Laboratory, Livermore, CA 94550, USA
- ⁴² Pacific Northwest National Laboratory, Richland, Washington 99354, USA
- ⁴³ University of Washington, Seattle, Washington 98195, USA
- ⁴⁴ Centro de Física de Materiales (CSIC-EHU), 20018 Donostia-San Sebastián, Euskadi, Spain
- ⁴⁵ BMW Group, Munich, Germany
- ⁴⁶ Princeton University, Princeton, New Jersey, USA
- ⁴⁷ Jülich Supercomputing Centre, Institute for Advanced Simulation,
Forschungszentrum Jülich, 52425 Jülich, Germany
- ⁴⁸ University of CA, Berkeley, CA 94720, USA
- ⁴⁹ National Energy Research Scientific Computing Center,
Lawrence Berkeley National Laboratory, Berkeley, CA, USA
- ⁵⁰ University of CA Los Angeles, Los Angeles, CA 90095, USA
- ⁵¹ JSR Corporation, 3-103-9, Tonomachi, Kawasaki-ku, Kawasaki, Kanagawa, 210-0821, Japan
- ⁵² Department of Physics, University of Wisconsin - Madison, Madison, WI 53706, USA
- ⁵³ IBM Research, Almaden Research Center, San Jose, CA 95120, USA
- ⁵⁴ IBM Quantum, IBM Research Europe - Dublin, IBM Technology Campus, Dublin 15, Ireland
- ⁵⁵ Algoritmiaq Ltd, Kanavakatu 3 C, FI-00160 Helsinki, Finland
- ⁵⁶ Korea Institute of Science and Technology Information, Daejeon 34141, Republic of Korea
- ⁵⁷ Robert Bosch LLC, Research and Technology Center, Sunnyvale, CA 94085, USA
- ⁵⁸ Department of Information and Computing Sciences, Utrecht University, NL
- ⁵⁹ CR&DL, Inc., USA
- ⁶⁰ Department of Physics, University of Central Florida, Florida 32816, USA
- ⁶¹ Applied Mathematics, Boeing Research & Technology, Huntsville, AL 35824
- ⁶² The University of Tokyo 7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan
- ⁶³ ExxonMobil Technology and Engineering Company, Annandale, NJ 08801, USA
- ⁶⁴ Sandia National Laboratories, Albuquerque, NM, USA
- ⁶⁵ NASA Ames, Mountain View, CA, USA
- ⁶⁶ Corporate Sector Research and Advanced Engineering,
Robert Bosch GmbH, Robert-Bosch-Campus 1, D-71272 Renningen, Germany
- ⁶⁷ University of Maryland, College Park, USA
- ⁶⁸ RIKEN Center for Quantum Computing (RQC), Wako, Saitama 351-0198, Japan
(Dated: September 20, 2024)

Computational models are an essential tool for the design, characterization, and discovery of novel materials. Computationally hard tasks in materials science stretch the limits of existing high-performance supercomputing centers, consuming much of their resources for simulation, analysis, and data processing. Quantum computing, on the other hand, is an emerging technology with the potential to accelerate many of the computational tasks needed for materials science. In order to do that, the quantum technology must interact with conventional high-performance computing in several ways: approximate results validation, identification of hard problems, and synergies in quantum-centric supercomputing. In this paper, we provide a perspective on how quantum-centric supercomputing can help address critical computational problems in materials science, the challenges to face in order to solve representative use cases, and new suggested directions.

I. INTRODUCTION

Materials science use cases are being investigated as some of the first applications of quantum computing to show practical quantum advantages. The strongest motivation for this belief is that a great deal of materials science problems are quantum in nature [1]. A computational advancement in the field will have consequences in many areas, from scientific exploration to industrial advancements and sustainability.

Currently, high-performance supercomputing centers dedicate a substantial amount of resources to computations in the materials science space [2–5]. A fundamental issue is that ab-initio computations of materials, in the full configuration interaction (CI) limit, suffer from an exponential or factorial growth of the computational resources in the system size [6–11]. Therefore, memory and processing times requirements rapidly become intractable, as the size of the quantum systems increase. To tame this dire scaling, various approximations are employed to reduce classical computational costs, however, even these either fail – for highly correlated materials – or become intractable for sufficiently large and general systems [12–23].

Quantum computers present an attractive alternative, since many quantum algorithms avoid incurring in the exponential memory overheads of classical computations of quantum matter. For practical use cases, quantum computers may be expected to operate embedded in classical high performance computing (HPC) environments, so-called *Quantum-centric Supercomputing (QCSC)*. Intensive classical processing is required before, after, and concurrently with these quantum computations, for a number of reasons: alleviating the workload of quantum computers, integration into existing classical HPC algorithms, signal extraction from noisy quantum devices, and exploiting fault-tolerant codes to the best extent.

In this paper, we motivate the need for quantum-centric supercomputing in materials science research and industrial development. We also identify the challenges ahead to achieve practical quantum advantage for materials science use cases and propose directions to address them. It is important to note that this work is intended as a perspective and not as an exhaustive review of the existing literature. More specifically, our discussion will be centered around addressing several particular aspects:

Key Algorithms. We identify quantum algorithms for materials science applications, in terms of practical applications and potential for quantum advantage.

Implications on the design of quantum-centric supercomputing architectures. We discuss the requirements that arise for quantum-centric supercomputing architectures. This includes evaluating the computational and operational demands, scalability, and integration hurdles when deploying these algorithms in coupled quantum-classical HPC environments.

Materials Science Use Cases. We highlight specific use cases in materials science where quantum and HPC

algorithms can be most effectively utilized. We select the algorithms and use cases according to three principles, which are necessary for quantum advantage:

1. the use case must be classically hard in some limit;
2. the use case and algorithm considered must be amenable to execution on a noisy or fault-tolerant quantum computer, depending on the scenario;
3. the use case represents an interesting problem in materials science.

Our discussion is organized as follows. Section II highlights some key existing algorithms relevant to materials science. Section III exposes classical computations in quantum-HPC workflows, keeping key algorithms and use cases as a reference. In Section IV, we detail of the challenges for hybrid quantum-classical workload management and quantum-HPC integration, and in Sec. V give a perspective on the main components of a programming model for such systems. Section VI summarizes the state-of-the-art in classical simulation of quantum circuits, which is a necessary crucial consideration in evaluating the potential for quantum advantage. Finally, Sec. VII gives an overview of potential use cases in materials science, bringing together considerations and threads from previous sections.

II. FUNDAMENTAL QUANTUM ALGORITHMS

In this section we summarize quantum algorithms that can be used for materials science applications. We consider algorithms that simulate quantum systems and mention quantum computing approaches to simulate partial differential equation (PDE). Although the simulation of quantum systems is the most natural application of quantum computers in the materials science space, we mention quantum algorithms for PDEs since they can be useful for some well-crafted use cases. Within the simulation category, we discuss the time dynamics, ground states, open quantum systems and finite-temperature properties.

Simulation of dynamics generated by a Hamiltonian of interest is one of the most natural applications for quantum computers. Quantum algorithms for approximating time evolution have been studied extensively, and so here briefly describe a few key methods. In recent years the development of *quantum signal processing* culminated in an algorithm called *qubitization* [24], the complexity of which is measured by the number of queries to an oracle that encodes the Hamiltonian. For a target evolution time t and an error tolerance ϵ , qubitization makes only $O(t + \log 1/\epsilon)$ queries to the oracle, which is asymptotically optimal. However, the construction of the Hamiltonian oracle generally involves ancillary qubits, making qubitization difficult to implement in near-term and early fault-tolerant devices.

In this regime, *Trotterization* [25, 26] (see also Sec. III B 1), the original quantum algorithm for simulating dynamics, remains one of the leading candidates for early applications, because its low overhead permits fine-grained tradeoffs between circuit depth and accuracy. Asymptotically, the gate complexity of the p -th order Trotterization scales as $O(t^{1+1/p}/\epsilon^{1/p})$, which can be made close to that of qubitization at large p . Additionally, the error of Trotterization depends on the commutativity of the terms in the Hamiltonian, allowing the algorithm to take advantage of the structure of the Hamiltonian. The final example we mention is the *HHKL* algorithm [27], which achieves nearly optimal scalings for Hamiltonians that are geometrically local (or even with power-law decaying interactions [28]) on some lattice. HHKL can be considered a hybrid between Trotterization and the high-accuracy methods that exploits commutativity between spatially-separated Hamiltonian terms.

Perhaps the most paradigmatic type of quantum simulation for materials is the approximation of ground states and low-energy states in general. While such problems are Quantum Merlin Arthur (QMA)-hard [29, 30] and thus believed to be intractable in worst cases, their broad utility across physics, chemistry, and materials science has made them the objects of intense study in the hope that for physically interesting cases, they may be manageable. The most heavily studied algorithm for simulating ground states in the near term is the Variational Quantum Eigensolver (VQE), e.g., [31–53]), which is based on optimizing a parameterized wavefunction prepared on the quantum computer, and is well-adapted to near term due to the possibility of implementing it with low-depth circuits. Similar considerations apply to the use of finite-depth quantum approximate optimization algorithms for finding the ground state of frustrated materials [54, 55]. However, VQE has a number of drawbacks, mostly related to the difficulties of establishing convergence of the optimization and of accurately measuring the target energy to be estimated, which can result in long runtimes or getting stuck in local minima.

If we consider fault-tolerant architectures, Quantum Phase Estimation (QPE) [56] and related algorithms (e.g., [57–59]) promise to offer ground state simulation with finite success probability, given access to initial reference states with sufficiently high overlaps with the true ground state. A good initial state assumption can be traded off by assumptions on the mixing time of a dissipative Lindbladian associated with the problem [60].

In the near term, quantum Krylov algorithms [61–81] and related techniques [82] have recently emerged as a promising family of ground state simulation methods since they fall somewhere between VQE and QPE: many variants [63, 64, 67–69, 71–73, 75–80, 83] have provable convergence [77, 78, 80, 81, 83] subject to similar assumptions to QPE, but the circuits they involve are only moderately more complex than those in VQE.

For some Hamiltonians, in particular spin and fermionic lattice models often studied in condensed mat-

ter and materials science, some variants of quantum Krylov methods could be amenable to noisy quantum devices. For a recent review of quantum Krylov algorithms, see [81].

Quantum computers also enable the simulation of open quantum system dynamics. Open systems more accurately model physical phenomenon in nature and their study is particularly relevant in fields such as condensed matter, material science, and quantum chemistry [84–91]. The simulation of open systems requires both coherent evolution under a Hamiltonian and dissipative processes that capture the interactions with the environment. Formally, the evolution of a state under an open-system dynamics is governed by the Lindblad quantum master equation [92].

Quantum algorithms for simulating Hamiltonian dynamics have been extended to simulate Lindbladian dynamics. A naive approach is to explicitly simulate the environment [93–95]. Trotterization can be used to simulate the dynamics of k -local Lindbladians [94] and non-Markovian systems in [96]. Simulation of sparse and non-local Lindbladian dynamics requires the implementation of sparse Stinespring isometries [97]. When Lindbladians can be written as a linear combination of Pauli operators, one can employ a variant of linear combination of unitaries [98]. Another promising approach is to implement imaginary-time evolution to approximate the dissipative dynamics [99]. Other emerging quantum algorithms to simulate open system dynamics include qubitization [100], wave matrix Lindbladization [101, 102], and partial probabilistic error cancellation [103].

In addition to estimating observables on ground states, expectations with respect to finite-temperature states are also critical for studying a variety of phenomena including phase transitions. In contrast to preparing ground states, which may encode a computationally intractable optimization problems [30], thermal state preparation is generally computationally easier as long as the temperature is sufficiently high. There is a range of quantum algorithms that aim to accelerate estimation of thermal properties of quantum systems. One approach [104], analogous to the ground-state projection method used by phase estimation, leverages Grover-assisted rejection sampling to either cool an infinite-temperature state to a given finite temperature [105], or to use two finite-temperature systems to simulate a single system at half the temperature [106].

A different approach for thermal state preparation involves simulating the physical process of thermalization. Early work in this direction involves a discrete-time Metropolis-Hastings algorithm over the eigenstates of the Hamiltonian that is guaranteed to converge on the thermal state [107]. Quantum techniques for achieving quadratic speedups in classical Markov Chain Monte Carlo have been adapted to these quantum analogues [108, 109]. More recent work involves the simulation of continuous-time quantum Markov processes known as Davies generators, a kind of open quantum system that

is regarded to emulate thermalization in nature [110]. A recent line of work aims to develop quantum algorithms that efficiently simulate these open systems despite their long-range interactions [111–113]. Some of these proposals aim for minimal quantum resources [60], and others explore connections between thermalization and the fundamental capabilities of quantum computation [114].

In addition to the Schrödinger equation, there are also applications of other PDEs in materials science, which include Maxwell’s equations for deriving electromagnetic properties, the heat equation, and PDEs involved in the characterization of structural properties. These frequently rely on the finite element method (FEM).

There are several proposals for quantum algorithms for PDEs, usually based on the quantum linear systems algorithm [115]. Boundary value problems analyzed via finite element methods often result in large linear systems of equations [116]. Their matrix elements can frequently be derived analytically which may allow sidestepping the need for a quantum random access memory (QRAM). Existing literature has shown that quantum algorithms for finite element methods exhibit polynomial speedups over their classical counterparts [117]. These speedups have been studied for applications in solving the heat equation [118] and electromagnetic structure problems [119].

There is also a growing body of literature on quantum algorithms for initial value problems of both linear [120, 121] and nonlinear PDEs [122]. These methods highlight that a variety of dynamics can be encoded into the Schrödinger equation [123], subject to certain restrictions pertaining to normalization [124] and chaos [125].

III. CLASSICAL INFORMATION PROCESSING IN QUANTUM-CLASSICAL WORKFLOWS

A. Mapping to qubits

Simulation of quantum systems is a promising use case for both noisy and fault-tolerant quantum computers. In order to simulate a physical quantum system, the governing physical Hamiltonian, for example, that of a fermionic or bosonic system, must be appropriately mapped to a qubit Hamiltonian such that the quantum computer faithfully simulates the physical system of interest. Note an appropriate mapping preserves the operator algebra: the physical operators that comprise the physical Hamiltonian and the qubit operators that comprise the qubit Hamiltonian must obey the same algebra. Fermions and bosons follow their respective algebras, commonly known as canonical anticommutation relations (CAR) and canonical commutation relations (CCR). Qubits are the so-called hard-core bosons [126].

Plentiful choices of the aforementioned kind of mappings exist as briefly described below. Which mapping to use may thus be determined by a practical consideration, such as the computational cost of the resulting

simulation. A good choice of mapping will likely depend on the specifics of a given physical Hamiltonian. In this subsection, we highlight some of the existing techniques and challenges around this problem.

Bosons have in common with qubits their commutation relations (bosonic wavefunctions are symmetric), but while a qubit has only two states, a bosonic mode has infinitely many. Any digital simulation of bosonic modes requires cutting off the dimension at some finite value, typically much larger than two. Once a cutoff is applied, however, the choice of mapping is essentially a decision about how to encode a nonnegative integer representing the occupation of the mode into qubits. One could for example use unary or binary representations, with the former yielding local representations of creation and annihilation operators with linear qubit cost in the occupation cutoff, and the latter yielding log-local representations of creation and annihilation operators with only logarithmic qubit cost in the occupation cutoff. This is a significant trade-off, but one that is at least well-understood — see for example [127–133].

Fermionic wavefunctions, on the other hand, are antisymmetric, which means that in second-quantization their creation and annihilation operators anticommute on different modes. Since spatially disjoint qubit operators commute, mapping fermions to qubits requires care. Also known as bosonization in the physics community (in the sense of qubits sharing the commutation relations of bosons), a collection of qubit operators can be used to embed a single fermionic creation or annihilation operator to preserve the fermionic algebra between the embedded fermionic operators. The first such mapping is the Jordan-Wigner (JW) transformation [134], and more recently this area has seen extensive research [50, 53, 135–145], with each of these references discussing new fermion-to-qubit mappings.

Some of alternative mappings may be related to the JW transformation by a unitary transformation since conjugation by a unitary preserves the entire algebra. Others, especially those designed for particular fermionic lattices, may preserve the algebra between the relevant operators only, implied by the given lattice. A subset of unitaries, such as Cliffords or linear-reversible circuits, can be used to describe other well-known mappings including Bravyi-Kitaev (BK) [135] or parity.

The art of finding and choosing a good mapping is centered around the computational resource savings. This is because different mappings can result in vastly different quantum circuits; e.g., with respect to the number of qubits or the number of quantum gates, when combined with both problem-specific and problem-agnostic circuit optimization techniques. In the near term, circuits will necessarily have limited width and depth due to noise, so providing a mapping that is efficient in both dimensions is crucial for reliable execution. Indeed, while the parity transform has been used to reduce the qubit counts [138] and the BK transform has been shown to admit the asymptotically optimal operator locality for

generic fermionic Hamiltonians [135], customized mappings that take advantage of the input problem specifics while working in concert with circuit compilation tools, e.g., [48, 146], have been reported to result in substantially smaller circuits [50, 53].

Other mappings that aim to exploit limited connectivity in fermionic lattices to localize qubit operators over the limited qubit connectivity offered by some quantum computers have also been explored [50, 142, 143, 145, 147, 148]. With very few exceptions, e.g., [144], most fermion-to-qubit mappings constitute mappings from products of Majorana operators to Pauli operators. Many interesting and useful configurations have been deployed within this framework. The space of mappings of this type is well defined [145], although not necessarily easy to search over. In this sense future work in constructing new useful mappings of this type constitutes a collective exercise in combinatorial optimization, which remains a useful and interesting activity, especially in the near-term regime where these details of efficiency can make or break an algorithm. The broader challenge for the field is finding ways to push beyond this paradigm, to either find new kinds of useful mappings that are not somehow equivalent to mapping Majoranas to Paulis, or to otherwise leverage our understanding of the landscape of mappings to some novel useful effect.

B. Hamiltonian encodings for dynamics simulation

The intractability of simulating dynamics of generic quantum systems using known classical methods has driven the development of quantum computation since its inception [149], and it is the branch of quantum simulation that is most strongly believed to offer a quantum advantage [25]. Hamiltonian dynamics, or the simulation of time evolution operator, e^{-iHt} , for a given Hamiltonian, H , is a core component in many quantum algorithms. Generally, the Hamiltonian is represented as a sum, $H = \sum_i H_i$, of operators which are diagonalizable using simple unitary transformations, $H_i = U_i D_i U_i^\dagger$; e.g., exact squares of one-electron fermionic operators for quantum chemistry calculations, weighted sum of fully commuting Pauli operators, or one-sparse operators, whose dynamics admit more straightforward simulation. The choice of these operators and their ordering, as well as the analysis of their sparsity and commutation relations, which is usually done classically, can greatly impact the efficiency of the algorithm [150]. Depending on the decomposition, a variety of well-known techniques, such as Trotterization and linear combination of unitaries, can then be used to compile the dynamics operator into fundamental gate operations.

1. Trotterization

The basic idea of Trotterization is to break down the evolution of a given quantum system into a series of smaller, more computationally manageable steps. For a given Hamiltonian operator $H = \sum_i H_i$, the first-order Trotter approximation is given by the product formula

$$\left(\prod_{k=1}^m e^{-iH_k \frac{t}{n}} \right)^n = e^{-iHt} + O\left(\frac{m^2 t^2}{n}\right). \quad (1)$$

As shown, the error vanishes for large n (number of Trotter steps) or small time. Higher-order product formulas are readily derivable but become increasingly complex to implement. In practice, the choice of which order, p , would best suit a particular application is application-specific. For example, the second-order Trotter-Suzuki formula

$$\left(\prod_{k=1}^m e^{-iH_k \frac{t}{2n}} \prod_{k=m}^1 e^{-iH_k \frac{t}{2n}} \right)^n = e^{-iHt} + O\left(\frac{m^3 t^3}{n^2}\right) \quad (2)$$

is a preferred choice for some lattice-gauge theory simulations [133, 151] due to its favorable asymptotic scaling in large parameters such as the number of levels used to simulate bosons, while admitting a relatively simple implementation detail. The fourth-order formula was shown to admit the best error bounds in intermediate-sized system simulations using several dozen qubits for the Heisenberg Hamiltonian [152, 153].

Trotterization introduces errors in the simulation as shown by the above examples. At the highest level, this error depends on the order p of the product formula used and scales asymptotically as $O(t^{p+1}/n^p)$ for n Trotter steps [25, 154]. However, this error is coming from non-commuting terms in the Hamiltonian decomposition, which suggests that tighter error bounds may be available when the specific structure of commutation is leveraged. In addition, in quantum computation, Trotterized time evolutions are always applied to some state, in which case one would ultimately care about the error in the resulting approximate time-evolved state rather than the operator error in the Trotter approximation itself. Research into providing tighter bounds and optimizing this error (e.g. by grouping and reordering of terms) is important and ongoing [26, 150, 155].

While Trotterization is most typically applied to Pauli decompositions of Hamiltonians, in principle, it can be applied to any decomposition into terms whose time evolutions can be simulated directly. For example, methods for simulation of sparse Hamiltonians usually involve decomposition into one-sparse terms [154, 156–159], and Trotterization of these is the approach for example in [154]. To our knowledge, there has been no systematic study of the Trotter error of these that goes beyond the asymptotic expressions obtained via tail bounds in [154].

In the case of Pauli decompositions

$$H = \sum_i c_i P_i \quad (3)$$

for Pauli operators P_i and real coefficients c_i , an upper bound on the number of steps can be given in terms of the ℓ_1 -norm of the coefficients, $\|H\|_1 = \sum_i |c_i|$. Such bounds, however, ignore commutativity between terms and tend to grossly overestimate the number of Trotter steps required to approximate the Hamiltonian evolution to a certain error. Indeed, in the trivial case where all terms in the Hamiltonian representation commute pairwise, a single Trotter step is enough to represent the Hamiltonian evolution exactly. A recent study on the theory of Trotter errors [26] provides more accurate bounds on the number of steps for a p th order formula in terms of norms of nested commutators $\sum_{i_1 \dots i_{p+1}} \| [H_{i_{p+1}}, [\dots [H_{i_2}, H_{i_1}] \dots]] \|$. Computing the precise commutator-scaling is exceptionally useful for reducing the number of Trotter steps, and, consequently, the quantum resources required, without sacrificing accuracy. It is, however, computationally expensive, as the computational cost, $O(N^{p+1})$, scales polynomially with the number of terms N in the Hamiltonian. Classical workflows (e.g. ordering of the Hamiltonian terms to reduce circuit depth) and their efficient implementation (e.g. leveraging high-performance parallel computing for accurate commutator-scaling evaluation) will play an important role in reducing the quantum resources required for practical applications and shortening the timeline to practical quantum advantage.

One final point to note is that the above discussion focuses on generic Hamiltonian simulation and the analysis of the Hamiltonian operator. In practice, especially when targeting early quantum applications, one is interested in the accurate evolution of a given Hamiltonian, for a given time, with a given initial quantum state. The details of choosing the appropriate Trotterization (and, more generally speaking, the Hamiltonian simulation approach) will be very much application-specific, and, when compared to results obtained by only analyzing the Hamiltonian operator, further reduction of the estimated resources will most certainly be achieved.

2. Linear Combination of Unitaries

Another approach to the simulation of e^{-iHt} involves block-encoding the Hamiltonian as a part of a larger unitary, $U(H)$. One of the most typical block-encodings is accomplished by writing H as a linear combination of unitaries, $H = \sum_n c_n U_n$, and building $U(H)$ using $U_P = \sum_n \mathcal{P}_n \otimes U_n$, where $\{\mathcal{P}_n\}$ is a complete set of orthogonal projectors, $\mathcal{P}_n \mathcal{P}_{n'} = \mathcal{P}_n \delta_{nn'}$ built in the ancilla qubit space. The computational cost of the algorithms based on linear combination of unitaries (LCU) depends on the ℓ_1 -norm of c_n coefficient vector, $\lambda = \sum_n |c_n|$ [160]. Thus, much effort has been directed to reducing the ℓ_1 -norm

by various selections of U_n operators [161–164]. It is also possible to show using triangle inequalities for norms that the lowest possible ℓ_1 -norm for LCU decompositions is $\lambda \geq (E_{\max} - E_{\min})/2$, where $E_{\max}/_{\min}$ are highest/lowest eigenvalues of the Hamiltonian in the entire fermionic Fock space or qubit Hilbert space [162]. Thus, in order to obtain ℓ_1 -norm lower than $(E_{\max} - E_{\min})/2$ one needs to modify the Hamiltonian.

Often in many-body problems, one is interested in eigenstates and eigenenergies in a particular symmetry subspace (e.g., with a fixed number of electrons for molecules). For quantum dynamics, if the initial state lies within a particular symmetry subspace, it continues to evolve there since the Hamiltonian preserves symmetries. In both cases, to reduce ℓ_1 -norm of the LCU decomposition, one can create an effective Hamiltonian that acts identically to the original in the particular symmetry subspace of interest, but has a lower spectral range, $(E_{\max} - E_{\min})/2$. In [163], this was done by employing the form of the effective Hamiltonian, $\hat{H}_{\text{eff}} = \hat{H} - \hat{O}_{1e}(\hat{N}_e - N_e \hat{I}) + c(\hat{N}_e^2 - N_e^2 \hat{I})$, where \hat{O}_{1e} is an arbitrary one-electron operator, \hat{N}_e is the operator of the number of electrons, N_e is the number of electrons in the symmetric subspace of interest, and c is a constant. Since finding the true spectral range of \hat{H}_{eff} is difficult, to optimize \hat{O}_{1e} and c in \hat{H}_{eff} one can use the simplest LCU decomposition in terms of Pauli products. \hat{O}_{1e} and c can be optimized to minimize the ℓ_1 -norm of Pauli product decomposition of \hat{H}_{eff} . As a direction for further improvement, simulating low-energy states or dynamics via LCU would benefit from either removing from consideration or shifting energies of highly excited states with the right number of electrons so that these states would not limit the ℓ_1 -norm of LCU.

Generally, the Pauli products (\hat{P}_k) are not the only unitaries that one can use for the LCU decomposition. A simple example of improving Pauli products is by grouping them into sets of anti-commuting Pauli products. If $\hat{A} = \sum_k c_k \hat{P}_k$, where c_k are real constants and $\{\hat{P}_k, \hat{P}_{k'}\} = 0$ then \hat{A} is proportional to the Hermitian unitary $\hat{R} = (1/\bar{c}) \sum_k c_k \hat{P}_k$, with $\bar{c} = \sqrt{\sum_k c_k^2}$. It is shown in Ref. [162] that such grouping of Pauli products in a qubit Hamiltonian will always reduce the ℓ_1 -norm with respect to the original Pauli product decomposition. There are other forms of unitaries that can be obtained from various fermionic tensor decompositions of the electronic Hamiltonians, for example, tensor hypercontraction or double factorization [161, 162]. To compare different LCU decompositions, one needs to account for three factors: 1) classical cost, 2) ℓ_1 -norm, 3) quantum resources for implementing (e.g., the number of ancilla qubits, T-gate count) of each time step. Note that ℓ_1 -norm mainly affects the size of the time step that one can take, while the third aspect concerns the quantum resources required for each time step.

C. Quantum measurement

Many quantum algorithms involve obtaining numerical values of observables, such as energies and dipole moments, as expectation values of their associated operators. Valuable use cases could lead to observables on 50 to 100 qubits, which when decomposed into Pauli operators can easily involve millions of terms; for instance, electronic structure Hamiltonians are characterized by approximately N^4 such Pauli products. Measuring all these Pauli products separately will take a prohibitively long time even on superconducting architectures, which have relatively rapid measurements. Grouping Pauli products is a natural approach to reduce the number of individual operators to measure. Here, we consider various grouping approaches to measure an expectation value of a single operator, like a Hamiltonian, for an arbitrary wavefunction. The generalization of this problem is measuring the expectation values of a set of operators for a set of wavefunctions. Such generalization poses additional optimization issues that can be important for some applications [165], but for the sake of brevity, we focus on the single operator, single wavefunction case for the remainder of this subsection.

All measurement schemes we describe can be seen as yielding estimators of the Hamiltonian expectation value by combining results of projective measurements of some Pauli operators. One can think of these as either measurements of generic Pauli operators or as measurements of the Ising Hamiltonian terms (i.e., diagonal Pauli operators) with respect to various states related by local basis transformations. Commonly examined physical Hamiltonians, like those for electronic structures, are not readily convertible to Ising form; for instance, through local unitary transformations. Therefore, they must be decomposed into simpler parts. A simple illustration of this idea is an observable $\hat{H} = \sum_n \hat{H}_n$ that is a sum of Hamiltonian fragments \hat{H}_n that can be transformed to the Ising form with some unitary rotations: $\hat{H}_n = \hat{U}_n^\dagger \hat{Z}_n \hat{U}_n$, where \hat{Z}_n is a polynomial of Pauli \hat{z} operators. Then

$$\langle \psi | \hat{H} | \psi \rangle = \sum_n \langle \psi | \hat{H}_n | \psi \rangle = \sum_n \langle \hat{U}_n \psi | \hat{Z}_n | \hat{U}_n \psi \rangle. \quad (4)$$

Hence,

$$\mathbb{E}(\hat{H}) = \sum_{n,k} \frac{Z_n^{(k)}}{M_n}, \quad (5)$$

where $Z_n^{(k)}$ are the results of measurements of \hat{Z}_n in the state $|\hat{U}_n \psi\rangle = \hat{U}_n |\psi\rangle$, and M_n are the numbers of measurements for each fragment.

Within this general scheme, a variety of expectation value estimation methods have been proposed. The methods can be categorized based on the following criteria:

1. Measurement Approach: Measurements within the same Hilbert space as \hat{H} versus in an extended

space using ancilla qubits (e.g., as for performing positive operator-valued measures);

2. Functional relations: Relations between estimators for measurable operators and that for the total Hamiltonian, $\mathbb{E}(\hat{H}) = f[\mathbb{E}(\hat{H}_n)]$ (e.g., linear combinations or more complex functional relations[166]);
3. Unitary Transformations: The types of unitary transformations that are needed to bring the measurable operators to the Ising form;
4. Estimator Error Scaling: The error scaling with the number of measurements is either $1/\sqrt{M}$ (shot-noise limit) or $1/M$ (quantum limit, i.e., Heisenberg scaling). Most near-term methods, which require low-depth measurement unitaries (i.e., \hat{U}_n) achieve only the $1/\sqrt{M}$ scaling. Achieving Heisenberg scaling requires much deeper circuits that are likely not feasible before fault-tolerance.

One of the estimation techniques that has become popular recently is classical shadow tomography [167]. In the framework described above, the unitary transformations \hat{U}_n are chosen randomly from the Clifford group (later, other groups were considered as well [168, 169]). Applying random Clifford transformations, $\{\hat{U}_k\}$, as an extension to a state preparation circuit allows one to effectively measure parts of the Hamiltonian that would be transformed to Ising forms \hat{Z}_k if $\{\hat{U}_k\}$ were applied to the Hamiltonian:

$$\hat{U}_k \hat{H} \hat{U}_k^\dagger = \hat{Z}_k + \hat{R}_k, \quad (6)$$

where \hat{R}_k is the non-Ising part). This shows that the information about the expectation value of \hat{H} is only obtained from the Ising parts because the expectation value of \hat{R}_k on any product state resulting from the wavefunction collapse is zero. This insight spurred various efforts to enhance the sampling convergence by biasing selection of $\{\hat{U}_k\}$ to those that increase the norm of $\{\hat{Z}_k\}$ parts [170–178]. A recent paper [179] has benchmarked these randomized measurement methods against a collection of molecular electronic structure Hamiltonians.

A more general view of this approach involves the use of adaptive, informationally incomplete (IC) Positive Operator Valued Measure (POVM) [180, 181]. This framework allows the optimization of the measurement for both the observable (e.g., H) and the state iteratively while collecting data. Moreover, the same mathematical framework can describe various physical implementations of the measurement [182] and naturally allows for the modelization of noise in the measurement [183]. Finally, it is worth mentioning that classical shadows and IC-POVMs allow for the estimation of multiple observables and for further post-processing, which can include noise mitigation methods [184].

To analyze the efficiency of various estimators one needs to take into account several factors: 1) classical

optimization cost for finding optimal measurable operators; 2) quantum resource overhead due to the need for extra unitary transformations (e.g., \hat{U}_n); and 3) measurement overhead due to the number of measurements needed to achieve a desired accuracy in the expectation value. In the near-term, the overhead related to additional unitaries is crucial for overall feasibility. For example, techniques using fermionic \hat{U}_n [185–188] could be less efficient than those using qubit-based Clifford transformations [189, 190] even though the former require lower numbers of measurements for a given accuracy. Additional options for mitigating circuit and measurement errors become possible using symmetries of measurable fragments. Selecting more symmetric fragments allows one to improve statistics by removing the results that violate expected symmetry constraints – e.g., the number of electrons – which provides an error mitigation protocol [191–193].

A crucial question for ongoing research is which types of physical Hamiltonians admit feasible measurement schemes in the near term. Spin models typically partition into a few natural measurement bases; e.g., a Heisenberg model typically contains terms of the form XX , YY , ZZ , and Z , which form three local measurement bases (all X , all Y , and all Z). On the other extreme, molecular electronic structure Hamiltonians have several terms that scale $\sim N^4$ for N orbitals. It is unclear whether any known measurement scheme can reduce the resulting extremely large measurement counts to something tractable on near-term devices for instances large enough to offer the potential for quantum advantage. Hence an important question is whether there exist intermediate problems where measurement techniques beyond simple grouping – as in spin models – can be leveraged to achieve tractable measurement counts. For fermionic Hamiltonians, this is closely connected to choices of fermion-to-qubit mapping, as discussed in Sec. III A. A potential example is the Fermi-Hubbard Hamiltonian, where it is possible to partition the Hamiltonian into 5 operators consisting of commuting Pauli products, and each operator is converted into Ising form by at most one layer of Bell measurement circuit (Controlled NOT gate (CNOT) + H) [194]. By estimating symmetries such as parity, these Bell measurement circuits allow one to mitigate circuit errors [195]. Another excellent study on the topic as well as other elements of the materials simulation stack can be found in Ref. [196].

D. Problem-level preprocessing and optimization

Achieving quantum advantage on pre-fault-tolerant quantum computers is likely to require optimization at all levels of the algorithmic stack. This includes optimization of the input problem itself, as well as optimization that takes the input instance into account. Since the family of possible input problems to quantum algorithms for materials science is too large for an exhaustive sur-

vey here, we instead provide two examples of this kind of problem-level preprocessing. Part of the goal of this is simply to emphasize that this step is almost certain to be a necessary part of any serious attempt to achieve quantum advantage on a noisy quantum computer.

Our first example is active space selection (see also Sec. VII C 2). A fundamental problem in chemistry and materials science is the eigenvalue problem, where the eigenstates of the Hamiltonian describing the system are to be solved. Diagonalizing the Hamiltonian of the full system is, however, too expensive and oftentimes unnecessary. In many systems of practical interest, e.g., point defects in wide bandgap semiconductors [197] and catalysts on surfaces or interfaces [198], electronic excited states of extended molecules [199], a certain subset of electrons and orbitals in the molecule or solid is more relevant to the problem than the rest. Therefore, a so-called active space can be defined, leading to an effective Hamiltonian. Solving this effective Hamiltonian (on a quantum computer) is much easier than solving the full Hamiltonian, while the essential properties of the system are retained. Generally speaking, defining active spaces requires identifying a chemically active site on the system. Efforts have been put into developing automated active space selection (see next paragraph), but a universal approach is so far lacking. In practice, prior knowledge of the system studied and the chemical intuition of the researcher is usually relied upon. For example, spin defect systems, e.g., the nitrogen-vacancy center in diamond and an electronic defect in MgO [200–203], are heterogeneous materials promising for the realization of quantum communication. They have a defect center hosted in a periodic solid, whose local environment is vital to its electronic, optical, and mechanical properties [204]. Therefore, a proper active space should include localized orbitals around the defect center [205, 206]. Other examples are pseudotetrahedral organometallic complexes containing chromium(IV) and aryl ligands which have been experimentally identified as promising molecular qubit candidates [207, 208].

Very often, local perturbations in extended systems such as spin defects described above or gas adsorption on solid-state catalysts are multiconfigurational, which means that several electronic states are degenerate in energy. The accurate description of such systems require multireference methods. Among the most common ones are the complete active space self-consistent field (CASSCF) method [209–211], density matrix renormalization group (DMRG) [212], multireference configuration interaction (MRCI) [213], and multireference perturbation theories. However, these methods scale poorly with system size and exponentially with the size of the active space and thus are impractical for use in extended systems.

An ab-initio quantum embedding technique, density matrix embedding theory (DMET) [214–217], and its periodic implementations provide a framework that allows users to effectively reduce an extended system to a

finite embedding subspace [218, 219]. The user selects a fragment space characterized by orbitals obtained by using orbital localization schemes. Electrons in such orbitals are generally entangled to the environment, so it is not possible to define a wave function in the fragment orbitals alone. However, the entangled part of the environment can be obtained approximately using the Schmidt decomposition of a single-determinant whole-system approximate wave function (usually the Hartree–Fock wave function) [216]. This involves the diagonalization of the environmental block, D_{env} , of the one-body reduced density matrix (RDM):

$$D_{\text{env}} = \mathbf{U}\lambda\mathbf{U}^\dagger \quad (7)$$

Here, λ represents a diagonal matrix containing the eigenvalues λ_i for $i = 0, 1, \dots, N_{\text{env}}$, and where N_{env} indicates the count of environmental orbitals. The columns of the unitary matrix \mathbf{U} , corresponding to non-zero and non-two λ_i values, identify the entangled bath orbitals. Localized fragment orbitals and entangled bath orbitals together define a subspace of the molecule, no larger than twice the size of the original fragment, which is unentangled at the level of the underlying single-determinantal level of theory and within which any quantum chemistry method can be applied by treating all other orbitals as part of a frozen core.

In the case of spin-defects, the fragment active space is usually selected by including orbitals centered around the defect, whereas in the case of gas adsorption, the orbitals centered around the gas molecules and the surface atoms they are bound to are considered. Methods like CAS-DMET and NEVPT2-DMET [201–203] utilize this framework and provide users the capability of modeling multireference ground and excited states arising from such local perturbations in extended solids with high accuracy. In CAS-DMET, the embedding subspace constructed using Eqn. (7) is subjected to a CASSCF calculation which further reduces the effective active space and accounts for the most relevant static correlation. The remaining dynamic correlation is captured by multireference perturbation theories like NEVPT2 [220–222] as is done in NEVPT2-DMET.

A more cost-effective method compared to NEVPT2-DMET to include electron correlation is DME-PDFT [223]. In this approach, the 1- and 2-RDMs generated from a CAS-DMET calculation are used to compute the densities and on-top pair densities, which are then used in a subsequent multiconfiguration pair-density functional theory (MC-PDFT) calculation [224–226]. MC-PDFT offers a method for integrating multiconfiguration wave function theory and density functional theory, allowing for the comprehensive treatment of both near-degeneracy correlation and dynamic correlation in strongly correlated systems. It proves to be a more cost-effective alternative to multireference perturbation theory, multireference configuration interaction, or multireference coupled cluster theory. Moreover, MC-

PDFT demonstrates superior accuracy for various properties compared to Kohn–Sham density functional theory. MC-PDFT computes the total energy as a functional of the electron density (ρ) and on-top pair density (Π) obtained from a multiconfiguration wave function.

The MC-PDFT energy is expressed as:

$$E_{\text{MC-PDFT}} = V_{NN} + \sum_{pq} h_{pq} D_{pq} + \frac{1}{2} \sum_{pqrs} g_{pqrs} D_{pq} D_{rs} + E_{ot}[\rho, \Pi], \quad (8)$$

where V_{NN} is the nuclear-nuclear repulsion energy, h_{pq} and g_{pqrs} are one- and two-electron integrals, D_{pq} are the elements of the one-electron RDM, and the on-top energy E_{ot} is a functional of the density ρ and the on-top pair-density Π . It has been used in combination with various multireference wave functions like the generalized active space wave function [227], DMRG [228, 229] and it has been used to compute the total energy of strongly correlated systems starting from the 2-electron RDM evaluated on the quantum device.

In the context of DMET, PDFT is especially attractive because the on-top energy is expressed in terms of electronic density on a quadrature grid using semi-local (and therefore linear-scaling) functionals which are agnostic to the origin of the underlying density matrices. This formalism thus models electron correlation of all electrons regardless of any underlying embedding methodology. Therefore, DME-PDFT is formally less sensitive to the size of the embedded fragment than methods such as NEVPT2-DMET, which can only model electron correlation within the embedded subspace of fragment and bath orbitals [223].

A promising step towards heterogeneous catalysis applications is the recent study of adsorption with memory-efficient DMET algorithms which, when combined with multireference electronic structure solvers, can be extended to the bond-breaking phenomenon on surfaces [230]. A computationally affordable alternative to CASSCF for modeling larger active spaces is the localized active space self-consistent field (LASSCF) method [231, 232]. LASSCF divides a given active space into localized non-interacting subspaces, connected through a mean field, and has the ability to capture strong electron correlation within these subspaces without encountering the computational cost associated with CASSCF.

The LASSCF wavefunction is represented as an antisymmetrized product of the Full Configuration Interaction (FCI) wavefunctions of these individual localized subspaces and the single determinantal wavefunction of the inactive space, which remains delocalized over the entire molecule. This is mathematically represented as:

$$\Psi_{\text{LAS}} = \bigwedge_K (\Psi_{A_K}) \wedge \Phi_D \quad (9)$$

where Ψ_{A_K} denotes the many-body (generally FCI) wavefunction of the K th localized subspace, and ϕ_D denotes the single-determinantal wavefunction of closed-shell occupied inactive orbitals. The energy of the system E_{LAS} is obtained through variational optimization and is expressed as:

$$E_{\text{LAS}} = \langle \psi_{\text{LAS}} | \hat{H} | \psi_{\text{LAS}} \rangle \quad (10)$$

where \hat{H} denotes the molecular Hamiltonian.

Although LASSCF treats strong correlations within these localized smaller active spaces, it only accounts for interactions between localized fragments at a mean-field level. Methods such as State Interaction (LASSI) on a classical computer [233] or the variational unitary coupled cluster singles and doubles (UCCSD) on a quantum computer (LAS-UCC) [234] model correlation between these fragments using LASSCF as a reference wave function. In the LAS-UCC method, the multiconfigurational subspace wavefunction Ψ_{A_K} for each fragment is prepared using either QPE or Direct Initialization (DI), as outlined by D’cunha *et al.* [235]. The inter-fragment correlations are then captured through a standard VQE optimization process using the Unitary Coupled Cluster (UCC) ansatz.

Despite these recent advances in computational methods for modeling electronic structures, significant research challenges exist, particularly involving the treatment of “larger” active spaces and the inherent poor scaling of multireference perturbation theories. Numerous research efforts have been directed towards the treatment of larger active spaces using classical electronic structure techniques like DMRG, selected CI, and quantum Monte Carlo methods. However, expanding the active space relies on physically motivated approximations to balance computational cost [7, 9, 236–238]. The basic philosophy in all these approaches is focusing on the most important configurations while eliminating less significant ones, based on physically motivated criteria. Quantum computing holds the promise of effectively modeling these larger active spaces accurately, especially as the number of available qubits increases and noise mitigation techniques become more powerful. Embedding techniques such as LAS-UCC may help accelerate this process. While perturbation theories scale poorly on classical hardware, recent studies on formulating quantum implementations of perturbation theories [239] highlight the potential for creating multireference quantum algorithms, including NEVPT2-DMET. With more qubits being available, the primary research challenge still remains to determine the optimal methods for quantum implementation that maintain a favorable balance between computational cost and the desired accuracy.

Our second example of problem-aware preprocessing is VQE ansatz design [33]. VQE (and its variants) is the most widely used algorithm on quantum computers for the eigenvalue problem, although its experimental applications to date have been limited to classically-simulable demonstrations; see, e.g., [31–47, 49, 51, 52]. The core of

the VQE algorithm is the parameterized quantum circuit $U(\vec{\theta})$, usually called the ansatz, whose output is a manifold of states that should contain an approximation to the ground state of the system. The ansatz can be evaluated according to two key attributes: expressivity and practicality. The former measures how well the ansatz can approximate the solution to the problem to a certain accuracy, while the latter takes circuit depth, the difficulty of optimizing those parameters, and the qubit connectivity of the underlying hardware into consideration. Quantum advantage in VQE requires identifying an ansatz that satisfies both criteria, i.e., it is a feasible circuit on a given device, it is trainable at scale on a classically hard problem, and it can represent a low-energy state with sufficient accuracy to surpass classical methods.

Two typical examples of ansatzes are the UCC ansatz [33, 36, 48, 240] and hardware efficient ansatzes [38]. The UCC is chemically-inspired and thus very expressive for most chemistry and materials, but usually entails a large number of gates and is likely beyond near-term hardware. The hardware-efficient ansatzes are problem-agnostic and designed specifically to use the native gates and connectivity of a given device, but may be impractical to optimize due to the barren plateau problem [241], and also may fail to be sufficiently expressive at low depth, lead to symmetry breaking, and non-smooth potential energy surfaces [242].

Efforts have been put into designing new ansatzes to strike a balance between these two dimensions and unify the advantages of UCC and hardware efficiency [243–246]. The qubit coupled-cluster ansatz [243, 247, 248] and ADAPT-VQE [246] are among such examples, where a pre-screening process is proposed to select certain Pauli operators to generate the gates in the ansatz, thereby reducing the circuit depth and number of parameters. These methods that trim the number of variational parameters are promising, but there is still a lack of evidence for the scalability of parameter training. An exception may be HMP2 [50], where second-order perturbation theory is employed to reduce the number of variational parameters, while providing real-time, quantum computation-based methodologies to identify additional Pauli operators to include in the VQE progression. These advancements indeed merit further research despite outstanding challenges due to the simplicity of variational quantum algorithm (VQA) in all other respects (particularly in circuit complexity) and the extensive success of classical variational algorithms.

Another challenge with ADAPT-VQE-like algorithms lies in the larger measurement overhead due to the selection of Pauli operators, which can however be mitigated by the use of IC-POVM measurements [249]. While it is important to design an ansatz structure that is most appropriate for a target problem, appropriately initializing the ansatz parameters is an equally critical task when it comes to the trainability of the resulting optimization. This is because random parameter initialization can be

extremely far from the optimal parameters, and the ability for VQE to converge to the optimal parameters (and thereby find the ground state) will heavily depend on the optimizer and the effects of noise. Identifying good classical initialization can therefore considerably improve the accuracy and speed of VQE convergence. One promising method is CAFQA [250]. The CAFQA ansatz is built with only Clifford gates and is therefore classically simulable in polynomial time. A discrete optimizer is used to search through different possible Clifford gate parameters until the optimizer converges to a set of parameters that minimizes the VQE expectation (to the best possible extent within the Clifford space). This ‘CAFQA’ state is then used to initialize traditional VQE. There is an opportunity to explore other forms of classical simulation to bootstrap VQE; near-Clifford circuits and tensor network approaches are promising directions [251, 252]. All these methods could benefit from problem-specific optimization. For instance, restricting the CAFQA Clifford search to a state space that is expected to have some overlap with the ground state can improve initialization speed and accuracy.

E. Classical data loading

At the highest level, quantum computation involves the following steps. First, a classical description for a specific problem is established. This description is input into a quantum circuit. After executing the quantum circuit, we retrieve classical measurement results as the solution to the problem. Within this workflow, an essential challenge is the data input. A robust quantum algorithm must explicitly detail this input technique and the associated hardware.

For some types of use cases, e.g., quantum simulation algorithms, the problems themselves are quantum mechanical in nature, and often this permits direct mappings to quantum circuits, such as those discussed in Sec. III A. In more general cases, including when the inputs are classical data, scientists who focus on theoretical complexity analysis often abstract such input methods as an ‘oracle’ and investigate how many times this oracle must be applied, using this as a basis to estimate the algorithm’s complexity. However, creating a physical realization of a quantum oracle is also a fundamental aspect of realizing quantum computation.

QRAM [253] serves as the physical embodiment of the quantum oracle. Unlike traditional RAM, where a single piece of data can be swiftly loaded into the central processor, QRAM allows for multiple classical data items to be loaded concurrently in superposition. This unique capability positions QRAM as a bridge between classical and quantum computing realms. More than that, QRAM can function as a rapid and versatile quantum oracle, effectively addressing the data-input bottleneck prevalent in numerous quantum applications. QRAM is also a fundamental assumption for many quantum al-

gorithms. For instance, many quantum machine learning algorithms [254, 255] may require QRAM as a prerequisite to potentially achieve quantum speedups, although some algorithms have been discovered to achieve the speedups without [256]. Currently, given the significant challenges in physically constructing QRAM, both theoretical and experimental research on QRAM is an active field of study.

For quantum simulation and material sciences, QRAM could primarily be used to provide fast, efficient, and reliable initial states and Hamiltonian encoding. QRAM architectures could provide useful interfaces, benefiting practical use cases of quantum simulation to large-scale, challenging problems in material science.

1. QRAM construction

QRAM could be considered as the following unitary,

$$|\psi_{\text{in}}\rangle = \sum_{i=0}^{N-1} \alpha_i |i\rangle^A |0\rangle^B \xrightarrow{\text{QRAM}} |\psi_{\text{out}}\rangle = \sum_{i=0}^{N-1} \alpha_i |i\rangle^A |x_i\rangle^B \quad (11)$$

starting from an input state $|\psi\rangle_{\text{in}}$ in superposition, one can get an output state $|\psi\rangle_{\text{out}}$. We assume that the data vector x_i has N components. In order to construct QRAM, quantum routers are important ingredients. One can write quantum routers as quantum circuits made by a sequence of controlled-SWAP gates. A QRAM can be assembled using quantum routers [257], as illustrated in Figure 1, where routers are organized in a binary tree formation. This basic QRAM configuration made by quantum routers is termed fanout QRAM, and it allows for data storage of $O(N)$ elements in a time complexity of $O(\log N)$.

On the other hand, the bucket-brigade architecture presented in [253] serves as a nuanced version of the fanout design, distinguished by dynamic routing of address qubits: address qubits are dynamically introduced into the tree during a query. If an address qubit comes across a router in the $|0\rangle$ or $|1\rangle$ state, it routes to the left or right, respectively. However, upon encountering a router in the $|W\rangle$ state (each router could incorporate a third state labeled $|W\rangle$, for “wait,” so it might be a qutrit system), the router and incident mode states are interchanged, setting the router’s state to $|0\rangle$ or $|1\rangle$ based on the initial incident address. Such designs are more robust to decoherence [253].

QRAM stands for a type of extremely fast quantum memory that may be better suited for handling large-scale data. On the other hand, when the amount of data is relatively small, one might use what’s known as Quantum Read-Only Memory (QROM) to upload quantum data, for tasks like simulating certain Hamiltonians [258]. Designs that combine QRAM and QROM are referred to as hybrid architectures [259]. Since QROM requires

$O(N \log N)$ time and $O(\log N)$ space, hybrid architectures might represent a space-time trade-off for specific hardware conditions (see Figure 1).

QRAM might potentially have numerous applications if it is realized on a large scale. Here we summarize some potential applications of QRAM.

- **Quantum algorithms.** Many quantum algorithms based on quantum machine learning [255, 260] might require QRAM as a quantum oracle for providing potential exponential quantum speedups. Refs. [261, 262] discuss how quantum algorithms could provide hardware requirements to QRAM architectures. Ref. [263] discusses how QRAM could bring potential enhancement for quantum simulation algorithms in a data center setup. For instance, one could design QRAM circuit for Hamiltonian simulation of quantum chemistry or quantum materials similar to [258], which is an efficient way for uploading Hamiltonian data towards quantum computers.
- **Quantum Data Center (QDC).** Nowadays, data centers are important businesses in the digital society, and it is natural to have a quantum version of data centers in the coming era. Refs. [263, 264] propose a theory of QDCs, where QRAM and quantum networks are combined to provide an elementary definition of QDC. These works show that QDCs with several applications in QRAM as fast quantum memories, will provide fast, secure, and accurate applications in quantum computing, communication, and sensing.
- **Quantum communication and sensing.** Quantum Private Queries [265] and the associated blind quantum computing [266] could be flagship applications of QRAM in quantum communications, where the users could ask data centers to provide specific services without privacy leakage. Moreover, [263] provides a quantum data compression algorithm based on variants of QRAM, an application helpful for distributive quantum sensing applied to, for instance, quantum telescopes [267].

Fault-tolerant designs of quantum random access memory are investigated by [268], showing that QRAM with full fault tolerance might be very challenging on a large scale. On the other hand, noisy QRAM might still be error-resilient for generic setups of decoherence and noisy channels [259]. In fact, though fanout QRAM designs might be vulnerable to decoherence, rendering them non-scalable, the bucket brigade QRAM architecture is known for its high noise resilience [253, 259]. Specifically, the infidelity of a query grows only logarithmically with memory size [259]. Yet, past studies leaned on specific noise models, raising questions about the practical advantages in real-world implementations. Ref. [259] delved into the effects of decoherence on QRAM in a comprehensive manner. A significant takeaway is the affirma-

tion that the logarithmic infidelity scaling persists across varied error channels, including challenges like depolarizing noise and coherent errors. The research identifies limited entanglement among the memory components as the chief reason behind this noise resilience. Interestingly, this insight also suggests potential architectural simplifications without compromising noise resilience, serving an interesting future research question. This understanding implies that with contemporary hardware, QRAM might be potentially realized in real-world noisy environments, facilitating high-fidelity queries even with early fault-tolerant technologies. Additionally, when quantum error correction is integrated, the bucket brigade architecture continues to show promise, enhancing hardware efficiency and fortifying against logical errors [253, 259].

2. Experimental efforts towards QRAM

Aside from theoretical studies, significant experimental progress has been made towards realizing QRAMs. In one of the original papers on QRAM [269], two different fanout designs are presented together with the bucket-brigade designs, quantum optical fanout and phase gate fanout, where hybrid systems of photons and trapped atoms are used. Ref. [270] realizes a quantum memory of 105 qubits carried by 210 memory cells in a macroscopic atomic ensemble, with demonstrations of storage of optical qubits into these memory cells and their readout, which might be a significant step towards full quantum random access memory with superpositions of addresses.

Refs. [259, 271, 272] propose a hardware-efficient QRAM design with Hybrid Quantum Acoustic Systems. More precisely, [259, 271, 272] introduce a methodology for quantum computation utilizing multi-modal quantum acoustic systems and, based on this framework, suggest a streamlined approach for QRAM construction. Quantum data is preserved in high-Q phonon modes. Interactions between these modes are deliberately crafted by applying off-resonant stimuli to a transmon qubit. Contrasted with previous suggestions that emphasize direct qubit excitation, it is claimed that their approach can significantly enhance gate fidelity, especially for enduring acoustic modes.

Ref. [273] introduces a QRAM construction approach using a photonic-integrated-circuit (PIC) architecture combined with solid-state memories. Additionally, [273] offers a novel scheme rooted in quantum teleportation, further expanding its application in quantum networks. Both of these methods effectively carry out the primary QRAM functions: quantum state transfer and quantum routing. These implementations are grounded in existing components, including electro-optic modulators, a Mach-Zehnder interferometer (MZI) network, and nanocavities linked to artificial atoms for spin-based memory operations.

Constructing a large-scale, fault-tolerant QRAM presents significant challenges, which could be summa-

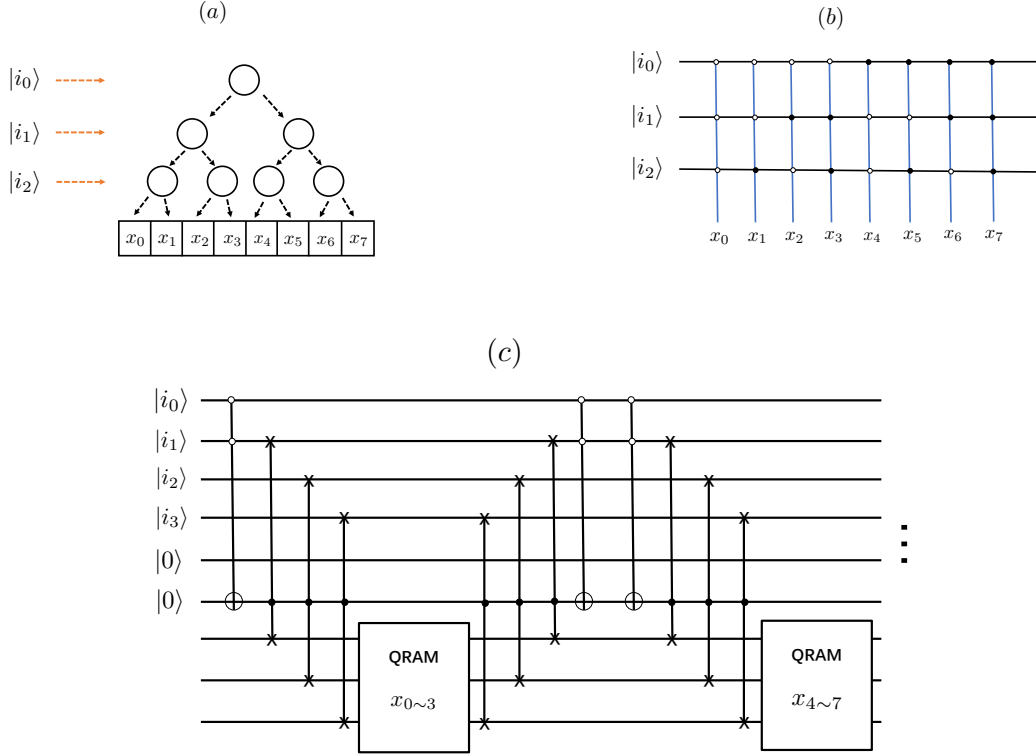


FIG. 1. (a) An illustration of QRAM made by quantum routers. (b) An illustration of QROM. (c) Hybrid QRAM-QROM architecture.

rized as the following.

- Scalability: The number of qubits required for a QRAM scales linearly with the data set size. As data sets grow, so does the qubit demand. Moreover, larger data sets necessitate more substantial QRAM hardware with higher circuit depth and more sophisticated designs.

There are various ways to address the scalability issue by hardware co-designs of QRAM. For instance, we can utilize Q^2 routers as foundational components for constructing QRAM (Here, Q^2 router refers to the router with both control and signal states being quantum states). This modular design allows us to assemble QRAM by interlinking various modules, as depicted in Figure 2. Alternatively, a 2-D integrated QRAM can be achieved using *H-tree* patterns, either through the teleportation protocol as outlined in [274] or by employing long wires (see Figure 3. Note that this design is beyond [274] by implementing long-range coupling wires, within the capability of IBM devices. Moreover, the design is compatible with a *single-layer* architecture without crossing coupling wires). For a scalable QRAM, extensive range connectivity might be bounded by some principles of physics as discussed in [275].

- Efficiency: To ensure QRAM’s efficiency and reliability, the quantum hardware must be both highly compact (fractal or self-similar structures in the design of QRAM with more efficient usages of space like [274]) and highly coherent. In addition to superconducting circuits, it is also intriguing to explore more compact phononic circuits for QRAM like [271].
- Error Correction: As QRAM size expands, errors become inevitable. Thus, strategies for reliable data loading despite these errors are essential. Standard quantum error correcting codes might be applicable to QRAM. However, developing quantum error correction techniques that are specifically adapted to QRAM [268, 276] remains an open research area. On the other hand, error resilience of designs of QRAM [259] might provide possible resolutions. Moreover, error resilience is a feature of bucket-brigade QRAM, which shows that errors only scale poly-logarithmically with the system size (see above). This indicates that QRAM can already be a useful application for some non-error-corrected quantum computers.
- Need for Specialized Solutions: Given its unique purpose and architecture, QRAM’s scalability issues should be approached with its specific needs

in mind. Unlike universal quantum computers, QRAM doesn't need a universal set of quantum operations, which could be an avenue for simplifying its architecture. Moreover, some classical optimization could be helpful to enable more sophisticated QRAM designs. For instance, [259] designs a noisy QRAM simulator in classical devices, which might be for further explorations of QRAM designs and simulation in a classical space.

F. Circuits, transpilation, and architecture

Just as in classical computing, every quantum computation that is to be carried out on a quantum computer needs to be broken down to a netlist of native operations that admit a direct implementation on a quantum computer. One of the most widely-used representation of such a netlist is the so-called quantum circuit, which is to be transpiled for a target quantum computer backend. Different quantum backend hardware may support different qubit-to-qubit connectivity, which limits the interactions between the qubits that may be implemented natively.

To this end, quantum applications in materials science are no exceptions. Indeed, quantum circuits for Heisenberg Hamiltonian simulation [152, 153], relevant for many-body localization and room-temperature superconductivity, and electronic structure calculations for small molecules [48, 50, 53], have been optimized for both pre-fault-tolerant and fault-tolerant settings before. Other examples include particle physics simulations [131, 133, 277], where fault-tolerant cost, including the number of non-Clifford quantum gates and qubit counts, have been minimized. In all these examples, it is critical that the resulting netlist of instructions to be carried out by a quantum computer is optimized, targeting *time to solution*: a key metric in all types of computation, including quantum. Case in point, if time to solution is of no matter, the computational advantages offered by quantum computing is moot and classical computing suffices, and a high-fidelity operation of a quantum computer would become unnecessary since repeated executions of computation can be performed until an acceptable answer is obtained. The important roles that circuits, transpilation, and architecture will continue to play in leveraging the computational power enabled by quantum computation then form the central topics of this section.

1. Connectivity and Architecture

We begin by highlighting key topics around qubit connectivity, a fundamental aspect of a quantum processing unit (QPU) that dictates the efficiency, scalability, and feasibility of the quantum devices. We first consider SWAP networks, highlighting their role in enabling

Hamiltonian simulation in quantum systems with limited connectivity. The discussion then moves to distributed quantum computing, emphasizing its importance in scaling quantum computing beyond the capabilities of current QPUs. Quantum circuit cutting is also examined for its practicality in decomposing complex quantum circuits into smaller segments, facilitating distributed processing.

a. SWAP networks One of the tools that can be leveraged to effectively implement Hamiltonian simulation primitives on hardware with minimal (linear) connectivity is the SWAP network [278–280]. Through a SWAP network, depicted in Figure 4, all $\binom{n}{2} \in O(n^2)$ pairs of (commuting) interactions between n qubits can be accomplished in just $O(n)$ circuit layers, even on a quantum computer architecture that admits nearest-neighbor qubit-to-qubit connectivity. With the exception of an all-to-all architecture that admits a single-instruction implementation of commuting, overlapping two-qubit gates [281–283], the asymptotic $O(n)$ scaling matches that of an architecture that allows for random-access two-qubit interactions that may be implemented in parallel so long as the interactions do not overlap. In the context of the Materials Science applications in this paper, any “dense” interactions with commuting terms can have a natural implementation on an all-to-all architecture (c.f. Refs. [50, 53, 133, 153]) or on a limited-connectivity architecture with SWAP networks. Reducing the overhead in moving the quantum information around for optimal implementation of a given set of two-qubit interactions remains an area of ongoing research.

b. Distributed/modular Just like classical HPC, distributed quantum computing is necessary for the continued scaling of quantum computing. While monolithic QPUs continue to make significant progress in both sizes and qubit qualities, there remains a huge gap between available QPUs and practical benchmarks. To make matters worse, developments from classical computing demonstrate that the computing and memory capacity gaps between large-scale industrial use cases and single computing cores only get larger, despite having more powerful monolithic cores. To combat this increasing gap, almost all industry-scale classical applications today such as Large Language Model (LLM) training, and even many personal-scale workloads such as intensive gaming, happen in parallel. In fact, distributed computing is only more imperative for quantum computing as single-core scalability poses significant engineering challenges.

There have been distributed quantum computing proposals from both hardware and software researchers. IBM, Quantinuum, and IonQ all announced plans to develop modular QPUs with interconnects to enable distributed quantum computing. Ref. [284] models hardware performances of potential superconducting modular QPUs. Ref. [285] proposes compilation algorithms for distributed QPUs with EPR pair connections. Ref. [286] proposes ClusterVQE which enables domain-specific optimizations to map applications from a target domain (here, VQA) to distributed QPUs. Further research

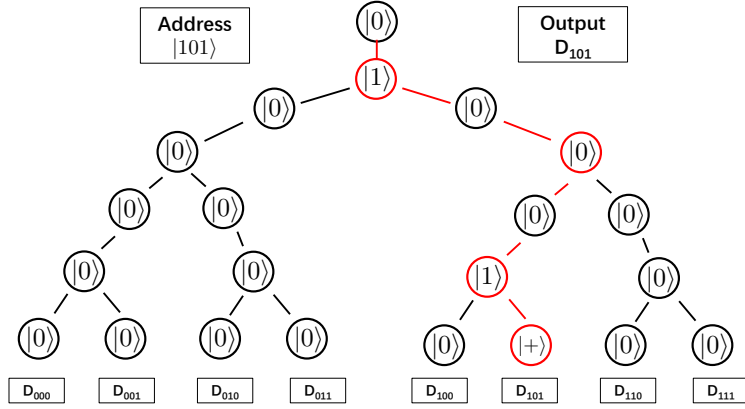


FIG. 2. QRAM as a binary tree of Q^2 routers. Here, Q^2 router refers to the router with both control and signal states being quantum states.




FIG. 3. (a) Self-similar fractal structure of the H-tree designs (see [274]). One can use the self-similar structure to create a size 2^{n+2} QRAM from 4 units of size 2^n QRAM. (b): An example of depth-4 tree with $2^4 = 16$ leaves.




FIG. 4. Example SWAP network for $n = 6$ qubits. In $O(n)$ steps, each of the $O(n^2)$ pair of qubits (colors) performs an interaction, even with just linear connectivity.

across the stack from applications to devices is necessary to enable distributed quantum computing with tolerable overheads.

c. Circuit cutting Quantum circuit cutting offers an orthogonal software path to enable distributed quantum computing to hardware modular QPUs. Cutting qubit wires [287] or quantum gates [288] breaks down a large quantum circuit into multiple smaller subcircuits.

Multiple less powerful QPUs run the smaller subcircuits in parallel. Classical computing eventually reconstructs the original circuit output. End-to-end implementations [289, 290] demonstrate cutting up to 200-qubit benchmarks to distribute onto multiple QPUs available nowadays, which are otherwise intractable for purely quantum computing or classical simulations. The primary challenges in applying quantum circuit cutting at practical scales stem from the demands of classical post-processing: firstly, the output length of an n -qubit circuit is 2^n , which rapidly strains classical memory and processing time for large quantum circuits; secondly, the classical post-processing effort scales exponentially with the number of required cuts, K , further constraining the runtime.

Despite its ability to execute previously intractable workloads, there is still a gap between the computational capacity of circuit cutting now and future practical quantum workloads. The first theory proof [287] proposed an exponential classical reconstruction cost. Instead, subsequent works [290] drastically reduce the overhead to sub-exponential but remain prohibitive in the worst case. Filling the gap calls for further developments in classical

reconstruction techniques, algorithm-backend co-design, and quantum computing cloud backend system adaptations [291].

In the next two paragraphs, we shift attention to elements beyond QPUs, addressing key developments in memory and qubit-state storage. We delve into QRAM and QROM, emphasizing how compiler strategies like the Gray Code can enhance memory efficiency. Further, we examine multimode cavity systems, an innovation in quantum memory, providing effective storage solutions for qubit states while also introducing distinctive integration and scalability challenges. How these memory elements are to be eventually integrated with a to-be-developed compiler to enable efficient quantum computing is an open area of research.

The QRAM and QROM memory primitives, discussed earlier in Sec. III E, can be optimally implemented with appropriate compiler techniques. As one example, consider the Gray Code, a sequence of the 2^n n -bit bitstrings such that only 1 bit changes between sequential neighbors, for instance, 000, 001, 011, 010, 110, 111, 101, 100. Integrating awareness of the Gray Code in a compiler can significantly reduce gate counts for memory implementations. Additional compilation techniques for these memory primitives, including routing and mapping, are discussed in prior work [274, 275, 292].

Multimode cavities with tens of modes and photon lifetimes on the order of milliseconds pose as efficient mediums to store multiple qubit states. When used as quantum memory or registers, they can allow for effectively extending the lifetimes of qubits. These cavities can be coupled with one physical qubit, to allow for swapping of a qubit state in and out of the cavity to perform multi-qubit operations with other qubits. While integrating multimode cavities reduces swap-distance to enable multi-qubit gates for a larger number of qubits, it also imposes serialization of swaps of qubit states from a single cavity posing additional compilation challenges. In terms of physical integration of these cavities to superconducting circuits, there remain challenges in terms of reducing 3D-cavity size or utilizing 2D resonators while maintaining long lifetimes of stored qubits [293].

2. Circuit optimization

a. Sequences of Pauli rotations Most standard implementations of Hamiltonian propagators rely on the implementation of a sequence of Pauli rotations, often coming from a Trotter expansion. In the near-term setting, a sequence $R_{P_k}(\theta_k) \cdots R_{P_1}(\theta_1)$ is usually implemented via an alternating sequence of Clifford circuits and single-qubit rotations. This prompts the natural problem of finding the best possible sequence of interleaved Clifford circuits and single-qubit layers implementing some goal sequence of Pauli rotations. In some cases, the ordering of the rotations may also be relaxed, leading to even further optimizations. Some work already tackles this

synthesis problem both for all-to-all architectures and for restricted coupling maps [48, 50, 53, 152, 153, 294–296]. This problem is tightly related to the *parity network synthesis problem* introduced in [297], for which many architecture-aware heuristics have been developed [298, 299]. It is still an open question to formalize the notion of Pauli networks in the presence of ancillas and measurements.

b. Feedforward One of the key emerging tools for efficient compilation is dynamic circuits that make mid-circuit measurements that feedforward into subsequent operations. While in principle any quantum circuit can be implemented without midcircuit measurements, due to the *principle of deferred measurement*, there are several scenarios where the dynamic approach significantly lowers circuit depth or circumvents topology constraints. For example, controlled rotation gates can be implemented with a smaller number of non-Clifford gates using feedforward [300, 301], which are beneficial for optimizing Heisenberg hamiltonian simulations [153]. Another example may be that, while a n -qubit GHZ state requires $O(n)$ layers with linear qubit connectivity (and $O(\log n)$ layers with all-to-all connectivity), the state can be produced in constant $O(1)$ depth with feedforward midcircuit measurements. More generally, any Clifford circuit can also be implemented in constant depth with midcircuit measurements on linear nearest-neighbor connectivity. In the context of materials science applications, this is useful for primitives such as the CNOT ladder used in Hamiltonian simulation. Provided with the already identified utility of feedforward in optimizing quantum circuits, further developments in the use of feedforward to optimize additional elements of quantum circuits are anticipated.

c. Pre-fault tolerant vs. Fault-tolerant The central task of a compiler is to translate a human-readable code to a machine code that can then be executed by a computer using its native operations. As such, the compiler plays a critical role in both the accessibility of the software development environment for the end users and the efficiency of the machine code that faithfully implements the code that the users want to execute using the least amount of computational resources. Specifically to the latter, an optimizing compiler can be developed, which takes into account various computational operations, possibly having varying cost [50, 146, 153]. An example applicable to quantum computing may be in the context of the pre-fault tolerant vs. fault-tolerant quantum computers, where the former typically enjoys faster and higher-fidelity single-qubit operations compared to entangling operations, whereas the latter admits significantly less expensive implementations of Clifford operations over non-Clifford counterparts, e.g., single-qubit T gates. A powerful compiler would then enable an end user to develop quantum software at ease – hopefully in most cases independent of the backend details – while behind-the-scenes custom-optimizing for different backend quantum computer hardware, including the avail-




FIG. 5. Two circuits implementing the same sequence $R_{YY}(\theta_3) \cdot R_{XZ}(\theta_2) \cdot R_{XX}(\theta_1)$.

ability of fault tolerance.

d. Approximate compilation Oftentimes, the freedom to implement target unitaries only *approximately* is a powerful tool that can lead to significant cost reductions. For example, this technique has enabled improvements in quantum volume [302], qudit SWAP gates [303], bosonic interactions [131, 133, 277], and even the leading proposal for Shor’s Algorithm [304], in which exact addition is replaced by approximate addition. Indeed, the widely-used quantum Fourier transform is also implemented approximately in practice [146, 305], and in most quantum simulation algorithms an integral component is implementing a time evolution operator approximately. Beyond approximating unitaries, approximate compilation has also been employed in conjunction with knowledge of the input state, where one quantum circuit may be replaced with another as long as their actions on the input state is similar (even if the actions on some other states might be different). Given the important role that approximate methods play in classical computing – consider for example numerical integrators or polynomial-time approximation scheme for NP-hard or #P-complete problems – fruitful research outcomes that leverage approximate compilation, broadly defined, are expected. Adding to the flavors of approximate compilation research works mentioned earlier, possible, non-limiting research avenues may include (c.f. [306–309] and references therein) continuing to quantify existing approximate methods in classical computing for a speed up (e.g., Grover-like or potentially more) or finding entirely novel approximation schemes for those problems known to be hard to solve for quantum computers (e.g., QMA-complete).

G. Real-time classical processing of quantum information (within coherence time)

In this model of hybrid computation, the classical part is embedded or coupled with the quantum program and runs simultaneously with the quantum operations within the coherence time. This can significantly reduce the data-movement between the classical and quantum processors while providing the desired flexibility and short timing in controlling the quantum state. With mid-circuit measurement [310, 311], classical computation can use previous quantum operation outcomes to profile the

quantum execution status, predict and even adjust the quantum operations for future iterations or steps. Some examples are quantum error correction (QEC) [312], progressive projection filtering [313, 314] and random walk phase estimation [315]. The key point is that the classical computation has to be performed while the quantum state is still coherent. Given the short coherence time for near-term devices, constraints on the latency of the classical computation can be stringent, possibly demanding a strongly coupled classical hardware accelerator – e.g., field programmable gate array (FPGA), graphics processing unit (GPU), or application-specific integrated circuit (ASIC) – physically in place where classical control is happening.

H. Error mitigation

The output of quantum devices is inevitably altered by noise throughout a computation. While QEC can theoretically be used to remove error from any computation once error rates of the physical devices are pushed below the fault-tolerance threshold, in the near term various techniques can be used to reduce the effect of noise on computation, generally referred to as quantum error mitigation (QEM) [43, 192, 316, 317]. When a noisy quantum device is used to compute expectation values of quantum observables, this in general results in a biased estimator. The goal of QEM is to minimize the bias and variance in expectation values from several different runs of the noisy quantum circuit using classical post-processing.

One method that can in principle recover the ideal noiseless expectation values is probabilistic error cancellation (PEC) [318]. PEC is based on the inversion of a learned noise model of the device through sampling from a distribution of noisy circuits related to the noise model. One can prove that the bias of the resulting estimator vanishes as the quality of the learned noise model improves (provided the noise does not change between the noise learning and the actual experiment), but the sampling overhead is exponential in the size of the circuit, with the base of the exponential being related to the noise rate. While PEC becomes challenging for larger experiments, various techniques can mitigate its cost such as the use of tensor-network-based post-processing [184] instead of the circuit sampling, which reduces the sampling overhead quadratically, at the cost of more classical

computational power. Moreover, the sampling overhead can be contained by improving the accuracy in estimating the noise factors.

An alternative approach to QEM is called Zero-Noise Extrapolation (ZNE). ZNE is based on the collection of noisy expectation values for various amplified values of the noise parameter. The expectation value at the zero-noise level is then extrapolated using a linear, polynomial, or exponential fit. ZNE can be implemented using analog (relying on pulse stretching [319]) or digital amplification methods (such as sub-circuit repetition [320–322]) that do not require a precise characterization of the noise, or using probabilistic error amplification (PEA) [323], which achieves the desired noise factor by randomly drawing circuits from a distribution implementing a rescaled Pauli noise model learned from the device. ZNE alleviates sampling overhead of PEC, however does not guarantee an unbiased estimator. Its success depends on a proper choice of the fitting method (methods for systematically choosing the most appropriate one have been proposed, e.g., [322]), amplification factors for the noise (e.g., too large amplification will destroy the signal, too small amplification will make the extrapolation less precise), as well as the amplification method itself. Recently, ZNE has also been extended to quantum error correction [324].

One more technique to boost quantum fidelity through post-processing is Quancorde [325], which uses Clifford “canary circuits” (which are classically simulable but also resemble the target application structure and thus suffer similar structural noise impact) to order an ensemble of devices or qubits/mappings approximately along the direction of increasing fidelity of the target application. One then estimates the correlation of measurement outcome probabilities with this ordering and uses this correlation to weight the noisy probability distribution; correct measurement outcomes are expected to have a higher correlation with the ensemble order, and thus their probabilities are boosted, while those of incorrect outcomes are suppressed.

QEM has become an integral part of quantum simulations with near-term devices and, together with the constant improvement of the hardware, has allowed for impressive experiments reaching the boundary of what is computable using brute-force classical calculations [323]. It will enable us to push the capabilities of the upcoming generations of hardware up to the transition to fault-tolerance, where it can still complement QEC for various use cases. One potential avenue towards increasing the effectiveness of QEM is to develop either applications that are well-adapted to certain types of error mitigation, or mitigation techniques that exploit information about the specific application.

IV. WORKLOAD MANAGEMENT

A. Overview

We cover here the different aspects of the management of heterogeneous workloads involving quantum and high-performance classical resources. Whereas both technologies have an extensive set of functional tools for resource access and management, scheduling, and data handling, there is still need for substantial work towards extending the elasticity of these tools to operate in a hybrid quantum-classical environment and to cater for the needs of the different user profiles.

We consider and discuss the specificities of different access models, with classical and quantum resources operating in co-located environments or remotely, and explore the impact that different integration timescales, given the disparity in speed and bandwidth between quantum and classical, have on such operational models.

Finally, we describe the current state of the middleware layer for quantum and classical integrations across many different vendors in the industry and give some brief overview of their challenges and involved complexities.

Computational workloads can be categorized broadly into two classes: (1) trivial (or “embarrassingly parallel”) tasks, where each task executes independently of all others using its own set of dedicated computational resources; and (2) non-trivial interoperating sets of tasks, potentially sharing an allocated set of computational resources, where the outcome of a subset of tasks can affect the behavior of another. For HPC systems to handle the various requirements of such workloads, sophisticated workload management systems (WMS) and job schedulers enable efficient resource allocation, initiate and monitor workload execution, and govern a queue of pending work requesting resources.

In QCSC systems, computational workloads can also be defined based on Quantum-HPC integration types, such as HPC-for-Quantum, Quantum-about-HPC, Quantum-in-HPC scenarios, involving interplay between quantum and classical tasks. Consistency in Quantum-x-HPC or HPC-x-Quantum seems desirable and possible. Quantum-about-HPC workload involves quantum tasks that require HPC resources for pre- and post-processing. In this integration type, the quantum tasks are independent of the HPC tasks, and the HPC resources are used to support the quantum tasks. Quantum-in-HPC workload involves HPC tasks that require quantum resources for acceleration, where the quantum tasks are tightly coupled with the HPC tasks, and the HPC resources are used to support the quantum tasks.

Finally, HPC-for-Quantum workload involves tight integration with quantum resources in real-time, within the coherence time of the quantum system III G. The classical resources are used to support the quantum tasks, and the classical tasks are tightly coupled with the quantum tasks [326]. Circuits with classical control often referred

as *dynamic circuits*.

As discussed above, integrating quantum systems into classical HPC facilities will add complexity to managing and scheduling quantum-based workloads due to the disparity between quantum and classical processing times, sometimes requiring dynamically allocating and deallocating resources within and among jobs.

In this section, we discuss in detail challenges related to classical-quantum workload management, job scheduling, queuing, and operation modes of these jobs. We review these concepts and current solutions in the classical setting, and give recommendations for on-premises and cloud-based integration of quantum computational resources with HPC systems.

We differentiate between (i) on-premises integration, which involves directly incorporating quantum computational resources within a local HPC system infrastructure, and (ii) cloud-based integration, which entails accessing quantum computational resources remotely through cloud services. Scenario (i) suits workloads requiring tight integration between quantum and classical resources (e.g., Quantum-in-HPC and HPC-for-Quantum workloads). Further, it is ideal for sensitive or proprietary workloads where data security and privacy are paramount, as the data does not leave the premises. Scenario (ii) is beneficial for workloads that can tolerate latency, e.g., Quantum-about-HPC workflows that offload encapsulated parts to a cloud quantum system.

B. Resource management

This section is structured into two main parts: exploring the resource model, which deals with the allocation and architecture of quantum computing resources, and job scheduling, which is essential for optimizing the performance and efficiency of quantum computing tasks.

1. Resource model

A resource model refers to the framework or method used to manage and allocate computational resources. Specifically, in this section, we explore the distinction and relationship between QPUs and qubits, the placement and location of quantum and classical resources and ownership models.

a. QPU vs qubits The introduction of QPUs marks a significant departure from the conventional approach to traditional HPC scaling. Instead of pursuing the conventional path of continually improving semiconductor processors for greater speed and density, or incorporating specialized processors like GPUs, and optimizing their coordinated use through parallelization, QPUs take a fundamentally different approach. While one might assume that QPUs merely serve as an additional component, they fundamentally alter the way data

is represented, employing the language of quantum circuits governed by the principles of quantum mechanics. This constitutes a distinct computational model, diverging from all that has come before it. However, it’s essential to note that QPUs are not designed for general-purpose computing tasks. Consequently, practical implementations necessitate the careful orchestration of both classical and quantum computation resources.

Of particular importance for this type of scheduling task is the fact that QPUs operate at very different time scales and bandwidths from standard HPC processes. One could argue that the quantum circuit represents the essence of QPUs as co-computing units with very specific tasks in a supercomputing environment. As the fundamental unit of quantum information, the quantum circuit is a computational routine consisting of coherent quantum operations, including unitary transformations and projective measurements, and concurrent classical computation. As such computational units, quantum circuits may require executions that span inordinate amounts of time for a single computation. For example, consider the fastest quantum hardware that exists today at scale, superconducting qubits, tackling the task of running a circuit with 100 qubits and 10^8 T gates. Using the surface code, and assuming a physical error rate of 10^{-4} and a syndrome cycle time of $1 \mu s$, a single execution of that circuit, considering both magic state distillation and consumption, can take on the order of tens of minutes and up to several hours, depending on the number of state distillation factories (and therefore physical qubits) used [327]. This circuit execution cannot be preempted, which is important for scheduling purposes.

The above scenario represents an ideal, still far from what quantum hardware can offer today. Currently, QPUs don’t run in fault tolerant ways, can only tackle much smaller circuits, and must leverage error mitigation solutions that result in large sampling overheads. In contrast with the example above, QPUs today are extremely non-fungible, each of them with its own noise profile, and this affects user preferences and scheduling complexity.

The introduction of this new type of resource opens a range of models of operation:

QPU-based Model: This model resembles the current GPU model, where the entire quantum device is treated as a unified resource unit. Due to its natural inflexibility, this approach has the downside of potential device underuse when jobs only target a subset of the QPU.

Qubit-Based Model: In this approach, the qubits themselves are treated as resources, allowing a QPU to be fully occupied simultaneously by a variety of jobs. In near-term devices, with qubits exhibiting different types and amount of noise, and in the presence of limited device connectivity, this approach has shortcomings, as there would be competition among jobs for the best or more suited resources within the QPU.

Shares Model: This introduces a pseudo-unit of resource definition where each “share” represents a part of the device, a specific period of time, or the circuit

size. Essentially, each share signifies a portion of the computational capacity of the resource. Can be treated as sub-category or QPU-based model [328].

b. Quantum resource location. In the landscape of HPC, the location of computing resources determines the efficiency, performance, and accessibility of computational tasks. Two primary modes of resource location, co-located and remote resources, offer distinct advantages and considerations for HPC facilities and users.

Co-located resources refer to computational resources physically situated within the same facility or data center as the user or workload management system. There are several advantages of having co-located quantum and classical resources. Minimal network latency and high-speed data transfer capabilities, which is especially advantageous for applications where access to large amounts of (low-level) data (like counts) is required. Full control and management of co-located resources, enabling customization and configuration based on specific requirements. Data security. Co-located resources location implements on-premises integration type (i) IV A.

Remote quantum resources, on the other hand, are located at a physical distance from HPC classical infrastructure, often in off-site data centers, cloud environments (cloud model integration type (ii) IV A), or partner organizations. Remote resources extend the reach and capabilities of HPC systems, but they come with unique characteristics. Elasticity is one of them, allowing for scaling up and down quantum resources based on need. Another characteristic is sharing. At this moment, it is a case when a quantum processor is shared among multiple users, which drives down the cost of execution, but also introduces variation in performance due to resource contention. Network overhead can impact data transfer rates and communication latency, especially in applications with large loads.

c. Quantum resource ownership. High-performance computing facilities generally serve large numbers of users with varying computational resource requirements for their computational workloads (jobs) III. Users are part of user groups and projects, which have assigned queues. A project has some resource allocation quota (e.g., 10^6 node hours), whereas queues are defined by criteria such as priority and the amount of computational resources available per job (e.g., n CPU cores, m GPUs) to a user of that queue. At the level of user groups and projects, the integration of quantum information processing devices with classical HPC computational resources is straightforward: access to a quantum resource could be granted to a given user assigned to a `qpu.user` group with project quota Q specifying the number of “quantum node” hours.

2. Job Scheduling

When introducing quantum resources, one could envisage *real-time* quantum nodes, where classical compu-

tations are performed within the coherence time of the qubits (e.g., for error correction studies), and *near-time* quantum nodes, where higher latency classical communication suffices (e.g., variational algorithms).

Common frameworks in use today within HPC facilities and cloud platforms include SLURM [329], TORQUE [330], Altair Grid Engine [331], and IBM Platform LSF [332], which provide highly scalable control over heterogeneous computer clusters and jobs submitted to them. To execute a workload, a user configures their job using a *submission script* based on the syntax of the associated scheduler. Submission scripts essentially define a job: project name, resource requirements (e.g., node count, CPU count, memory size, and wall clock limit), tasks to execute, input and output data paths, etc. Note the governance of an HPC system (see Sec. IV D) could require all input data be readily available, e.g., on a scratch area accessible only by the user. These scripts are submitted from a login (head) node to the WMS or job scheduler, where they are queued and finally executed on one or more compute nodes.

A queued user job is executed in the order of the project’s or queue’s priority as well as resource availability; for instance, a higher-priority job whose resource requirements are met will be executed before a lower-priority job. Resources can be exclusive or shared, so if `job_x` does not need all cores on the node(s) to which it is assigned, `job_y` could be assigned to the remaining cores. However, resources allocated to a given job are fixed and, therefore, remain allocated for the duration of the job’s lifetime (and thus charging against the user/project quota), regardless of whether or not they are performing useful work.

Extensions to existing job schedulers could include quantum-specific configuration settings, e.g., number of circuits to be executed and depths, from which, along with architecture-dependent factors such as data loading, gate application, and readout times, runtime estimates can be made. However, while separating classical and quantum configuration spaces may be amenable to near-term applications (i.e., wherein classical and quantum operate cooperatively but have loose latency requirements, as in VQE), mid-term quantum-classical applications requiring lower latencies likely will necessitate tighter integration between the two types of resources.

Additionally, quantum-centric supercomputing workload scheduling must also consider the coupling between the quantum and classical tasks and assign and co-locate tasks to the resources accordingly. Different quantum-classical coupling types represent workloads with unique requirements and challenges, such as the lack of a unified standard for accessing QPUs and expressing hybrid workloads, and the need for parallelization of quantum workloads across nodes, cores, and accelerators. For example, classical tasks tightly coupled to quantum tasks, e.g., for the HPC-for-Quantum integration type, must be co-allocated and assigned to nearby QPUs and GPUs. Additionally, scheduling of QPU jobs must be prioritized

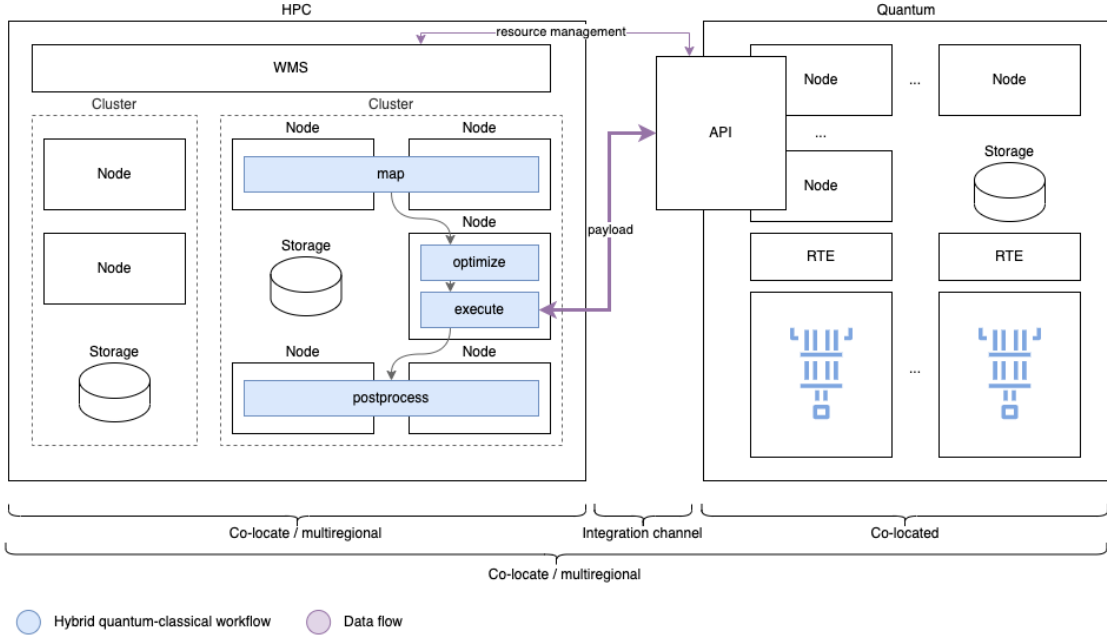


FIG. 6. Quantum-centric supercomputing integration overview. Integration channel can be tight or loosely coupled, which will affect throughput and latency of transmitted payload. Real-time compute is possible on co-located quantum side; near-time compute can be executed on quantum and HPC sides; long-time compute should happen on HPC side. Depending on implementation WMS can schedule tasks on quantum side through API or direct access to classical nodes of quantum side.

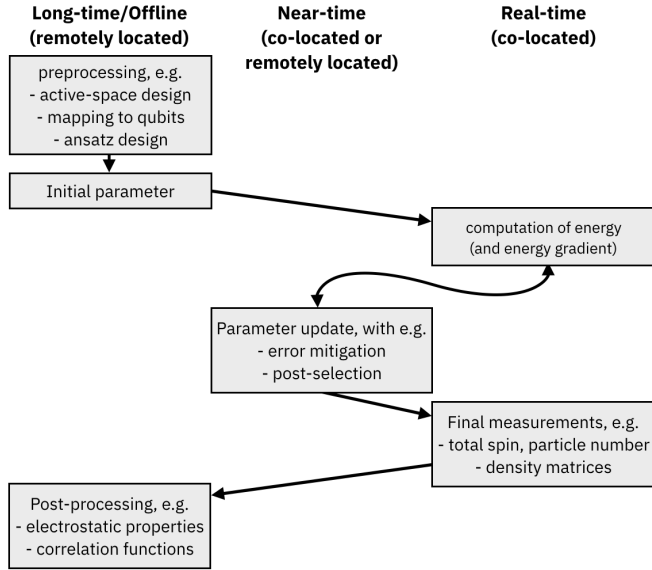


FIG. 7. Integration between classical (HPC) and quantum computing resources exemplified by the variational quantum eigensolver. The steps of the calculation are represented by gray blocks, connected by arrows describing the flow of operations and arranged left/center/right for operation that require “long-time/near-time/real-time” interaction between HPC and quantum computers (see also main text).

to ensure the optimal utilization of the QPUs while minimizing the time-to-solution and energy-to-solution [326] (section IV B 2 a).

a. *Queueing* HPC environments typically offer shared resources that are made available to a large user base. To ensure the equitable and efficient distribution of computing resources, HPC systems implement queuing mechanisms.

One of the aspects of execution management in HPC is the presence of multiple queues, each configured with specific priorities and constraints. These queues are designed to categorize and prioritize computational tasks based on the specific needs and requirements of the users or projects.

Queues can have resource constraints. Resource constraints may include limitations on the number of nodes, processors, memory, or specific software specifications available for jobs in each queue. These constraints help balance the allocation of resources across different user groups and projects while ensuring that the computing environment remains stable and efficient. As QPUs are “rare” resources, it is beneficial to increase the priority of the “quantum” queue while maintaining constraints on execution time and classical resources. It might be also beneficial to increase the priority of classical jobs connected to the quantum queue.

Because QPUs are a new and currently limited resource type, it is beneficial to specify a separate queue for quantum jobs.

Furthermore, as quantum computing technology evolves and the number of available QPUs increases, HPC administrators will need to adapt the queuing system to accommodate this growth. This might involve not only adjusting the priority of the quantum queue but also dynamically redistributing workloads to make optimal use of the expanding quantum infrastructure. Balancing the allocation of quantum and classical resources will remain a complex task, requiring constant monitoring and fine-tuning to ensure efficient utilization of the entire HPC facility. In this context, the queuing system’s flexibility and adaptability will be crucial in meeting the evolving demands of quantum computing within the HPC environment.

Cloud queues. There is a practice nowadays of cloud vendors are providing access to quantum services that are executing quantum payload. In order to provide fair allocation of resources, cloud providers are implementing their own queuing mechanisms behind cloud endpoints. It is important to understand that cloud vendors have different job queuing and priority implementations, and all might employ a shared resource ownership model IV B 1 c.

Qiskit Runtime service fair share queue. Access to quantum resources are divided between hubs, groups and projects. A *Hub* represents the top level of an organization such as an academic, industry, or research partner. A *Group* represents a mid-level structure to which access shares can be allocated by the hub for one or more collections of users (projects). A *Project* represents the base-level construct to which shares are allocated from the overarching group, and to which users are directly assigned.

For each group and project, the duration of the scheduling period is used to convert shares into an equivalent amount of time that an instance would receive under ideal conditions. The ratio between the time used and the shares equivalent time is used as the basis for scheduling jobs.

Many HPC job scheduling systems employ a so-called “fair-share algorithm” designed to distribute computational resources among multiple jobs in a manner that is considered fair. The definition of “fair” can vary significantly depending on the context and the specific requirements of the system. The Qiskit Runtime fair-share algorithm takes into consideration how shares of resources are distributed across groups and projects to determine job prioritization [333].

b. Cloud-bursting model Similar to a classical cloud bursting [334], we propose a Quantum cloud bursting model in which an HPC facility offloads quantum computation to cloud vendors when there are insufficient quantum resources within its local infrastructure.

As with any approach, there are pros and cons of quantum cloud bursting. Quantum resources may be expensive to maintain locally. In such cases, offloading workloads to cloud vendors with readily available QPUs can significantly decrease maintenance. Certain applications

require proprietary or specialized software. Acquiring and managing licenses for such software can be both costly and time-consuming. Cloud bursting allows organizations to leverage cloud vendors’ pre-configured environments, providing access to the required software without local setup and licensing. It also presents challenges related to data transfer and increased requirements for security.

Figure 6 is showing how quantum cloud bursting model can be implemented via integration of Quantum cloud API and workload management system. Table I shows a minimal API specification that cloud vendors should implement to unlock better integration of quantum and classical resources.

An illustration is given in Figure 7, using the variational quantum eigensolver as an example. Classical preprocessing (including mapping to qubits, choice of the variational ansatz, and parameter initialization) is carried out on HPC resources located remotely from the quantum computer (“long-time” interaction). The execution of quantum circuits may require the use of real-time communication between classical and quantum units through real-time electronics of RTE (see also Figure 6). Real-time interaction is required e.g. to implement dynamical circuits used for symmetry verification/enforcing, monitored quantum dynamics, and quantum error correction. The update of variational parameters, which requires the calculation of energies and energy gradients typically in conjunction with readout/gate error mitigation and post-selection operations, may be carried out on HPC resources located remotely from the quantum computer or co-located with it (“near-time” interaction).

c. Quantum execution runtime prediction Quantum execution runtime, determined by circuit complexity (number and type of gates, layers), can be predicted using historical data [335]. Despite potential overhead from error mitigation and circuit transpilation, accurate predictions are achievable. However, dynamic circuits with their non-deterministic workload structures pose challenges.

Runtime prediction is crucial for effective workload management, allowing quantum resources to be optimally utilized. Incorporating runtime predictions into scheduling enhances priorities and overall workload efficiency, showcasing the valuable role of prediction in quantum computing.

Orchestrating workflows that involve the integration of both HPC and quantum systems presents a set of unique challenges, particularly in the design and management of the payload, which encapsulates the instructions and data for these complex tasks. One of the primary challenges lies in achieving interoperability between classical and quantum computing environments. Quantum algorithms and classical algorithms have different computing models and requirements. Ensuring that the payload can effectively bridge this gap, translating between classical and quantum instructions and data formats, is a

Endpoint	Request Body	Response	Description
<i>GET allocations/</i>	allocation_ids: list(str) fields: list(str)	allocations: list(dict)	Fetch resource allocations listed by id.
<i>POST allocations/</i>	specification: dict	allocation_id: str	This allocates a set of resources from a resource request specification.
<i>DEL allocations/</i>	allocation_id: str	status: bool	Destroy a resource allocation and return resources to the pool of available resources.
<i>GET resources/</i>	specification: dict fields: list(str)	resources: list(dict)	Fetch the current state of resources on a system. This should report status of anything that a user may be able to specify for their job placement, including QPU status and software stack status
<i>GET tasks/</i>	task_ids: list(str) fields: list(str)	tasks: list(dict)	If a list of tasks are passed in, provide information on currently running tasks on the system. If no tasks passed in, provide information on all running tasks. This may constrain the results for the executing user so that it shows only their own jobs if the user is otherwise unprivileged.
<i>POST tasks/</i>	task_information: dict allocation_id: str batch_id: str	allocation_id: str	Launch a task.
<i>DEL tasks/</i>	task_ids: list(str) signal: int	status: bool	Send a signal to tasks for termination.

TABLE I. Minimal quantum cloud bursting API specification.

non-trivial task. This translation layer must not only accommodate the distinct computational paradigms but also manage the conversion of data types and formats, all while maintaining computational efficiency.

Optimizing resource allocation is another significant challenge in the payload for orchestrating quantum-centric supercomputing workflows. Balancing the computational load between HPC and quantum systems and efficiently using both types of resources is a complex task. It involves determining when to offload tasks to the quantum processor and when to keep them on classical hardware based on factors such as quantum system availability, workload characteristics, and quantum hardware constraints. Additionally, managing the data movement between these systems while minimizing latency and maximizing throughput adds further complexity to the payload design. To address these challenges, advanced workload scheduling and resource management strategies, as this paper mentioned, are required, leveraging real-time monitoring and adaptive decision-making to ensure that HPC and quantum systems work in harmony to deliver optimal performance for complex scientific and computational tasks.

C. Middleware

Quantum computing is evolving towards modular architectures comprising multiple QPUs coupled to classical computing nodes. Moreover, emerging applications that aim to benefit from quantum acceleration involve significant classical and quantum computational components, such as different phases of data collection and streaming, parameter tuning, simulation, and analysis. Typically, such quantum and classical coupling depend upon interaction between application components within and outside the coherence time of the quantum system, i.e., on tight- and medium-coupling across multiple quantum computing units (QPUs) or loose-coupling in a workflow application [326] (see section IV A). Hence, there is a need to design middleware systems that can facilitate the efficient understanding and interplay between quantum and classical components in an end-to-end workflow. Regarding quantum-centric supercomputing systems, middleware should also leverage well-established high-performance computing abstractions and must be compatible with existing HPC software stacks for managing hybrid workloads, tasks, and resources to integrate quantum computing into HPC systems seamlessly.

This section initially describes existing hybrid plat-

forms, runtime management, and quantum workflow frameworks instrumental in realizing quantum-centric supercomputing middleware systems. Further, we envision a conceptual middleware built upon the quantum-centric supercomputing integration limitations of existing middleware systems.

a. Existing middleware: Recently, several tools and frameworks emerged to manage quantum and classical tasks and resources efficiently. XACC is a system-level approach to integration of quantum and conventional processing, whereby quantum kernels are offloaded using hardware-agnostic constructs [336]. QCOR is a high-level language specification and associated tooling for hybrid programming [337, 338]. NVIDIA developed the CUDA Quantum [339] to integrate classical and quantum computing devices. It supports programming hybrid quantum-classical applications with optimized control and communication between quantum processors and classical tasks. It is also integrated with CUDA libraries for accelerating and scaling quantum simulations and classical HPC computations across distributed multi-node, and multi-GPU architectures.

There also exists hybrid quantum-classical runtime systems, such as Qiskit runtime [340] and Braket Jobs [341], that can be integrated into quantum software frameworks (e.g., PennyLane [342]) and platforms. The Qiskit runtime system provides primitives for defining, scheduling, and optimizing near-time quantum-classical workloads and enables their execution both synchronously and asynchronously. On the contrary, Amazon Braket Jobs are limited to a proprietary cloud environment and execute quantum-classical tasks as braket jobs with on-demand priority access to QPUs on a variety of quantum hardware (e.g., IonQ, Rigetti, etc.). There are also emerging middleware software stacks; for instance, QCG and QCG-PilotJob deliver highly efficient services and access tools for remote job management in large-scale computing environments, including HPC and quantum environments, by adding and extending capabilities to existing queuing systems (e.g. SLURM) [343] [344].

To support end-to-end quantum workflows and the interplay between quantum and classical components, several workflow orchestration frameworks exist, such as Orquestra [345], Covalent [346] and Tierkreis [347]. Orquestra and Covalent frameworks provide a quantum development and execution environment for quantum and quantum-inspired workflows with support for deployment, scaling, and parallelizing workflows on classical and quantum processors. Tierkreis instead utilizes a data flow-based programming model to orchestrate hybrid quantum workflows and provides a runtime environment allowing for concurrent and asynchronous execution. Recently, another middleware tool called IBM Quantum Serverless [348] emerged, which combined with Qiskit Patterns, a mechanism to build quantum workflows at scale, allow users to execute hybrid workloads on cloud or on-premise infrastructure.

b. Conceptual gap and vision: The intersection of quantum and HPC presents unique resource management and scheduling challenges, heterogeneity in handling different quantum hardware types, and thus, a high application development complexity. Effective resource management in a quantum-centric supercomputing middleware system requires sophisticated resource allocation mechanisms and scheduling algorithms that can dynamically allocate quantum and classical resources and manage complex workloads. This allocation must consider quantum processors’ distinctive computational capabilities and limitations. Different quantum technologies, such as superconducting qubits or trapped ions, have varied characteristics and requirements. Thus, the complexity of developing applications for such systems is substantial.

Despite significant advancements in integrating quantum and classical computing, current tools and frameworks have limitations. One major constraint is their nature as point solutions, often tailored to specific quantum providers or hardware. This specificity can limit their applicability and flexibility, making it challenging to adapt these solutions to different quantum computing environments or hardware platforms. Furthermore, many of these tools focus on narrow types of workloads. For instance, some are optimized for near-time variational algorithms. Others are designed for workflows, which, while effective in those contexts, may not efficiently handle real-time and near-time applications. As a result, tools and approaches are fragmented, posing challenges for developers seeking a unified, versatile platform for diverse quantum-centric supercomputing tasks.

D. Governance

1. Security

Similar to classical infrastructure, some of the security topics that must be addressed when developing hybrid quantum classical environments:

- Authentication and authorization mechanisms to restrict access to authorized users or entities.
- Fine-Grained permission to ensure that users can only access the resources they are entitled to.
- Data encryption protocols for both data in transit and data at rest. Explore the use of post-quantum cryptography to secure data.
- Secure data transfer between classical and quantum components in a workflow, preventing data interception or tampering during the transition.
- Access monitoring to detect and respond to unauthorized access attempts or unusual behaviors.

- Auditing and logging to track job submissions, resource allocations, and user interactions, aiding in post-incident analysis and compliance.
- Data privacy of user data and compliance with data protection regulations such as GDPR [349] or HIPAA [350], especially when dealing with sensitive data in hybrid computations.
- Threat detection to proactively identify and respond to potential security incidents.
- Resource isolation: In the event of a security breach, ensure that affected resources are isolated and investigated to prevent further damage or unauthorized access.

2. Policies

Storage Policies: Policies will define a data lifecycle that incorporates measures for archival, backup, and the secure erasure of data to safeguard intellectual property and sensitive information.

Storage Quotas: Storage quotas are allocated to projects, reflecting the anticipated volume of data and the project’s duration. Requests for adjustments will be entertained to these quotas in response to significant changes in data needs, subject to a review process and availability of resources.

Data Lifecycle Management: Procedures for data storage, archival, backup, and deletion. Data retention periods will be determined based on the data’s significance, sensitivity, and ongoing relevance, ensuring essential data is preserved and non-critical data does not occupy valuable storage indefinitely.

Queue Policies: Ideally, a queue scheduling policy will efficiently prioritize jobs based on estimated runtime, research urgency, and users’ historical consumption of resources, while also incorporating backfilling, to minimize idle resources. Regular reviews of queue performance should be undertaken to validate and adjust our prioritization strategies.

Scheduled Access: An advance reservation system ideally would be available for projects with foreseeable intensive computing needs, ensuring the availability of critical HPC and quantum resources for time sensitive projects. Such reservation requests would need to provide technical justification and be subject to a merit review process.

Job Time Limits: Reasonable runtime limits must be set for various job types to ensure equitable system access, as well as providing predictability in scheduling for users waiting for system access.

Job Interruptions Handling: Users should be notified in advance of system maintenance where possible. The capacity to pause jobs during system maintenance or unexpected downtime, and resume jobs after would be ideal,

minimizing the impact on research continuity and computational efficiency. Users should be encouraged to implement best practices for checkpointing, and saving work.

Data Use Agreements (DUA): Users are expected to engage in data use agreements that explicitly outline conditions for access, use, and sharing of data.

Responsibilities of Data Users: Users are tasked with upholding data confidentiality and integrity, in compliance with security protocols.

Data Sharing and Collaboration: Establish conditions under which data can be shared with external parties, paying special attention to the transfer of data between classical and quantum computing frameworks.

Termination and Renewal: The terms for the termination and renewal of DUAs will be clearly communicated. In the event of termination, data will be managed according to pre-established procedures, ensuring the secure disposition or transfer of data, with particular care for sensitive quantum data.

Breach of Agreement: Outline the consequences for any breach of DUAs, recognizing the potential for increased severity in the context of a hybrid HPC-Quantum environment.

Learning and Resources: Given that both HPC and quantum computing require specialized knowledge and skills, providing a robust array of training programs and resources encourages both individual and cross-domain competency.

User Feedback: Intelligent incorporation of user feedback ensures effective alignment with user needs and expectations, and aiding in identifying bugs and issues. User feedback provides a window into how quantum and HPC resources are actually being utilized versus initial assumptions. User feedback can also identify an accessibility or security issues, and inform what specific training and educational resources are most needed. Given the pace of evolution of the field of quantum-centric supercomputing, frequent and timely inclusion of user feedback is essential in guiding future development. A structured approach to collect user feedback involves regular forums, surveys, and user support help tickets to ensure we capture the user experience comprehensively.

Compliance and Audit: Importance of maintaining rigorous compliance with regulatory standards and internal policies. Our protocols will undergo periodic reviews to align with legal and industry benchmarks. Regular internal audits will be conducted to ensure adherence to our established policies and to identify areas for improvement. These audits will also assess the precision of resource allocation and the efficacy of our data management practices.

V. PROGRAMMING MODELS

A. Introduction

A programming model describes and provides the mechanisms for instructing how a computer system is to perform a computation, based on the execution models of specific computing hardware and programming language. These can usually be separated because any underlying execution model for a given hardware platform may be made transparent in any given language through calls to functions through abstractions. For the purposes of this paper, we will adopt this common viewpoint and focus on language-independent programming models supporting developers in several roles working across a layered execution model. Likewise, we will assume independence of programming paradigm, allowing for its choice to be driven by the details of an individual project and its requirements.

A programming model for a heterogeneous computing system combining both classical and quantum elements can be described in terms of four layers, illustrated in Fig. 8, comprising an application’s execution and categories of developers who write code describing computational procedures across these layers.

Working up from the bottom, Layer-1, which we call the Hardware layer, is where hardware-dependent operations occur on specific devices. On the quantum side, the programming model at Layer-1 provides access to processors and other electronics, software, and firmware which directly control the quantum device, and on the classical side it provides processor-specific instructions according to its Instruction Set Architecture (ISA), itself an abstraction of the processor hardware.

Layer-2, which we call the Computation layer, is where low level computational flows and information are described. For quantum this means constructing and optimizing circuits which are then mapped onto specific quantum hardware by compilation (and transpilation), operating on quantum information and interfaces to classical information. On the classical side this means functions optimized for specific underlying processors including CPUs (of varying ISAs), GPUs via a manufacturer’s API, or other environments for specialized processors such as inference engines, processors aimed at AI problems, or custom FPGAs.

The Library or Function layer, Layer-3, is where computational capabilities in Layer-2 are abstracted into functions aimed at particular problem domains and the hardware-independent computations supporting them. These are used by application developers without necessarily being concerned with the details of the underlying execution model.

Finally, Layer-4, which we call the Application layer, is where domain-specific application developers construct programs and workflows that perform computations to solve problems relevant to domain experts, expressed in domain-relevant terms, without necessarily being con-




FIG. 8. Programming model layers. (L1) Hardware layer: electronics, firmware V E. (L2) Computational layer: circuits V C V F. (L3) Library layer: functions. (L4) Application layer: domain specific applications.

cerned with the detailed mechanisms of those computations.

B. Related Work

It’s useful and instructive to consider popular existing programming models and how they relate to this layering. Apache Spark [351] is a programming model for large-scale data processing whose base programming language is Java; SYCL [352], CUDA [353], and High Performance ParalleX (HPX) [354] are programming models developed in idiomatic C/C++, spanning Layer-2 and Layer-3.

There do exist domain-specific programming libraries such as OpenFermion [355], which facilitate the mapping of electron-structure problems to quantum circuits. While certainly useful, it is arguable whether this approach represents a complete programming model per se rather than a suite of helper routines for a specific workflow (module specification \rightarrow integral generation \rightarrow circuit mapping) at Layer-2.

SimuQ [356] defines a domain-specific language for modeling arbitrary Hamiltonians, compiles these to an abstract *analog* instruction set, then lowers that to a sequence of pulses rather than to a quantum circuit at Layer-2. Programming-model support for the direct expression of Hamiltonians may be beneficial in Layer-3 in defining functions for building materials-science applications. One challenge, however, is performance scalability of the software stack. Peng et al. [356] report SimuQ com-

pilation times that grow rapidly with the number of sites in the Hamiltonian as well as a number of Hamiltonians that fail to map to the underlying quantum computer due to hardware constraints.

Message Passing Interface (MPI) [357] and Open Multi-Processing (OpenMP) [358] are open-standard specifications that aim generally for the programming of parallelized computations on heterogeneous classical computing platforms, and are examples operations spanning Layers 2 and 3.

Programming models can be directives-based, decorated, or provide extensions to their associated programming languages. In each case, a compiler capable of translating code written using a given programming model is needed to target specific backend hardware at Layer-1. A programming model incorporating semantics and execution model for quantum computing systems will naturally extend existing models targeting classical hardware. Importantly, a quantum computing-enabled programming model should provide abstractions such that applications can be developed independently from the underlying computing model or hardware; that is, at Layer-4 they should be agnostic with respect to both classical and quantum hardware. In this way, application developers can write software that can be executed on arbitrary hardware using functions in libraries at Layer-3, with possible customizations in Layer-2 if needed, and a backend compiler to translate the higher-level code (including language extensions, etc.) to the target machine code.

An initial prototype of such a programming model for the development of quantum computing applications is XACC/qcor [359, 360], an extensible language and classical-quantum compilation framework for Rigetti, IBM, and D-Wave quantum systems. XACC follows the co-processor model, akin to SYCL and CUDA, treating the quantum system as an computational accelerator, similar how a general purpose graphics processing unit (GPGPU) computing is used. This work paved the way for NVIDIA’s open-source CUDA Quantum platform [361] which extends their existing CUDA framework, providing support for programming QPUs in addition to CPUs and CUDA GPUs.

C. Execution Model

An execution model describes the behavior of an application and is exposed via a programming model. The hardware, or device, architectures help to define the execution models based on their capabilities and functionalities, and the programming model presents an abstract implementation of the underlying capabilities to enable general (or expected) purpose computation. In classical computing, a well known classification of architectures includes those found in Flynn’s taxonomy [362]. In this section, we complement Flynn in the context of quantum computation and how this classification can be incor-

porated into a programming model for hybrid classical-quantum applications.

Flynn’s taxonomy is laid out in two dimensions: the number of instruction streams or programs and the number of data streams. The single instruction, single data (SISD) model represents classical sequential machines following the von Neumann architecture. In the immediate term of quantum computing, until QPU partitioning or parallel QPU processing is possible, we foresee that SISD will be the predominant architecture. Programming languages exposing this execution model include C, C++, and Fortran, which are dominant languages within HPC. Parallel quantum computations, wherein, e.g., a large QPU can be partitioned into topologically equivalent subprocessors, single instruction, multiple data (SIMD) will be beneficial; SIMD capabilities can then be exploited to parallelize quantum computations. A recent popular addition to the SIMD family is GPUs, which manipulate multiple data through the same set of instructions executed across many processors in lock-step, known as single instruction, multiple threads (SIMT) [363].

Multiple instruction, multiple data (MIMD) is generally accepted to represent parallelism. Again, supporting an architecture within a programming model depends on the underlying hardware capabilities which are exposed to the user. Modern languages such as Erlang [364], Go [365], or Python [366], and recent revisions of C++ [367] present built-in architecture-dependent concurrency models, and parallelism is most often exploited either via extensions to traditional HPC programming models such as CUDA [368], OpenMP [369], OpenMPI [370], or via careful allocation of resources and explicit control flow declaration known as “workload orchestration”. Ray [371] is an example of a parallel workload orchestrator.

In our layered vision of programming models, at Layer-1 (i.e., Hardware Layer), classical hardware shows instances of several kinds of execution models: instruction-level parallel (superscalar) processors, instances of MIMD and GPUs, address SIMD and SIMT; Tensor Processing Units, and Data Processing Units support unloading specific workloads from CPUs, ASICs are examples of programmable technologies for specialized SISD; and novel applications can be explored through the use of reconfigurable hardware such as FPGAs. On the quantum side, general-purpose quantum computers can essentially be considered SIMD machines with a computational memory in which each qubit has both independent and combined processing capabilities.

At Layer-2 (Computation Layer), the hardware architecture is abstracted into the programming model’s execution model. This, in turn, will be reflected by the programming language subject to compilation. Language-specific statements control the underlying hardware through a human-friendly syntax and developer ergonomics, presenting an abstract SISD model. It is through the careful coordination that access to paral-

lelism and other execution models is granted to higher layers. In quantum computing, although there exist programming languages capturing quantum semantics (e.g., OpenQASM 3 [372], Q# [373], and Quipper[374]), the control flow of quantum circuits, circa dynamic circuits, is simple enough to be captured as data and to be translated into hardware instructions mostly directly.

Layer-3 (Library Layer) takes advantage of Layer-2’s execution model libraries to access actual programming models beyond SISD variants. This layer increases the level of abstraction into libraries, frameworks, and software development kits (SDK), introducing concepts agnostic to the hardware execution model, and tools for interfacing with the abstract execution model. A classical example, in AI engineering, TensorFlow [375] exploits the concept of a graph, and implements a variant of the data-flow execution model, where operations are executed as soon as the input operands are ready. Tools such as TensorBoard [376] allow for the inspection and debugging of graphs. Similarly, on the quantum side, the Qiskit SDK [377] introduces concepts from quantum computing such as quantum registers, gates, and measurements, and captures the circuit execution model as a directed acyclic graph (DAG). The circuit and DAG visualizers [378] allow for the exploration of these quantum program representations.

Finally, Layer-4 makes use of these libraries, frameworks, and SDKs to implement domain-specific applications. Applications at this level act as coordinators, directing the workloads needed to solve specific problems. Within the quantum computing ecosystem, Quantum Serverless [291] can coordinate an execution workload through many clouds and quantum computing providers and thus, understand concepts coming from both the classical and quantum worlds. At this uppermost level, developers make use of the most suitable programming model based on trade-offs between speed, solution quality, and resource cost.

D. Quantum Kernels

In the context of classical computing, invocation of accelerator kernels can happen in several ways: (1) host (CPU) code; (2) strictly via other kernels (device-side, kernel launch); or (3) host or kernel code. Scheduling of host code and accelerator kernels similarly can be performed in three ways: (1) implicitly, if a programming model provides runtime mechanisms for managing a task DAG; (2) asynchronously, through non-blocking calls that return immediately after being called; or (3) explicitly, via synchronization points.

A programming model should enable quantum kernels (circuits, possibly with classical control flow) to execute in parallel, synchronously and asynchronously with respect to the kernels themselves and classical computations in order to cover a wide variety of use cases. This type of execution model is akin to that of classical accel-

erator (e.g., GPU) programming wherein asynchronous accelerated kernels can be desirable but data dependencies permit only synchronous execution. However, a major difference between purely classical and classical-quantum hybrid execution are the latencies involved; for instance, classical accelerators interface via high-speed buses whereas today’s free and publicly available quantum devices are generally cloud-based.

The rate at which quantum computations can be performed (i.e., the analog to classical CPU clock speed) is set by the qubit technology which range from $O(1)$ Hz (photonic qubits) to $O(10^6)$ Hz (superconducting qubits). Programming models must provide kernel scheduling mechanisms, or runtime, for orchestrating devices where latencies can range from $O(10^{-6})$ seconds to $O(1)$ minutes or greater. In cases where classical and quantum computations are decoupled — i.e., independent classical and quantum information processing with no communication between processes is needed — no orchestration would be required from the programming model. Indeed, classical and quantum routines or applications can be developed in separate languages and executed independently. In this section, we are concerned with ‘near-time’ hybrid computations wherein classical and quantum resources work in unison to complete a task.

Scheduling and execution of quantum kernels can be handled by a queuing system within the programming model. Similar to the SYCL programming model [352], a queue could have a one-to-one correspondence with a device; multiple queues would be necessary to work with multiple devices. Once a queue has associated to it a quantum device, that device can then be queried directly from the queue to obtain device-specific information, e.g., number of qubits and supported features. A quantum extension to the queue model would include sampling functionality to specify the number of executions (or shots) of a given kernel. Furthermore, a queue would capture the necessary semantics for execution, data input and output, etc., and provide within its API synchronous and asynchronous capabilities.

In principle, quantum accelerator kernels should be indistinguishable from classical accelerator kernels. That is, the semantics should be such that both classical and quantum kernel can expressed in the same way. For example, classical kernels generally define an “index space” over which an instance of the kernel executes. Each kernel instance executes the same code but may operate on different data or traverse a different execution pathway through its code. This mechanism, coupled to embarrassingly parallel workloads, is what empowers GPUs in today’s HPCs. For a quantum device kernel, the concept of an index space can be extended to include device (qubit) partitioning wherein, e.g., a single circuit instance can be mapped multiple times to the topology of a large device and hence execute in parallel.

In addition to partitioning, various quantum “back-ends” (i.e., control electronics, classical computational devices, etc.) will likely have different technical capabil-

ities; e.g., control-flow, dynamical decoupling schemes, etc. Such capabilities could be configurable so a developer may choose the schemes most appropriate for their kernel. Therefore, an API specification should permit extensibility for vendors to expose all desired functionality for their backend devices via a common plugin mechanism, allowing for interoperability between different vendor library implementations.

In HPC, vendor libraries are typically written in low-level systems languages (Layer-2) such as C and C++ which are the de facto for applications targeting a set of heterogeneous architectures on current HPC platforms. A programming model built atop these would require language extensions for allow the expressiveness necessary to write circuit applications. Language extensions also allow for a single-source programming model such that development would be streamlined and eliminate the interplay of multiple languages (and thus compilers, interpreters, etc.). Moreover, this makes classical control-flow within quantum kernels more natural for developers. An alternative to a single large-scale application is an admixture of distinct executables whose orchestration is handled by a workflow management system (WFMS), see [379]. The WFMS is effectively a wrapper around multiple executables — that may or may not have intercommunication — and steers their execution based on a predefined DAG or logic that enables dynamism of the overall flow.

E. Memory Model

Quantum processors operate at a very particular level in terms of memory and, in fact, a quantum processor can be regarded as a computational memory. In contrast with typical classical formalism where operations usually happen outside dedicated memory, quantum computing operates on the data input and the lifetime of quantum information is limited to the duration of the circuit carrying out the computation, with some notable exceptions like inherently quantum memories (QRAM and QROM). These quantum memories, hitherto never realized, would naturally span layers 1 and 2, and could play a fundamental role in classical data loading for some algorithms in which the implementation of the quantum oracle is critical to retain any level of quantum advantage [380]. Section ?? provides a thorough discussion on this subject.

On a more classical standpoint, and similar to the concept of shared memory in HPC architectures, one could envision a model whereby some memory is shared at the QPU controller level with other processing units (GPUs, CPUs) in order to exploit different aspects of the quantum pre- and post-processing tasks, as well as fast loading of quantum circuits or parts of circuits.

F. Data Types

The integration of classical computing resources with QPUs involves necessarily some common interface for data consumption that bridges largely different DSLs. A seamless workflow that exploits both CPUs/GPUs/xPUs and QPUs will include steps to transform data types as needed by each unit. It is worthwhile to differentiate between data interpreted as instructions and other type of data like metadata or data carrying computing information. Within the scope of programming models, the quantum nature of the information existing in the QPU during the computation is unimportant. Any interaction of the QPU with supporting or co-computing units before, during, or after the quantum computation occurs largely in the form of classical information. In fact, both the inputs and outputs for quantum processors are naturally classical except for a few notable exceptions (QRAM being one example, as discussed in Section III E). The natural outputs from QPUs are probability distributions, to which we have a level of access limited by our ability to sample them. Certain quantum algorithms like quantum phase estimation or Shor’s algorithm produce distributions so peaked that one or only a few samples are enough to provide the answer. Other algorithmic approaches based on heuristic methods, however, produce output distributions from which the value of certain quantum observable is computed by repeated sampling. This value can either represent the answer to the quantum subroutine, or it can be passed back to other -classical- computing resources to generate a new input to the QPU.

Inputs to quantum processors are, at the highest level of abstraction (Layer-2), a series of operations, that can be both unitary and non-unitary, and which can be represented by a quantum circuit. These quantum circuits can then take many levels of specificity throughout the transpilation and compilation chains, from the multilevel representation that a quantum assembly language like OpenQASM provides to the binaries used by the classical control hardware to implement the quantum gates on the quantum processor. These gates, or instructions, are parameterized to get instanced by data from the algorithm input. The gates are ultimately expressed as physical pulses of energy that interact with the quantum hardware (Layer-1). Thus, these pulses ultimately contain both instruction data and computational data. It is at higher levels of abstraction, however, where HPC could possibly help with circuit optimization, transpilation, and routing, working with circuit representations involving data types digestible by the DSL of the different processing units.

VI. CLASSICAL SIMULATION OF QUANTUM SYSTEMS

A. Why use classical simulators

For near-term quantum computers, one might question the necessity of classical simulators to emulate quantum computations. After all, the promise of quantum computing lies in its ability to perform certain tasks exponentially faster than classical computers. However, classical simulators remain invaluable for several compelling reasons [251, 252, 381–386].

1. **Resource Efficiency:** Quantum computer time is a limited and valuable resource. When developing and testing new quantum circuits or algorithms, researchers often need to execute them numerous times to validate their functionality and robustness. Running experiments on a real quantum machine can be time-consuming and costly, particularly when waiting for quantum computer availability after each circuit modification. Classical simulators offer an efficient alternative, allowing researchers to rapidly iterate through experiments without waiting for quantum resources.
2. **Noise and Error Analysis:** Real quantum machines are susceptible to noise and errors due to environmental factors, making it challenging to control and maintain the desired quantum state fidelity. Classical simulators provide a controlled environment for introducing and analyzing various noise scenarios. Researchers can simulate quantum noises to assess how robust their algorithms are under different conditions, as well as explore error correction techniques. This is a crucial step in building fault-tolerant quantum systems.
3. **Scalability:** Current quantum machines have limitations in terms of the number of qubits and the noise levels they exhibit. Classical simulators, on the other hand, can be adapted to simulate quantum systems with a larger number of qubits, enabling researchers to explore complex quantum algorithms and states that are beyond the capabilities of existing quantum hardware.
4. **Versatility:** Classical simulators offer flexibility regarding what they can simulate. Researchers can use them to capture "snapshots" of quantum states during computations, perform measurements, and make decisions based on measurement outcomes. Additionally, they can obtain more comprehensive information, such as the density operator of the system, rather than just specific measurement results.
5. **Efficiency Trade-Offs:** Building an efficient quantum simulator involves addressing trade-offs between computational efficiency and the range of

quantum states it can represent. The choice of simulator depends on the specific needs of the quantum computation being emulated.

6. **Connectivity:** When executing a circuit on quantum hardware, the transpilation step involves mapping logical qubits to physical ones, in a fashion that may be one-to-many. This is due to the constraint of the physical connection of qubits, which is not usually all-to-all on true QPUs. When simulating, all qubits can be entangled with all others, thus keeping down the 'physical' qubits needed and easing the computational burden of solving the minor-embedding problem.
7. **Pre-optimization:** Variational algorithms can be expensive and require many iterations of a quantum circuit. Pre-optimization with classical simulators can be used to find approximate circuit parameters that can be used or refined with quantum hardware. [251, 252, 385]
8. **HPC-assisted quantum computing:** There is a way in which HPC can substantially help to extend the depth of quantum circuits, as demonstrated in the recently developed method called "Operator Backpropagation." [386] The idea is to compute a part of a quantum circuit on a quantum device, and compute the second part of a circuit classically by using a quantum circuit simulator on an HPC system, and then stitch the results together.

In summary, classical simulators allow for efficient development, testing, and analysis of quantum circuits, algorithms, and states, while also providing a controlled environment for exploring quantum noise and errors. While quantum computers hold immense promise, classical simulators remain a vital component of the quantum computing toolkit, enabling progress and innovation in the field. Furthermore, simulators will likely continue to serve as yet another type of circuit execution environment among multi-node compute clusters for executing distributed workloads, which will be adept at handling (sub-)circuits that fall within their scope.

B. Quantum circuits simulators

In this section, we describe different types of classical quantum circuit simulators. In particular, we describe their capabilities in terms of what the maximum size systems (in both qubits and gates) that can be simulated using HPC. The advantages and disadvantages of quantum circuit simulators are also discussed, as well as use cases.

Before we start, we note that not every quantum circuit is difficult to simulate on a classical computer. Trivially, a quantum circuit that carries out a purely classical computation is not difficult to simulate on a classical

computer. Indeed, some of the quantum circuits that are not difficult to simulate include Clifford circuits [387–389] and the quantum Fourier transform [390, 391], to name a couple. Some quantum circuits are known to be classically simulable for specific input states as well [389]. If the simulation does not need to produce the full spectrum of the output state of a given quantum computation but rather a sampling of such, as it would be the case for obtaining a final, classical result out of a quantum computer upon measurements, the difficulty of the classical simulation can dramatically decrease [392]. Further decrease in the difficulty is obtained if the simulation is to mimic an error-prone quantum computer executing a quantum circuit [393]. We further note in passing that simulation of a quantum circuit that exhibits sparse couplings between densely connected sets of qubits can be more amenable to classical simulations [394], using methodologies not unlike circuit cutting discussed in Sec. III F 1.

1. State-vector simulators

State-vector based simulation is to simulate the operations of applying a series of unitary operators $U_{m-1} \cdots U_1 U_0$ to the state-vector representation of the quantum states $|\psi\rangle = \sum_{i=0}^{2^n-1} \alpha_i |i\rangle$, where n is the number of qubits and m is the number of operations or gates. Typically, a complex-valued double or single precision floating-point vector $\vec{\alpha}$ of size 2^n is used to store the coefficients α_i , which costs 16×2^n bytes of memory for classical simulation. U_i with $i \in [0, m-1]$ is a 2×2 (for one-qubit gate) or 4×4 (for two-qubit gate) complex matrix. It has been shown that an arbitrary quantum circuit can be decomposed into 1-qubit and 2-qubit gates [395]. In fact, most quantum devices internally run 1-qubit or 2-qubit basis gates. For example, IBM adopts 1-qubit gate X, SX, RZ, ID and 2-qubit gate CX (recently ECR) as the basis gates. To apply a gate U , the operation is $|\psi\rangle \rightarrow U|\psi\rangle$. For 1-qubit U applying on qubit q in a quantum register, $\vec{\alpha}$ is updated through the following expression where $s_i = \lfloor i/2^q \rfloor 2^{q+1} + (i \% 2^q)$ for every integer $i \in [0, 2^{n-1} - 1]$:

$$\begin{bmatrix} \alpha_{s_i} \\ \alpha_{s_i+2^q} \end{bmatrix} \rightarrow U_{2 \times 2} \cdot \begin{bmatrix} \alpha_{s_i} \\ \alpha_{s_i+2^q} \end{bmatrix} \quad (12)$$

Regarding 2-qubit unitary gate U applying on qubit p and q (assuming $p < q$ without losing generality), $\vec{\alpha}$ is updated through:

$$\begin{bmatrix} \alpha_{s_i} \\ \alpha_{s_i+2^p} \\ \alpha_{s_i+2^q} \\ \alpha_{s_i+2^p+2^q} \end{bmatrix} \rightarrow U_{4 \times 4} \cdot \begin{bmatrix} \alpha_{s_i} \\ \alpha_{s_i+2^p} \\ \alpha_{s_i+2^q} \\ \alpha_{s_i+2^p+2^q} \end{bmatrix} \quad (13)$$

where $s_i = \lfloor \lfloor i/2^p \rfloor / 2^{q-p-1} \rfloor 2^{q+1} + (\lfloor i/2^p \rfloor \% 2^{q-p-1}) 2^{p+1} + (i \% 2^p)$ for every integer $i \in [0, 2^{n-2} - 1]$.

Therefore, state-vector based quantum numerical simulation is to perform a sequence of 2×2 or 4×4 operations

Eq. (12) and Eq. (13) over the large state-vector coefficient array of complex numbers in a classical system.

State vector simulators hold the full coefficients of pure quantum states within the classical system's memory, which scales exponentially with the number of qubits. Consequently, the state vector is often distributed across many nodes of a classical HPC [396–399]. Such simulators however exhibit linear scaling with respect to the circuit depth [394]. As such, state vector simulations are sensitive to the number of qubits but much less to the gate count or circuit depth. For example, a recent work shows that a low-energy nuclear quantum circuit with 115 million gates on 21-qubits can be effectively simulated within one hour using a GPU for state vector simulation [400].

We note that various forms of approximate state-vector simulators can be tailored to problem specific applications. For example, recent truncated state vector simulators have been used to simulate VQE for various chemistry problems on 64 qubits without significant loss of accuracy or the need for large computing resources [252, 401]. Combining the state vector simulators with Feynman path summation can trade the circuit depth with the number of qubits or memory usage [394], rendering some hard to simulate circuits according to a naive state vector simulator to be simulable due to orders of magnitude improvement in their simulability [394, 402].

2. Density matrix simulators

When dealing with a mixed system where noise is present, a pure state can no longer provide sufficient information about the system. In this case, a mixed state corresponds to a statistical ensemble, or probabilistic mixture of pure states, can be used to describe the condition where a system is entangled with another system, such as the environment from which the noise is imposed. A mixed state is represented by a density matrix or a density operator, which is defined by choosing the basis in the underlying space. The density matrix is given by:

$$\rho = \sum p_s |\psi_s\rangle \langle \psi_s|$$

where ρ represents the fraction of the ensemble of each pure state and is generally an unknown real value. A density matrix contains all the information of a quantum system, allowing the calculation of the probabilities of the outcomes of any measurement performed.

Compared to state-vector simulation, the density operator requires the conservation of 4^n coefficients, where n is the number of eigenstates for each pure state, i.e., the number of qubits. Therefore, the memory cost of a density matrix simulation is 2^n times that of a pure state simulation using state-vector. The system also evolves according to the operator or gate sequence. When a par-

ticular gate is applied, the density matrix evolves as follows [403]:

$$\rho(t) = G(t)\rho(0)G(t)^\dagger \quad (14)$$

where $G(t)$ is the gate at time t , and $G(t)^\dagger$ is its adjoint. The evolution of a density matrix is more complicated than that of a state vector, considering the size of 4^n and the extra adjoint operator per gate. Although $G(t)$ by itself is a unitary operator, it becomes a general matrix and is not necessarily unitary in the presence of noise. Depending on the noise model, the evolution can become an ensemble of evolutions of corresponding channels, each described by a Kraus operator.

The noisy density matrix quantum numerical simulation is to compute ρ_{out} for a n -qubit quantum system, density matrix quantum circuit simulation is to compute ρ_{out} for a n -qubit quantum register, given initial state ρ_{in} and m non-unitary transformations G_0, G_1, \dots, G_{m-1} :

$$\rho_{\text{out}} = G_{m-1} \cdots (G_1(G_0\rho_{\text{in}}G_0^\dagger)G_1^\dagger) \cdots G_{m-1}^\dagger \quad (15)$$

where G and ρ are $2^n \times 2^n$ matrices. Due to noise, G is not necessarily a unitary matrix. G^\dagger is the adjoint of G verifying $G^\dagger = (G^*)^T$. As real quantum devices typically use 1-qubit or 2-qubit gates representing as 4×4 or 16×16 matrices, to obtain the matrix G_i with a $2^n \times 2^n$ size, Kronecker product or tensor product is used with the identity matrix I for the other qubits.

3. Tensor network simulators

Quantum tensor networks can be used for the simulation of quantum circuits [404–410]. These simulators leverage the mathematical framework of tensor networks, which are graphical representations of quantum states and operations, as networks of interconnected tensors. Specifically, the nodes in these networks correspond to tensors, which encode gate operations, while the edges represent indices along which tensors are contracted, reflecting the entanglement and interactions between qubits.

The strength of tensor network simulators lies in their ability to compactly represent quantum states that would otherwise require a prohibitive amount of memory. This is particularly valuable for simulating shallow-depth quantum circuits with a large number of qubits, which are challenging for traditional state-vector simulators. These simulators are still limited by the amount of available distributed memory on a supercomputer. For the current generation of supercomputers, tensor network simulators can run quantum circuits up to approximately 200 qubits without approximations for certain quantum circuits [406]. Larger simulations are possible with a truncated bond dimension as well as other applications with different resource requirements [251].

There are different types of tensor network simulators: Matrix Product States (MPS), Projected Entan-




FIG. 9. Different types of tensor network simulators: (a) Matrix Product States (MPS), (b) Tree Tensor Networks (TTN), (c) Projected Entangled Pair States (PEPS), and (d) Multi-scale Entanglement Renormalization Ansatz (MERA).

gled Pair (PEPS), Tree Tensor Networks (TTN), Multi-scale Entanglement Renormalization Ansatz (MERA), and Continuous Variable Tensor Network (CVTN) to name a few as shown on Figure 9. MPS is arguably one of the most popular simulators, and it is particularly useful for simulating one-dimensional quantum systems. In MPS, the tensors are arranged in a chain, with each tensor representing the state of a qubit and the connections between them capturing the nearest-neighbor entanglement. One of the MPS applications is an efficient tensor network simulation of lossy Gaussian boson sampling experiments, which has been recently demonstrated [411, 412].

Each type of tensor network has its specific applications, advantages, and limitations, and the choice of which to use depends on the application, size, and structure of a quantum circuit:

- MPS simulators excel in one-dimensional quantum systems with short-range interactions. They are efficient for simulating ground and excited states, as well as dynamical properties, due to their ability to capture entanglement in a scalable manner with limited entanglement entropy.
- PEPS are generalizations of MPS to higher dimensions, making them suitable for two-dimensional quantum systems. They are adept at handling both short- and long-range interactions, but their computational complexity increases significantly with the system’s size and entanglement.
- TTN simulators are structured in a hierarchical, tree-like manner, offering efficient computation for certain quantum circuits, especially those with hierarchical or layered structures. They are partic-

ularly useful for simulating states that exhibit hierarchical entanglement patterns and for providing insights into quantum many-body systems.

- MERA is designed for critical systems with long-range entanglement. It excels in representing ground states of quantum many-body systems near criticality, offering insights into scaling and universality in quantum phase transitions.
- CVTN simulators are tailored for quantum systems with continuous variables, like quantum fields or modes of light. They are adept at handling systems where particle number isn't conserved and are crucial in studying non-Gaussian states and processes in quantum optics and field theories.

Tensor network simulators continue to evolve, with ongoing research aimed at increasing their efficiency, scalability, and applicability to a broader range of quantum computing tasks.

4. Open system Lindblad quantum simulators

Open-system Lindblad quantum simulators are specifically designed to model the dynamics of quantum systems interacting with external environments. The applications include, for example, simulating quantum magnetism, topological materials, quantum phase transitions, and electron transportation. Unlike closed quantum systems, open systems are subject to environmental influences that lead to non-unitary processes such as decoherence and dissipation. These simulators leverage the Lindblad master equation, which blends the unitary evolution dictated by the system's Hamiltonian with the non-unitary aspects resulting from environmental interactions. This approach allows for a comprehensive simulation of open quantum systems, capturing the complexity of quantum noise and the environmental effects. Recently, a novel approach, called noisy quantum gates, has been proposed, as a classical simulation of the Lindblad dynamics [413]. It is based on integrating the noise into the gates, rather than keeping gates and noise as two separate dynamics, an approach can be generalized to non-Markovian dynamics by using colored noises.

In general, there are a few ways to formulate, but the most common is to use the density matrix formalism, which is adept at representing mixed states, a critical aspect when dealing with open systems. These simulators provide researchers with the flexibility to define specific system-environment interactions, making them a versatile tool across various quantum research domains, especially in material science applications. Its ability to accurately model quantum noise induced by the environment is an especially valuable aspect of it. Overall, the open-system Lindblad quantum simulators stand as a very valuable tool in quantum research, enabling a deeper understanding of the environmental impacts on quantum

systems and aiding in the advancement of practical quantum applications.

C. Overview of classical simulators

In summary, we described four main types of quantum circuit simulators. In the context of material science, the choice of a quantum circuit simulator depends on the specific characteristics of the system under study and the phenomena of interest. Here is how each type of simulator can be used:

State-vector simulators are ideal for systems that can be accurately represented by pure quantum states. In material science, they are particularly useful for studying the evolution of quantum states under Hamiltonians with relatively few degrees of freedom. They excel in scenarios where the full quantum state needs to be tracked, such as in the simulation of small, isolated quantum systems or systems with limited entanglement.

Density matrix simulators are well suited for studying systems where mixed states are prevalent, which includes most real-world scenarios in material science. They can handle decoherence and other noise effects, making them suitable for simulating open quantum systems or systems under non-ideal conditions. Density matrix simulators are ideal for investigating phenomena in quantum materials where environmental interactions play a significant role.

Different types of tensor network simulators have different use cases. MPS and TTN are powerful in simulating one-dimensional and certain hierarchical quantum systems, respectively. They can efficiently model systems with short-range interactions, which are common in material science. These simulators are especially useful for studying ground state properties and low-energy excitations in materials. PEPS and MERA are more suited for higher-dimensional systems. PEPS can handle both short- and long-range interactions in two-dimensional materials, making them valuable for exploring complex quantum materials, such as high-temperature superconductors. MERA is particularly effective in studying critical phenomena and phase transitions in materials.

Open system Lindblad quantum simulators are designed to handle non-unitary evolution, which is typical in open quantum systems, where the system is in contact with an external environment. In material science, they are crucial for studying decoherence, dissipation, and thermalization processes in materials. Lindblad simulators are particularly relevant for investigating quantum materials and devices operating under realistic, non-ideal conditions, where environmental interactions cannot be ignored. They are essential for understanding the behavior of materials in quantum information processing and quantum sensing applications, where control and mitigation of decoherence are critical.

Using HPC is critical for simulating large quantum circuits. Ultimately, supercomputers can simulate rela-

tively small quantum circuits because of the exponential requirement of available distributed memory. The next generation of quantum simulators will probably use small quantum computers to simulate very large quantum circuits. The idea is to use small quantum computers to perform tasks that are inherently quantum in nature, such as the contraction of intermediate multi-dimensional tensors, which are central to tensor network simulations. This idea is closely related to Feynman’s proposal to use quantum computers to simulate quantum systems [414]. One way is to use Harrow-Hassidim-Lloyd (HHL) algorithm for the contraction of very high-dimension tensors, which is currently a bottleneck of classical tensor network simulators. The HHL algorithm can be adapted to perform tensor contractions by encoding the tensors as linear systems. This could potentially revolutionize tensor network simulations by dramatically reducing the computational complexity of these contractions. This approach promises a scalable pathway for quantum simulation, as improvements in quantum hardware, such as increased qubit counts and enhanced coherence times, would directly translate into an increased capacity for simulating larger and more complex quantum circuits.

VII. USE CASE IDENTIFICATION

A. Criteria for use case identification

In this Section, we identify use cases relevant to real-world applications [415]. In such identification, we will rely on the following criteria:

1. **Potential impact:** selecting real-world use cases requires identifying research topics at the frontier of materials science and engineering that could benefit from the deployment of quantum computing for solving one or more of the discovery steps involved. Representative examples include semiconductor integration and processing, electronics, and optoelectronics applications. Also, materials use cases in carbon capture and storage, catalysis, energy storage, materials replacement, and recycling are being intensely researched, along with pharmaceutical and biochemical applications. In some cases, impact arises from insight into fundamental phenomena, while in other cases it comes from screening a large number of candidate materials and selecting the highest-performing candidate.
2. **Physical scale:** typically, realistic materials use cases involve computational workflows consisting of various simulation steps with length and time scales covering more than 10 orders of magnitude. Therefore, predicting a material’s performance, from the molecular scale to the process scale, requires a broad range of computational methodologies including classical and quantum mechanical molecular dynamics simulations as well as the thermody-

amic and bulk continuum models. Recently, artificial intelligence and machine-learning techniques have been added to the physics-based simulations toolkit for applications including generative materials design, high-throughput materials screening, and accelerated simulation.

3. **Complexity:** application-guided and performance-oriented computational discovery work streams feature complex materials, where important phenomena occur in volumes often consisting of thousands of molecular building blocks with millions of atoms. These are challenging to represent adequately in any computational simulation environment. A real-world discovery use case requires, therefore, the identification of a minimum representative materials volume for conducting the performance analysis. This representation can be as small as the active space within a molecule, a surface at which a physical or chemical process occurs, or as large as an entire composite wing of an airplane.

In principle, quantum computing could be applied in each discovery step with the potential of enhancing computational accuracy, reducing computational time, or delivering any other improvement of the materials discovery outcome. For selecting the proper application of quantum computing in a real-world use case, it is however important to first identify and compare the difficulty of the classical and quantum simulation tasks and to ensure that proper data representations and algorithms can be developed for performing the discovery task on a quantum computing system. To that end, in our analysis of use cases, we will consider:

1. the relevance to real-world problems
2. the hardness of simulating the use case under consideration on classical computers
3. the hardness of simulating the use case under consideration on quantum computers (near-term and fault-tolerant)
4. the hardness of simulating quantum circuits involved in the previous point on quantum hardware and with classical simulators
5. the possibility of using classical HPC and quantum computers in concert

B. Overview of use cases

A central task of quantum theory, in the fields of quantum condensed matter physics and quantum chemistry, consists of describing the properties of interacting atomic, molecular, amorphous, and crystalline systems. Apart from advancing our fundamental understanding of quantum mechanics, a generic solution to the quantum problem would allow progress in application fields ranging




FIG. 10. Schematics overview of use cases in materials discovery. A central task of quantum theory is describing the properties of molecules and solids (blue). This goal can be achieved by solving for the electronic Schrodinger equation to compute e.g. ground-state, excited-states, and non-equilibrium dynamical properties (vertical arrow). Due to computational limitations, one can introduce approximations to the fundamental electronic structure problem, that reduce the size of the quantum problem within the existing computational budget (red). Accurate solutions of the electronic Schrodinger equation are the starting point (green) for important applications in materials science, including vibrational structure calculations, coarse-grained models (e.g. corrosion, deformation), and materials discovery (e.g. catalysis, metamaterials).

from materials discovery (better magnets, solar cells, catalysts, or qubit hardware) to drug design.

The required operation count and the number of quantum degrees of freedom required for the accurate ab initio description of such systems puts them far outside of the accessible regime on current or near-term quantum hardware, except for very small systems. These requirements were elucidated by many authors including Reiher et al [416], Tubman et al [417], Elfving et al [418] and Goings et al [419], who considered the application of quantum phase estimation to challenging instances of the electronic structure problem (respectively nitrogenase cofactor, chromium dimer, and cytochrome P450) on fault-tolerant devices.

It is therefore advantageous to consider approximations to the fundamental electronic structure problem that reduce the size of the quantum problem both in terms of qubit and operation requirements (see Figure 10, left).

Two complementary approaches are presently pursued. First, the embedding route identifies a set of strongly quantum parts of the problem (typically “orbitals” or “sites”), extracts the relevant physics to an auxiliary quantum problem, and solves this problem using a quantum processor. The remainder of the system is then treated classically with polynomial scaling approximations. The exact limit is recovered by enlarging the embedding subspace or increasing the accuracy of the classical method.

Second, the model systems route approximates the

original quantum system by a minimal set of appropriately chosen effective low-energy degrees of freedom. In this scheme, the exact limit is recovered in a “bottom-up” approach by gradually adding degrees of freedom until the original model is recovered. A typical example of this approach is given by the Heisenberg spin model, which (in the appropriate limit) captures the low-energy limit of the so-called t-J model, which itself is an approximation of the low-energy physics of the Hubbard model.

For both of these approximation approaches, simulation on quantum computers will become feasible much earlier than the simulation of the full quantum electronic structure problem. Realistic embedding approaches for materials typically result in compact but dense quantum problems (requiring many operations on few qubits), whereas model systems generally yield larger but sparse systems (requiring many qubits but fewer operations). The sparsity allows us to take advantage of devices with limited qubit connectivity with a reasonable operation count.

A central question concerns the physical observables under study. Traditional electronic structure mainly investigates ground states and their properties, such as the “band structure” or the ground-state energy as a function of nuclear coordinates. Out of these states, information about structure, vibrational modes, and mechanical properties is inferred.

However, experiments are performed at non-zero temperatures and, additionally, probe aspects of the excitation spectra. Photoemission spectroscopy, for instance, obtains the single-particle excitation spectrum and the gap size [420]. Nuclear magnetic resonance and magnetic neutron spectroscopy obtain nuclear spin excitations, and optical conductivity measurements probe two-particle excitations related to a current-current correlation function. Standard ground state methodology such as density functional theory (DFT) requires additional approximations to predict such observables. Famously the DFT “band gap” is not a reliable proxy for the gap in semiconductors [421–423]. Quantum (or classical) computing methodologies that intrinsically contain excitation information present a valuable addition to our materials toolkit.

In many experimental setups, a weak external field (such as an incident light beam) probes the physics of the material under study. A central assumption is that this field leaves the material essentially unperturbed, such that linear response theory provides an appropriate description for its physics. However, once fields become strong, the system itself acquires an explicit time dependence which may often be modeled as a quench (an instantaneous change of the Hamiltonian such as a very short but very strong light pulse), a steady state (the time-evolution of a system in a translation-invariant long-time limit, such as a system with an applied voltage after initial transients have decayed), or a periodic drive (such as a continuous laser field).

While our classical toolkit for simulating ground-state

physics is rather mature, the study of excitations within linear response is much more demanding and often requires high-performance computing resources. The study of time-dependent phenomena in real materials is still in its infancy. These science domains therefore offer promising application fields for future quantum computers.

Reliable solutions to the electronic structure problem also allow us to define the nuclear Schrodinger equation, which in turn can be solved to determine molecular vibrations, and coarse-grained models, whose solution allows us to characterize complex phenomena including material stress, corrosion, and cracking (see Figure 10, right).

C. Relevant use cases

1. Ground-state electronic structure

One of the goals of electronic structure is to determine the ground and low-lying eigenstates of the Hamiltonian of a molecule, often within the Born-Oppenheimer approximation, i.e., for a fixed set of nuclear positions \mathbf{R} . Determining the main features of the potential energy surface, i.e. the electronic energy as a function of nuclear positions $E(\mathbf{R})$, is key to understanding chemical reactivity, reaction rates, and product distributions. In particular, accurate estimation of reaction barrier heights (i.e. within ~ 1 kcal/mol from experiments) is important for many applications, such as catalyst development, battery electrode modeling, and manufacturing process design, and a known challenge for electronic structure methods and in particular for DFT [423].

As an example, in the design of manufacturing processes, the desired reaction(s) should dominate, e.g., have lower reaction barrier heights than undesired reactions. One way to arrange this is through reactant selection, however, there may be many possible reactants. Experimental screening to identify the best reactant candidates is often time, cost, and labor-intensive due to various factors. Reactants may be expensive, hazardous, or rare. Computational screening methods offer the possibility of screening candidate reactant barrier heights orders of magnitude faster, safer, and cheaper than experimental methods. However, current state-of-the-art classical computational methods are often insufficiently accurate to compare reactant barrier heights.

a. Static and dynamic correlation Within electronic structure, wavefunctions are often classified as single-reference and multi-reference. In the former case, a single electronic configuration dominates the target wavefunction. While the ground states of the majority of molecules have a single-reference character, one often encounters a multi-reference character in molecular excited states, at stretched bond geometries, in transition-metal chemistry, and in the presence of small HOMO-LUMO gap (e.g. large π -conjugated systems).

Single-reference problems can be accurately solved with various quantum chemistry methods for classical

computers, from DFT methods (that can routinely treat $O(10^3)$ atoms [424]) to high-level wavefunction methods such as coupled-cluster theory (that can attain kcal/mol accuracy on systems of $O(10)$ atoms [425]). Extending quantum chemistry methods to multi-reference situations is an active and challenging research field where, notwithstanding steady progress, the accuracy attainable in molecules with more than a $O(1)$ atoms is significantly lower than in the single-reference case.

Important examples of multi-reference quantum-chemical problems are found in transition-metal chemistry [7, 426]. Transition-metal oxides are an important class of materials that are used in catalysts, semiconductors, pigments, and many other applications. Benchmarking studies of DFT for bond dissociation enthalpies of transition metal oxides and their cations demonstrate mean absolute errors of 30 kJ/mol (7 kcal/mol) or greater [427], identifying the energetics of reactions involving transition metals as a challenge for DFT [427]. Examples of challenging instances of the electronic structure problem connected with transition-metal chemistry are:

1. Active sites of enzymes containing transition metals. These often involve multiple coupled transition metals, exemplified by the systems in Figure 11:
 - the Fe2S2 and Fe4S4 clusters of ferredoxins [428], where multi-reference character can be investigated in (30e,20o) and (36e,54o) active spaces respectively [429, 430]
 - the iron-molybdenum cofactor (FeMo-cofactor) of nitrogenase, which catalyzes the 6-electron reduction of N_2 to ammonium in biological nitrogen fixation [416, 431], and can be modeled by a (113e,76o) active space [429, 432]
 - the P-cluster of nitrogenase [433] and oxygen-evolving complex of photosystem II [434, 435]. In the context of nitrogenase, the hypothesized function of the P-cluster is to mediate electron transfer from the Fe-cluster to the FeMo-cofactor where nitrogen reduction takes place. The relevant states are the resting state P^N , the one-electron oxidized state P^{1+} , and the two-electron oxidized state P^{OX} , which can be modeled by (114e,73o), (117e,75o) and (120e,77o) active spaces respectively.

Active sites of enzymes containing transition metals present highly complex multi-reference quantum chemistry problems motivated by biological applications. Combined theoretical and experimental studies, primarily at the level of DFT, have illustrated many structural and electronic features of such active sites [436–438]. On the other hand, to interpret aspects of experimental spectroscopy, it is necessary to achieve a detailed understanding of the interplay between spin-coupling and delocalization between metals, which in turn




FIG. 11. The Fe_2S_2 (a) and Fe_4S_4 (b) clusters of ferredoxins, the P-cluster of nitrogenase in the resting state P^{N} (c), and the FeMo-cofactor of nitrogenase (d). Orange, yellow, teal, red, blue, gray, and white sticks denote Fe, S, Mo, O, N, C, and H atoms respectively.

requires accurate multi-reference quantum chemistry [428, 433, 435, 439, 440].

2. Nanocatalysts and surface catalysts containing transition metals. Simulating the mechanism of action of synthetic heterogeneous catalysts is challenging, as in the case of enzymes. DFT predictions of quantities such as the adsorption energy of small molecules are unreliable [441, 442]. While not all such problems are expected to be multi-reference in character, even the single-reference modeling of such chemistry, at a level significantly beyond DFT, is currently challenging or impossible. In addition, multi-reference effects are expected to play a role in certain catalysts, such as transition metal oxides, or at intermediate geometries in reaction pathways [441, 443–445].

For all these systems, the basic quantum simulation problem is the computation of the ground-state (or low-energy eigenstate) of the electronic Hamiltonians, and the basic metric is whether quantum algorithms yield more accurate energies (for comparable computational resources) than their classical counterparts. Proof-of-principle demonstrations can be carried out in the active spaces listed above. However, to describe the actual chemistry of these systems, one should improve over active-space calculations by (i) treating larger numbers of electrons and orbitals, (ii) accounting for environmental, dynamical, and solvent effects, and (iii) accounting for dynamical electronic correlation. While multi-reference problems are a known computational challenge, even single-reference problems become hard to model classically when large numbers of electrons and orbitals are involved [446].

b. Relativistic effects A discussion of the advantages and present challenges posed by hybrid quantum-classical approaches to electronic structure calculations

begs mention of the treatment of relativistic effects. Accurate modeling of the ground and excited states of many molecules, such as atomic gold and uranium dimers, are examples of systems whose precise computational treatment requires the inclusion of relativistic effects in the Hamiltonian. Without this inclusion, a fundamental precision limit to quantum computation remains in place [447].

c. Correlated electronic structure in solids There is some overlap in methods and ideas between the electronic structure problem in materials, e.g. crystalline solids, and quantum chemistry. When electron-electron interactions are weak, the low-energy properties of a material can be described by computing the band structure using DFT. However, in some materials (called strongly correlated), the electron-electron interaction fundamentally changes the behavior and the effective non-interacting ansatz underpinning DFT is no longer appropriate. Typical examples are Mott insulators, which appear as conductors in band-structure theory but are insulating due to electron-electron interaction. While Mott insulators are now well understood, many other phenomena in strongly correlated systems lack a justification in terms of an underlying microscopic mechanism and a quantitative and predictive theory of the associated physics. Some important examples include:

- *High-temperature superconductivity.* The mechanism driving superconductivity is not yet fully elucidated. Two regimes close to superconductivity, the pseudogap and strange-metal phase remain elusive [448]. The nature and relation to the superconducting phase of both are not understood. The strange-metal phase features behavior inconsistent with a simple weakly interacting metal even at high energies, hence the name of “non-Fermi liquid phase”, and has motivated a whole area of research on exotic metallic systems [449]. On the other hand, the pseudogap phase features several competing ordering tendencies [450], which are very challenging to resolve in numerical methods, because most methods naturally favor a particular ordering pattern, so discriminating physical effects from method biases is very difficult.
- *Frustrated spin systems* realize a host of high non-trivial phases, in particular topological and gapless spin liquids [451]. They have historically been the test bed for computational methods such as tensor networks and variational methods and, as such, appear as good test cases for quantum simulations.

2. Embedding

As seen in the previous Section, solving the electronic structure problem in situations that are relevant to real-world applications, e.g. catalysis, requires significant quantum resources. That example illustrates how the

goals of relevance to real-world applications and ease of execution on quantum devices are in conflict. However, there are two ways of making feasible progress toward the solution of challenging instances of the electronic structure problem: quantum embedding theories and effective low-energy models. These are not only useful steps toward the electronic structure problem, but also very compelling applications for current and near-term devices. Examples and applications are described in the remainder of this and following Sections.

The idea of quantum embedding theories stems from the fact that in chemistry, physics, and materials science, phenomena of interest often take place within a small region, motivating the separation of the whole system into a so-called active space surrounded by a host environment. The electronic structure problem is solved separately for the environment at a low level of theory and for the active space at a higher, i.e., more accurate and more expensive, level of theory. As an active space only accounts for a small portion of the whole system with a greatly reduced Hilbert space, simulating it using quantum chemistry methods becomes feasible on quantum computers and emulators. Examples of active spaces include point defects in semiconductors and insulators and active sites of catalysts at interfaces or on surfaces.

Various quantum embedding theories have been proposed, categorized by the level of theory used to describe different parts of the system.

- Density-based embedding theories define an active space in real space and partition the density of the system into the active space and the environment. The latter is described at the DFT level, and the former is described either using DFT with a higher-level exchange-correlation functional (DFT-in-DFT) or with wavefunction-based quantum chemistry methods (WF-in-DFT) [452–455].
- Density matrix embedding theories (DMET) [456–465], as introduced previously, also define an active space in real space. The active space is treated at a higher level of theory. The interaction with the environment is accounted for through a set of bath orbitals obtained from a calculation of the full system at a lower level of theory, with an additional one-body term that ensures the density matrices at both levels of theory coincide.
- Embedding theories based on Green’s function define an active space by a set of single-particle electronic states. The non-local, frequency-dependent self-energy of the active space is expressed as a sum of terms evaluated at high and low levels of theory, with an additional term to remove the double counting. Examples of such theories include dynamical mean-field theory embedding (DMFT) (DMFT+DFT [466, 467] or DMFT+GW [468–473]), self-energy embedding theory (SEET) [474–476], and quantum defect embedding theory (QDET) [477–480]. They mainly

differ in the choice of the high-level and low-level theories and the partitioning of the total self-energy into the active space and the environment.

In practice, a quantum embedding calculation typically starts with a low-level calculation of the full system, the results of which are subsequently employed to determine parameters of the many-body effective Hamiltonian of the active space. In QDET, for example, the effective Hamiltonian comprises a one-body term and a two-body term. In the one-body term, an exact double counting correction has been derived for the case where G_0W_0 is the low-level theory. In the two-body term, the effect of the environment is incorporated through an effective screening. The QDET Hamiltonian can be computed without any explicit summation over empty states, enabling its application to large systems. Leveraging GPU-accelerated classical supercomputers, it is feasible to simulate systems consisting of over a thousand atoms and active spaces consisting of over a hundred orbitals. The effective Hamiltonian, written in second quantization, can be diagonalized on classical computers using FCI, or on quantum computers [248, 478, 481–487] by mapping it to qubits and quantum gates and finding its lowest eigenstates using a suitable quantum eigensolver. Challenges remain in facilitating the communication of data from classical (super)computers to quantum computers, accommodating the effective Hamiltonian with the available number of gates and circuit depth, and mitigating noises present on near-term computers which may lead to unphysical results. Intriguingly, quantum computers themselves are employed to investigate the properties of spin defects in solids. These defects hold promise as prospective candidates for the implementation of improved quantum computers, thus forming a positive feedback loop at the intersection of computational materials science and quantum computing.

Additionally, CAS-DMET and NEVPT2-DMET described before have been effectively used to examine the adsorption and emission spectra of the negatively charged nitrogen-vacancy (NV^{-1}) and the neutral silicon vacancy (SiV^0) defects in diamond, as well as the neutral oxygen vacancy (OV^0) defect in both bulk and surface magnesium oxide, as reported in recent literature. [201–203] In particular, for the NV^{-1} defect in diamond, NEVPT2-DMET with a significantly reduced embedding space and a complete active space (CAS) configuration of 10 electrons in 9 orbitals (denoted as 10e, 9o) accurately predicts a triplet-triplet excitation energy of 2.31 eV. This is in close agreement with the experimentally observed value of 2.18 eV. Additionally, it predicts an optically inactive singlet-singlet transition at 1.02 eV, closely approximating the experimental value of 1.26 eV. For the neutral OV^0 defect in MgO bulk, NEVPT2-DMET, with a CAS configuration of 2e, 5o, estimates the excitation energy to be 5.24 eV, aligning well with the experimental absorption maximum of 5.03 eV. Studies of the optically allowed $S_0 \rightarrow S_1$ transition for the OV^0 defect on the MgO surface have also been conducted using CAS-

DMET, NEVPT2-DMET, and DME-PDFT methods. However, varying experimental results ranging from 1-5 eV have been reported, which makes direct comparison difficult. Using MP2 and CCSD solvers for the DMET embedding subspace, the adsorption energies for CO on MgO have been calculated using the largest embedding subspaces, and these have been found to be within 1.2 kcal/mol of the non-embedding reference values. [230] So far, the LASUCC and LASSI methods have primarily been tested on strongly correlated molecular systems and has been utilized to study the dissociation of $(\text{H}_2)_2$ into two H_2 molecules, the dissociation of the two double bonds in trans-butadiene, and the J -coupling in tris- $(\mu$ -hydroxo)-bridged chromium compound model [233, 234]. The accuracy of these studies has been confirmed through excellent agreement with the results from corresponding CASCI calculations performed using LAS orbitals.

3. Model systems

Low-energy models introduce an effective Hamiltonian describing the low-energy degrees of freedom of a quantum system. The effective Hamiltonian is then solved for the ground- and low-energy excited states. How to propose a low-energy effective Hamiltonian, and how to propose reliable and precise parameters for that Hamiltonian, are important questions that affect the predictive power of that model for real-world applications. On the other hand, accurate solutions of a given model allow us to capture trends in physical properties as the free parameters of the model vary (e.g. temperature, interaction strength, number of electrons). Therefore, the value of a model as it relates to real-world application lies both in its definition and in its computational solution: the physical properties predicted by a model should be equally robust to perturbations in model parameters and solutions of the Schrödinger equation, for robust inferences about the physical properties of a real system to be drawn from the use of effective low-energy models.

Although band structure and Fermi liquid theory are used as the primary models for quantum particles in materials, these descriptions can fail to capture actual physical properties, for example in high- T_c cuprates and other transition-metal compounds like spin liquids. This offers a challenge in how to improve their performance for such systems. Various models have been suggested for these systems, such as the Emery and Heisenberg models reviewed below. However, the ability of such models to capture actual physics of interest is again often unclear, even though they have been widely studied and have greatly advanced our understanding of correlated electron physics. Furthermore, these models can themselves be challenging to solve.

A typical approach to connecting ab initio electron-level models and large-scale models is downfolding using DFT. In this method, DFT bands are modeled by projecting them onto localized Wannier functions, which

yields one-particle hopping terms [488]. Interactions are added by modeling screening, for example using constrained DFT or random phase approximations (RPA). Finally, since the original DFT bands also already incorporated interactions, this double-counting must be corrected. This approach and related downfolding techniques continue to be actively developed and broadly used [489–494].

In the remainder of this Section, we review some important model systems, namely spin Hamiltonians and lattice models of correlated electrons. While the relationship between these models and real materials is very delicate, both due to the choice of the effective Hamiltonian and its parameterization, studying computationally the phase diagram of these models as a function of their free parameters can reveal trends and phase transitions that are relevant for the physics of real materials. Furthermore, it constitutes a promising research area for simulations on current and near-term quantum devices.

a. Spin Hamiltonians Spin Hamiltonians are frequently used to model the magnetic properties of materials as well as other systems where there is local correlation or frustration. Consider a lattice with n sites occupied by spin- s particles that are allowed to interact with each other. In such a lattice, there is an allowed number of $2s + 1$ microstates at each lattice site, giving a total of $(2s + 1)^n$ spin configurations for a lattice with n sites. At an infinitely high temperature, spins point randomly in all directions, agitated by thermal motion. On the other hand, as the temperature approaches $T = 0\text{K}$, in many situations spins display the tendency to become part of a long-range magnetically ordered state under the effect of the interaction between them. The two extreme limits are antiferromagnetic order, where adjacent spins point in opposite directions, and ferromagnetic order, where all spins point in the same direction. If we consider paramagnetism (namely the absence of magnetic order in the absence of an externally applied magnetic field) as analogous to an ideal gas (where particles interact weakly with each other), and magnetic order as analogous to a solid (where particles form a static and regularly repeating lattice), it is natural to inquire if there exists an analogous liquid of spins.

Spin liquids [495–500] are exotic materials where spins do not order at any finite temperature but interact strongly with each other, unlike in a paramagnet where spins are weakly coupled to each other. Normally, strong spin interactions result in a single spin configuration (or a small subset thereof) having lower energy than the others, thereby driving magnetic order. In a spin liquid, interactions conspire to make all spin configurations (or a large subset thereof) nearly equivalent in energy, thereby evading magnetic order. In a spin liquid state, the ground state consists of a superposition of spin configurations that is not dominated by a single configuration (or a small subset thereof). A spin liquid thus does not possess magnetic order, as individual spins fluctuate, but possesses strong entanglement in the ground state,

so that spins are not randomly mutually oriented.

The key ingredient to obtain a ground state that is a superposition of a large number of spin configurations is to have strong interaction between neighboring spins and have a large number of configurations with near-identical energies, i.e. to frustrate the spins. This often arises due to the underlying geometric frustration of the lattice in which spins sit and/or the nature of the interaction between spins. Furthermore, particles with higher spin s tend to show more classical behavior as opposed to quantum behavior, because the energy barrier between different microstates scales as s^2 , which strongly suppresses quantum fluctuations. Therefore, a spin as low as possible, e.g. $s = 1/2$, greatly favors spin liquid behavior.




FIG. 12. Left: crystal structure of α -RuCl₃, illustrating the van der Waals gap between the RuCl₆ honeycomb layers (with Ru/Cl atoms represented as gray/green circles) and magnetic Ru³⁺ ions aligned as in the zigzag phase (red arrows). Right: schematic representation of the two-dimensional Heisenberg-Kitaev model with spins represented as white circles and $x/y/z$ bonds as green/blue/red links respectively. Quantum simulations to date have focused on 1 and 2 hexagons, corresponding to 6 and 10 spins respectively [501].

A concrete example of a material of interest as a proximate spin liquid [497] is ruthenium trichloride, α -RuCl₃ (see Fig. 12). This transition metal halide has a crystal structure made up of stacked honeycomb layers of edge-sharing RuCl₆ octahedra. Plumb et al [502] found that spin-orbit coupling in this material is substantial, leading to a $j = 1/2$ description of the Ru³⁺ valence electrons. Since this material is built up with edge-sharing RuCl₆ octahedra, it is believed [503, 504] that its low-energy excitations are parametrized by the Kitaev-Heisenberg Hamiltonian

$$H_{kh} = \sum_{\langle ij \rangle_\gamma} J \vec{S}_i \cdot \vec{S}_j + K S_{i\gamma} S_{j\gamma}, \quad (16)$$

where $\langle ij \rangle_\gamma$ denotes summation over the bonds $\gamma \in \{x, y, z\}$. This feature makes α -RuCl₃ a material of great interest in the ongoing search for a Kitaev spin-liquid ground state [505–507]. Although α -RuCl₃ orders magnetically at low temperature with zigzag magnetic order [507, 508], this material has shown some signatures of spin-liquid physics, such as a broad continuum of magnetic excitations identified in both Raman scattering [509] and inelastic neutron scattering measurements [507]. Therefore, the degree of similarity to the Kitaev

model, as well as the interpretation of spectroscopic and thermal measurements in terms of modifications to the Kitaev model, is currently intensely debated. We remark that, in the literature [510], the Kitaev-Heisenberg Hamiltonian is augmented by additional couplings,

$$H = H_{kh} + \sum_{\langle ij \rangle_\gamma} \Gamma (S_{i\alpha} S_{j\beta} + S_{i\beta} S_{j\alpha}) + \sum_{\langle ij \rangle_\gamma} \Gamma' (S_{i\gamma} S_{j\alpha} + S_{i\gamma} S_{j\beta} + S_{i\alpha} S_{j\gamma} + S_{i\beta} S_{j\gamma}), \quad (17)$$

with $\{\alpha, \beta, \gamma\} = \{x, y, z\}, \{z, x, y\}, \{y, z, x\}$ for $\gamma = z, y, x$ respectively. Several parameter configurations have been proposed to describe the Hamiltonian H , ranging from first-principle methods [510, 511] to phenomenological analyses [507, 510]. Recent literature has focused on $J = -1.53$, $K = -24.4$, $\Gamma = 5.25$, $\Gamma' = -0.95$ [501]. The coupling that is believed to be the leading one is the (negative) Kitaev term $K < 0$. The off-diagonal exchange term $\Gamma > 0$ is also believed to be significant, while the ferromagnetic exchange J is believed to be subleading [512]. However, it should be noted that the parameters of H vary quite significantly between the studies [510, 513]. The exactly solvable Kitaev model corresponds to $J = \Gamma = \Gamma' = 0$, and it is argued that, in the parameter regime of α -RuCl₃, both the excitations and the heat capacity show echoes of the two kinds of Majorana fermions that exist at the solvable point [514, 515].

The combination of rich physical behavior and challenging ground- and thermal-state preparation makes quantum spin liquid models like the Kitaev-Heisenberg Hamiltonian approximating α -RuCl₃ interesting and challenging problems for classical and quantum simulations. The presence of 2-local spin-1/2 Hamiltonians makes them compelling targets for simulations on quantum devices, focused on the calculation of specific heats and dynamical structure factors relevant to understanding the low-energy excitation spectrum of the Hamiltonian.

b. Lattice models A central feature of correlated electron materials is the competition between different inhomogeneous orders. Important examples are metal-insulator [516], paramagnet-ferromagnet [517], and conductor-superconductor phase transitions [518, 519]. Lattice models of correlated electrons serve as the canonical microscopic physical models for the computational description of such competition phenomena in materials. A paradigmatic example of a fermionic lattice model is the one-band Hubbard model [520–523],

$$\hat{H} = -t \sum_{\langle ij \rangle_\sigma} (\hat{a}_{i\sigma}^\dagger \hat{a}_{j\sigma} + \text{h.c.}) + U \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow}. \quad (18)$$

In Eq. (18) where \hat{a}/\hat{a}^\dagger denote the usual fermion annihilation/creation operators, $\hat{n} = \hat{a}^\dagger \hat{a}$ is the number operator, t and U are the kinetic and repulsion energies, and $\langle ij \rangle$ denote summation over nearest neighbors sites ij on a lattice (e.g. the 2D square lattice). The

Hubbard model is one of the simplest models of interacting fermions, but despite its simplicity, it exhibits a wide range of correlated electron behavior including interaction-driven metal-insulator transitions, superconductivity, and magnetism. The precise behavior depends delicately on parameters, creating an interesting challenge for theory and computation. Exact solutions are available for one-dimensional [524] and infinite-dimensional cases [525]. High-temperature series expansions provide numerically exact results but only for temperatures too high to be relevant for physically interesting situations [526]. In general dimensions at relevant temperatures, only approximate numerical solutions are available, including wavefunction-based, diagrammatic, and embedding methods [527]. These numerical methods have been steadily developed and carefully benchmarked [527, 528], providing illuminating insights into the physics of this model.




FIG. 13. (a) schematic view of a [CuO₂]₄ cell in a cuprate CuO₂ plane. Cu/O atoms are represented with orange/red circles and orbitals by blue and green contour plots (blue/green for positive/negative values). (b) definition of the three-band Hubbard model parameters (the curve connectors represent the kinetic energy coefficients).

One concrete example of a fermionic lattice model that can be connected to a technological application is the three-band Hubbard model in two dimensions, also known as the Emery model [529]. This model is generally believed to contain the essential properties of high-T_c cuprate superconductors that arise from the interplay between the copper d_{x²-y²} and oxygen p_{x,y} orbitals in the CuO₂ layers (see Figure 13). It is defined, in the so-called “phase representation”, by the Hamiltonian

$$\begin{aligned}
 \hat{H} = & t_{pd} \sum_{\langle ij \rangle \sigma} (\hat{d}_{i\sigma}^\dagger \hat{p}_{j\sigma} + \text{h.c.}) + t_{pp} \sum_{\langle ij \rangle \sigma} (\hat{p}_{i\sigma}^\dagger \hat{p}_{j\sigma} + \text{h.c.}) \\
 & - \Delta_{pd} \sum_{i\sigma} \hat{n}_{i\sigma}^d \\
 & + U_d \sum_i \hat{n}_{i\uparrow}^d \hat{n}_{i\downarrow}^d + U_p \sum_i \hat{n}_{i\uparrow}^p \hat{n}_{i\downarrow}^p + V_{pd} \sum_{\langle ij \rangle \sigma \tau} \hat{n}_{i\sigma}^p \hat{n}_{j\tau}^p
 \end{aligned} \quad (19)$$

where $\langle \rangle$ denotes nearest neighbors, $d_{i\sigma}^{(\dagger)}/p_{i\sigma}^{(\dagger)}$ destroys (creates) a hole with spin σ on the Cu[3d]/O[2p] orbital, $\hat{n}_{i\sigma}^d/\hat{n}_{i\sigma}^p$ is the number of holes on a Cu[3d]/O[2p] orbital

with spin σ , and the charge transfer gap Δ_{pd} is defined as the difference between the O[2p] and Cu[3d] orbital energies.

On classical computers, the Emery model has been investigated with several numerical methods, including direct simulations of lattices by exact diagonalization of up to 24 sites [530], quantum Monte Carlo and DMRG of 80 and 144 sites [531, 532], and embedding methods like density matrix embedding theory with 12 sites [533]. These simulations have been primarily directed at simulating ground-state spin-spin, and *d*-wave correlation functions, to establish e.g. the presence of stripe or *d*-wave superconducting order [530–533].

Due to the complexity of the model, unlike in the one-band case, a consensus on much of the physics has yet to be reached, providing an opportunity for high-accuracy quantum computing algorithms to elucidate the properties of high-T_c cuprate superconductors by simulating lattices of O(100) sites.

The Emery model is often simplified considering a one-band Hubbard model. While much of the physics of high-T_c cuprate superconductors, e.g. *d*-wave pairing, density waves, pseudogap phase, stripe order, is described by the one-band Hubbard model Hamiltonian, there are still important reasons to go beyond the one-band picture to study the original three-band model directly. For instance, (a) some important physics may be lost in the reduction to the one-band approximation (such as a role for the oxygen degrees of freedom in the pseudogap phase [534]), (b) near degeneracies of competing states seen in the one-band case [528] may be resolved with the additional degrees of freedom of the three-band model, and (c) the three-band model retains the atomic structure of the CuO₂ layer and thus has a direct link to the structure of real materials as well as experimental measurements of orders at the atomic scale [535–537].

D. Excited states

Many experiments on condensed-matter systems probe dynamical properties, rather than equilibrium properties. For example, material properties are often explored through scattering experiments, such as neutron and photon scattering or angle-resolved photoemission spectroscopy (ARPES) [538], which probe the structure factor and spectral function respectively [539, 540]. In these experiments, a system is initialized at equilibrium at inverse temperature β under a Hamiltonian \hat{H} , and a time-dependent perturbation $\hat{V}(t)$ is applied. The response of the system is characterized by linear response theory, computing frequency-dependent correlation functions of the form

$$C_{AB}(\omega) = \int_{-\infty}^{\infty} \frac{dt}{2\pi} e^{i\omega t} \text{Tr}[\hat{\rho}(\beta) e^{it\hat{H}} \hat{A} e^{-it\hat{H}} \hat{B}]. \quad (20)$$

where \hat{A}, \hat{B} are suitable operators, e.g. components of the dipole moment operator and density operators at two

different spatial points in UV/visible light and neutron scattering experiments respectively. It should be noted that $C_{AB}(\omega)$ can be computed knowing the eigenpairs of the unperturbed Hamiltonian \hat{H} and/or its time evolution operator $e^{-it\hat{H}}$. In other words, by virtue of the linear response approximation, the perturbation $\hat{V}(t)$ does not appear in the function $C_{AB}(\omega)$.

Going beyond spectral properties, the non-equilibrium real-time dynamics of quantum systems have increasingly come into focus, both because of experiments that can probe quantum dynamics at atomic scales, and because of fundamental interest in studying the equilibration of quantum systems, which serves as a bridge between the theories of quantum mechanics and statistical mechanics. Experimental setups now exist that can probe ultra-fast dynamics in materials including, for example, free-electron lasers [541, 542] as well as other pulsed laser systems. These allow the application of experimental techniques, such as pump-probe spectroscopy [543, 544], to provide novel insights into the behavior of correlated quantum systems. Within non-equilibrium real-time dynamics, one is often interested in the calculation of time-dependent expectation values, $B(t) = \text{Tr}[\hat{\rho}(\beta)\hat{U}(t,0)^\dagger\hat{B}\hat{U}(t,0)]$. The time-evolution operator $\hat{U}(t,0)$ under $\hat{H} + \hat{V}(t)$ can be written as a time-ordered exponential, yielding the following expression for $B(t)$, based on a Keldysh contour integral,

$$B(t) = \sum_{nm=0}^{\infty} i^{n-m} \int_0^t dr_n \dots dr_1 \int_0^t ds_m \dots ds_1 \quad (21)$$

$$\text{Tr}[\hat{X}(r_n \dots r_1)^\dagger \hat{B} \hat{X}(s_m \dots s_1)]$$

where $r_1 < \dots < r_n$, and $s_1 < \dots < s_m$, and

$$\hat{X}(s_m \dots s_1) = \hat{U}_0(t, s_m) \hat{V}(s_m) \hat{U}_0(s_m, s_{m-1}) \dots \quad (22)$$

$$\dots \hat{U}_0(s_2, s_1) \hat{V}(s_1) \hat{U}_0(s_1, 0),$$

with $\hat{U}_0(t, t') = e^{-i(t-t')\hat{H}}$. Linear-response frequency-dependent correlation functions and non-equilibrium real-time correlation functions are very relevant physical quantities, they are challenging to compute on classical computers because they require access to the time dynamics of a system (or equivalently to its excited states), and they are very suitable tasks for fault-tolerant quantum computers, due to their ability to simulate complex unitary transformations like time evolution operators at polynomial cost and with controllable approximations only.

E. Applications of electronic structure

1. Vibrational structure calculations

Obtaining accurate vibrational spectra of molecules is a costly task on conventional computers. While uncovering the electronic structure of molecules stands as a

fundamental challenge in quantum chemistry and material design, to truly make an impact in both scientific research and practical applications, it is vital to go beyond the electronic structure. This requires creating a kinetic model that relies on a deep understanding of a molecule's vibrational structure. Knowing a molecule's vibrational structure enables the prediction of thermodynamic properties that are key in many fields, such as atmospheric science, catalysis, and fuel combustion modeling. Although classical computers often handle simulation of the electronic structure of small molecules reasonably well, they struggle with calculating vibrational structures beyond the harmonic approximation, even for small molecules. Computational challenges emerge when higher-order terms are involved due to deviations from harmonicity and also the interplay between different bosonic modes. This can be described, e.g. by quartic force field Hamiltonian [545]

$$H_{\text{anharm}} = \frac{1}{2} \sum_i^M \omega_i (q_i^2 + p_i^2) + \sum_{\{ijk\}} h_{ijk} q_i q_j q_k +$$

$$+ \sum_{\{ijkl\}} h_{ijkl} q_i q_j q_k q_l + \dots,$$

where the first term indicates the harmonic approximation and the remaining terms are the anharmonic corrections, with h being the force constants.

Applications of quantum computing in calculating molecular vibrational structures have not been extensively explored, see the references in [415, 546, 547]. While there have been a few proposals for both near and long-term quantum algorithms to address the vibrational structure problem, the implementation on currently available devices requires developing efficient quantum algorithms for near-term devices as well as a better understanding of possible optimizations to minimize the qubit numbers and circuit depth. Recent studies on quantum resources required for vibrational structure calculations on quantum computers indicate that the combined resources needed for achieving quantum advantage in vibrational structure problems might be lower than those for electronic structure problems [547, 548]. The vibrational qubit Hamiltonian's Pauli strings are more localized, making them potentially easier and faster to simulate compared to those in the electronic qubit Hamiltonian. It is noted that the number of Pauli strings simultaneously executed in a quantum circuit during a single Trotter step is also larger for the vibrational structure problem compared to the electronic one.

Enhancing vibrational structure calculations can be achieved through two possible approaches. One is centered on refining the accuracy of the electronic potential energy surface (PES) using quantum computers. This method proves effective when dealing with rigid molecules, especially those where the vibrational structure is well-described by harmonic approximations or classical algorithms. The second approach focuses on el-

evating vibrational structure calculations through quantum computing, assuming a sufficiently accurate electronic energy surface. This becomes particularly crucial in characterizing floppy molecules and capturing anharmonic effects. One pertinent example of the latter case, with industrial significance, is polyynes molecules. These organic compounds, characterized by alternating carbon single and triple bonds, are challenging to stabilize. They belong to high symmetry point groups and feature numerous silent modes, undetectable through infrared or Raman spectroscopy. A precise depiction of their low-frequency vibrations is crucial for understanding their reactivity. Polyynes play a pivotal role in floating catalytic chemical vapor deposition and the emerging technology of carbon nanotubes, ultimately contributing to reducing the carbon footprint. Precise numerical results on the quantum computational cost for polyyne molecules have been provided in the recent study [547]. Another recent study has explored the first approach, examining how a hybrid quantum and classical computation could improve the vibrational structure simulations [549]. For a model lithium hydride system, the PES is computed on quantum hardware, and then classical computation is used to both fit analytic forms to the PES and from those generate the vibrational energy levels that can be used to predict thermodynamic properties. This hybrid algorithm was demonstrated with actual calculations on quantum hardware, though the authors point out that extensions to systems with more degrees of freedom may require innovations to avoid the combinatorial complexity of sampling the entire space of molecular deformations.

One more classic example of a multiscale problem that couples electronic, vibrational, and transport phenomena is photosynthesis. The core process is the conversion of electromagnetic energy to a stored chemical form, which involves photon absorption plus exciton transmission, transduction, and dissipation. This is all done by biological systems with optimized electronic-vibronic coupling to move energy, with estimated photochemical yields for some subsystems of up to 85% [550]. It has been shown recently that these systems may even include alternate pathways to dissipate energy and avoid damage to the key molecular components [551]. Multi-scale modeling of photosynthesis has been a long-standing area of interest in biological simulation [552]. Simplified molecular simulations can provide a model for vibrational modes, but they need to be combined with exciton transfer calculations as well as detailed simulations of reactive chemistry to create the full picture. The exciton transfer calculations are a good example of a quantum process with a relatively simple Hamiltonian, as the number and geometry of donor and acceptor species are known. Fluctuations in their geometry, however, require coupling to a simpler simulation model. The last chemical step then requires accurate quantum chemical modeling of complex systems. In photosystem II, for example, the key chemical step is using the excited electron for water oxidation.

This happens in a small active site complex of ten mixed magnesium, calcium, and oxygen atoms, and the process is still incompletely understood [553].

Photosynthetic complexes, modeled as open quantum systems interacting with a vibrational environment, have been the focus of various analytical and numerical studies. However, the computational resources required by most classical methods increase exponentially as the number of particles in the simulation grows. This challenge intensifies significantly for open quantum systems with structured environments, such as photosynthetic complexes. While numerically exact solutions are achievable for only small systems (fewer than 20 particles) and under restrictions concerning bath modes, the exploration of more complex open quantum systems remains a challenging computational barrier. This limitation is illustrated in Figure 14, indicating the current upper limits of computational capabilities on classical computers for simulating the dynamics of such open quantum systems. As it is shown, there is a trade-off between the complexity of the spectral density and the system size [554]. Various proposals have surfaced for simulating these complex systems on quantum hardware, including gate-based simulations, a hybrid quantum-classical approach, as well as analog quantum simulators, see the references in [555], for example. Nevertheless, conducting large-scale simulations on existing quantum hardware poses a challenge, due to the limitations, e.g., circuit depth, in currently available quantum devices, so near-term progress will likely depend on hybrid quantum-classical algorithms or better optimization and resource reduction techniques.

2. Coarse-grained simulations

The commercialization of new, advanced materials is often a cost-, time-, and labor-intensive process. Typical times to bring a new material to market are on the order of 10-20 years from initial research to first use [556]. The long, expensive commercialization process is partially due to the experimental nature of materials development, e.g., many material samples must be physically prepared and then evaluated. Challenges in the process include the sample preparation, which for new materials has a steep learning curve, with many iterations of learning required to identify the best methods to create high-quality samples for evaluation. The precursors and by-products for sample preparation may be expensive, hard to obtain, and/or unsafe, thus requiring specialized equipment to handle. Further, evaluation of the samples may be laborious with many data points to be taken under carefully controlled conditions. Finally, more often than not, once a sample is finally prepared and evaluated, the conclusion is that the material does not meet the desired specifications, and the entire process must be repeated for a different material candidate.

The materials commercialization process may be accelerated by conducting material exploration computa-




FIG. 14. The shaded region in the graph illustrates the estimated sizes of treatable systems when simulating Frenkel exciton Hamiltonians using current classical supercomputing resources. Three photosynthetic systems are presented: the Fenna-Matthews-Olson complex of Green-Sulfur Bacteria, the Light Harvesting I and II complexes of Purple Bacteria, and Photosystem II of higher plants. The simulation was conducted utilizing the hierarchical equations of motion approach on 64 AMD Opteron cores, utilizing a total of 250 GB of RAM. Image from [554].

tionally. New materials may be evaluated quickly and efficiently by virtually modeling a material and predicting its properties. The role of experimentation is moved from exploration to confirmation of the computational results.

In computational materials discovery, the process often has two steps. The first identifies or generates material candidates, which may be done, for example, with machine learning methods. In the second step, the material candidates are filtered by computationally predicting their material properties. Methods to compute material properties are discussed here.

One type of material property in demand is those that are both light and strong. The main modeling for these concerns two main areas. The first is materials discovery, in which computation is used to test predicted materials' stability. This includes stability against spontaneous transition to alternative undesirable phases with the application of small amounts of pressure or temperature change; or indeed stability against phase separation, in which constituent atoms cluster together and do not form compounds or layers as desired. The computational tests for this include perturbing atomic positions and seeing minute energy changes in order to infer reaction barriers. These are DFT calculations using often nudged elastic band techniques. Additional tests are done by calculating the phonon spectra. A state-of-the-art set of such tests appears in [557], pages 6-8.

The second main area of modeling for advanced ma-

terials is defect studies: crack propagation and corrosion are two key areas. Both of these defects can be either across a surface in the form of failure of protective coatings and surface cracks; or deep into a surface, whether a crack or deep pit corrosion. Failure of mechanical properties occurs on all length scales, starting at the atomic and spreading to macroscopic, visible effects on the scale or centimeters or longer. For this, the main modeling approach is a type of embedding theory as introduced above. The essence of embedding techniques is a set of calculations that are done at much higher accuracy, usually at the cost of both time and resources, in a key important region chosen as the active space. The results of the embedding calculations then need to be matched to those done at a coarser level, a challenging task to do well, and which difficulty depends to some degree on the type of embedding done, see the paragraph below. For example, if the type of embedding is by calculating quantum wave functions in the active space, connecting to classical more macroscopic descriptions of phonon and other vibrational levels can be a very challenging task. However, if instead the active space consists of precise calculations of atomic positions at zero temperature, then these can be connected to those at finite temperature by a process of ramping up vibrations, and this can work quite well. When the defect to be modeled, such as a crack, occurs at multiple length scales, then each scale working up from the smallest to the largest can be considered the embedded active space of the one that comes at a larger scale, and matching must be done at each stage. This is termed multi-scale quantum embedding.

Both for crack propagation and defect modeling such as corrosion of the key material or failure of a coating, often first-principles molecular dynamics is also used, with DFT used to calculate the forces and energies. These calculations are used both to test the matching at the different scales and also as a study tool in itself, to try to infer mechanisms for materials failure. With a crack, the action can spread very quickly from the atomic to the macroscopic, if forces are strong enough, and MD is key for measures of the speed of such scale transitions. For other kinds of defects such as wear and corrosion, processes typically happen on a much slower time scale, and deeper calculations can be done at the atomic level to study the origins of the failure. However, in almost all cases, these, too, are done with quantum embedding methods of detailed fully quantum calculations, often extending to full-configurational chemistry, for the active regions.

Other types of material challenges are in handling creep, a slow deformation when a material is exposed to long-term stress (pressure, force, twist). These types of slower macroscopic changes are usually treated with differential equations, which have optimal algorithms in quantum computing, particularly if the problem can be linearized. Additional observables such as Young's modulus and Poisson's ratio measure the properties of materials and can reveal defects if the material is not up to




FIG. 15. The schematic structure of the planar multi-layer in the transparent radiative cooling metamaterial that can be mapped into a binary vector (from [562]).

the expected level. Methods to compute these properties are also on a macroscopic scale and use both differential and integral equations.

When materials are exposed to magnetic or electrical fields, we term these open quantum problems. Systems in magnetic or electrical fields are known as driven systems, and those at finite temperatures as dissipative systems. Thus one can have a driven, dissipative system. These types of systems are usually handled with lattice models, as discussed in a previous Subsection.

3. An example use case: quantum metamaterials

The above use cases show how quantum and classical computing can effectively work together to answer key questions in materials design. As noted in our criteria, such modeling must be relevant to real-world problems, challenging to simulate classically, include feasibly solvable sub-problems on quantum computers, and potentially involve classical HPC and quantum computers in concert. The overall evolution of simulation in materials discovery generally follows the three phases of phenomenological simulation for insight, property prediction for candidate screening, and finally direct optimization of materials for target properties. To illustrate this further, we explore an example use case of metamaterials in more detail.

The design of metamaterials, characterized by their unique properties like negative refractive index and selective spectral reflection, holds immense promise for diverse real-world applications [558–561]. These artificially structured materials, functioning at sub-wavelength scales, offer a novel path to enhance the performance of conventional electronic devices. This improvement, in turn, has the potential to revolutionize energy-saving in fields such as thermophotovoltaics, heating, ventilation, and air conditioning systems. Consequently, the pursuit of optimal metamaterial designs has become an area of intense research interest.

The properties of metamaterials hinge not only on their inherent material characteristics but also on their geometrical parameters. Therefore, determining the ideal geometrical parameters is fundamental to achieving high-performance metamaterials. However, the challenge arises from the multi-dimensional nature of these parameters, creating vast, non-convex design spaces. As such, finding the optimal metamaterial structures through intuition, trial-and-error, exhaustive searches, or random approaches becomes a complex and laborious task.

This leads to essential questions in metamaterial design: How can we identify high-performance metamaterials that are close to the global optimal design? How can we optimally design metamaterials with complex geometric parameters that lead to significantly large design spaces? To address these questions, a novel strategy is proposed - an active learning algorithm that harnesses the power of machine learning, quantum computing, and HPC to efficiently design metamaterials, even in cases with complex target properties and expansive search spaces. Metamaterials, often consisting of discretized design concepts, necessitate effective solutions for combinatorial optimization problems. Quantum computing, exploiting the unique principles of quantum physics like superposition, entanglement, and tunneling, offers exponential computational speedup over classical computing [563].

To solve optimization problems using quantum computing processors, an objective function is formulated as a Hamiltonian, describing the relationship between metamaterial states and their properties. Factorization machine, a supervised machine learning model, can effectively capture the relationship between input variables and output values, providing model parameters such as bias, linear terms, and quadratic terms. These parameters can be directly utilized to construct a Hamiltonian matrix [564]. Consequently, an iterative algorithm that combines factorization machines [565] and quantum computing can efficiently discover optimal metamaterial designs.

Despite the efficiency of quantum computing in solving Hamiltonian objective functions, current quantum devices have a limited number of qubits. Thus, their utility for numerical simulations is restricted. Parallel computing, particularly HPC, holds the potential to provide substantial acceleration for such tasks. Therefore, the integration of machine learning, quantum computing, and HPC emerges as a powerful approach for the efficient design of complex metamaterials. One more question emerges: How can continuous variables be optimized using a quantum computing-based algorithm, which typically returns binary variables? To address this limitation, another quantum computing-based algorithm is proposed [566], capable of simultaneously optimizing discrete and continuous spaces, making it a more versatile tool in metamaterial design.

VIII. CONCLUSIONS AND OUTLOOK

This work can be summarized as a perspective on quantum computing in the field of materials science. Rather than an exhaustive review of the existing literature, this paper aims at pointing out the challenges for a quantum-centric supercomputing architecture to achieve quantum advantage for materials science use cases. For that, we have presented a comprehensive view of topics at different level of abstraction, ranging from hybrid quantum-classical architectural problems to the final use cases. We believe this is the major contribution of this paper.

The different sections have highlighted the challenges at various levels. The first two sections closely tie together, listing fundamental algorithms for simulation of quantum systems and specific hurdles that arise in implementing them at scale both for noisy quantum devices and fault-tolerant architectures, with an emphasis on the classical workloads. Then, we moved our focus closer to the architecture design arguing in details about the difficulties that are faced when optimizing quantum and classical workloads coming from hybrid algorithms. We discussed more general topics around quantum-centric supercomputing, namely, workload management and programming models, and how extensions to existing classical computing frameworks can help to address quantum computing device requirements. Before tying up all the threads in the section about use cases, we have discussed how classical simulations of quantum systems can be used both for approximate verification and to identify hard use cases. Finally, we have provided criteria for the identification of good use cases in materials science, drawing upon the previous sections. Based on these criteria, we have discussed a few exemplary cases, representative of the variety of applications of materials science in mind.

One major takeaway from this perspective is to suggest how we can lay the grounds to think about quantum computing together with classical high-performance computing, in quantum-centric supercomputing centers. Materials science provides a great setting for use cases, which have potential for quantum advantage in scientific

and industrial applications.

ACKNOWLEDGMENTS

Y.A. acknowledges support from the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357 at Argonne National Laboratory. A.D.M., M.G, and S.V. are supported by CERN through the CERN Quantum Technology Initiative (CERN QTI). A.F.I. acknowledges financial support from the Natural Sciences and Engineering Council of Canada (NSERC). The work at the DIPC was funded by the Gipuzkoa Provincial Council (project QUAN-000021-01), the European Union (project NextGenerationEU/PRTR-C17.I1), as well as by the IKUR Strategy under the collaboration agreement between Ikerbasque Foundation and DIPC on behalf of the Department of Education of the Basque Government. Y.A., G.G., L.G., A.M. M.H. and B.H. acknowledge funding support from the Next Generation Quantum Science and Engineering (Q-NEXT), supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. A.L. and T.S.H. acknowledge support from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center (QSC). W.A.d.J. acknowledges funding support from the Quantum Systems Accelerator (QSA), supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. N.M.T. acknowledges support from by U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-Design Center for Quantum Advantage under Contract No. DE-SC0012704 (C2QA). L.G., A.M. and M.H. acknowledge partial support from the NSF QuBBE Quantum Leap Challenge Institute (Grant No. NSF OMA-2121044). The work at the Oak Ridge National Laboratory used resources of the Oak Ridge Leadership Computing Facility which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

-
- [1] R. P. Feynman, Simulating physics with computers, *International Journal of Theoretical Physics* **21** (1982).
 - [2] Argonne Leadership Computing Facility, Materials Science, <https://www.alcf.anl.gov/domains/materials-science> (2023).
 - [3] B. Austin *et al.*, Nersc-10 workload analysis, https://portal.nersc.gov/project/m888/nersc10/workload/N10_Workload_Analysis.latest.pdf (2020).
 - [4] C. Connor, *Institutional Computing for Open Science and Engineering*, 2023 IS&T Capability Review Presentation LA-UR-23-26091 (Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 2023).
 - [5] Donostia International Physics Center Supercomputing Center, Supercomputing Center, <https://dipc.ehu.eus/en/supercomputing-center> (2023), [Online; accessed 8-December-2023].
 - [6] A. Szabo and N. S. Ostlund, *Modern quantum chemistry: introduction to advanced electronic structure theory* (Courier Corporation, 2012).
 - [7] N. M. Tubman, J. Lee, T. Y. Takeshita, M. Head-Gordon, and K. B. Whaley, A deterministic alternative to the full configuration interaction quantum monte carlo method, *The Journal of Chemical Physics* **145**,

- 044112 (2016), <https://doi.org/10.1063/1.4955109>.
- [8] J. B. Schriber and F. A. Evangelista, Adaptive configuration interaction for computing challenging electronic excited states with tunable accuracy, *Journal of Chemical Theory and Computation* **13**, 5354 (2017), pMID: 28892621, <https://doi.org/10.1021/acs.jctc.7b00725>.
- [9] N. M. Tubman, C. D. Freeman, D. S. Levine, D. Hait, M. Head-Gordon, and K. B. Whaley, Modern approaches to exact diagonalization and selected configuration interaction with the adaptive sampling ci method, *Journal of Chemical Theory and Computation* **16**, 2139 (2020).
- [10] C. Mejuto-Zaera, N. M. Tubman, and K. B. Whaley, Dynamical mean field theory simulations with the adaptive sampling configuration interaction method, *Phys. Rev. B* **100**, 125165 (2019).
- [11] J. Olsen, P. Jørgensen, and J. Simons, Passing the one-billion limit in full configuration-interaction (fci) calculations, *Chemical Physics Letters* **169**, 463 (1990).
- [12] E. Solomonik, D. Matthews, J. R. Hammond, J. F. Stanton, and J. Demmel, A massively parallel tensor contraction framework for coupled-cluster computations, *Journal of Parallel and Distributed Computing* **74**, 3176 (2014).
- [13] J. Kim, A. D. Baczewski, T. D. Beaudet, A. Benali, M. C. Bennett, M. A. Berrill, N. S. Blunt, E. J. L. Borda, M. Casula, D. M. Ceperley, S. Chiesa, B. K. Clark, R. C. Clay, K. T. Delaney, M. Dewing, K. P. Esler, H. Hao, O. Heinonen, P. R. C. Kent, J. T. Krogel, I. Kylänpää, Y. W. Li, M. G. Lopez, Y. Luo, F. D. Malone, R. M. Martin, A. Mathuriya, J. McMinis, C. A. Melton, L. Mitas, M. A. Morales, E. Neuscamman, W. D. Parker, S. D. P. Flores, N. A. Romero, B. M. Rubenstein, J. A. R. Shea, H. Shin, L. Shulenburger, A. F. Tillack, J. P. Townsend, N. M. Tubman, B. V. D. Goetz, J. E. Vincent, D. C. Yang, Y. Yang, S. Zhang, and L. Zhao, Qmcpack: an open source ab initio quantum monte carlo package for the electronic structure of atoms, molecules and solids, *Journal of Physics: Condensed Matter* **30**, 195901 (2018).
- [14] D. B. Williams-Young, N. M. Tubman, C. Mejuto-Zaera, and W. A. de Jong, A parallel, distributed memory implementation of the adaptive sampling configuration interaction method, *The Journal of Chemical Physics* **158**, 214109 (2023), https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0148650/17914342/214109_1.5.0148650.pdf.
- [15] M. Valiev, E. Bylaska, N. Govind, K. Kowalski, T. Straatsma, H. Van Dam, D. Wang, J. Nieplocha, E. Apra, T. Windus, and W. de Jong, Nwchem: A comprehensive and scalable open-source solution for large scale molecular simulations, *Computer Physics Communications* **181**, 1477 (2010).
- [16] K. Foyevtsova, J. T. Krogel, J. Kim, P. R. C. Kent, E. Dagotto, and F. A. Reboredo, Ab initio quantum monte carlo calculations of spin superexchange in cuprates: The benchmarking case of Ca_2CuO_3 , *Phys. Rev. X* **4**, 031003 (2014).
- [17] M. Holzmann, R. C. Clay, M. A. Morales, N. M. Tubman, D. M. Ceperley, and C. Pierleoni, Theory of finite size effects for electronic quantum monte carlo calculations of liquids and solids, *Phys. Rev. B* **94**, 035126 (2016).
- [18] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi, Towards an exact description of electronic wavefunctions in real solids, *Nature* **493**, 365 (2013).
- [19] J. J. Eriksen, T. A. Anderson, J. E. Deustua, K. Ghanem, D. Hait, M. R. Hoffmann, S. Lee, D. S. Levine, I. Magoulas, J. Shen, N. M. Tubman, K. B. Whaley, E. Xu, Y. Yao, N. Zhang, A. Alavi, G. K.-L. Chan, M. Head-Gordon, W. Liu, P. Piecuch, S. Sharma, S. L. Ten-no, C. J. Umrigar, and J. Gauss, The ground state electronic energy of benzene, *The Journal of Physical Chemistry Letters* **11**, 8922 (2020), pMID: 33022176, <https://doi.org/10.1021/acs.jpcclett.0c02621>.
- [20] C.-C. Chang, B. M. Rubenstein, and M. A. Morales, Auxiliary-field-based trial wave functions in quantum monte carlo calculations, *Phys. Rev. B* **94**, 235144 (2016).
- [21] V. c. v. Vlček, W. Li, R. Baer, E. Rabani, and D. Neuhauser, Swift *gw* beyond 10,000 electrons using sparse stochastic compression, *Phys. Rev. B* **98**, 075107 (2018).
- [22] W. A. De Jong, E. Bylaska, N. Govind, C. L. Janssen, K. Kowalski, T. Müller, I. M. Nielsen, H. J. van Dam, V. Veryazov, and R. Lindh, Utilizing high performance computing for chemistry: parallel computational chemistry, *Physical Chemistry Chemical Physics* **12**, 6896 (2010).
- [23] N. M. Tubman, J. L. DuBois, R. Q. Hood, and B. J. Alder, Prospects for release-node quantum Monte Carlo, *The Journal of Chemical Physics* **135**, 184109 (2011).
- [24] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [25] S. Lloyd, Universal quantum simulators, *Science* , 1073 (1996).
- [26] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [27] J. Haah, M. Hastings, R. Kothari, and G. H. Low, Quantum algorithm for simulating real time evolution of lattice hamiltonians, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (2018) pp. 350–360.
- [28] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, Locality and digital quantum simulation of power-law interactions, *Phys. Rev. X* **9**, 031006 (2019).
- [29] A. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*, Graduate Studies in Mathematics, Vol. 47 (American Mathematical Society, Providence, RI, 2002).
- [30] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local hamiltonian problem, *SIAM Journal on Computing* **35**, 1070 (2006).
- [31] J. Du, N. Xu, X. Peng, P. Wang, S. Wu, and D. Lu, Nmr implementation of a molecular hydrogen quantum simulation with adiabatic state preparation, *Phys. Rev. Lett.* **104**, 030502 (2010).
- [32] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, A. Aspuru-Guzik, and A. G. White, Towards quantum chemistry on a quantum computer, *Nature Chemistry* **2**, 106 (2010).
- [33] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum

- processor, *Nature communications* **5**, 4213 (2014).
- [34] Y. Wang, F. Dolde, J. Biamonte, R. Babbush, V. Bergholm, S. Yang, I. Jakobi, P. Neumann, A. Aspuru-Guzik, J. D. Whitfield, and J. Wrachtrup, Quantum simulation of helium hydride cation in a solid-state spin register, *ACS Nano* **9**, 7769 (2015).
- [35] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, Scalable quantum simulation of molecular energies, *Phys. Rev. X* **6**, 031007 (2016).
- [36] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure, *Physical Review A* **95**, 020501 (2017).
- [37] S. Paesani, A. A. Gentile, R. Santagati, J. Wang, N. Wiebe, D. P. Tew, J. L. O'Brien, and M. G. Thompson, Experimental bayesian quantum phase estimation on a silicon photonic chip, *Phys. Rev. Lett.* **118**, 100503 (2017).
- [38] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *nature* **549**, 242 (2017).
- [39] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum chemistry calculations on a trapped-ion quantum simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [40] R. Santagati, J. Wang, A. A. Gentile, S. Paesani, N. Wiebe, J. R. McClean, S. Morley-Short, P. J. Shadbolt, D. Bonneau, J. W. Silverstone, D. P. Tew, X. Zhou, J. L. O'Brien, and M. G. Thompson, Witnessing eigenstates for quantum simulation of hamiltonian spectra, *Science Advances* **4**, 10.1126/sciadv.aap9646 (2018), <https://advances.sciencemag.org/content/4/1/eaap9646.full>.
- [41] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, Cloud quantum computing of an atomic nucleus, *Phys. Rev. Lett.* **120**, 210501 (2018).
- [42] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, Self-verifying variational quantum simulation of lattice models, *Nature* **569**, 355 (2019).
- [43] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature* **567**, 491 (2019).
- [44] M. Ganzhorn, D. J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, *et al.*, Gate-efficient simulation of molecular eigenstates on a quantum computer, *Physical Review Applied* **11**, 044092 (2019).
- [45] R. Sagastizabal, X. Bonet-Monroig, M. Singh, M. A. Rol, C. Bultink, X. Fu, C. Price, V. Ostroukh, N. Muthusubramanian, A. Bruno, *et al.*, Experimental error mitigation via symmetry verification in a variational quantum eigensolver, *Physical Review A* **100**, 010302 (2019).
- [46] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, and R. C. Pooser, Quantum chemistry as a benchmark for near-term quantum computers, *npj Quantum Information* **5**, 1 (2019).
- [47] S. E. Smart and D. A. Mazziotti, Quantum-classical hybrid algorithm using an error-mitigating n-representability condition to compute the mott metal-insulator transition, *Physical Review A* **100**, 022517 (2019).
- [48] Y. Nam, J.-S. Chen, N. C. Pienti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, K. M. Hudek, A. M. Ducore, M. Keenan, S. M. Kreikemeier, J. Mizrahi, P. Solomon, M. Williams, J. D. Wong-Campos, D. Moehring, C. Monroe, and J. Kim, Ground-state energy estimation of the water molecule on a trapped-ion quantum computer, *npj Quantum Information* **6**, 33 (2020).
- [49] J. Liu, F. Tacchino, J. R. Glick, L. Jiang, and A. Mezzacapo, Representation learning via quantum neural tangent kernels, *PRX Quantum* **3**, 030323 (2022).
- [50] Q. Wang, M. Li, C. Monroe, and Y. Nam, Resource-optimized fermionic local-hamiltonian simulation on a quantum computer for quantum chemistry, *Quantum* **5**, 509 (2021).
- [51] J. Liu, K. Najafi, K. Sharma, F. Tacchino, L. Jiang, and A. Mezzacapo, Analytic theory for the dynamics of wide quantum neural networks, *Physical Review Letters* **130**, 150601 (2023).
- [52] H. Zheng, Z. Li, J. Liu, S. Strelchuk, and R. Kondor, Speeding up learning quantum states through group equivariant convolutional quantum ansätze, *PRX Quantum* **4**, 020327 (2023).
- [53] Q. Wang, Z.-P. Cian, M. Li, I. L. Markov, and Y. Nam, Ever more optimized simulations of fermionic systems on a quantum computer, *arXiv preprint arXiv:2303.03460* (2023).
- [54] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe, Quantum approximate optimization of the long-range ising model with a trapped-ion quantum simulator, *Proceedings of the National Academy of Sciences* **117**, 25396 (2020), <https://www.pnas.org/doi/pdf/10.1073/pnas.2006373117>.
- [55] P. C. Lotshaw, H. Xu, B. Khalid, G. Buchs, T. S. Humble, and A. Banerjee, Simulations of frustrated ising hamiltonians using quantum approximate optimization, *Philosophical Transactions of the Royal Society A* **381**, 20210414 (2023).
- [56] A. Y. Kitaev, Quantum measurements and the abelian stabilizer problem, *arxiv preprint, arXiv:quant-ph/9511026* (1995).
- [57] L. Lin and Y. Tong, Near-optimal ground state preparation, *Quantum* **4**, 372 (2020).
- [58] L. Lin and Y. Tong, Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers, *PRX Quantum* **3**, 010318 (2022).
- [59] Y. Dong, L. Lin, and Y. Tong, Ground-state prepara-

- tion and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices, *PRX Quantum* **3**, 040305 (2022).
- [60] Z. Ding, C.-F. Chen, and L. Lin, Single-ancilla ground state preparation via lindbladians (2023), arXiv:2308.15676 [quant-ph].
- [61] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states, *Phys. Rev. A* **95**, 042308 (2017).
- [62] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of molecular spectra on a quantum processor with an error-resilient algorithm, *Phys. Rev. X* **8**, 011021 (2018).
- [63] R. M. Parrish and P. L. McMahon, Quantum filter diagonalization: Quantum eigendecomposition without full quantum phase estimation, arXiv 10.48550/arXiv.1909.08925 (2019), arXiv:1909.08925 [quant-ph].
- [64] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nature Physics* **16**, 205 (2020).
- [65] T. Takeshita, N. C. Rubin, Z. Jiang, E. Lee, R. Babush, and J. R. McClean, Increasing the representation accuracy of quantum simulations of chemistry without extra quantum resources, *Phys. Rev. X* **10**, 011004 (2020).
- [66] W. J. Huggins, J. Lee, U. Baek, B. O'Gorman, and K. B. Whaley, A non-orthogonal variational quantum eigensolver, *New Journal of Physics* **22**, 073009 (2020).
- [67] N. H. Stair, R. Huang, and F. A. Evangelista, A multireference quantum krylov algorithm for strongly correlated electrons, *Journal of Chemical Theory and Computation*, *Journal of Chemical Theory and Computation* **16**, 2236 (2020).
- [68] M. Urbanek, D. Camps, R. Van Beeumen, and W. A. de Jong, Chemistry on quantum computers with virtual quantum subspace expansion, *Journal of Chemical Theory and Computation* **16**, 5425 (2020).
- [69] J. Cohn, M. Motta, and R. M. Parrish, Quantum filter diagonalization with compressed double-factorized hamiltonians, *PRX Quantum* **2**, 040352 (2021).
- [70] N. Yoshioka, H. Hakoshima, Y. Matsuzaki, Y. Tokunaga, Y. Suzuki, and S. Endo, Generalized quantum subspace expansion, *Phys. Rev. Lett.* **129**, 020502 (2022).
- [71] K. Seki and S. Yunoki, Quantum power method by a superposition of time-evolved states, *PRX Quantum* **2**, 010333 (2021).
- [72] C. L. Cortes and S. K. Gray, Quantum krylov subspace algorithms for ground- and excited-state energy estimation, *Phys. Rev. A* **105**, 022417 (2022).
- [73] K. Klymko, C. Mejuto-Zaera, S. J. Cotton, F. Wudarski, M. Urbanek, D. Hait, M. Head-Gordon, K. B. Whaley, J. Moussa, N. Wiebe, W. A. de Jong, and N. M. Tubman, Real-time evolution for ultracompact hamiltonian eigenstates on quantum hardware, *PRX Quantum* **3**, 020323 (2022).
- [74] U. Baek, D. Hait, J. Shee, O. Leimkuhler, W. J. Huggins, T. F. Stetina, M. Head-Gordon, and K. B. Whaley, Say no to optimization: A nonorthogonal quantum eigensolver, *PRX Quantum* **4**, 030307 (2023).
- [75] N. V. Tkachenko, Y. Zhang, L. Cincio, A. I. Boldyrev, S. Tretiak, and P. A. Dub, Quantum davidson algorithm for excited states, arXiv 10.48550/arXiv.2204.10741 (2022), arXiv:2204.10741 [quant-ph].
- [76] G. Lee, D. Lee, and J. Huh, Sampling error analysis in quantum krylov subspace diagonalization, arXiv 10.48550/arXiv.2307.16279 (2023), arXiv:2307.16279 [quant-ph].
- [77] Z. Zhang, A. Wang, X. Xu, and Y. Li, Measurement-efficient quantum krylov subspace diagonalisation, arXiv 10.48550/arXiv.2301.13353 (2023), arXiv:2301.13353 [quant-ph].
- [78] W. Kirby, M. Motta, and A. Mezzacapo, Exact and efficient Lanczos method on a quantum computer, *Quantum* **7**, 1018 (2023).
- [79] Y. Shen, K. Klymko, J. Sud, D. B. Williams-Young, W. A. de Jong, and N. M. Tubman, Real-Time Krylov Theory for Quantum Computing Algorithms, *Quantum* **7**, 1066 (2023).
- [80] W. Kirby, Analysis of quantum krylov algorithms with errors, arXiv:2401.01246 (2024).
- [81] M. Motta, W. Kirby, I. Liepuoniute, K. J. Sung, J. Cohn, A. Mezzacapo, K. Klymko, N. Nguyen, N. Yoshioka, and J. E. Rice, Subspace methods for electronic structure simulations on quantum computers 10.48550/arXiv.2312.00178 (2023), arXiv:2312.00178 [quant-ph].
- [82] Y. Shen, D. Camps, A. Szasz, S. Darbha, K. Klymko, D. B. Williams-Young, N. M. Tubman, and R. V. Beeumen, Estimating eigenenergies from quantum dynamics: A unified noise-resilient measurement-driven approach, arXiv 10.48550/arXiv.2306.01858 (2023), arXiv:2306.01858 [quant-ph].
- [83] E. N. Epperly, L. Lin, and Y. Nakatsukasa, A theory of quantum subspace diagonalization, *SIAM Journal on Matrix Analysis and Applications* **43**, 1263 (2022).
- [84] B. Olmos, I. Lesanovsky, and J. P. Garrahan, Facilitated spin models of dissipative quantum glasses, *Physical Review Letters* **109**, 020403 (2012).
- [85] V. May and O. Kühn, *Charge and energy transfer dynamics in molecular systems* (John Wiley & Sons, 2023).
- [86] A. Nitzan, *Chemical dynamics in condensed phases: relaxation, transfer and reactions in condensed molecular systems* (Oxford university press, 2006).
- [87] M. J. Kastoryano, F. G. Brandão, A. Gilyén, *et al.*, Quantum thermal state preparation, arXiv preprint arXiv:2303.18224 (2023).
- [88] Z. Ding, L. Lin, *et al.*, Single-ancilla ground state preparation via lindbladians, arXiv preprint arXiv:2308.15676 (2023).
- [89] T. S. Cubitt, Dissipative ground state preparation and the dissipative quantum eigensolver, arXiv preprint arXiv:2303.11962 (2023).
- [90] J. Hubisz, B. Sambasivam, and J. Unmuth-Yockey, Quantum algorithms for open lattice field theory, *Physical Review A* **104**, 052420 (2021).
- [91] A. W. Schlimgen, K. Head-Marsden, L. M. Sager, P. Narang, and D. A. Mazziotti, Quantum simulation of the lindblad equation using a unitary decomposition of operators, *Physical Review Research* **4**, 023216 (2022).
- [92] G. Lindblad, On the generators of quantum dynamical semigroups, *Communications in Mathematical Physics*

- 48**, 119 (1976).
- [93] H. Wang, S. Ashhab, and F. Nori, Quantum algorithm for simulating the dynamics of an open quantum system, *Physical Review A* **83**, 062317 (2011).
- [94] M. Kliesch, T. Barthel, C. Gogolin, M. Kastoryano, and J. Eisert, Dissipative quantum church-turing theorem, *Physical review letters* **107**, 120501 (2011).
- [95] Z. Ding, X. Li, and L. Lin, Simulating open quantum systems using hamiltonian simulations, arXiv preprint arXiv:2311.15533 (2023).
- [96] R. Sweke, M. Sanz, I. Sinayskiy, F. Petruccione, and E. Solano, Digital quantum simulation of many-body non-markovian dynamics, *Physical Review A* **94**, 022317 (2016).
- [97] A. M. Childs and T. Li, Efficient simulation of sparse markovian quantum dynamics, arXiv preprint arXiv:1611.05543 (2016).
- [98] R. Cleve and C. Wang, Efficient quantum algorithms for simulating lindblad evolution, arXiv preprint arXiv:1612.09512 (2016).
- [99] H. Kamakari, S.-N. Sun, M. Motta, and A. J. Minnich, Digital quantum simulation of open quantum systems using quantum imaginary-time evolution, *PRX Quantum* **3**, 010320 (2022).
- [100] M. Pocrnic, D. Segal, and N. Wiebe, Quantum simulation of lindbladian dynamics via repeated interactions, arXiv preprint arXiv:2312.05371 (2023).
- [101] D. Patel and M. M. Wilde, Wave matrix lindbladization i: Quantum programs for simulating markovian dynamics, *Open Systems & Information Dynamics* **30**, 2350010 (2023).
- [102] D. Patel and M. M. Wilde, Wave matrix lindbladization ii: General lindbladians, linear combinations, and polynomials, *Open Systems & Information Dynamics* **30**, 2350014 (2023).
- [103] J. D. Guimarães, J. Lim, M. I. Vasilevskiy, S. F. Huelga, and M. B. Plenio, Noise-assisted digital quantum simulation of open systems using partial probabilistic error cancellation, *PRX Quantum* **4**, 040329 (2023).
- [104] D. Poulin and P. Wocjan, Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer, *Physical Review Letters* **103**, 10.1103/physrevlett.103.220502 (2009).
- [105] A. N. Chowdhury and R. D. Somma, Quantum algorithms for gibbs sampling and hitting-time estimation (2016), arXiv:1603.02940 [quant-ph].
- [106] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman, M. Greiner, H. Pichler, and P. Hayden, Quantum virtual cooling, *Physical Review X* **9**, 10.1103/physrevx.9.031013 (2019).
- [107] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, Quantum metropolis sampling, *Nature* **471**, 87–90 (2011).
- [108] M.-H. Yung and A. Aspuru-Guzik, A quantum-quantum metropolis algorithm, *Proceedings of the National Academy of Sciences* **109**, 754–759 (2012).
- [109] P. Wocjan and K. Temme, Szegedy walk unitaries for quantum maps, *Communications in Mathematical Physics* **402**, 3201–3231 (2023).
- [110] E. Davies, *Quantum Theory of Open Systems* (Academic Press, 1976).
- [111] P. Rall, C. Wang, and P. Wocjan, Thermal state preparation via rounding promises, *Quantum* **7**, 1132 (2023).
- [112] C.-F. Chen, M. J. Kastoryano, F. G. S. L. Brandão, and A. Gilyén, Quantum thermal state preparation (2023), arXiv:2303.18224 [quant-ph].
- [113] C.-F. Chen, M. J. Kastoryano, and A. Gilyén, An efficient and exact noncommutative quantum gibbs sampler (2023), arXiv:2311.09207 [quant-ph].
- [114] C.-F. Chen, H.-Y. Huang, J. Preskill, and L. Zhou, Local minima in quantum systems (2023), arXiv:2309.16596 [quant-ph].
- [115] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Physical review letters* **103**, 150502 (2009).
- [116] Copyright, in *The Finite Element Method in Engineering (Fifth Edition)*, edited by S. S. Rao (Butterworth-Heinemann, Boston, 2011) fifth edition ed., p. iv.
- [117] A. Montanaro and S. Pallister, Quantum algorithms and the finite element method, *Phys. Rev. A* **93**, 032324 (2016).
- [118] N. Linden, A. Montanaro, and C. Shao, Quantum vs. classical algorithms for solving the heat equation, *Communications in Mathematical Physics* **395**, 601 (2022).
- [119] J. Zhang, F. Feng, and Q. J. Zhang, Quantum method for finite element simulation of electromagnetic problems, in *2021 IEEE MTT-S International Microwave Symposium (IMS)* (2021) pp. 120–123.
- [120] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM Journal on Computing* **46**, 1920 (2017), <https://doi.org/10.1137/16M1087072>.
- [121] D. Fang, L. Lin, and Y. Tong, Time-marching based quantum solvers for time-dependent linear differential equations, *Quantum* **7**, 955 (2023).
- [122] H. Krovi, Improved quantum algorithms for linear and nonlinear differential equations, *Quantum* **7**, 913 (2023).
- [123] S. Jin, N. Liu, and Y. Yu, Quantum simulation of partial differential equations: Applications and detailed analysis, *Phys. Rev. A* **108**, 032603 (2023).
- [124] D. An, J.-P. Liu, D. Wang, and Q. Zhao, A theory of quantum differential equation solvers: limitations and fast-forwarding (2023), arXiv:2211.05246 [quant-ph].
- [125] D. Lewis, S. Eidenbenz, B. Nadiga, and Y. Subaşı, Limitations for quantum algorithms to solve turbulent and chaotic systems (2023), arXiv:2307.09593 [quant-ph].
- [126] L.-A. Wu and D. Lidar, Qubits as parafermions, *Journal of Mathematical Physics* **43**, 4506 (2002).
- [127] R. D. Somma, Quantum computation, complexity, and many-body physics, arXiv preprint arXiv:quant-ph/0512209 (2005).
- [128] A. Macridin, P. Spentzouris, J. Amundson, and R. Harnik, Electron-phonon systems on a universal quantum computer, *Phys. Rev. Lett.* **121**, 110504 (2018).
- [129] S. McArdle, A. Mayorov, X. Shan, S. Benjamin, and X. Yuan, Digital quantum simulation of molecular vibrations, *Chem. Sci.* **10**, 5725 (2019).
- [130] N. P. D. Sawaya and J. Huh, Quantum algorithm for calculating molecular vibronic spectra, *The Journal of Physical Chemistry Letters* **10**, 3586 (2019).
- [131] A. Kan and Y. Nam, Lattice quantum chromodynamics and electrostatics on a universal quantum computer, arXiv preprint arXiv:2107.12769 (2021).
- [132] M. Kreshchuk, W. M. Kirby, G. Goldstein, H. Beau-

- chemin, and P. J. Love, Quantum simulation of quantum field theory in the light-front formulation, *Phys. Rev. A* **105**, 032418 (2022).
- [133] A. Kan and Y. Nam, Simulating lattice quantum electrodynamics on a quantum computer, *Quantum Science and Technology* **8**, 015008 (2022).
- [134] P. Jordan and E. Wigner, Über das paulische äquivalenzverbot., *Z. Phys.* **47**, 631 (1928).
- [135] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, *Annals of Physics* **298**, 210 (2002).
- [136] F. Verstraete and J. I. Cirac, Mapping local hamiltonians of fermions to local hamiltonians of spins, *Journal of Statistical Mechanics: Theory and Experiment* **2005**, P09012 (2005).
- [137] J. T. Seeley, M. J. Richard, and P. J. Love, The bravyi-kitaev transformation for quantum computation of electronic structure, *The Journal of Chemical Physics* **137**, 224109 (2012).
- [138] S. Bravyi, J. M. Gambetta, A. Mezzacapo, and K. Temme, Tapering off qubits to simulate fermionic hamiltonians, *arXiv preprint*, arXiv:1701.08213 (2017), arXiv:1701.08213 [quant-ph].
- [139] K. Setia and J. D. Whitfield, Bravyi-Kitaev Superfast simulation of electronic structure on a quantum computer, *The Journal of Chemical Physics* **148**, 164104 (2018), https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.5019371/13715475/164104_1_online.pdf.
- [140] M. Steudtner and S. Wehner, Fermion-to-qubit mappings with varying resource requirements for quantum simulation, *New Journal of Physics* **20**, 063010 (2018).
- [141] M. Steudtner, *Methods to simulate fermions on quantum computers with hardware limitations*, Ph.D. thesis, Leiden University (2019).
- [142] C. Derby, J. Klassen, J. Bausch, and T. Cubitt, Compact fermion to qubit mappings, *Phys. Rev. B* **104**, 035118 (2021).
- [143] C. Derby and J. Klassen, A compact fermion to qubit mapping part 2: Alternative lattice geometries, *arXiv preprint*, arXiv:2101.10735 (2021), arXiv:2101.10735 [quant-ph].
- [144] W. Kirby, B. Fuller, C. Hadfield, and A. Mezzacapo, Second-quantized fermionic operators with polylogarithmic qubit and gate complexity, *PRX Quantum* **3**, 020351 (2022).
- [145] R. W. Chien and J. Klassen, Optimizing fermionic encodings for both hamiltonian and hardware, *arXiv preprint*, arXiv:2210.05652 10.48550/arXiv.2210.05652 (2021), arXiv:2210.05652 [quant-ph].
- [146] Y. Nam, N. J. Ross, Y. Su, A. M. Childs, and D. Maslov, Automated optimization of large quantum circuits with continuous parameters, *npj Quantum Information* **4**, 23 (2018).
- [147] A. Miller, Z. Zimborás, S. Knecht, S. Maniscalco, and G. García-Pérez, Bonsai algorithm: Grow your own fermion-to-qubit mappings, *PRX Quantum* **4**, 030314 (2023).
- [148] Y.-A. Chen and Y. Xu, Equivalence between fermion-to-qubit mappings in two spatial dimensions, *PRX Quantum* **4**, 010326 (2023).
- [149] R. P. Feynman, *Simulating physics with computers*, in *Feynman and computation* (CRC Press, 2018) pp. 133–153.
- [150] L. A. Martínez-Martínez, T.-C. Yen, and A. F. Izmaylov, Assessment of various hamiltonian partitionings for the electronic structure problem on a quantum computer using the trotter approximation, *Quantum* **7**, 1086 (2023).
- [151] Z. Shangnan, Quantum semi-supervised learning with quantum supremacy, *arXiv preprint* arXiv:2110.02343 (2021).
- [152] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proceedings of the National Academy of Sciences* **115**, 9456 (2018).
- [153] Y. Nam and D. Maslov, Low-cost quantum circuits for classically intractable instances of the hamiltonian dynamics simulation problem, *npj Quantum Information* **5**, 44 (2019).
- [154] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Efficient quantum algorithms for simulating sparse hamiltonians, *Communications in Mathematical Physics* **270**, 359 (2007).
- [155] A. Tranter, P. J. Love, F. Mintert, N. Wiebe, and P. V. Coveney, Ordering of trotterization: Impact on errors in quantum simulation of electronic structure, *Entropy* **21**, 10.3390/e21121218 (2019).
- [156] D. W. Berry and A. M. Childs, Black-box hamiltonian simulation and unitary implementation, *Quantum Info. Comput.* **12**, 29 (2012).
- [157] D. W. Berry, A. M. Childs, and R. Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters, in *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on* (IEEE, 2015) pp. 792–809.
- [158] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating hamiltonian dynamics with a truncated taylor series, *Physical review letters* **114**, 090502 (2015).
- [159] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization, *Quantum* **3**, 208 (2019).
- [160] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Info. Comput.* **12**, 901–924 (2012).
- [161] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX Quantum* **2**, 030305 (2021).
- [162] I. Loaiza, A. Marefat Khah, N. Wiebe, and A. F. Izmaylov, Reducing molecular electronic hamiltonian simulation cost for linear combination of unitaries approaches, *Quantum Sci. Technol.* **8**, 035019 (2023).
- [163] I. Loaiza and A. F. Izmaylov, Block-invariant symmetry shift: Preprocessing technique for second-quantized hamiltonians to improve their decompositions to linear combination of unitaries, arXiv:2304.13772 (2023).
- [164] I. Loaiza and A. F. Izmaylov, Majorana tensor decomposition: A unifying framework for decompositions of fermionic hamiltonians to linear combination of unitaries, arXiv:2311.XXX (2023).
- [165] S. Choi and A. F. Izmaylov, Measurement optimization techniques for excited electronic states in near-term quantum computing algorithms, *J. Chem. Theory Comput.* (arXiv:2302.11421) **19**, 3184 (2023).
- [166] G. Wang, D. E. Koh, P. D. Johnson, and Y. Cao, Minimizing estimation runtime on noisy quantum comput-

- ers, *PRX Quantum* **2**, 010346 (2021).
- [167] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [168] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [169] B. O’Gorman, Fermionic tomography and learning (2022), [arXiv:2207.14787 \[quant-ph\]](#).
- [170] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nature Physics* **16**, 1050 (2020).
- [171] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient estimation of Pauli observables by derandomization, *Phys. Rev. Lett.* **127**, 030503 (2021).
- [172] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum hamiltonians with locally-biased classical shadows, *Comm. Math. Phys.* **391**, 951 (2022).
- [173] C. Hadfield, Adaptive Pauli shadows for energy estimation, [arXiv:2105.12207 \(2021\)](#), 2105.12207.
- [174] J. M. Lukens, K. J. Law, and R. S. Bennink, A Bayesian analysis of classical shadows, *npj Quantum Information* **7**, 1 (2021).
- [175] A. Shlosberg, A. J. Jena, P. Mukhopadhyay, J. F. Haase, F. Leditzky, and L. Dellantonio, Adaptive estimation of quantum observables, *Quantum* **7**, 906 (2023).
- [176] D. E. Koh and S. Grewal, Classical shadows with noise, *Quantum* **6**, 776 (2022).
- [177] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, *Nature Reviews Physics* **5**, 9 (2023).
- [178] S. Hillmich, C. Hadfield, R. Raymond, A. Mezzacapo, and R. Wille, Decision diagrams for quantum measurements with shallow circuits, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE Computer Society, Los Alamitos, CA, USA, 2021) pp. 24–34.
- [179] A. Dutt, W. Kirby, R. Raymond, C. Hadfield, S. Sheldon, I. L. Chuang, and A. Mezzacapo, Practical benchmarking of randomized measurement methods for quantum chemistry hamiltonians, [arXiv 10.48550/arXiv.2312.07497 \(2023\)](#), [arXiv:2312.07497 \[quant-ph\]](#).
- [180] G. García-Pérez, M. A. Rossi, B. Sokolov, F. Tacchino, P. K. Barkoutsos, G. Mazzola, I. Tavernelli, and S. Maniscalco, Learning to measure: Adaptive informationally complete generalized measurements for quantum algorithms, *PRX Quantum* **2**, 040342 (2021).
- [181] A. Acharya, S. Saha, and A. M. Sengupta, Shadow tomography based on informationally complete positive operator-valued measure, *Phys. Rev. A* **104**, 052418 (2021).
- [182] L. E. Fischer, D. Miller, F. Tacchino, P. K. Barkoutsos, D. J. Egger, and I. Tavernelli, Ancilla-free implementation of generalized measurements for qubits embedded in a qudit space, *Phys. Rev. Res.* **4**, 033027 (2022).
- [183] A. Glos, A. Nykänen, E.-M. Borrelli, S. Maniscalco, M. A. C. Rossi, Z. Zimborás, and G. García-Pérez, Adaptive povm implementations and measurement error mitigation strategies for near-term quantum devices (2022), [arXiv:2208.07817 \[quant-ph\]](#).
- [184] S. Filippov, M. Leahy, M. A. C. Rossi, and G. García-Pérez, Scalable tensor-network error mitigation for near-term quantum computing (2023), [arXiv:2307.11740 \[quant-ph\]](#).
- [185] W. J. Huggins, J. R. McClean, N. C. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 1 (2021).
- [186] M. Motta, E. Ye, J. R. McClean, Z. Li, A. J. Minnich, R. Babbush, and G. K.-L. Chan, Low rank representations for quantum simulation of electronic structure, *npj Quantum Information* **7**, 83 (2021).
- [187] T.-C. Yen and A. F. Izmaylov, Cartan subalgebra approach to efficient measurements of quantum observables, *PRX Quantum* **2**, 040320 (2021).
- [188] S. Choi, I. Loaiza, and A. F. Izmaylov, Fluid fermionic fragments for optimizing quantum measurements of electronic hamiltonians in the variational quantum eigensolver, *Quantum* **7**, 889 (2023).
- [189] S. Choi, T.-C. Yen, and A. F. Izmaylov, Improving quantum measurements by introducing “ghost” pauli products, *J. Chem. Theory Comput.* ([arXiv:2208.06563](#)) **18**, 7394 (2022).
- [190] T.-C. Yen, A. Ganeshram, and A. F. Izmaylov, Deterministic improvements of quantum measurements with grouping of compatible operators, non-local transformations, and covariance estimates, *npj Quantum Inf* **9** (2023).
- [191] Z. Cai, Quantum error mitigation using symmetry expansion, *Quantum* **5**, 548 (2021).
- [192] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O’Brien, Quantum error mitigation (2023), [arXiv:2210.00921 \[quant-ph\]](#).
- [193] J. Cohn, M. Motta, and R. M. Parrish, Quantum filter diagonalization with compressed double-factorized hamiltonians, *PRX Quantum* **2**, 040352 (2021).
- [194] C. Cade, L. Mineh, A. Montanaro, and S. Stanisic, Strategies for solving the fermi-hubbard model on near-term quantum computers, *Phys. Rev. B* **102**, 235122 (2020).
- [195] Z. Cai, Resource estimation for quantum variational simulations of the hubbard model, *Phys. Rev. Appl.* **14**, 014059 (2020).
- [196] L. Clinton, T. Cubitt, B. Flynn, F. M. Gambetta, J. Klassen, A. Montanaro, S. Piddock, R. A. Santos, and E. Sheridan, *Towards near-term quantum simulation of materials* (2022), [arXiv:2205.15256 \[quant-ph\]](#).
- [197] G. Wolfowicz, F. J. Heremans, C. P. Anderson, S. Kanai, H. Seo, A. Gali, G. Galli, and D. D. Awschalom, Quantum guidelines for solid-state spin defects, *Nature Reviews Materials* **6**, 906 (2021).
- [198] T. P. Gujarati, M. Motta, T. N. Friedhoff, J. E. Rice, N. Nguyen, P. K. Barkoutsos, R. J. Thompson, T. Smith, M. Kagele, M. Brei, *et al.*, Quantum computation of reactions on surfaces using local embedding, *npj Quantum Information* **9**, 88 (2023).
- [199] A. Sarkar, M. R. Hermes, C. J. Cramer, J. S. Anderson, and L. Gagliardi, Understanding Antiferromagnetic and Ligand Field Effects on Spin Crossover in a Triple-Decker Dimeric Cr(II) Complex, *Journal of the American Chemical Society* **145**, 22394 (2023).
- [200] J. R. Maze, A. Gali, E. Togan, Y. Chu, A. Trifonov, E. Kaxiras, and M. D. Lukin, Properties of nitrogen-

- vacancy centers in diamond: the group theoretic approach, *New Journal of Physics* **13**, 025025 (2011).
- [201] A. Mitra, H. Q. Pham, R. Pandharkar, M. R. Hermes, and L. Gagliardi, Excited states of crystalline point defects with multireference density matrix embedding theory, *J. Phys. Chem. Lett.* **12**, 11688 (2021).
- [202] S. Haldar, A. Mitra, M. R. Hermes, and L. Gagliardi, Local excitations of a charged nitrogen vacancy in diamond with multireference density matrix embedding theory, *J. Phys. Chem. Lett.* **14**, 4273 (2023).
- [203] S. Verma, A. Mitra, Y. Jin, S. Haldar, C. Vorwerk, M. R. Hermes, G. Galli, and L. Gagliardi, Optical Properties of Neutral F Centers in Bulk MgO with Density Matrix Embedding, *Journal of Physical Chemistry Letters*, 7703 (2023).
- [204] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. Hollenberg, The nitrogen-vacancy colour centre in diamond, *Physics Reports* **528**, 1 (2013).
- [205] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Computational Materials* **6**, 85 (2020).
- [206] N. Sheng, C. Vorwerk, M. Govoni, and G. Galli, Green's function formulation of quantum defect embedding theory, *Journal of Chemical Theory and Computation* **18**, 3512 (2022).
- [207] A. Sauza-de la Vega, R. Pandharkar, G. D. Strocio, A. Sarkar, D. G. Truhlar, and L. Gagliardi, Multiconfiguration Pair-Density Functional Theory for Chromium(IV) Molecular Qubits, *JACS Au* **2**, 2029 (2022).
- [208] S. L. Bayliss, D. W. Laorenza, P. J. Mintun, B. D. Kovos, D. E. Freedman, and D. D. Awschalom, Optically addressable molecular spins for quantum information processing, *Science* **370**, 1309 (2020).
- [209] B. O. Roos, P. R. Taylor, and P. E. Siegbahn, A complete active space scf method (casscf) using a density matrix formulated super-ci approach, *Chem. Phys.* **48**, 157 (1980).
- [210] P. E. M. Siegbahn, J. Almlöf, A. Heiberg, and B. O. Roos, The complete active space scf (casscf) method in a newton-raphson formulation with application to the hno molecule, *J. Chem. Phys.* **74**, 2384 (1981), <https://doi.org/10.1063/1.441359>.
- [211] P. Siegbahn, A. Heiberg, B. Roos, and B. Levy, A comparison of the super-CI and the newton-raphson scheme in the complete active space SCF method, *Phys. Scr.* **21**, 323 (1980).
- [212] S. R. White, Density matrix formulation for quantum renormalization groups, *Physical Review Letters* **69**, 2863 (1992).
- [213] R. J. Buenker and S. D. Peyerimhoff, Individualized configuration selection in CI calculations with subsequent energy extrapolation, *Theoretica chimica acta* **35**, 33 (1974).
- [214] G. Knizia and G. K.-L. Chan, Density matrix embedding: A simple alternative to dynamical mean-field theory, *Phys. Rev. Lett.* **109**, 186404 (2012).
- [215] G. Knizia and G. K.-L. Chan, Density matrix embedding: A strong-coupling quantum embedding theory, *J. Chem. Theory Comput.* **9**, 1428 (2013), <https://doi.org/10.1021/ct301044e>.
- [216] S. Wouters, C. A. Jiménez-Hoyos, Q. Sun, and G. K.-L. Chan, A practical guide to density matrix embedding theory in quantum chemistry, *J. Chem. Theory Comput.* **12**, 2706 (2016), <https://doi.org/10.1021/acs.jctc.6b00316>.
- [217] S. Wouters, C. A. Jiménez-Hoyos, and G. K.L. Chan, Five years of density matrix embedding theory, in *Fragmentation* (John Wiley & Sons, Ltd, 2017) Chap. 8, pp. 227–243, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119129271.ch8>
- [218] H. Q. Pham, M. R. Hermes, and L. Gagliardi, Periodic electronic structure calculations with the density matrix embedding theory, *Journal of Chemical Theory and Computation* **16**, 130 (2019).
- [219] Z.-H. Cui, T. Zhu, and G. K.-L. Chan, Efficient implementation of ab initio quantum embedding in periodic systems: Density matrix embedding theory, *Journal of Chemical Theory and Computation* **16**, 119 (2020).
- [220] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J.-P. Malrieu, Introduction of n-electron valence states for multireference perturbation theory, *J. Chem. Phys.* **114**, 10252 (2001), <https://doi.org/10.1063/1.1361246>.
- [221] C. Angeli, S. Borini, M. Cestari, and R. Cimiraglia, A quasidegenerate formulation of the second order n-electron valence state perturbation theory approach, *J. Chem. Phys.* **121**, 4043 (2004), <https://doi.org/10.1063/1.1778711>.
- [222] N-electron valence state perturbation theory: A fast implementation of the strongly contracted variant, *Chem. Phys. Lett* **350**, 297 (2001).
- [223] A. Mitra, M. R. Hermes, and L. Gagliardi, Density Matrix Embedding Using Multiconfiguration Pair-Density Functional Theory, *Journal of Chemical Theory and Computation* **10.1021/acs.jctc.3c00247** (2023).
- [224] G. Li Manni, R. K. Carlson, S. Luo, D. Ma, J. Olsen, D. G. Truhlar, and L. Gagliardi, Multiconfiguration Pair-Density Functional Theory, *J. Chem. Theory Comput.* **10**, 3669 (2014).
- [225] L. Gagliardi, D. G. Truhlar, G. Li Manni, R. K. Carlson, C. E. Hoyer, and J. L. Bao, Multiconfiguration pair-density functional theory: A new way to treat strongly correlated systems, *Acc. Chem. Res.* **50**, 66 (2017), <https://doi.org/10.1021/acs.accounts.6b00471>.
- [226] C. Zhou, M. R. Hermes, D. Wu, J. J. Bao, R. Pandharkar, D. S. King, D. Zhang, T. R. Scott, A. O. Lykhin, L. Gagliardi, and D. G. Truhlar, Electronic structure of strongly correlated systems: recent developments in multiconfiguration pair-density functional theory and multiconfiguration nonclassical-energy functional theory, *Chemical Science* **13**, 7685 (2022).
- [227] S. Ghosh, C. J. Cramer, D. G. Truhlar, and L. Gagliardi, Generalized-active-space pair-density functional theory: an efficient method to study large, strongly correlated, conjugated systems, *Chemical Science* **8**, 2741–2750 (2017).
- [228] P. Sharma, V. Bernales, S. Knecht, D. G. Truhlar, and L. Gagliardi, Density matrix renormalization group pair-density functional theory (dmrg-pdft): singlet–triplet gaps in polyacenes and polyacetylenes, *Chemical Science* **10**, 1716–1723 (2019).
- [229] P. Sharma, D. G. Truhlar, and L. Gagliardi, Magnetic coupling in a tris-hydroxo-bridged chromium dimer occurs through ligand mediated superexchange in conjunction with through-space coupling, *Journal of the American Chemical Society* **142**, 16644–16650 (2020).

- [230] A. Mitra, M. R. Hermes, M. Cho, V. Agarawal, and L. Gagliardi, Periodic Density Matrix Embedding for CO Adsorption on the MgO(001) Surface, *J. Phys. Chem. Lett.* **13**, 7483 (2022).
- [231] M. R. Hermes and L. Gagliardi, Multiconfigurational Self-Consistent Field Theory with Density Matrix Embedding: The Localized Active Space Self-Consistent Field Method, *Journal of Chemical Theory and Computation* **15**, 972 (2019).
- [232] M. R. Hermes, R. Pandharkar, and L. Gagliardi, Variational Localized Active Space Self-Consistent Field Method, *J. Chem. Theory Comput.* **16**, 4923 (2020).
- [233] R. Pandharkar, M. R. Hermes, C. J. Cramer, and L. Gagliardi, Localized Active Space-State Interaction: a Multireference Method for Chemical Insight, *Journal of Chemical Theory and Computation* **18**, 6557 (2022).
- [234] M. Otten, M. R. Hermes, R. Pandharkar, Y. Alexeev, S. K. Gray, and L. Gagliardi, Localized Quantum Chemistry on Quantum Computers, *Journal of Chemical Theory and Computation* **18**, 7205 (2022).
- [235] R. D’Cunha, M. Otten, M. R. Hermes, L. Gagliardi, and S. K. Gray, State preparation in quantum algorithms for fragment-based quantum chemistry, *arXiv abs/2305.18110* (2023).
- [236] A. A. Holmes, N. M. Tubman, and C. J. Umrigar, Heat-Bath Configuration Interaction: An Efficient Selected Configuration Interaction Algorithm Inspired by Heat-Bath Sampling, *Journal of Chemical Theory and Computation* **12**, 3674 (2016).
- [237] C. J. Stein and M. Reiher, autoCAS: A Program for Fully Automated Multiconfigurational Calculations, *Journal of Computational Chemistry* **40**, 2216 (2019).
- [238] G. H. Booth, A. J. W. Thom, and A. Alavi, Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space, *The Journal of Chemical Physics* **131**, 54106 (2009).
- [239] J. Li, B. A. Jones, and S. Kais, Toward perturbation theory methods on a quantum computer, *Science Advances* **9**, eadg4576 (2023).
- [240] A. Anand, P. Schleich, S. Alperin-Lea, P. W. Jensen, S. Sim, M. Díaz-Tinoco, J. S. Kottmann, M. Degroote, A. F. Izmaylov, and A. Aspuru-Guzik, A quantum computing view on unitary coupled cluster theory, *Chemical Society Reviews* **51**, 1659 (2022).
- [241] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature Communications* **9**, 4812 (2018).
- [242] R. D’Cunha, T. D. Crawford, M. Motta, and J. E. Rice, Challenges in the use of quantum computing hardware-efficient ansätze in electronic structure theory, *The Journal of Physical Chemistry A* **127**, 3437 (2023).
- [243] I. G. Ryabinkin, T.-C. Yen, S. N. Genin, and A. F. Izmaylov, Qubit coupled-cluster method: A systematic approach to quantum chemistry on a quantum computer, *J. Chem Theory Comput.* **14**, 6317 (2018).
- [244] V. E. Elfving, M. Millaruelo, J. A. Gámez, and C. Gogolin, Simulating quantum chemistry in the seniority-zero space on qubit-based quantum computers, *Physical Review A* **103**, 032605 (2021).
- [245] M. Motta, K. J. Sung, K. B. Whaley, M. Head-Gordon, and J. Shee, Bridging physical intuition and hardware efficiency for correlated electronic states: the local unitary cluster jastrow ansatz for electronic structure, *Chem. Sci.* **14**, 11213 (2023).
- [246] H. L. Tang, V. Shkolnikov, G. S. Barron, H. R. Grimley, N. J. Mayhall, E. Barnes, and S. E. Economou, qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor, *PRX Quantum* **2**, 020310 (2021).
- [247] I. Ryabinkin, R. Lang, S. Genin, and A. Izmaylov, Iterative qubit coupled cluster approach with efficient screening of generators, *J. Chem. Theory Comput.* **16**, 1055 (2020).
- [248] B. Huang, N. Sheng, M. Govoni, and G. Galli, Quantum simulations of fermionic Hamiltonians with efficient encoding and ansatz schemes, *Journal of Chemical Theory and Computation* **19**, 1487 (2023).
- [249] A. Nykänen, M. A. C. Rossi, E.-M. Borrelli, S. Maniscalco, and G. García-Pérez, Mitigating the measurement overhead of adapt-vqe with optimised informationally complete generalised measurements (2023), *arXiv:2212.09719 [quant-ph]*.
- [250] G. S. Ravi, P. Gokhale, Y. Ding, W. M. Kirby, K. N. Smith, J. M. Baker, P. J. Love, H. Hoffmann, K. R. Brown, and F. T. Chong, Cafqa: A classical simulation bootstrap for variational quantum algorithms (2023), *arXiv:2202.12924 [quant-ph]*.
- [251] A. Khan, B. K. Clark, and N. M. Tubman, Pre-optimizing variational quantum eigensolvers with tensor networks, *arXiv e-prints*, *arXiv:2310.12965* (2023), *arXiv:2310.12965 [quant-ph]*.
- [252] J. W. Mullinax and N. M. Tubman, Large-scale sparse wavefunction circuit simulator for applications with the variational quantum eigensolver, *arXiv e-prints*, *arXiv:2301.05726* (2023), *arXiv:2301.05726 [quant-ph]*.
- [253] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum random access memory, *Physical review letters* **100**, 160501 (2008).
- [254] S. Aaronson, Read the fine print, *Nature Physics* **11**, 291 (2015).
- [255] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [256] P. Niroula and Y. Nam, A quantum algorithm for string matching, *npj Quantum Information* **7**, 37 (2021).
- [257] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
- [258] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding electronic spectra in quantum circuits with linear t complexity, *Physical Review X* **8**, 041015 (2018).
- [259] C. T. Hann, G. Lee, S. Girvin, and L. Jiang, Resilience of quantum random access memory to generic noise, *PRX Quantum* **2**, 020311 (2021).
- [260] P. Wittek, *Quantum machine learning: what quantum computing means to data mining* (Academic Press, 2014).
- [261] H. Jiang, Z.-J. M. Shen, and J. Liu, Quantum computing methods for supply chain management, in *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)* (IEEE, 2022) pp. 400–405.
- [262] J. Liu, H. Zheng, M. Hanada, K. Setia, and D. Wu, Quantum power flows: From theory to practice, *arXiv preprint arXiv:2211.05728* (2022).
- [263] J. Liu, C. T. Hann, and L. Jiang, Data centers with quantum random access memory and quantum net-

- works, *Physical Review A* **108**, 032610 (2023).
- [264] J. Liu and L. Jiang, Quantum data center: Perspectives, arXiv preprint arXiv:2309.06641 (2023).
- [265] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum private queries, *Physical review letters* **100**, 230502 (2008).
- [266] V. Giovannetti, L. Maccone, T. Morimae, and T. G. Rudolph, Efficient universal blind quantum computation, *Physical review letters* **111**, 230501 (2013).
- [267] D. Gottesman, T. Jennewein, and S. Croke, Longer-baseline telescopes using quantum repeaters, *Physical review letters* **109**, 070503 (2012).
- [268] O. Di Matteo, V. Gheorghiu, and M. Mosca, Fault-tolerant resource estimation of quantum random-access memories, *IEEE Transactions on Quantum Engineering* **1**, 1 (2020).
- [269] V. Giovannetti, S. Lloyd, and L. Maccone, Architectures for a quantum random access memory, *Physical Review A* **78**, 052310 (2008).
- [270] N. Jiang, Y.-F. Pu, W. Chang, C. Li, S. Zhang, and L.-M. Duan, Experimental realization of 105-qubit random access quantum memory, *npj Quantum Information* **5**, 1 (2019).
- [271] C. T. Hann, C.-L. Zou, Y. Zhang, Y. Chu, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, Hardware-efficient quantum random access memory with hybrid quantum acoustic systems, *Physical review letters* **123**, 250501 (2019).
- [272] C. Hann, *Practicality of Quantum Random Access Memory*, Ph.D. thesis, Yale University (2021).
- [273] K. C. Chen, W. Dai, C. Errando-Herranz, S. Lloyd, and D. Englund, Scalable and high-fidelity quantum random access memory in spin-photon networks, *PRX Quantum* **2**, 030319 (2021).
- [274] S. Xu, C. T. Hann, B. Foxman, S. M. Girvin, and Y. Ding, Systems architecture for quantum random access memory, arXiv preprint arXiv:2306.03242 (2023).
- [275] Y. Wang, Y. Alexeev, L. Jiang, F. T. Chong, and J. Liu, Fundamental causal bounds of quantum random access memories, arXiv preprint arXiv:2307.13460 (2023).
- [276] S. Jaques and A. G. Rattew, Qram: A survey and critique, arXiv preprint arXiv:2305.10310 (2023).
- [277] A. F. Shaw, P. Lougovski, J. R. Stryker, and N. Wiebe, Quantum algorithms for simulating the lattice schwinger model, *Quantum* **4**, 306 (2020).
- [278] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum simulation of electronic structure with linear depth and connectivity, *Physical review letters* **120**, 110501 (2018).
- [279] T. Tomesh, P. Gokhale, E. R. Anschuetz, and F. T. Chong, Coreset clustering on small quantum computers, *Electronics* **10**, 1690 (2021).
- [280] A. Hashim, R. Rines, V. Omole, R. K. Naik, J. M. Kreikebaum, D. I. Santiago, F. T. Chong, I. Siddiqi, and P. Gokhale, Optimized swap networks with equivalent circuit averaging for qaoa, *Physical Review Research* **4**, 033028 (2022).
- [281] N. Grzesiak, R. Blümel, K. Wright, K. M. Beck, N. C. Pimenti, M. Li, V. Chaplin, J. M. Amini, S. Debnath, J.-S. Chen, *et al.*, Efficient arbitrary simultaneously entangling gates on a trapped-ion quantum computer, *Nature communications* **11**, 2963 (2020).
- [282] N. Grzesiak, A. Maksymov, P. Niroula, and Y. Nam, Efficient quantum programming using ease gates on a trapped-ion quantum computer, *Quantum* **6**, 634 (2022).
- [283] S. Bravyi, D. Maslov, and Y. Nam, Constant-cost implementations of clifford operations and multiply-controlled gates using global interactions, *Physical Review Letters* **129**, 230501 (2022).
- [284] J. Ang, G. Carini, Y. Chen, I. Chuang, M. A. DeMarco, S. E. Economou, A. Eickbusch, A. Faraon, K.-M. Fu, S. M. Girvin, *et al.*, Architectures for multinode superconducting quantum computers, arXiv preprint arXiv:2212.06167 (2022).
- [285] A. Wu, Y. Ding, and A. Li, Collcomm: Enabling efficient collective quantum communication based on epr buffering, arXiv preprint arXiv:2208.06724 (2022).
- [286] Y. Zhang, L. Cincio, C. F. A. Negre, P. Czarnik, P. J. Coles, P. M. Anisimov, S. M. Mniszewski, S. Tretiak, and P. A. Dub, Variational quantum eigensolver with reduced circuit complexity, *npj Quantum Information* **8**, 10.1038/s41534-022-00599-z (2022).
- [287] T. Peng, A. W. Harrow, M. Ozols, and X. Wu, Simulating large quantum circuits on a small quantum computer, *Physical review letters* **125**, 150504 (2020).
- [288] C. Piveteau and D. Sutter, Circuit knitting with classical communication, *IEEE Transactions on Information Theory* (2023).
- [289] W. Tang, T. Tomesh, M. Suchara, J. Larson, and M. Martonosi, Cutqc: using small quantum computers for large quantum circuit evaluations, in *Proceedings of the 26th ACM International conference on architectural support for programming languages and operating systems* (2021) pp. 473–486.
- [290] W. Tang and M. Martonosi, Scaleqc: A scalable framework for hybrid computation on quantum and classical processors, arXiv preprint arXiv:2207.00933 (2022).
- [291] IBM, Introducing quantum serverless, a new programming model for leveraging quantum and classical resources, <https://research.ibm.com/blog/quantum-serverless-programming> (2021), accessed: Nov 6, 2023.
- [292] O. Di Matteo, A. McCoy, P. Gysbers, T. Miyagi, R. Woloshyn, and P. Navrátil, Improving hamiltonian encodings with the gray code, *Physical Review A* **103**, 042405 (2021).
- [293] S. Chakram, A. E. Oriani, R. K. Naik, A. V. Dixit, K. He, A. Agrawal, H. Kwon, and D. I. Schuster, Seamless high-q microwave cavities for multimode circuit quantum electrodynamics, *Physical review letters* **127**, 107701 (2021).
- [294] A. Cowtan, W. Simmons, and R. Duncan, A generic compilation strategy for the unitary coupled cluster ansatz, arXiv preprint arXiv:2007.10515 (2020).
- [295] G. Li, Y. Shi, and A. Javadi-Abhari, Software-hardware co-optimization for computational chemistry on superconducting quantum processors, in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)* (IEEE, 2021) pp. 832–845.
- [296] S. Martiel and T. G. de Brugière, Architecture aware compilation of quantum circuits via lazy synthesis, *Quantum* **6**, 729 (2022).
- [297] M. Amy, P. Azimzadeh, and M. Mosca, On the controlled-not complexity of controlled-not-phase circuits, *Quantum Science and Technology* **4**, 015002 (2018).
- [298] V. Vandaele, S. Martiel, and T. G. de Brugière,

- Phase polynomials synthesis algorithms for nisq architectures and beyond, *Quantum Science and Technology* **7**, 045027 (2022).
- [299] A. Meijer-van de Griend and R. Duncan, Architecture-aware synthesis of phase polynomials for nisq devices, *Electronic Proceedings in Theoretical Computer Science* (2020).
- [300] C. Jones, Novel constructions for the fault-tolerant tofoli gate, 2013, *Phys. Rev. A* **87**, 022328.
- [301] Y. Nam, Y. Su, and D. Maslov, Approximate quantum fourier transform with $\mathcal{O}(n \log(n))$ t gates, *NPJ Quantum Information* **6**, 26 (2020).
- [302] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, Validating quantum computers using randomized model circuits, *Physical Review A* **100**, 032328 (2019).
- [303] C. Campbell, F. T. Chong, D. Dahl, P. Frederick, P. Goiporia, P. Gokhale, B. Hall, S. Issa, E. Jones, S. Lee, *et al.*, Superstaq: Deep optimization of quantum programs, arXiv preprint arXiv:2309.05157 (2023).
- [304] C. Gidney and M. Ekerå, How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits, *Quantum* **5**, 433 (2021).
- [305] D. Coppersmith, An approximate fourier transform useful in quantum factoring, arXiv preprint quant-ph/0201067 (2002).
- [306] E. Campbell, A. Khurana, and A. Montanaro, Applying quantum algorithms to constraint satisfaction problems, *Quantum* **3**, 167 (2019).
- [307] S. Pabst and Y. Nam, A quantum algorithm for network reliability, arXiv preprint arXiv:2203.10201 (2022).
- [308] R. Shaydulin, C. Li, S. Chakrabarti, M. DeCross, D. Herman, N. Kumar, J. Larson, D. Lykov, P. Minssen, Y. Sun, *et al.*, Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem, arXiv preprint arXiv:2308.02342 (2023).
- [309] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, *et al.*, Quantum algorithms: A survey of applications and end-to-end complexities, arXiv preprint arXiv:2310.03011 (2023).
- [310] A. D. Córcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Mineev, J. M. Chow, and J. M. Gambetta, Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits, *Physical Review Letters* **127**, 100501 (2021).
- [311] L. C. Govia, P. Jurcevic, C. J. Wood, N. Kanazawa, S. T. Merkel, and D. McKay, A randomized benchmarking suite for mid-circuit measurements, *New Journal of Physics* (2022).
- [312] N. C. Brown, J. P. Campora III, C. Granade, B. Heim, S. Wernli, C. Ryan-Anderson, D. Lucchetti, A. Paetznick, M. Roetteler, K. Svore, *et al.*, Advances in compilation for quantum hardware—a demonstration of magic state distillation and repeat-until-success protocols, arXiv preprint arXiv:2310.12106 (2023).
- [313] Y. Ge, J. Tura, and J. I. Cirac, Faster ground state preparation and high-precision ground energy estimation with fewer qubits, *Journal of Mathematical Physics* **60** (2019).
- [314] I. Stetcu, A. Baroni, and J. Carlson, Projection algorithm for state preparation on quantum computers, *Physical Review C* **108**, L031306 (2023).
- [315] C. Granade and N. Wiebe, Using random walks for iterative phase estimation, arXiv preprint arXiv:2208.04526 (2022).
- [316] A. Maksymov, J. Nguyen, Y. Nam, and I. Markov, Enhancing quantum computer performance via symmetrization, arXiv preprint arXiv:2301.07233 (2023).
- [317] S. Stein, N. Wiebe, Y. Ding, J. Ang, and A. Li, Q-beep: Quantum bayesian error mitigation employing poisson modeling over the hamming spectrum, in *Proceedings of the 50th Annual International Symposium on Computer Architecture* (2023) pp. 1–13.
- [318] E. Van Den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse pauli-lindblad models on noisy quantum processors, *Nature Physics* (2023).
- [319] Y. Kim, C. J. Wood, T. J. Yoder, S. T. Merkel, J. M. Gambetta, K. Temme, and A. Kandala, Scalable error mitigation for noisy quantum circuits produces competitive expectation values, *Nature Physics* **19**, 752–759 (2023).
- [320] K. Landsman, M. Keesan, and C. Monroe, Toward convergence of effective-field-theory simulations on digital quantum computers, *Physical Review A* **100**, 062319 (2019).
- [321] T. Gurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, Digital zero noise extrapolation for quantum error mitigation, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2020) pp. 306–316.
- [322] R. Majumdar, P. Rivero, F. Metz, A. Hasan, and D. S. Wang, Best practices for quantum error mitigation with digital zero-noise extrapolation, arXiv preprint arXiv:2307.05203 (2023).
- [323] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature* **618**, 500 (2023).
- [324] M. A. Wahl, A. Mari, N. Shammah, W. J. Zeng, and G. S. Ravi, Zero noise extrapolation on logical qubits by scaling the error correction code distance (2023), arXiv:2304.14985 [quant-ph].
- [325] G. S. Ravi, J. M. Baker, K. N. Smith, N. Earnest, A. Javadi-Abhari, and F. Chong, Boosting quantum fidelity with an ordered diverse ensemble of clifford canary circuits (2022), arXiv:2209.13732 [quant-ph].
- [326] N. Saurabh, S. Jha, and A. Luckow, A conceptual architecture for a quantum-hpc middleware, in *2023 IEEE International Conference on Quantum Software (QSW)* (IEEE, 2023) pp. 116–127.
- [327] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, *Quantum* **3**, 128 (2019).
- [328] Ibm quantum documentation: shares and administration., <https://docs.quantum-computing.ibm.com/run/fair-share-queue#shares-and-administration> (), accessed: 2023-11-20.
- [329] A. B. Yoo, M. A. Jette, and M. Grondona, Slurm: Simple linux utility for resource management, in *Workshop on job scheduling strategies for parallel processing* (Springer, 2003) pp. 44–60.
- [330] Torque, <https://hpc-wiki.info/hpc/Torque>, accessed: 2023-11-20.
- [331] Altair grid engine: Distributed resource management

- and optimization, <https://altair.com/grid-engine>, accessed: 2023-11-20.
- [332] Ibm spectrum lsf suites, <https://www.ibm.com/products/hpc-workload-management>, accessed: 2023-11-20.
- [333] Ibm quantum documentation: fair-share queue., <https://docs.quantum-computing.ibm.com/run/fair-share-queue> (), accessed: 2023-11-20.
- [334] T. Bicer, D. Chiu, and G. Agrawal, A framework for data-intensive computing with cloud bursting, in *2011 IEEE International Conference on Cluster Computing* (2011) pp. 169–177.
- [335] G. S. Ravi, K. N. Smith, P. Murali, and F. T. Chong, Adaptive job and resource management for the growing quantum cloud (2022), [arXiv:2203.13260](https://arxiv.org/abs/2203.13260) [quant-ph].
- [336] A. J. McCaskey, D. I. Lyakh, E. F. Dumitrescu, S. S. Powers, and T. S. Humble, Xacc: a system-level software infrastructure for heterogeneous quantum-classical computing*, *Quantum Science and Technology* **5**, 024002 (2020).
- [337] T. M. Mintz, A. J. McCaskey, E. F. Dumitrescu, S. V. Moore, S. Powers, and P. Lougovski, Qcor: A language extension specification for the heterogeneous quantum-classical model of computation, *J. Emerg. Technol. Comput. Syst.* **16**, 10.1145/3380964 (2020).
- [338] A. Mccaskey, T. Nguyen, A. Santana, D. Claudino, T. Kharazi, and H. Finkel, Extending c++ for heterogeneous quantum-classical computing, *ACM Transactions on Quantum Computing* **2**, 10.1145/3462670 (2021).
- [339] NVIDIA, NVIDIA CUDA Quantum: The platform for hybrid quantum-classical computing, <https://developer.nvidia.com/cuda-quantum> (2023).
- [340] A. Cross, The ibm q experience and qiskit open-source quantum computing software, in *APS March meeting abstracts*, Vol. 2018 (2018) pp. L58–003.
- [341] D. Poccia, Introducing amazon braket hybrid jobs – set up, monitor, and efficiently run hybrid quantum-classical workloads, <https://aws.amazon.com/blogs/aws/introducing-amazon-braket-hybrid-jobs-set-up-monitor-and-optimizing-compilation-for-proceedings-of-classical-workloads/> (2021).
- [342] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri, *et al.*, Pennylane: Automatic differentiation of hybrid quantum-classical computations, [arXiv:1811.04968](https://arxiv.org/abs/1811.04968) (2018).
- [343] PSNC, PSNC QCG HPC & Quantum middleware platform, <https://qcg.psnc.pl> (2023).
- [344] B. Bosak, T. Piontek, P. Karlshoefler, E. Raffin, J. Lakhilili, and P. Kopta, Verification, validation and uncertainty quantification of large-scale applications with qcg-pilotjob (2021) pp. 495–501.
- [345] Zapata Computing, Orquestra: A platform for hybrid quantum-classical computing, Zapata Computing, <https://www.zapatacomputing.com/orquestra-platform/> (2023).
- [346] Covalent, Covalent: Open source workflow orchestration for heterogenous computing, Covalent, <https://www.covalent.xyz/> (2023).
- [347] S. Sivarajah, L. Heidemann, A. Lawrence, and R. Duncan, Tierkreis: A dataflow framework for hybrid quantum-classical computing (2022).
- [348] I. Faro, I. Sitdikov, D. Valinas, F. Fernandez, C. Codella, and J. Glick, Middleware for quantum: An orchestration of hybrid quantum-classical systems, in *2023 IEEE International Conference on Quantum Software (QSW)* (IEEE Computer Society, Los Alamitos, CA, USA, 2023) pp. 1–8.
- [349] European Parliament and Council of the European Union, Regulation (EU) 2016/679 of the European Parliament and of the Council.
- [350] Centers for Medicare & Medicaid Services, The Health Insurance Portability and Accountability Act of 1996 (HIPAA), Online at <http://www.cms.hhs.gov/hipaa/> (1996).
- [351] The Apache Software Foundation, *SparkR: R Front End for 'Apache Spark'* (2023), <https://www.apache.org> <https://spark.apache.org>.
- [352] R. Reyes, G. Brown, R. Burns, and M. Wong, Sycl 2020: More than meets the eye, in *Proceedings of the International Workshop on OpenCL, IWOCCL '20* (Association for Computing Machinery, New York, NY, USA, 2020).
- [353] J. Sanders and E. Kandrot, *CUDA by Example: An Introduction to General-Purpose GPU Programming*, 1st ed. (Addison-Wesley Professional, 2010).
- [354] T. Heller, P. Diehl, Z. Byerly, J. Biddiscombe, and H. Kaiser, Hpx – an open source c++ standard library for parallelism and concurrency (2023), [arXiv:2401.03353](https://arxiv.org/abs/2401.03353) [cs.DC].
- [355] J. R. McClean, N. C. Rubin, K. J. Sung, I. D. Kivlichan, X. Bonet-Monroig, Y. Cao, C. Dai, E. S. Fried, C. Gidney, B. Gimby, P. Gokhale, T. Häner, T. Hardikar, V. Havlíček, O. Higgott, C. Huang, J. Izaac, Z. Jiang, X. Liu, S. McArdle, M. Neeley, T. O’Brien, B. O’Gorman, I. Ozfidan, M. D. Radin, J. Romero, N. P. D. Sawaya, B. Senjean, K. Setia, S. Sim, D. S. Steiger, M. Staudtner, Q. Sun, W. Sun, D. Wang, F. Zhang, and R. Babbush, OpenFermion: The electronic structure package for quantum computers, *Quantum Science and Technology* **5**, 034014:1 (2020).
- [356] Y. Peng, J. Young, P. Liu, and X. Wu, SimuQ: A framework for programming quantum hamiltonian simulation with circuit compilation, *Proceedings of ACM Languages & Programming* **8**, 2425 (2024).
- [357] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra, *The MPI Core*, 2nd ed., Vol. 1 (The MIT Press, Cambridge, Massachusetts).
- [358] L. Dagum and R. Menon, OpenMP: An industry-standard API for shared-memory programming, *IEEE Computational Science and Engineering* **5**, 46.
- [359] A. J. McCaskey, D. I. Lyakh, E. F. D. itrescu, S. S. Powers, and T. S. Humble, XACC: A system-level software infrastructure for heterogeneous quantum-classical computing, *Quantum Science and Technology* **5**, 024002:1 (2020).
- [360] A. Hayashi, A. Adams, J. Young, A. McCaskey, E. Dumitrescu, V. Sarkar, and T. M. Conte, Enabling multi-threading in heterogeneous quantum-classical programming models, in *2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (2023) pp. 509–516.
- [361] J.-S. Kim, A. McCaskey, B. Heim, M. Modani, S. Stanwyck, and T. Costa, CUDA Quantum: The platform for integrated quantum-classical computing, in *60th ACM/IEEE Design Automation Conference (DAC)* (San Francisco, California, USA, 2023) pp. 1–4.
- [362] M. Flynn, Very high-speed computing systems, *Pro-*

- ceedings of the IEEE **54**, 1901 (1966).
- [363] E. Lindholm, J. R. Nickolls, S. F. Oberman, and J. Montrym, Nvidia tesla: A unified graphics and computing architecture, *IEEE Micro* **28** (2008).
- [364] *HOPL III: Proceedings of the third ACM SIGPLAN conference on History of programming languages* (Association for Computing Machinery, New York, NY, USA, 2007).
- [365] *Go Programming Language*, accessed: 2023-11.
- [366] Welcome to Python.org — python.org, <https://www.python.org/>, [Accessed 09-02-2024].
- [367] Standard library header `jexecution.h` (C++17) - cppreference.com, accessed: 2024-02.
- [368] CUDA Toolkit - Free Tools and Training — NVIDIA Developer, accessed: 2024-02.
- [369] OpenMP Architecture Review Board (2023), accessed: 2023-11.
- [370] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R. H. Castain, D. J. Daniel, R. L. Graham, and T. S. Woodall, Open MPI: Goals, concept, and design of a next generation MPI implementation, in *Proceedings, 11th European PVM/MPI Users' Group Meeting* (Budapest, Hungary, 2004) pp. 97–104.
- [371] Productionizing and scaling Python ML workloads simply — Ray — ray.io, <https://www.ray.io/>, [Accessed 09-02-2024].
- [372] A. W. Cross *et al.*, OpenQASM 3: A broader and deeper quantum assembly language, <https://arxiv.org/abs/2104.14722> (2022), <https://arxiv.org/abs/2104.14722>.
- [373] Microsoft, Q#: Enabling Scalable Quantum Computing and Development with a High-Level DSL, <https://learn.microsoft.com/en-us/azure/quantum/> (2024), <https://learn.microsoft.com/en-us/azure/quantum/user-guide/>.
- [374] P. Selinger *et al.*, Quipper: A Scalable Quantum Programming Language, <http://www.mathstat.dal.ca/~selinger/quipper/> (2013), <https://arxiv.org/abs/1304.3390>.
- [375] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, <https://www.tensorflow.org/> (2015).
- [376] TensorFlow Team, *Tensorboard: Tensorflow's visualization toolkit* (2023), accessed: 2023-02-14.
- [377] IBM, Qiskit: An Open-source Framework for Quantum Computing, <https://qiskit.org/> (2023), <https://www.ibm.com/quantum/qiskit>.
- [378] Qiskit Development Team, *Dag drawer* (2023), accessed: 2023-02-14.
- [379] A. Al-Saadi, D. H. Ahn, Y. Babuji, K. Chard, J. Corbett, M. Hategan, S. Herbein, S. Jha, D. Laney, A. Merzky, T. Munson, M. Salim, M. Titov, M. Turilli, and J. M. Wozniak, Exaworks: Workflows for exascale (2021), [arXiv:2108.13521 \[cs.DC\]](https://arxiv.org/abs/2108.13521).
- [380] M. Mosca, Quantum algorithms (2008), [arXiv:0808.0369 \[quant-ph\]](https://arxiv.org/abs/0808.0369).
- [381] T. Jones, A. Brown, I. Bush, and S. C. Benjamin, Quest and high performance simulation of quantum computers, *Scientific Reports* **9**, 10736 (2019).
- [382] A. Li, B. Fang, C. Granade, G. Prawiroatmodjo, B. Heim, M. Roetteler, and S. Krishnamoorthy, Sv-sim: Scalable pgas-based state vector simulation of quantum circuits, in *SC21: International Conference for High Performance Computing, Networking, Storage and Analysis* (2021) pp. 1–14.
- [383] Y. A. Danylo Lykov, Importance of diagonal gates in tensor network simulations, [arXiv https://doi.org/10.48550/arXiv.2106.15740](https://doi.org/10.48550/arXiv.2106.15740) (2021).
- [384] D. A. Fedorov, Y. Alexeev, S. K. Gray, and M. Otten, Unitary Selective Coupled-Cluster Method, *Quantum* **6**, 703 (2022).
- [385] U. Baek, D. Hait, J. Shee, O. Leimkuhler, W. J. Huggins, T. F. Stetina, M. Head-Gordon, and K. B. Whaley, Say NO to Optimization: A Nonorthogonal Quantum Eigensolver, *PRX Quantum* **4**, 030307 (2023), [arXiv:2205.09039 \[quant-ph\]](https://arxiv.org/abs/2205.09039).
- [386] Laying the Groundwork for Quantum Powered Use Cases, https://www.youtube.com/watch?v=Hd-IGvuARfE&ab_channel=IBMRResearch (2023), [Video].
- [387] D. Gottesman, The heisenberg representation of quantum computers, *arXiv preprint quant-ph/9807006* (1998).
- [388] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Physical Review A* **70**, 052328 (2004).
- [389] H. J. Garcia and I. L. Markov, Simulation of quantum circuits via stabilizer frames, *IEEE Transactions on Computers* **64**, 2323 (2014).
- [390] D. Aharonov, Z. Landau, and J. Makowsky, The quantum fft can be classically simulated, *arXiv preprint quant-ph/0611156* (2006).
- [391] N. Yoran and A. J. Short, Efficient classical simulation of the approximate quantum fourier transform, *Physical Review A* **76**, 042321 (2007).
- [392] S. Hillmich, I. L. Markov, and R. Wille, Just like the real thing: Fast weak simulation of quantum computation, in *2020 57th ACM/IEEE Design Automation Conference (DAC)* (IEEE, 2020) pp. 1–6.
- [393] S. Hillmich, A. Zulehner, R. Kueng, I. L. Markov, and R. Wille, Approximating decision diagrams for quantum circuit simulation, *ACM Transactions on Quantum Computing* **3**, 1 (2022).
- [394] A. Fatima and I. L. Markov, Faster schrödinger-style simulation of quantum circuits, in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE, 2021) pp. 194–207.
- [395] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Physical review A* **52**, 3457 (1995).
- [396] T. Häner and D. S. Steiger, 0.5 petabyte simulation of a 45-qubit quantum circuit, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (ACM, 2017) p. 33.
- [397] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, T. Magerlein, E. Solomonik, and R. Wisnieff, Breaking the 49-qubit barrier in the simulation of quantum circuits, *arXiv preprint arXiv:1710.05867* (2017).
- [398] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff, Leveraging secondary storage to simulate deep 54-qubit sycamore circuits, *arXiv preprint arXiv:1910.09534* (2019).
- [399] A. Li, B. Fang, C. Granade, G. Prawiroatmodjo, B. Heim, M. Roetteler, and S. Krishnamoorthy, Sv-sim: scalable pgas-based state vector simulation of quantum

- circuits, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2021) pp. 1–14.
- [400] A. Li, A. Baroni, I. Stetcu, and T. S. Humble, Deep quantum circuit simulations of low-energy nuclear states, arXiv preprint arXiv:2310.17739 (2023).
- [401] M. R. Hirsbrunner, D. Chamaki, J. W. Mullinax, and N. M. Tubman, Beyond MP2 initialization for unitary coupled cluster quantum circuits, arXiv e-prints , arXiv:2301.05666 (2023), arXiv:2301.05666 [quant-ph].
- [402] I. L. Markov, A. Fatima, S. V. Isakov, and S. Boixo, Massively parallel approximate simulation of hard quantum circuits, in *2020 57th ACM/IEEE Design Automation Conference (DAC)* (IEEE, 2020) pp. 1–6.
- [403] A. Li, O. Subasi, X. Yang, and S. Krishnamoorthy, Density matrix quantum circuit simulation via the bsp machine on modern gpu clusters, in *Sc20: international conference for high performance computing, networking, storage and analysis* (IEEE, 2020) pp. 1–15.
- [404] I. L. Markov and Y. Shi, Simulating quantum computation by contracting tensor networks, *SIAM Journal on Computing* **38**, 963 (2008).
- [405] D. Lykov and Y. Alexeev, Importance of diagonal gates in tensor network simulations (2021), arXiv:2106.15740 [quant-ph].
- [406] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, Tensor network quantum simulator with step-dependent parallelization, arXiv preprint arXiv:2012.02430 (2020).
- [407] D. Lykov, A. Chen, H. Chen, K. Keipert, Z. Zhang, T. Gibbs, and Y. Alexeev, Performance evaluation and acceleration of the qtensor quantum circuit simulator on gpus, in *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)* (IEEE, 2021) pp. 27–34.
- [408] D. Lykov, R. Schutski, A. Galda, V. Vinokur, and Y. Alexeev, Tensor network quantum simulator with step-dependent parallelization, in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2022) pp. 582–593.
- [409] W. Berquist, D. Lykov, M. Liu, and Y. Alexeev, Stochastic approach for simulating quantum noise using tensor networks, in *2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)* (IEEE, 2022) pp. 107–113.
- [410] M. Shah, X. Yu, S. Di, D. Lykov, Y. Alexeev, M. Becchi, and F. Cappello, Gpu-accelerated error-bounded compression framework for quantum circuit simulations, in *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)* (IEEE, 2023) pp. 757–767.
- [411] M. Liu, C. Oh, J. Liu, L. Jiang, and Y. Alexeev, Simulating lossy gaussian boson sampling with matrix-product operators, *Physical Review A* **108**, 052604 (2023).
- [412] C. Oh, M. Liu, Y. Alexeev, B. Fefferman, and L. Jiang, Tensor network algorithm for simulating experimental gaussian boson sampling, arXiv:2306.03709 (2023).
- [413] G. Di Bartolomeo, M. Vischi, F. Cesa, R. Wixinger, M. Grossi, S. Donadi, and A. Bassi, Noisy gates for simulating quantum computers, *Phys. Rev. Res.* **5**, 043210 (2023).
- [414] A. Hey, *Feynman and computation* (CRC Press, 2018).
- [415] N. P. Sawaya, D. Marti-Dafcik, Y. Ho, D. P. Tabor, D. Bernal, A. B. Magann, S. Premaratne, P. Dubey, A. Matsuura, N. Bishop, W. A. de Jong, S. Benjamin, O. D. Parekh, N. Tubman, K. Klymko, and D. Camps, HamLib: A library of Hamiltonians for benchmarking quantum algorithms and hardware, arXiv e-prints , arXiv:2306.13126 (2023), arXiv:2306.13126 [quant-ph].
- [416] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, *Proceedings of the national academy of sciences* **114**, 7555 (2017).
- [417] N. M. Tubman, C. Mejuto-Zaera, J. M. Epstein, D. Hait, D. S. Levine, W. Huggins, Z. Jiang, J. R. McClean, R. Babbush, M. Head-Gordon, and K. B. Whaley, Postponing the orthogonality catastrophe: efficient state preparation for electronic structure simulations on quantum devices, arXiv e-prints , arXiv:1809.05523 (2018), arXiv:1809.05523 [quant-ph].
- [418] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. Bochevarov, How will quantum computers provide an industrially relevant computational advantage in quantum chemistry?, arXiv preprint arXiv:2009.12472 (2020).
- [419] J. J. Goings, A. White, J. Lee, C. S. Tautermann, M. Degroote, C. Gidney, T. Shiozaki, R. Babbush, and N. C. Rubin, Reliably assessing the electronic structure of cytochrome p450 on today’s classical computers and tomorrow’s quantum computers, *Proceedings of the National Academy of Sciences* **119**, e2203533119 (2022).
- [420] C. Mejuto-Zaera, G. Weng, M. Romanova, S. J. Cotton, K. B. Whaley, N. M. Tubman, and V. Vlček, Are multi-quasiparticle interactions important in molecular ionization?, *The Journal of Chemical Physics* **154**, 121101 (2021).
- [421] M. Shishkin and G. Kresse, Self-consistent *gw* calculations for semiconductors and insulators, *Phys. Rev. B* **75**, 235102 (2007).
- [422] J. P. Perdew, Density functional theory and the band gap problem, *International Journal of Quantum Chemistry* **28**, 497 (1985).
- [423] A. J. Cohen, P. Mori-Sánchez, and W. Yang, Challenges for density functional theory, *Chemical reviews* **112**, 289 (2012).
- [424] R. O. Jones, Density Functional Theory: Its Origins, Rise to Prominence, and Future, *Rev. Mod. Phys.* **87**, 897 (2015).
- [425] M. Sparta and F. Neese, Chemical Applications Carried Out by Local Pair Natural Orbital Based Coupled-Cluster Methods, *Chem. Soc. Rev.* **43**, 5032 (2014).
- [426] D. Hait, N. M. Tubman, D. S. Levine, K. B. Whaley, and M. Head-Gordon, What levels of coupled cluster theory are appropriate for transition metal systems? a study using near-exact quantum chemical values for 3d transition metal binary compounds, *Journal of Chemical Theory and Computation* **15**, 5370 (2019), pMID: 31465217, <https://doi.org/10.1021/acs.jctc.9b00674>.
- [427] K. A. Moltved and K. P. Kepp, Performance of density functional theory for transition metal oxygen bonds, *ChemPhysChem* **20**, 3210 (2019).
- [428] S. Sharma, K. Sivalingam, F. Neese, and G. K.-L. Chan, Low-Energy Spectrum of Iron–Sulfur Clusters Directly from Many-Particle Quantum Mechanics, *Nat. Chem.* **6**, 927 (2014).
- [429] <https://github.com/zhendongli2008> (2020).
- [430] Z. Li and G. K.-L. Chan, Spin-projected matrix product states: Versatile tool for strongly correlated systems, J.

- Chem. Theory Comput. **13**, 2681 (2017).
- [431] L. Noodleman, T. Lovell, T. Liu, F. Himo, and R. A. Torres, Insights into Properties and Energetics of Iron-Sulfur Proteins from Simple Clusters to Nitrogenase, *Curr. Opin. Chem. Biol.* **6**, 259 (2002).
- [432] Z. Li, J. Li, N. S. Dattani, C. Umrigar, and G. K. Chan, The electronic complexity of the ground-state of the fmo cofactor of nitrogenase as relevant to quantum simulations, *J. Chem. Phys.* **150** (2019).
- [433] Z. Li, S. Guo, Q. Sun, and G. K.-L. Chan, Electronic Landscape of the P-Cluster of Nitrogenase as Revealed Through Many-Electron Quantum Wavefunction Simulations, *Nat. Chem.* **11**, 1026–1033 (2019).
- [434] C. W. Cady, R. H. Crabtree, and G. W. Brudvig, Functional Models for the Oxygen-Evolving Complex of Photosystem II, *Coord. Chem. Rev.* **252**, 444 (2008).
- [435] Y. Kurashige, G. K.-L. Chan, and T. Yanai, Entangled Quantum Electronic Wavefunctions of the Mn_4CaO_5 Cluster in Photosystem II, *Nat. Chem.* **5**, 660 (2013).
- [436] P. E. Siegbahn, Structures and Energetics for O_2 Formation in Photosystem II, *Acc. Chem. Res.* **42**, 1871 (2009).
- [437] K. M. Lancaster, M. Roemelt, P. Ettenhuber, Y. Hu, M. W. Ribbe, F. Neese, U. Bergmann, and S. DeBeer, X-ray Emission Spectroscopy Evidences a Central Carbon in the Nitrogenase Iron-Molybdenum Cofactor, *Science* **334**, 974 (2011).
- [438] J. Batool, T. Hahn, and M. R. Pederson, Magnetic Signatures of Hydroxyl- and Water-Terminated Neutral and Tetra-Anionic Mn_{12} -Acetate, *J. Comput. Chem.* **40**, 2301 (2019).
- [439] V. G. Chilkuri, S. DeBeer, and F. Neese, Ligand Field Theory and Angular Overlap Model Based Analysis of the Electronic Structure of Homovalent Iron-Sulfur Dimers, *Inorg. Chem.* **59**, 984 (2020).
- [440] L. Cao, O. Caldararu, and U. Ryde, Protonation and Reduction of the FeMo Cluster in Nitrogenase Studied by Quantum Mechanics/Molecular Mechanics (QM/MM) Calculations, *J. Chem. Theory Comput.* **14**, 6653 (2018).
- [441] L. Schimka, J. Harl, A. Stroppa, A. Grüneis, M. Marsman, F. Mittendorfer, and G. Kresse, Accurate Surface and Adsorption Energies from Many-Body Perturbation Theory, *Nat. Mater.* **9**, 741 (2010).
- [442] M. Capdevila-Cortada, Z. Lodziana, and N. López, Performance of DFT+U Approaches in the Study of Catalytic Materials, *ACS Catalysis* **6**, 8370 (2016).
- [443] J. K. Nørskov, F. Abild-Pedersen, F. Studt, and T. Bligaard, Density Functional Theory in Surface Chemistry and Catalysis, *Proc. Natl. Acad. Sci.* **108**, 937 (2011).
- [444] A. M. Wodtke, Electronically Non-Adiabatic Influences in Surface Chemistry and Dynamics, *Chem. Soc. Rev.* **45**, 3641 (2016).
- [445] J. K. Nørskov, T. Bligaard, J. Rossmeisl, and C. H. Christensen, Towards the Computational Design of Solid Catalysts, *Nat. Chem.* **1**, 37 (2009).
- [446] J. VandeVondele, U. Borstnik, and J. Hutter, Linear Scaling Self-Consistent Field Calculations with Millions of Atoms in the Condensed Phase, *J. Chem. Theory Comput.* **8**, 3565 (2012).
- [447] T. F. Stetina, A. Ciavarella, X. Li, and N. Wiebe, Simulating effective qed on quantum computers, *Quantum* **6**, 622 (2022).
- [448] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, From Quantum Matter to High-Temperature Superconductivity in Copper Oxides, *Nature* **518**, 179 (2015).
- [449] S.-S. Lee, Recent Developments in Non-Fermi Liquid Theory, *Annu. Rev. Condens. Matter Phys.* **9**, 227 (2018).
- [450] E. Fradkin, S. A. Kivelson, and J. M. Tranquada, Colloquium: Theory of Intertwined Orders in High-Temperature Superconductors, *Rev. Mod. Phys.* **87**, 457 (2015).
- [451] L. Balents, Spin Liquids in Frustrated Magnets, *Nature* **464**, 199 (2010).
- [452] M. Rossmannek, P. K. Barkoutsos, P. J. Ollitrault, and I. Tavernelli, Quantum hf/dft-embedding algorithms for electronic structure calculations: Scaling up to complex molecular systems, *The Journal of Chemical Physics* **154** (2021).
- [453] M. Rossmannek, F. Pavošević, A. Rubio, and I. Tavernelli, Quantum embedding method for the simulation of strongly correlated systems on quantum computers, *The Journal of Physical Chemistry Letters* **14**, 3491 (2023).
- [454] Y. Gao, F. Albrecht, I. Rončević, I. Etdedgui, P. Kumar, L. Scriven, K. Christensen, S. Mishra, L. Righetti, M. Rossmannek, I. Tavernelli, H. Anderson, and L. Gross, On-surface synthesis of a doubly anti-aromatic carbon allotrope: Cyclo[16]carbon [10.26434/chemrxiv-2023-w124b](https://doi.org/10.26434/chemrxiv-2023-w124b) (2023).
- [455] I. Rončević, F. J. Leslie, M. Rossmannek, I. Tavernelli, L. Gross, and H. Anderson, Aromaticity reversal induced by vibrations in cyclo[16]carbon [10.26434/chemrxiv-2023-8kc5t](https://doi.org/10.26434/chemrxiv-2023-8kc5t) (2023).
- [456] G. Knizia and G. K.-L. Chan, Density matrix embedding: A simple alternative to dynamical mean-field theory, *Physical Review Letters* **109**, 186404 (2012).
- [457] G. Knizia and G. K.-L. Chan, Density matrix embedding: A strong-coupling quantum embedding theory, *Journal of Chemical Theory and Computation* **9**, 1428 (2013).
- [458] S. Wouters, C. A. Jiménez-Hoyos, Q. Sun, and G. K.-L. Chan, A practical guide to density matrix embedding theory in quantum chemistry, *Journal of Chemical Theory and Computation* **12**, 2706 (2016).
- [459] H. Q. Pham, V. Bernalles, and L. Gagliardi, Can density matrix embedding theory with the complete activate space self-consistent field solver describe single and double bond breaking in molecular systems?, *Journal of Chemical Theory and Computation* **14**, 1960 (2018).
- [460] M. R. Hermes and L. Gagliardi, Multiconfigurational self-consistent field theory with density matrix embedding: The localized active space self-consistent field method, *Journal of Chemical Theory and Computation* **15**, 972 (2019).
- [461] H. Q. Pham, M. R. Hermes, and L. Gagliardi, Periodic electronic structure calculations with the density matrix embedding theory, *Journal of Chemical Theory and Computation* **16**, 130 (2020).
- [462] Z.-H. Cui, T. Zhu, and G. K.-L. Chan, Efficient implementation of ab initio quantum embedding in periodic systems: Density matrix embedding theory, *Journal of Chemical Theory and Computation* **16**, 119 (2020).
- [463] B. T. G. Lau, G. Knizia, and T. C. Berkelbach, Regional embedding enables high-level quantum chemistry for surface science, *The Journal of Physical Chemistry Letters* **12**, 1104 (2021).

- [464] Z.-H. Cui, H. Zhai, X. Zhang, and G. K.-L. Chan, Systematic electronic structure in the cuprate parent state from quantum many-body simulations, *Science* **377**, 1192 (2022).
- [465] A. Mitra, M. R. Hermes, and L. Gagliardi, Density matrix embedding using multiconfiguration pair-density functional theory, *Journal of Chemical Theory and Computation* **19**, 3498 (2023).
- [466] V. I. Anisimov, A. I. Poteryaev, M. A. Korotin, A. O. Anokhin, and G. Kotliar, First-principles calculations of the electronic structure and spectra of strongly correlated systems: Dynamical mean-field theory, *Journal of Physics: Condensed Matter* **9**, 7359 (1997).
- [467] A. I. Lichtenstein and M. I. Katsnelson, Ab initio calculations of quasiparticle band structure in correlated systems: LDA++ approach, *Physical Review B* **57**, 6884 (1998).
- [468] P. Sun and G. Kotliar, Extended dynamical mean-field theory and GW method, *Physical Review B* **66**, 085120 (2002).
- [469] S. Biermann, F. Aryasetiawan, and A. Georges, First-principles approach to the electronic structure of strongly correlated systems: Combining the GW approximation and dynamical mean-field theory, *Physical Review Letters* **90**, 086402 (2003).
- [470] S. Biermann, Dynamical screening effects in correlated electron materials—a progress report on combined many-body perturbation and dynamical mean field theory: GW+DMFT, *Journal of Physics: Condensed Matter* **26**, 173202 (2014).
- [471] L. Boehnke, F. Nilsson, F. Aryasetiawan, and P. Werner, When strong correlations become weak: Consistent merging of gw and dmft, *Physical Review B* **94**, 201106 (2016).
- [472] S. Choi, A. Kutepov, K. Haule, M. van Schilfhaarde, and G. Kotliar, First-principles treatment of Mott insulators: Linearized QSGW+DMFT approach, *npj Quantum Materials* **1**, 1 (2016).
- [473] F. Nilsson, L. Boehnke, P. Werner, and F. Aryasetiawan, Multitier self-consistent GW+EDMFT, *Physical Review Materials* **1**, 043803 (2017).
- [474] T. N. Lan and D. Zgid, Generalized self-energy embedding theory, *The Journal of Physical Chemistry Letters* **8**, 2200 (2017).
- [475] D. Zgid and E. Gull, Finite temperature quantum embedding theories for correlated systems, *New Journal of Physics* **19**, 023047 (2017).
- [476] A. A. Rusakov, S. Iskakov, L. N. Tran, and D. Zgid, Self-energy embedding theory (SEET) for periodic systems, *Journal of Chemical Theory and Computation* **15**, 229 (2019).
- [477] H. Ma, N. Sheng, M. Govoni, and G. Galli, First-principles studies of strongly correlated states in defect spin qubits in diamond, *Physical Chemistry Chemical Physics* **22**, 25522 (2020).
- [478] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, *npj Computational Materials* **6**, 85 (2020).
- [479] H. Ma, N. Sheng, M. Govoni, and G. Galli, Quantum embedding theory for strongly correlated states in materials, *Journal of Chemical Theory and Computation* **17**, 2116 (2021).
- [480] N. Sheng, C. Vorwerk, M. Govoni, and G. Galli, Green's function formulation of quantum defect embedding theory, *Journal of Chemical Theory and Computation* **18**, 3512 (2022).
- [481] I. Rungger, N. Fitzpatrick, H. Chen, C. H. Alderete, H. Apel, A. Cowtan, A. Patterson, D. M. Ramo, Y. Zhu, N. H. Nguyen, E. Grant, S. Chretien, L. Wossnig, N. M. Linke, and R. Duncan, Dynamical mean field theory algorithm and experiment on quantum computers (2020), [arXiv:1910.04735 \[quant-ph\]](https://arxiv.org/abs/1910.04735).
- [482] T. Keen, T. Maier, S. Johnston, and P. Lougovski, Quantum-classical simulation of two-site dynamical mean-field theory on noisy quantum hardware, *Quantum Science and Technology* **5**, 035001 (2020).
- [483] Y. Kawashima, E. Lloyd, M. P. Coons, Y. Nam, S. Matsuura, A. J. Garza, S. Johri, L. Huntington, V. Senicourt, A. O. Maksymov, J. H. V. Nguyen, J. Kim, N. Alidoust, A. Zaribafiyani, and T. Yamazaki, Optimizing electronic structure simulations on a trapped-ion quantum computer using problem decomposition (2021), [arXiv:2102.07045 \[quant-ph\]](https://arxiv.org/abs/2102.07045).
- [484] Y. Yao, F. Zhang, C.-Z. Wang, K.-M. Ho, and P. P. Orth, Gutzwiller hybrid quantum-classical computing approach for correlated materials, *Physical Review Research* **3**, 013184 (2021).
- [485] J. Tilly, P. V. Sriluckshmy, A. Patel, E. Fontana, I. Rungger, E. Grant, R. Anderson, J. Tennyson, and G. H. Booth, Reduced density matrix sampling: Self-consistent embedding and multiscale electronic structure on current generation quantum computers, *Physical Review Research* **3**, 033230 (2021).
- [486] C. Vorwerk, N. Sheng, M. Govoni, B. Huang, and G. Galli, Quantum embedding theories to simulate condensed systems on quantum computers, *Nature Computational Science* **2**, 424 (2022).
- [487] B. Huang, M. Govoni, and G. Galli, Simulating the electronic structure of spin defects on quantum computers, *PRX Quantum* **3**, 010339 (2022).
- [488] E. Pavarini, I. Dasgupta, T. Saha-Dasgupta, O. Jepsen, and O. Andersen, Band-structure trend in hole-doped cuprates and correlation with t_c max, *Physical review letters* **87**, 047003 (2001).
- [489] K. Haule, Exact double counting in combining the dynamical mean field theory and the density functional theory, *Physical review letters* **115**, 196403 (2015).
- [490] F. Aryasetiawan, M. Imada, A. Georges, G. Kotliar, S. Biermann, and A. Lichtenstein, Frequency-dependent local interactions and low-energy effective models from electronic structure calculations, *Physical Review B* **70**, 195104 (2004).
- [491] H. O. Jeschke, F. Salvat-Pujol, and R. Valentí, First-principles determination of heisenberg hamiltonian parameters for the spin-1 2 kagome antiferromagnet zncu₃(oh)₆cl₂, *Physical Review B* **88**, 075106 (2013).
- [492] S. Zhou and D. Ceperley, Construction of localized wave functions for a disordered optical lattice and analysis of the resulting hubbard model parameters, *Physical Review A* **81**, 013402 (2010).
- [493] S. R. White, Numerical canonical transformation approach to quantum many-body problems, *The Journal of chemical physics* **117**, 7472 (2002).
- [494] T. Yanai and G. K. Chan, Canonical transformation theory for multireference problems, *The Journal of chemical physics* **124** (2006).
- [495] L. Savary and L. Balents, Quantum spin liquids: a review, *Reports on Progress in Physics* **80**, 016502 (2016).

- [496] Y. Zhou, K. Kanoda, and T.-K. Ng, Quantum spin liquid states, *Reviews of Modern Physics* **89**, 025003 (2017).
- [497] J. R. Chamorro, T. M. McQueen, and T. T. Tran, Chemistry of quantum spin liquids, *Chemical Reviews* **121**, 2898 (2020).
- [498] J. A. Sears, Y. Zhao, Z. Xu, J. W. Lynn, and Y.-J. Kim, Phase diagram of α -rucl 3 in an in-plane magnetic field, *Physical Review B* **95**, 180411 (2017).
- [499] A. Jahin, A. C. Y. Li, T. Iadecola, P. P. Orth, G. N. Perdue, A. Macridin, M. S. Alam, and N. M. Tubman, Fermionic approach to variational quantum simulation of kitaev spin models, *Phys. Rev. A* **106**, 022434 (2022).
- [500] A. C. Li, M. S. Alam, T. Iadecola, A. Jahin, J. Job, D. M. Kurkcuoglu, R. Li, P. P. Orth, A. B. Özgüler, G. N. Perdue, *et al.*, Benchmarking variational quantum eigensolvers for the square-octagon-lattice kitaev model, *Physical Review Research* **5**, 033071 (2023).
- [501] R. N. Tazhigulov, S.-N. Sun, R. Haghshenas, H. Zhai, A. T. Tan, N. C. Rubin, R. Babbush, A. J. Minnich, and G. K.-L. Chan, Simulating models of challenging correlated molecules and materials on the sycamore quantum processor, *PRX Quantum* **3**, 040318 (2022).
- [502] K. Plumb, J. Clancy, L. Sandilands, V. V. Shankar, Y. Hu, K. Burch, H.-Y. Kee, and Y.-J. Kim, α -rucl 3: A spin-orbit assisted mott insulator on a honeycomb lattice, *Physical Review B* **90**, 041112 (2014).
- [503] A. Kitaev, Anyons in an exactly solved model and beyond, *Annals of Physics* **321**, 2 (2006).
- [504] G. Jackeli and G. Khaliullin, Mott insulators in the strong spin-orbit coupling limit: from heisenberg to a quantum compass and kitaev models, *Physical review letters* **102**, 017205 (2009).
- [505] A. Koitzsch, C. Habenicht, E. Müller, M. Knupfer, B. Büchner, H. C. Kandpal, J. Van Den Brink, D. Nowak, A. Isaeva, and T. Doert, J eff description of the honeycomb mott insulator α -rucl 3, *Physical Review Letters* **117**, 126403 (2016).
- [506] J. Nasu, J. Knolle, D. L. Kovrizhin, Y. Motome, and R. Moessner, Fermionic response from fractionalization in an insulating two-dimensional magnet, *Nature Physics* **12**, 912 (2016).
- [507] A. Banerjee, C. Bridges, J.-Q. Yan, A. Aczel, L. Li, M. Stone, G. Granroth, M. Lumsden, Y. Yiu, J. Knolle, *et al.*, Proximate kitaev quantum spin liquid behaviour in a honeycomb magnet, *Nature materials* **15**, 733 (2016).
- [508] J. A. Sears, M. Songvilay, K. Plumb, J. Clancy, Y. Qiu, Y. Zhao, D. Parshall, and Y.-J. Kim, Magnetic order in α -rucl 3: A honeycomb-lattice quantum magnet with strong spin-orbit coupling, *Physical Review B* **91**, 144420 (2015).
- [509] L. J. Sandilands, Y. Tian, K. W. Plumb, Y.-J. Kim, and K. S. Burch, Scattering continuum and possible fractionalized excitations in α -rucl 3, *Physical review letters* **114**, 147201 (2015).
- [510] P. Maksimov and A. Chernyshev, Rethinking α -rucl 3, *Physical Review Research* **2**, 033011 (2020).
- [511] R. Yadav, N. A. Bogdanov, V. M. Katukuri, S. Nishimoto, J. Van Den Brink, and L. Hozoi, Kitaev exchange and field-induced quantum spin-liquid states in honeycomb α -rucl3, *Scientific reports* **6**, 37925 (2016).
- [512] S. M. Winter, A. A. Tsirlin, M. Daghofer, J. van den Brink, Y. Singh, P. Gegenwart, and R. Valentí, Models and materials for generalized kitaev magnetism, *Journal of Physics: Condensed Matter* **29**, 493002 (2017).
- [513] S. M. Winter, Y. Li, H. O. Jeschke, and R. Valentí, Challenges in design of kitaev materials: Magnetic interactions from competing energy scales, *Physical Review B* **93**, 214431 (2016).
- [514] M. Gohlke, R. Verresen, R. Moessner, and F. Pollmann, Dynamics of the kitaev-heisenberg model, *Physical review letters* **119**, 157203 (2017).
- [515] P. Laurell and S. Okamoto, Dynamical and thermal magnetic properties of the kitaev spin liquid candidate α -rucl3, *npj quantum materials* **5**, 2 (2020).
- [516] M. Imada, A. Fujimori, and Y. Tokura, Metal-insulator transitions, *Reviews of modern physics* **70**, 1039 (1998).
- [517] N. A. Spaldin, *Magnetic materials: fundamentals and applications* (Cambridge university press, 2010).
- [518] J. Orenstein and A. Millis, Advances in the physics of high-temperature superconductivity, *Science* **288**, 468 (2000).
- [519] E. Dagotto, Correlated electrons in high-temperature superconductors, *Reviews of Modern Physics* **66**, 763 (1994).
- [520] J. Hubbard, Electron correlations in narrow energy bands. ii. the degenerate band case, *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **277**, 237 (1964).
- [521] J. Kanamori, Electron correlation and ferromagnetism of transition metals, *Progress of Theoretical Physics* **30**, 275 (1963).
- [522] M. C. Gutzwiller, Effect of correlation on the ferromagnetism of transition metals, *Physical Review Letters* **10**, 159 (1963).
- [523] D. P. Arovas, E. Berg, S. A. Kivelson, and S. Raghu, The hubbard model, *Annual review of condensed matter physics* **13**, 239 (2022).
- [524] E. H. Lieb and F.-Y. Wu, Absence of mott transition in an exact solution of the short-range, one-band model in one dimension, *Physical Review Letters* **20**, 1445 (1968).
- [525] W. Metzner and D. Vollhardt, Correlated lattice fermions in $d = \infty$ dimensions, *Physical review letters* **62**, 324 (1989).
- [526] J. Oitmaa, C. Hamer, and W. Zheng, *Series expansion methods for strongly interacting lattice models* (Cambridge University Press, 2006).
- [527] J. P. LeBlanc, A. E. Antipov, F. Becca, I. W. Bulik, G. K.-L. Chan, C.-M. Chung, Y. Deng, M. Ferrero, T. M. Henderson, C. A. Jiménez-Hoyos, *et al.*, Solutions of the two-dimensional hubbard model: Benchmarks and results from a wide range of numerical algorithms, *Physical Review X* **5**, 041041 (2015).
- [528] B.-X. Zheng, C.-M. Chung, P. Corboz, G. Ehlers, M.-P. Qin, R. M. Noack, H. Shi, S. R. White, S. Zhang, and G. K.-L. Chan, Stripe order in the underdoped region of the two-dimensional hubbard model, *Science* **358**, 1155 (2017).
- [529] V. Emery, Theory of high- T_c superconductivity in oxides, *Physical Review Letters* **58**, 2794 (1987).
- [530] M. Greiter and R. Thomale, No evidence for spontaneous orbital currents in numerical studies of three-band models for the cuo planes of high-temperature superconductors, *Physical review letters* **99**, 027005 (2007).
- [531] E. W. Huang, C. B. Mendl, S. Liu, S. Johnston, H.-C. Jiang, B. Moritz, and T. P. Devereaux, Numerical evi-

- dence of fluctuating stripes in the normal state of high- t_c cuprate superconductors, *Science* **358**, 1161 (2017).
- [532] E. Vitali, H. Shi, A. Chiciak, and S. Zhang, Metal-insulator transition in the ground state of the three-band hubbard model at half filling, *Physical Review B* **99**, 165116 (2019).
- [533] Z.-H. Cui, C. Sun, U. Ray, B.-X. Zheng, Q. Sun, and G. K.-L. Chan, Ground-state phase diagram of the three-band hubbard model from density matrix embedding theory, *Physical Review Research* **2**, 043259 (2020).
- [534] B. Fauqué, Y. Sidis, V. Hinkov, S. Pailhes, C. Lin, X. Chaud, and P. Bourges, Magnetic order in the pseudogap phase of high- t_c superconductors, *Physical Review Letters* **96**, 197001 (2006).
- [535] M. S. Hybertsen, M. Schlüter, and N. E. Christensen, Calculation of coulomb-interaction parameters for La_2CuO_4 using a constrained-density-functional approach, *Physical Review B* **39**, 9028 (1989).
- [536] R. L. Martin, Electronic localization in the cuprates, *Physical Review B* **53**, 15501 (1996).
- [537] W. Hanke, M. L. Kiesel, M. Aichhorn, S. Brehm, and E. Arrighoni, The 3-band hubbard-model versus the 1-band model for the high- t_c cuprates: Pairing dynamics, superconductivity and the ground-state phase diagram, *The European Physical Journal Special Topics* **188**, 15 (2010).
- [538] A. Damascelli, Probing the Electronic Structure of Complex Systems by ARPES, *Physica Scripta* **T109**, 61 (2004).
- [539] T. J. Zuehlsdorff, A. Montoya-Castillo, J. A. Napoli, T. E. Markland, and C. M. Isborn, Optical spectra in the condensed phase: Capturing anharmonic and vibronic features using dynamic and static approaches, *The Journal of Chemical Physics* **151**, 074111 (2019).
- [540] Z. Wiethorn, K. Hunter, T. Zuehlsdorff, and A. Montoya-Castillo, Beyond the Condon limit: Condensed phase optical spectra from atomistic simulations, *arXiv e-prints*, arXiv:2310.04333 (2023), arXiv:2310.04333 [physics.chem-ph].
- [541] B. D. Patterson *et al.*, Coherent Science at the Swiss-FEL X-Ray Laser, *New J. Phys.* **12**, 035012 (2010).
- [542] S. P. Weathersby *et al.*, Mega-Electron-Volt Ultrafast Electron Diffraction at SLAC National Accelerator Laboratory, *Rev. Sci. Instrum* **86**, 073702 (2015).
- [543] M. C. Fischer, J. W. Wilson, F. E. Robles, and W. S. Warren, Invited Review Article: Pump-Probe Microscopy, *Rev. Sci. Instrum* **87**, 031101 (2016).
- [544] M. Buzzi, M. Först, R. Mankowsky, and A. Cavalleri, Probing Dynamics in Quantum Materials with Femtosecond X-rays, *Nat. Rev. Mater* **3**, 299 (2018).
- [545] P. J. Ollitrault, A. Baiardi, M. Reiher, and I. Tavernelli, Hardware efficient quantum algorithms for vibrational structure calculations, *Chem. Sci.* **11**, 6842 (2020).
- [546] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, G. Salton, S. Wang, and F. G. S. L. Brandão, Quantum algorithms: A survey of applications and end-to-end complexities (2023), arXiv:2310.03011 [quant-ph].
- [547] D. Trenev, P. J. Ollitrault, S. M. Harwood, T. P. Gujarati, S. Raman, A. Mezzacapo, and S. Mostame, Refining resource estimation for the quantum computation of vibrational molecular spectra through trotter error analysis (2023), arXiv:2311.03719 [quant-ph].
- [548] N. P. D. Sawaya, F. Paesani, and D. P. Tabor, Near- and long-term quantum algorithmic approaches for vibrational spectroscopy, *Phys. Rev. A* **104**, 062419 (2021).
- [549] S. T. Stober, S. M. Harwood, D. Trenev, P. K. Barkoutsos, T. P. Gujarati, and S. Mostame, Considerations for evaluating thermodynamic properties with hybrid quantum-classical computing work flows, *Phys. Rev. A* **105**, 012425 (2022).
- [550] B. Demmig-Adams and W. Adams, Photosynthesis and partitioning — photoinhibition, in *Encyclopedia of Applied Plant Sciences*, edited by B. Thomas (Elsevier, Oxford, 2003) pp. 707–714.
- [551] J. S. Higgins, L. T. Lloyd, S. H. Sohail, M. A. Allodi, J. P. Otto, R. G. Saer, R. E. Wood, S. C. Massey, P.-C. Ting, R. E. Blankenship, and G. S. Engel, Photosynthesis tunes quantum-mechanical mixing of electronic and vibrational states to steer exciton energy transfer, *Proceedings of the National Academy of Sciences* **118**, e2018240118 (2021).
- [552] T. Ritz, A. Damjanović, and K. Schulten, The quantum physics of photosynthesis, *ChemPhysChem* **3**, 243 (2002).
- [553] N. Cox, D. A. Pantazis, and W. Lubitz, Current understanding of the mechanism of water oxidation in photosystem ii and its relation to xfel data, *Annual Review of Biochemistry* **89**, 795 (2020), PMID: 32208765, <https://doi.org/10.1146/annurev-biochem-011520-104801>.
- [554] S. Mostame, J. Huh, C. Kreisbeck, A. J. Kerman, T. Fujita, A. Eisfeld, and A. Aspuru-Guzik, Emulation of complex open quantum systems using superconducting qubits, *Quantum Information Processing* **16**, 44 (2016).
- [555] B. Jaderberg, A. Eisfeld, D. Jaksch, and S. Mostame, Recompile-enhanced simulation of electron-phonon dynamics on ibm quantum computers, *New Journal of Physics* **24**, 093017 (2022).
- [556] N. Science and T. C. (US), *Materials genome initiative for global competitiveness* (Executive Office of the President, National Science and Technology Council, 2011).
- [557] D.-B. Long, N. V. Tkachenko, Q. Feng, X. Li, A. I. Boldyrev, J. Yang, and L.-M. Yang, Two-dimensional bimetal-embedded expanded phthalocyanine monolayers: a class of multifunctional materials with fascinating properties, arXiv preprint arXiv:2307.05237 (2023).
- [558] D. R. Smith, J. B. Pendry, and M. C. K. Wiltshire, Metamaterials and negative refractive index, *Science* **305**, 788 (2004).
- [559] S. A. Cummer, J. Christensen, and A. Alù, Controlling sound with acoustic metamaterials, *Nature Reviews Materials* **1**, 16001 (2016).
- [560] K. Bertoldi, V. Vitelli, J. Christensen, and M. van Hecke, Flexible mechanical metamaterials, *Nature Reviews Materials* **2**, 17066 (2017).
- [561] X. Tian, P. M. Lee, Y. J. Tan, T. L. Y. Wu, H. Yao, M. Zhang, Z. Li, K. A. Ng, B. C. K. Tee, and J. S. Ho, Wireless body sensor networks based on metamaterial textiles, *Nature Electronics* **2**, 243 (2019).
- [562] S. Kim, W. Shang, S. Moon, T. Pastega, E. Lee, and T. Luo, High-performance transparent radiative cooler designed by quantum computing, *ACS Energy Letters* **7**, 4134 (2022).
- [563] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, Evidence for quantum annealing with more than one

- hundred qubits, *Nature Physics* **10**, 218 (2014).
- [564] S. Rendle, Factorization machines, in *2010 IEEE International Conference on Data Mining* (2010) pp. 995–1000.
- [565] K. Kitai, J. Guo, S. Ju, S. Tanaka, K. Tsuda, J. Shiomi, and R. Tamura, Designing metamaterials with quantum annealing and factorization machines, *Phys. Rev. Res.* **2**, 013319 (2020).
- [566] S. Izawa, K. Kitai, S. Tanaka, R. Tamura, and K. Tsuda, Continuous black-box optimization with an ising machine and random subspace coding, *Phys. Rev. Res.* **4**, 023062 (2022).