Snowmass White Paper:
Quantum Computing Systems and Software for High-energy Physics Research[*]

Travis S. Humble,[†] Andrea Delgado, Raphael Pooser, Christopher Seck, Ryan Bennink, Vicente Leyton-Ortega, C.-C. Joseph Wang, Eugene Dumitrescu, Titus Morris, Kathleen Hamilton, Dmitry Lyakh, Prasanna Date, Yan Wang, Nicholas A. Peters, Katherine J. Evans, and Marcel Demarteau
Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA

Alex McCaskey
NVIDIA, Santa Clara CA, USA

Thien Nguyen
Quantum Brilliance, The Australian National University, Canberra, Australia

Susan Clark and Melissa Reville
Sandia National Laboratories, Albuquerque, New Mexico, USA

Alberto Di Meglio, Michele Grossi, and Sofia Vallecorsa
CERN - European Organization for Nuclear Research, Geneva, Switzerland

Kerstin Borras,[‡] Karl Jansen, and Dirk Krücker,
Deutsches Elektronen-Synchrotron - DESY, Hamburg and Zeuthen, Germany
(Dated: March 15, 2022)

Quantum computing offers a new paradigm for advancing high-energy physics research by enabling novel methods for representing and reasoning about fundamental quantum mechanical phenomena. Realizing these ideals will require the development of novel computational tools for modeling and simulation, detection and classification, data analysis, and forecasting of high-energy physics (HEP) experiments. While the emerging hardware, software, and applications of quantum computing are exciting opportunities, significant gaps remain in integrating such techniques into the HEP community research programs. Here we identify both the challenges and opportunities for developing quantum computing systems and software to advance HEP discovery science. We describe opportunities for the focused development of algorithms, applications, software, hardware, and infrastructure to support both practical and theoretical applications of quantum computing to HEP problems within the next 10 years.

I. MOTIVATION

The aim of high-energy physics (HEP) is to understand matter at the most fundamental level. The current understanding of the building blocks of the universe and their interactions is embodied in the so-called standard model (SM) of particle physics. In this model, fundamental particles are quantum mechanical entities with fixed quantum numbers such as mass, spin, electric, and color charge, and are regarded as excitations of relativistic quantum fields. While the field of HEP has enjoyed many decades of progress, a long-standing challenge for both theorists and experimentalists in HEP is that the properties and behaviors of perturbative quantum fields are mathematically complex, present infrared and ultraviolet divergence challenges, and expansions converge slowly when the coupling constant is large. Furthermore, the time evolution of final state particles and their hadronization is particularly complex and are difficult to accurately simulate on digital computers, even on high performance computer (HPC) systems.

Meanwhile, a new approach to computing that explicitly leverages the complexities of quantum mechanics for computational purposes has been in development. In the last few years, prototype quantum computers capable of performing small computations using a few dozen quantum bits ("qubits") have become widely accessible, prompting a burgeoning field of research in applications of near-term and future quantum computers. Quantum computing (QC) approaches have been developed and demonstrated for applications in many disciplines

including chemistry, materials science, nuclear physics, and HEP. However, the field of QC is far from mature and the potential of QC for HEP has only begun to be explored.

In this paper, we highlight some of the significant opportunities and challenges for the development of quantum computing systems and associated software to deliver unprecedented capabilities for HEP discovery science. In Sec. II, we survey current approaches to applying QC to HEP. In Sec. III, we describe the leading technical challenges and gaps facing the adoption of quantum computing for HEP research problems, while in Sec. IV, we respond with long-term research priorities that will address these gaps in the coming decade.

II. BACKGROUND

Quantum computing may be viewed as part of the larger subject of quantum information science and technology (QIST), which concerns the use of quantum systems to store, transmit, and process information [1]. QIST involves the precise isolation and control of well-defined quantum particles and fields. These quantum physical systems can be used to simulate naturally occurring quantum phenomena (such as quantum chromodynamics (QCD)) or to solve computational tasks unrelated to physical phenomena (such as clustering or regression on a dataset). As will be discussed later in this section, key concepts in quantum computing such as superposition, entanglement, and computational complexity have also provided new, information-oriented ways to address fundamental questions in HEP.

For the most part, efforts to apply QC in HEP to date have consisted of adapting established general-purpose quantum algorithms and tools to HEP-specific tasks. While most demonstrations have involved only toy problems, there have been a few demonstrations in which the quantum computations were competitive with more traditional methods. Some examples include the simulation of a full $SU(2)$ gauge theory on a noisy, intermediate-scale quantum (NISQ) device [2], the training of large-scale quantum machine learning (QML) models for supernova classification [3], and the generation of synthetic detector data [4].

In the United States, much of the work to date exploring the potential of quantum information science for HEP has been through the Department of Energy (DOE) QuantISED program. QuantISED was launched in 2018 as part of the DOE Office of Science (SC) QIS initiative following a series of community roundtables, pilot studies, and [5, 6]. The program has relied upon interdisciplinary collaboration between HEP and QIS researchers and partnerships between DOE laboratories, university consortia, and hardware vendors such as IBM, Rigetti, and Google. Major topic areas addressed by QuantISED include:

- **Cosmos and qubits:** theoretical connections between cosmic physics and qubit systems that can be studied in laboratory settings. In particular, the study of quantum gravity in the context of quantum information has yielded new insights. For example, principles of computational complexity and quantum error correction have helped scientists investigate the interiors of black holes [7] and whether gravity emerges from entanglement [8].

- **Foundational theory:** formulations of gauge theories amenable to simulation on quantum computers [9].

- **Quantum simulation:** algorithms for simulating quantum mechanical systems including scalar quantum field theories [10], nuclear physics [11], and other many-body systems [12].

- **Quantum computing:** quantum-enhanced machine learning and data analysis for HEP experiments. Developments in this area include: quantum machine learning algorithms [3, 13, 14] and optimization algorithms for the analysis of collider data and quantum simulation of field theories [15].

- **Quantum sensors:** sensors leveraging quantum information concepts and technologies to yield new detection capabilities. One advancement in this area is quantum-enhanced searches for dark matter [16, 17].

- **Quantum technology:** advances in QIST and its integration with HEP technology. Efforts in this topic include quantum internet [18], noise mitigation techniques for near-term quantum devices [19], analysis of the technological requirements for HEP programs [20], and others.

While there is not space to describe all of these efforts in detail, recent work relating black hole thermodynamics and quantum information theory exemplifies the kinds of fruitful connections being developed between QIS and HEP. A long-standing problem in particle physics is the so-called "black-hole information paradox", which concerns the apparent loss of information that occurs when a particle disappears inside a black hole. The particle's information chaotically mixes with all the other matter and information inside the black hole, generating quantum entanglement between distant regions and seemingly making it impossible to retrieve. Recent landmark work [21, 22] suggests that the information passing the event horizon of a black hole is not lost forever but is eventually released.

In the field of quantum information, this process of apparent information loss through widespread entanglement is known as scrambling. A quantum processor can directly measure the scrambling-induced spread of initially localized information via the decay of out-of-time-ordered correlation (OTOC) functions. In [23], the authors realize scrambling behavior in a multi-qutrit sys-

tem. In [24], a seven-qubit circuit executed on an ion-trap quantum computer enabled the authors to bound the scrambling-induced decay of the OTOC measurement experimentally. These two experiments pave the way for laboratory investigations of the black hole information paradox, something that would hardly be imaginable without QC.

III. CHALLENGES

The emerging field of QC faces many technical challenges related to the development of algorithms, software, hardware, and infrastructure to support HEP research. In this section, we identify how these existing research challenges influence on-going efforts to solve key problems in HEP research using QC. In Sec. III A, we survey algorithms and applications which are known to apply to HEP challenge problems, and we identify several key gaps in extending these methods beyond proof of principle efforts. In Sec. III B, we review how existing computational tools, including quantum programming methods and numerical simulators, have influenced the accessibility and adoption of QC for the HEP community. In Sec. III C, we briefly summarize early efforts in developing testbeds for experimental QC with a focus on HEP application areas. Then, in Sec. III D, we discuss some of the broader infrastructure challenges that face adoption of QC by the HEP community.

A. Algorithms and Applications

The Hamiltonian formulation of lattice quantum field theories allows us to use the tools from quantum computation to solve problems in strongly correlated quantum systems by understanding the dynamics in terms of quantum circuits or representing the action of a measurement as a projection in a Hilbert space [25, 26]. Different approaches in the quantum computation domain for quantum field theories cover subatomic many-body physics [27, 28], real-time model dynamics of lattice gauge theories [29, 30], self-verifying variational quantum simulation of the lattice Schwinger model [31], Non-abelian SU(2) lattice gauge theories on superconducting devices [32], optical abelian and simulations of non-abelian gauge theories on ultra-cold atoms platforms [33, 34], and simulation of Z(2) lattice gauge theories with dynamical fermionic matter [35, 36].

In most of these applications, the limited number of qubits and gate fidelities available in QC hardware implementations represent bottlenecks for a proper comparison with classical algorithms and, therefore, any assessment of quantum advantage. On the other hand, the application of QCD with the lattice gauge model [37, 38] represents an exciting challenge on the real-device implementation that requires a novel design following quantum device limitations, state-of-the-art classical simula-

tors handling a decent number of qubits to benchmark the quantum computation algorithms with existing classical approaches, and the exploration of gravity/QFT duality [26].

From an algorithmic point of view, several challenges address the exploration of variational quantum algorithms in a higher dimension, design of quantum algorithms beyond the existing hybrid classical-quantum format with the combination of different strategies to solve problems in HEP, exploitation of quantum singular-value-decomposition (SVD) methods for unfolding problems in HEP, and on-hardware implementation of quantum algorithms to reconstruct Feynman path integrals.

1. Digital and Analog Simulations

Interestingly, digital, analog, and mixed digital-analog models of quantum computation have been considered within the context of lattice gauge theories [39, 40]. In the digital realm the problem of a *gauge-invariant* problem formulation is already a non-trivial matter with quantum-error-correction-like protocols [41] and manifestly gauge-invariant formulations [42, 43] having been considered. After overcoming the first hurdle of problem formulation the ever-growing set of quantum algorithms, including quantum phase estimation and linear combinations of unitaries or block encoding, can be used to compute the various quantities of interest appearing in the target theory [43–45]. These algorithms seem suitable for fault tolerant computation rather than near term quantum computational models.

Analog approaches have long been taken with synthetic gauge fields in cold atom simulators being a prime example of analog simulation [46]. Analog realizations of gauge invariant encodings, such loop-string-hadron encoding in cold atom systems [47], have also appeared. Lastly, with some modest abstraction one popular near term approach has been to leverage the natural device physics in order to augment the conventional digital gate-sets with analog processes [48]. For example, select bosonic vibrational modes can be used to simulate $U(1)$ gauge variables in a resource efficient manner [49], e.g., using trapped ion systems [50].

2. Variational Methods

Generative modeling learns how to generate a distribution by analyzing a known data. The model itself may then be used to draw additional samples representative of the distribution. With a quantum processor, preparing and sampling from arbitrary distributions can be directly implemented, and there is no need to compute partition functions which may be intractable. Additionally, using quantum systems has an attractive scaling characteristic: a model built on n qubits can potentially prepare and

sample from a distribution over 2^n states, where each state is a length-n binary bitstring.

In quantum machine learning there is a broad field of research which includes using classical machine learning algorithms for studying and simulating quantum systems [51–53] and using quantum algorithms to implement analogues of classical machine learning algorithms. Many machine learning applications in HEP leverage generative modeling [54–57], and the leading model is the generative adversarial network (GAN). Quantum GAN models [58–60] have been developed that utilize a trainable quantum circuit model for the generator network, yet adversarial training can still encounter mode collapse as with classical models. However, there are other generative models that can be constructed with quantum circuits including quantum Boltzmann machines [61, 62] and quantum circuit Born machines (and related Ising Born Machines) [63, 64].

Near-term demonstrations of QML can reach comparable performance with, but have not yet demonstrated an advantage over, classical machine learning models. Classical machine learning and deep learning can utilize large, distributed computing clusters to implement models with millions of parameters. Coupled with parallelization, these training workflows dwarf the current state of the art in quantum machine learning that uses variational, hybrid algorithms. However, as the size and capabilities of available quantum devices improve, larger quantum models can be built and executed. While near-term demonstrations and applications may not outperform classical ML models, their development is essential for understanding how quantum models learn.

3. Noise Mitigation

The design space associated with quantum generative models, and parameterized quantum models in general, is large. Best practices for building parameterized quantum circuit models, in order to effectively prepare and sample from these distributions, remains an open question [65]. Many heuristic approaches abound including: using sparse parameterizations (single rotation gates instead of arbitrary three angle decompositions) reduces training overhead, at the expense of trainability; and using fewer entangling and two-qubit operations can reduce susceptibility to hardware noise and improve transpilation for sparse hardware connectivity graphs, but at the expense of the correlations that can be modeled.

Mitigating hardware noise requires efficient noise characterization [66, 67]. Many methods of error mitigation that rely on regression modeling (e.g. [68]) are focused on mitigating the effect of noise on expectation values. For generative modeling and algorithms that use the distribution over all basis states, matrix-based error mitigation methods have been incorporated into variational training workflows [69], and scalable error mitigation methods are an active area of research [70, 71]. These methods are

needed to assess the stability of the quantum devices as well as reproducibility of the applications [72].

4. Quantum Machine Learning

Modern machine learning (ML) techniques, including deep learning, are rapidly being applied, adapted, and developed for HEP applications. ML is currently used across the different areas of particle physics from collider physics [73–75] to the study of the cosmos [76, 77] and quantum gravity. ML is also present throughout the various stages of HEP studies from the simulation of hypothesized particles [78], in the parameterization of cross-sections calculated for sensitivity analyses [79], to the analysis of experimental data [73–75]. Most of these applications have replaced non-ML methods and techniques and can process experimental data very efficiently. These methods and techniques have even been deployed in FPGAs [80] and integrated into the data acquisition systems for fast selection of information.

More recently, applications at the intersection of quantum computing and machine learning have been explored to analyze experimental data. The hope is to adapt or develop algorithms that can efficiently process HEP data. One example includes algorithms to construct physics objects amenable to analysis from the signals generated in a particle detector–i.e., the clustering of detector hits into so-called tracks for reconstructing a particle's trajectory [81–88] or tracks and calorimeter energy depositions into jets [89–91]. Furthermore, quantum-assisted algorithms have been explored in unsupervised learning settings to classify jets according to their origin (b-tagging) [92], generative tasks [93–95], and the selection of events or interactions along with background suppression [3, 13, 96–108]. In particular, generative models have been explored extensively as an alternative for the simulation of particle interactions and the detector's response to such interactions [4, 94].

Leveraging quantum computers for machine learning in HEP has several advantages. In general, quantum computers can operate more efficiently in high-dimensional tensor product spaces because of quantum superposition, which is vital for analyzing large-scale and complex HEP data patterns. Moreover, they can leverage quantum tunneling to train machine learning models more efficiently [109–112]. In some cases, quantum-assisted techniques have been shown to train on less data than their classical counterparts. This might be particularly useful in the context of data augmentation in simulation tasks, where using a generative model can reduce the resources allocated for traditional ML methods.

Nonetheless, even though the applications of QML to HEP data are many, there are several challenges associated with the trainability and deployment of these models on NISQ devices. For example, the size and complexity of HEP datasets requires an efficient encoding into quantum states, resulting in a large overhead in

pre-processing. Many training workflows serially encode training data and paralellization requiring either multiple devices or large quantum devices that can run multiple circuits with no cross-talk or interference between circuits. Furthermore, there is little understanding on how to tailor ansatz/circuit design for specific applications and a need for error correction/mitigation techniques that can be scalably constructed and incorporated into variational training workflows.

B. Software and Compilers

1. Compiling

Improvements in quantum computers, and their subsequent enabling of increasingly nontrivial simulations of fermionic systems, have pushed the development of novel encoding of fermionic degrees of freedom. The commonly used mappings from these degrees of freedom to that of spin (i.e., the Jordan-Wigner, Bravyi-Kitaev, or parity mappings) in general lead to largely non-local spin Hamiltonians and unnecessarily large qubit numbers per fermionic mode mapped. These properties drastically compound circuit and measurement complexity of simulations. Towards countering these expenses, several works have shown both general and system specific mappings that outperform Jordan-Wigner and others in terms of locality of spin Hamiltonians as well as the number of fermionic modes per qubit [113, 114].

Although these economic mappings are useful, their utility in the fault-tolerant era may be overshadowed by instead using *logical* fermionic mappings [115–118]. These avoid the overhead associated with fermionic-to-qubit mapping followed by the cost of error corrected codes by directly mapping the underlying fermionic symmetries to error correcting code spaces. While the benefit is obvious in terms of qubit and gate overhead, such embeddings require couching fundamental fermionic operations in terms of fault tolerant operations, which will likely encourage simple, more physically motivated implementations of algorithms and compiling of code.

Explorations of such mappings, and their relative quantum resource efficiency for gauge-invariant HEP problems of interest remain an open, but important, first step. There have been several successful HEP problem implementations, generally focusing on how to truncate gauge degrees of freedom while simultaneously satisfying gauge constraints efficiently [42, 119]. This has paved the way, similarly to fermionic mappings, for methods of how to embed HEP simulations natively in error detection and correction schemes [41, 120]. These methods have so far proven to be very specific to the target models and its symmetries. Thus the endeavor to find general, optimal mappings will be essential in the coming years for scaling up near term HEP simulations into the fault tolerant era.

C. Experimental Testbeds

The current state of the art for testbed devices consists of quantum processors that contain dozens of qubits. These machines are capable of circuit depths in the 100s of gates and up to 1000s of gates in the most advanced systems. Alternately, algorithms for simulating HEP problems of interest have been implemented on a comparably smaller number of qubits and circuit depths.

Given the current level of noise and imperfection in practical realizations of quantum computing devices, efficient numerical simulators capable of running on large-scale heterogeneous (HPC) platforms will continue to play an important role in characterization, verification, and validation of quantum devices, as well as analysis of the quantum algorithms intended for those devices. Probing the potential for quantum advantage will require extensive research and development in classical algorithms for quantum computing simulations in order to make a true boundary where quantum wins over classical.

An ability to utilize existing and future classical hardware efficiently is another dimension of optimization on the classical computing side. Furthermore, the HEP-related quantum simulators will require generalization to the qudit-based representation of quantum devices and algorithms. In order to better understand and cope with the noise and decoherence, analog-level simulations of open quantum device dynamics will become vital for their performance optimization. These simulations will require significant HPC resources coupled with highly optimized computer codes. Yet another challenge is related to loading classical (or quantum) data to quantum devices–a problem that can also benefit from advanced quantum-inspired classical algorithms coupled to large-scale HPC resources.

D. Instrument and Data Networks

Historically, HEP computing has been performed on sizeable, purpose-built computing systems. These began as single-site computing facilities but have evolved into the distributed computing grids that we use today. Therefore, integration with current heterogeneous computing resources is needed to accelerate the adoption of quantum computing technologies within the HEP community. For example the Fermilab HEP cloud is a portal to diverse computing resources on local clusters, campus farms, grid resources, commercial clouds, HPC centers, and quantum computing resources.

IV. PRIORITIES

A. Algorithms and Applications

With ever increasing numbers of qubits into the hundreds or thousands, while remaining in the NISQ era,

quantum algorithm research for HEP applications to improve/overtake classical methods is an important research focus for the next ten years. Quantum computing for HEP provides a great challenge and important intellectual stimulus for overcoming practical quantum information challenges with big data. In addition, quantum computing for simulating lattice gauge theories provides both motivation and a framework for interdisciplinary research towards developing special-purpose digital and analog quantum simulators, and ultimately scalable universal quantum computers.

Another potential area that may yield quantum advantage involves processing the large data sets resulting from detected events. For instance, simulations and on-hardware implementations for the unfolding problem can exploit both quantum generative modeling as well as singular value decomposition. Recent advances also enabled light-front encoding of relativistic fields and subsequent simulation with variational methods [121, 122]. Encoding techniques allow for large simulations to be mapped onto relatively few qubits via the Lie algebraic method to reduce simulation complexity via gauge invariant and sub-sector encodings pertaining to a symmetry [123], while computing the quantum gradient on the Reimannian manifold allows for increased circuit depths when complexity reductions have been maximized. Approaches beyond the current hybrid quantum-classical paradigm, which combine efficient quantum algorithms with classical distributed algorithms, may also benefit HEP applications.

1. Variational Algorithms

One advantage to using quantum models for statistical learning tasks is that with n qubits, one can build a parameterized model that can efficiently prepare arbitrary distributions over 2^n states. The development of QML models for HEP applications can provide insight into how data can be efficiently transformed in high-dimensional probability spaces. The challenges for these approaches to quantum computing can be organized into two areas: designing circuit ansatz models that can efficiently transform data into high-dimensional quantum Hilbert spaces, and building scalable training methods that can optimize these models in noisy landscapes.

Many state of the art quantum algorithms in the HEP application space that can be deployed on near-term hardware are variational algorithms: these are hybrid workflows that utilize quantum and classical processors to train parameterized quantum models. Training a variational model is a non-convex optimization problem, and the training task can be NP-hard [124]. In classical machine learning, gradient-based training methods can be scaled to large-scale applications in non-convex and convex optimization, and efficient training of deep learning models has been facilitated through the use of backpropagation. Analogous methods for variational workflows

have been developed [125, 126], but have not been demonstrated yet on hardware.

Currently, gradient-based and gradient-free optimization are the leading approaches to training. Gradient-based training is predicted to improve convergence [127], yet developing gradient-based optimization routines that can effectively incorporate queue-based access to quantum hardware should be a priority for near-term research. As the number of parameters in a circuit increases, and the number of qubits in the register increases, effective scheduling of the experiments that are needed in order to execute one step of an optimization method is needed. Hardware access which allows for job batching (where a set of experiments sent by a single user are executed and results are returned as a single group) can facilitate efficient training. With batched jobs, it may be possible to execute all circuits needed for one optimization step, either evaluation of the loss function gradient (e.g. for gradient descent [128–130]), executing a direct search (e.g. for particle swarm [131] or coordinate descent [4]), or for gradient approximation (e.g. using finite differences or SPSA [132]). Additionally, a larger number of shots are needed to efficiently sample from n-qubit states.

2. Noisy Algorithms

The challenges of effective state preparation, circuit training in noisy landscapes are not mutually exclusive. One observed effect of hardware noise is the loss landscape flattening and the existence of "barren plateaus" [133, 134]. In a flattened landscape, the efficacy of gradient-based training will be reduced. However, the impact of barren plateaus can be mitigated with circuit ansatz and cost function design [135–137].

There are a diverse set of error mitigation methods developed for short depth circuits [138]. But highly expressible variational models may not be short depth and variational training is not guaranteed to prepare well-verified states. An alternate approach to error mitigation is to use matrix-based methods where sparse linear filters are used to reduce spurious counts in prepared states. Matrix-based error mitigation methods can be incorporated into variational training methods as a data post-processing state, as the added cost of preparing the linear filter.

3. Machine Learning Algorithms

The known challenges for quantum machine learning (QML) in the HEP application space can be organized into three areas: embedding data into quantum Hilbert spaces, training quantum models in noisy landscapes, and working with quantum data. Currently, the capabilities in classical distributed computing make it unlikely that a quantum advantage will be observed with classical data, or classical representations of quantum data, using

QML models deployed on near-term hardware with < 100 qubits. However, the development of QML models for HEP applications can provide insight into how data can be efficiently transformed in high-dimensional probability spaces.

These three challenges are not mutually exclusive. For example the challenges of training quantum models are encountered irrespective of the data source (classical data, classical representation of quantum data, or quantum data). Examples of how challenges of embedding data can manifest: finding low-dimensional representations of classical data into qubit registers, or high-fidelity transmission of quantum data from sensors to qubits used for QML. In addition to these, it is important in QML to explore purely quantum approaches for both training and inferencing, i.e. approaches that are not hybrid in nature (such as variational approaches). Working with the NISQ-era machines in the near-term, it is also important to explore quantum error-correction and error-mitigation subroutines specific to QML circuits used in HEP. From a logistical point of view, accessing quantum computers requires better scheduling policies and remote access protocols. Lastly, today's quantum computers are noisy and error-prone—this severely limits our ability to work with large-scale data reliably. In this regard, it is important to build fault tolerant quantum computers that are also scalable.

B. Software and Compilers

As discussed in Sec. III A and Sec. IV A, hybrid quantum algorithms, such as those based on variational principles, have emerged as a promising candidate for quantum computational advantage in HEP areas, especially in the Noisy Intermediate-Scale Quantum (NISQ) era [139]. The iterative nature of these algorithms necessitates a tight integration of quantum hardware with classical computing resources. For instance, the execution time of quantum circuits that involves a variational parameter loop has been shown to be dominated (more than 90% [140]) by classical computing procedures for compilation, control, and data transfer. As quantum computing capabilities mature, software infrastructures for compilation and control will necessarily need to improve in order to utilize the full computing power of heterogeneous quantum-classical resources for handling practical applications.

We envision the need for system-level and hardware-agnostic compiler toolchains akin to those we have today for classical computing - extensible, modular systems with unified intermediate representations that enable a wide array of optimization and code-generation techniques for multiple source languages or target architectures [141]. The most important aspect of any compilation toolchain design is the Intermediate Representation (IR), which is how code is represented in the compiler. A standardized, robust, and forward-looking quantum-classical IR is, therefore, an essential step towards ensuring full interoperability within the quantum software ecosystem, promoting contributions from broad communities, and enabling optimization and transformation of quantum programs that we discuss in detail in the later point.

For instance, multiple community-led efforts [142, 143] in the field have coalesced around the idea of leveraging existing classical compiler infrastructures, like the LLVM, for quantum computing. The approach ensures the low-level coupling of quantum computation with classical computing resources and also supports integration with classical languages and tools readily available in those compiler toolchains. For example, the Quantum Intermediate Representation (QIR) Alliance [143] is an open-source community within the Linux Foundation focusing on enabling an LLVM-based IR specification for quantum computation (QIR), as well as developing tooling around this unified representation. Oak Ridge National Laboratory is a founding member of this organization along with several industry players, including Microsoft, Quantinuum, Rigetti, and Quantum Circuits Inc. Importantly, by leveraging LLVM for quantum compilations we could also benefit from its powerful tools, especially the Multi-Level Intermediate Representation (MLIR) [144] library, which is designed for heterogeneous hardware and domain-specific languages analogous to that of quantum-classical computing paradigm. In this regard, we have successfully demonstrated a prototype compiler toolchain, QCOR [145], capable of compiling the novel OpenQASM version 3 [146] quantum programming language to LLVM IR adhering to the QIR specification adopting the multi-stage, progressive lowering approach of MLIR. This compilation technology is specifically relevant to many HEP application use cases whereby the algorithmic procedure is best captured by high-level domain-specific descriptions, such as analog and digital quantum simulation.

An essential feature of compiler toolchains is the ability to apply transformations on the IR to improve execution quality, e.g., code's runtime or the computing resources required, while preserving the semantics of the input program. In compiler technology, such transformations are often called *passes* since they can be assembled into a pipeline of passes each of which processes the IR and applies its transformation rule.

In the NISQ era of quantum computing, compiler passes that implement circuit simplification and efficient hardware topology mapping are indispensable to any quantum software stack. A broad variety of techniques have been developed quantum circuit simplifications at the gate level, such as those based on ZX-calculus [147, 148], template-based [149] or peep-hole [150] optimizations. Similarly, various circuit-to-hardware placement methods have been developed to minimize the number of qubit swapping operations [151, 152] required or to maximize circuit execution fidelity by mapping the circuit to best-performing qubits [153].

For dynamical Hamiltonian simulation algorithms that are particularly relevant to HEP applications, techniques such as advanced term splitting [154] and ordering [155, 156] or algebraic circuit compression [157] can also be incorporated as optimization passes targeting the high-level description of the quantum algorithms where algorithmic semantics of the program are expressed. In this respect, a quantum compiler infrastructure that supports a hierarchy of abstractions, such as the MLIR, is highly desirable since it can retain high-level algorithmic information in the IR which compiler optimizers can reason [158].

As quantum computers evolve into fault-tolerant machines, the associated software infrastructure will also need to be upgraded in order to handle quantum error correction protocols [159]. Specifically, we envision that quantum programs would always be expressed in terms of logical operations whose corresponding physical qubit level instructions incorporating a target-dependent quantum error correction code are compiled by a compiler software infrastructure [160]. In this regard, quantum compilers leveraging classical computing technology, such as LLVM, would be best suited to performing this logical-physical transformation thanks to their ability to incorporate a variety of classical computing resources required by the error decoding and correction protocol, e.g., minimum weight perfect matching [161] or maximum likelihood decoding [162] for surface code [163]. Furthermore, as quantum programming languages move towards this fault-tolerant paradigm whereby qubit measurement-based control flow is intertwined with that of conventional computing for classical data, we can leverage the vast amount of IR optimization capability from the LLVM infrastructure to handle common classical optimization passes like function inlining or loop unrolling. While these optimization passes are intrinsically classical, they help unlock many quantum-related simplification patterns which would otherwise be obscure to the quantum optimizer due to the complex nature of the IR tree with control flows.

1. Encodings

The HEP-specific encodings discussed in section II are an important first step to solving gauge-invariant HEP problems of interest in a resource-efficient way. Explorations of such mappings, and their relative quantum resource efficiency for gauge-invariant HEP problems of interest remain an open, but important first step. There have been several successful HEP problem implementations, generally focusing on how to truncate gauge degrees of freedom while simultaneously satisfying gauge constraints efficiently [42, 119]. This has paved the way, similarly to fermionic mappings, for methods of how to actually embed HEP simulations natively in error detection and correction schemes [41, 120]. These methods have so far proven to be very specific to the target mod-

els and its symmetries. Thus the endeavor to find general, optimal mappings will be essential in the coming years for scaling up near term HEP simulations into the fault tolerant era. Comparing these methods to general fault tolerant encodings for universal QC will yield insight into the efficacy of building custom machines dedicated to HEP problems which implement custom encodings in hardware rather than reconfigurable devices capable of universal operations.

2. Numerical Simulators

Since numerical simulators will continue to play a vital role in quantum device verification and validation as well as in quantum algorithm analysis [164, 165], the relevant research priorities are concerned with scaling numerical simulation techniques towards larger NISQ devices as well as improving their ability to faithfully reproduce the behavior of actual quantum hardware via more accurate noise models, in particular proper modeling of the multi-qubit cross-talk [166]. In turn, devising more accurate noise models will require further research and development in scalable approaches for modeling open quantum system dynamics coupled with either a Markovian or non-Markovian bath. Further research related to simulations of multi-level qudit systems at their native Hamiltonian level will be necessary for assessing their computational power as compared to the qubit-based quantum devices. Pulse-level simulations, optimal quantum control, and quantum gate implementation optimization will form another direction of simulation heavy research and development efforts. More efficient numerical tensor-algebraic techniques coupled with advances in classical machine learning will be required for scalable characterization of larger NISQ devices. On the engineering side, all these newly devised numerical simulation techniques will have to be implemented in an efficient manner in order to fully exploit the computational power of Exascale HPC platforms which are becoming widely available worldwide [159, 167–169].

C. Experimental Hardware

1. Testbeds

It is in the interest of the HEP community to build specialized devices to address the above application needs. For example, multiple platforms offer analog-level operations to users (ie, pulse controls). In many scenarios, access to lower level, analog operations allows for higher accuracy in quantum simulation problems. Further, hybrid analog-digital simulations have shown increased accuracy in field theory problems for scattering calculations [170]. Currently, HEP programs - especially those in the QuantISED program in the US - utilize testbeds around the world, including the DOE-supported testbeds consisting

of trapped ion and superconducting devices, to solve prototypical problems of interest in the field. For example, the information scrambling problem [23] has been demonstrated on superconducting hardware [171], while variational methods for solving problems using the light-field encoding has been mapped onto trapped ion platforms [172]. Uncovering the most useful aspects of the testbeds in these demonstrations will help us discover the most important ingredients for testbeds that support HEP goals. In particular, the "close to the metal", analog controls were a defining factor in these demonstrations.

In the NISQ era, fully analog systems will also be useful for implementing complex Hamiltonians for which digital, prefault-tolerant platforms do not support sufficient circuit depth. Hamiltonians for problems of interest can be mapped to trapped ion systems' motional mode degrees of freedom, or the spin-spin coupling present in neutral atom, quantum dot, or other platforms which support both analog spin-spin coupling and tunable transverse fields. In the future, when digital platforms move closer to fault tolerance, direct fermionic (HEP problem-specific) encodings should be supported in custom testbeds. These encodings ensure that the resources required in a given testbed required to represent a specific HEP problem are far fewer than that required for general purpose, universal quantum simulation. A custom, fault tolerant fermionic code may not be suitable to simulate Boson sampling, for example, but it would be ideal for simulating scattering in fermionic field theories.

2. Analog Simulators

While digital quantum computers scale to the extent needed for fault tolerant quantum computations, *analog* quantum simulators are beginning to push state of the art, for example probing topological spin liquids on a programmable device at a scale not previously observed on neutral atom-based simulators [173–175]. Analog neutral atomic physics supplements photonic, superconducting, trapped-ion, and many more analog modalities which currently exist. There has also been recent research interest in exploring digital-analog computational modalities (in which hard operations are done in an analog fashion and digital operations play a complimentary supporting role, e.g. providing local basis transformations) and their application to HEP (and related condensed-matter) models as discussed in the previous section. Despite this promise, developing current analog methodologies to solve substantially more complex HEP domain problems is a research avenue in its own right.

3. HEP-dedicated Testbeds

For HEP, an open and relevant question is whether quantum computers can efficiently simulate quantum field theories (QFT). QFT encompasses all fundamental interactions, possibly excluding gravity. The simulation of QFT has HEP applications related to event generators for quantum chromodynamics (QCD) simulations of nuclear matter. Outside HEP, it might find applications in the simulation of strongly coupled theories and the characterization of the computational complexity of quantum states.

Nonetheless, the simulation of quantum field theories is not a trivial task due to the sign problem. In physics, the sign problem is encountered in calculations of the properties of a quantum mechanical system with a large number of strongly interacting fermions or QFTs involving a non-zero density of strongly interacting fermions. Since the particles are strongly interacting, perturbation theory is inapplicable, and one is forced to use brute-force numerical methods. The sign problem is one of the major unsolved problems in the physics of many-particle systems, limiting progress in nuclear physics, preventing the ab initio calculation of properties of nuclear matter, and limiting our understanding of nuclei and neutron stars. In quantum field theory, it prevents the use of lattice QCD [176] to predict the phases and properties of quark matter [177].

An exciting and promising alternative to the simulation of QFT is based on the idea of adiabatic quantum computation (AQC) [178, 179]. In general, AQC is as powerful as universal quantum computation when a non-stoquastic Hamiltonian is used [180], which is the case of Hamiltonians dealing with the sign problem. One of the critical aspects of AQC is quantum annealing [181], which is a metaheuristic algorithm to solve combinatorial optimization problems by changing the parameters adiabatically or even non-adiabatically. Measurements on currently available quantum annealers are done only on the standard computational basis [182]; thereby, the Hamiltonian is stoquastic.

Thus, a quantum testbed that fully realizes or exploits the power of AQC is needed to show that quantum speedup can be achieved when non-stochastic Hamiltonians relevant to HEP are used. Some models experiencing first-order and second-order phase transitions using non-stoquastic Hamiltonians have been proposed and have demonstrated a quantum speedup [183]. This natural and efficient way to simulate QFT on a quantum testbed will eliminate the need to reformulate HEP-relevant Hamiltonians (such as QCD) to fit a finite-dimensional Hilbert space amenable to circuit implementations.

Formulating the future of QCD studies as relying on a successful implementation of AQC provides an alternative to collider experiments, where large statistics are needed to understand rare processes. If we can simulate the QCD Hamiltonian we can study some observables in the context of information theory. For example, understanding the probabilistic parton distribution in terms of entangling entropy [184]. In other words, in a collider experiment we have access to a probability distribution by collapsing the wave function to one of the states, by

using a quantum-mechanical simulation we have access to the full density matrix.

D. Infrastructures, Platforms and Data Networks

1. Software Development Frameworks, Benchmarks, and Optimisation

The HEP community has developed over the years a broad range of algorithms and methods for different applications, from optimization to simulation and machine learning. Initial work on developing and assessing quantum equivalents of common algorithms have produced valuable insight [185]. However, the performance of NISQ systems is still limited and full advantage is expected only once reasonable fault tolerant devices will become available. To fully exploit the potential of quantum algorithms, software and hardware systems should be designed as part of co-development efforts across the HEP and quantum computing science and engineering communities. Experience in the classical domain shows that innovation in hardware technology, readily available to software and algorithm experts, can boost the development of new software stacks and algorithms that, in turn, produce critical feedback to hardware development in a virtuous cycle. Examples include hardware-aware compiler optimization, design of application-aware hardware architectures (i.e. systems dedicated to accelerating quantum simulations) and so on.

A fundamental component of measuring the progress and impact of quantum algorithms is the establishment of open benchmarking frameworks and platforms where different combinations of algorithms, software and hardware architecture can be tested, and the results published as a means of establishing baselines for further development. Such benchmarks can form the base for the development of common software tools and interfaces with an increasing level of abstraction and platform neutrality as typical today for classical computing resources.

As an example, in 2021, CERN launched the ABAQUS project with the objective to build an open, easily accessible platform for researchers to become familiar with different quantum computing architectures and tools, produce or get access to benchmarks over different metrics, and build a community database of knowledge of the applicability of quantum algorithms to HEP workloads. There are specific aspects that need to be investigated and studied. For example, there are easier simulations that are only polynomially hard, such as tensor networks with limited entanglement or limited number of non-Clifford gates that scale polynomially in the number of Clifford gates, but exponentially in the number of non-Clifford gates. It needs to be understood how to adapt and employ these simplifications.

2. Heterogeneous Computing Networks

When dealing with the quantum computing paradigm, the role of infrastructure deserves a particular attention. Quantum computing should not over time require separate computing infrastructures, application platforms, or languages, but must be integrated into the existing cloud-based physical and software infrastructure of today and enable the deployment of hybrid, heterogeneous workloads where the quantum processing units (QPUs) are part of a broad range of accelerators. Presently classical resources can reliably store, manage and process huge quantities of data, while quantum devices are expected to efficiently explore high-dimensional spaces to extract insights and identify optimal answers.

The development and deployment of such hybrid infrastructures is currently happening across a series of steps. As current NISQ resources are limited in availability and efficiency, experience is built by simulating quantum algorithms and building noise mitigation strategies on testbeds of classical resources of increasing scale. Since memory requirements double for every qubit added in the simulation and quickly hit an exponential wall, simulation of large scale systems can be performed by means of distributed and parallelised computation across HPC clusters of CPU/GPU. Novel architectural frameworks to integrate quantum capabilities APIs in cloud architectures are also emerging [186].

At the same time, it is critical to start building paths to architecture optimization of hybrid resources, workload splitting and hybrid algorithms, scheduling, data flows, and circuits transpilation processes. Looking at hybrid models of variational classifiers, reinforcement-learning optimisation algorithms, quantum-classic GAN models for simulation, the infrastructure is evolving rapidly to make the classic and quantum interaction more efficient, defining containerised modules to be tested on distributed HPC environments and loaded on real quantum computers. This will need to be achieved by integrating HPC and quantum computing strategies as part of national and international initiatives.

One example of how such integration may be addressed is the Quantum Computing User Program (QCUP) at the Oak Ridge Leadership Computing Facility, a US Department of Energy user facility operated by Oak Ridge National Laboratory [187]. QCUP offers access to a diversity of quantum computing systems following merit-based review of user proposal for demonstrate scientific applications. In addition, these users are capable of integrating their quantum computing workflows alongside conventional HPC systems. As quantum computing technology matures, this setting will offer a natural means by which to integrate quantum computing and conventional HPC systems for purposes of advanced computation.

3. Quantum Data Networks

A particular aspect of future Quantum/HPC infrastructures is the role of quantum data networks. Today quantum computing and quantum communication are usually handled as related by separate fields of applications of quantum technologies. However, quantum infrastructures, time and frequency distribution networks, and the development of quantum data technologies will play a fundamental and transformative role in physics experiments at different energy regimes (for example standard reference signals for anti-matter and low-energy experiments, such as ASACUSA and ALPHA, using high-precision laser spectroscopy [188]).

Many national and international initiatives are being deployed to build future quantum networks and the HEP community must remain an active part of this development as it was at the beginning of the century in the development of distributed grid and cloud networks.

4. International Collaborations and Co-Development

The HEP community is distributed over the whole globe and the collaboration across institutes has always played a fundamental role in pushing the limits of science and technology. For example the development and operation of the Worldwide LHC Computing Grid (WLCG) was and still is instrumental for joint progress and employing all brain power wherever located in the world.

The same approach needs to be taken in the development of future quantum infrastructures, Synergies between the Snowmass exercise and equivalent discussions in the European Particle Physics Strategy and other similar initiatives will be critical. Joint R&D projects, testbeds, progressive development of common frameworks and tools across US Quantum Information Science institutes, the CERN Quantum Technology Initiative [189], DESY QUANTUM, and many other national and international initiatives in EU and other countries will accelerate the development and adoption of quantum computing across the HEP community. In the same way, attention should be given to co-development not only across HEP, but also between HEP institutes, industry, and other scientific disciplines, such as Astrophysics, Cosmology, Earth Observation, Climate research, Condensed Matter, Chemistry, Biology, and many more. Often very similar challenges appear in different disciplines and can be solved with one solution.

ACKNOWLEDGEMENTS

[1] T. S. Humble, H. Thapliyal, E. Munoz-Coreas, F. A. Mohiyaddin, and R. S. Bennink, Quantum computing circuits and devices, IEEE Design & Test **36**, 69 (2019).

[2] Y. Atas, J. Zhang, R. Lewis, A. Jahanpour, J. F. Haase, and C. A. Muschik, Su(2) hadrons on a quantum computer (2021).

[3] E. Peters, J. Caldeira, A. Ho, S. Leichenauer, M. Mohseni, H. Neven, P. Spentzouris, D. Strain, and G. N. Perdue, Machine learning of high dimensional data on a noisy quantum processor, npj Quantum Information **7**, 10.1038/s41534-021-00498-9 (2021).

[4] A. Delgado and K. E. Hamilton, Unsupervised quantum circuit learning in high energy physics (2022), arXiv:2203.03578 [quant-ph].

[5] Report of the doe roundtable held february 25, 2016 (2016).

[6] E. Farhi, S. Jordan, P. Hayden, M. Lukin, J. Maldacena, J. Preskill, P. Shor, J. Taylor, and C. Williams, Grand challenges - science.osti.gov (2015).

[7] I. Kim, E. Tang, and J. Preskill, The ghost in the radiation: robust encodings of the black hole interior, Journal of High Energy Physics **2020**, 10.1007/jhep06(2020)031 (2020).

[8] A. Periwal, E. S. Cooper, P. Kunkel, J. F. Wienand, E. J. Davis, and M. Schleier-Smith, Programmable interactions and emergent geometry in an array of atom clouds, Nature **600**, 630 (2021).

[9] M. Carena, H. Lamm, Y.-Y. Li, and W. Liu, Lattice renormalization of quantum simulations, Phys. Rev. D **104**, 094519 (2021).

[10] K. Yeter-Aydeniz, E. F. Dumitrescu, A. J. McCaskey, R. S. Bennink, R. C. Pooser, and G. Siopsis, Scalar quantum field theories as a benchmark for near-term quantum computers, Physical Review A **99**, 032306 (2019).

[11] K. Yeter-Aydeniz, R. C. Pooser, and G. Siopsis, Practical quantum computation of chemical and nuclear energy levels using quantum imaginary time evolution and lanczos algorithms, npj Quantum Information **6**, 1 (2020).

[12] K. Yeter-Aydeniz, G. Siopsis, and R. C. Pooser, Scattering in the ising model with the quantum lanczos algorithm, New Journal of Physics **23**, 043033 (2021).

[13] S. L. Wu, J. Chan, W. Guan, S. Sun, A. Wang, C. Zhou, M. Livny, F. Carminati, A. Di Meglio, A. C. Y. Li, and et al., Application of quantum machine learning using the quantum variational classifier method to high

energy physics analysis at the LHC on IBM quantum computer simulator and hardware with 10 qubits, Journal of Physics G: Nuclear and Particle Physics **48**, 125003 (2021).

[14] S. L. Wu, S. Sun, W. Guan, C. Zhou, J. Chan, C. L. Cheng, T. Pham, Y. Qian, A. Z. Wang, R. Zhang, and et al., Application of quantum machine learning using the quantum kernel algorithm on high energy physics analysis at the LHC, Physical Review Research **3**, 10.1103/physrevresearch.3.033221 (2021).

[15] B. Nachman, D. Provasoli, W. A. de Jong, and C. W. Bauer, Quantum algorithm for high energy physics simulations, Physical Review Letters **126**, 10.1103/physrevlett.126.062001 (2021).

[16] K. Wurtz, B. M. Brubaker, Y. Jiang, E. P. Ruddy, D. A. Palken, and K. W. Lehnert, A cavity entanglement and state swapping method to accelerate the search for axion dark matter (2021).

[17] A. V. Dixit, S. Chakram, K. He, A. Agrawal, R. K. Naik, D. I. Schuster, and A. Chou, Searching for dark matter with a superconducting qubit, Phys. Rev. Lett. **126**, 141302 (2021).

[18] R. Valivarthi, S. I. Davis, C. Peña, S. Xie, N. Lauk, L. Narváez, J. P. Allmaras, A. D. Beyer, Y. Gim, M. Hussein, G. Iskander, H. L. Kim, B. Korzh, A. Mueller, M. Rominsky, M. Shaw, D. Tang, E. E. Wollman, C. Simon, P. Spentzouris, D. Oblak, N. Sinclair, and M. Spiropulu, Teleportation systems toward a quantum internet, PRX Quantum **1**, 020317 (2020).

[19] B. Nachman, M. Urbanek, W. A. de Jong, and C. W. Bauer, Unfolding quantum computer readout noise, npj Quantum Information **6**, 10.1038/s41534-020-00309-7 (2020).

[20] E. J. Gustafson and H. Lamm, Toward quantum simulations of F_2 gauge theory without state preparation, Phys. Rev. D **103**, 054507 (2021).

[21] A. Almheiri, N. Engelhardt, D. Marolf, and H. Maxfield, The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole, J. High Energ. Phys. **2019** (12), 63.

[22] G. Penington, Entanglement wedge reconstruction and the information paradox, J. High Energ. Phys. **2020** (9), 2.

[23] M. Blok, V. Ramasesh, T. Schuster, K. O'Brien, J. Kreikebaum, D. Dahlen, A. Morvan, B. Yoshida, N. Yao, and I. Siddiqi, Quantum Information Scrambling on a Superconducting Qutrit Processor, Physical Review X **11**, 021010 (2021), publisher: American Physical Society.

[24] K. A. Landsman, C. Figgatt, T. Schuster, N. M. Linke, B. Yoshida, N. Y. Yao, and C. Monroe, Verified quantum information scrambling, Nature **567**, 61 (2019).

[25] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. Van Acoleyen, F. Verstraete, U.-J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, Simulating lattice gauge theories within quantum technologies, The European Physical Journal D **74**, 165 (2020).

[26] A. J. Buser, H. Gharibyan, M. Hanada, M. Honda, and J. Liu, Quantum simulation of gauge theory via orbifold lattice, Energ. Phys. **34**, 2021 (2021).

[27] S. P. Jordan, K. S. Lee, and J. Preskill, Quantum algorithms for quantum field theories, Science **336**, 1130 (2012).

[28] H.-H. Lu, N. Klco, J. M. Lukens, T. D. Morris, A. Bansal, A. Ekström, G. Hagen, T. Papenbrock, A. M. Weiner, M. J. Savage, et al., Simulations of subatomic many-body physics on a quantum frequency processor, arXiv:1810.03959 (2018).

[29] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller, et al., Real-time dynamics of lattice gauge theories with a few-qubit quantum computer, Nature **534**, 516 (2016).

[30] C. Muschik, M. Heyl, E. Martinez, T. Monz, P. Schindler, B. Vogell, M. Dalmonte, P. Hauke, R. Blatt, and P. Zoller, U(1) Wilson lattice gauge theories in digital quantum simulators, New J. Phys. **19**, 103020 (2017).

[31] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, et al., Self-verifying variational quantum simulation of the lattice Schwinger model, arXiv:1810.03421 (2018).

[32] A. Mezzacapo, E. Rico, C. Sabín, I. Egusquiza, L. Lamata, and E. Solano, Non-abelian SU(2) lattice gauge theories in superconducting circuits, Phys. Rev. Lett. **115**, 240502 (2015).

[33] L. Tagliacozzo, A. Celi, A. Zamora, and M. Lewenstein, Optical abelian lattice gauge theories, Ann. Phys. **330**, 160 (2013).

[34] L. Tagliacozzo, A. Celi, P. Orland, M. Mitchell, and M. Lewenstein, Simulation of non-abelian gauge theories with optical lattices, Nat. Comm. **4**, 2615 (2013).

[35] E. Zohar, A. Farace, B. Reznik, and J. I. Cirac, Digital quantum simulation of Z_2 lattice gauge theories with dynamical fermionic matter, Phys. Rev. Lett. **118**, 070501 (2017).

[36] E. Zohar, A. Farace, B. Reznik, and J. I. Cirac, Digital lattice gauge theories, Physical Review A **95**, 023604 (2017).

[37] A. Kan and Y. Nam, Lattice quantum chromodynamics and electrodynamics on a universal quantum computer (2021), arXiv:2107.12769 [quant-ph].

[38] T. F. Stetina, A. Ciavarella, X. Li, and N. Wiebe, Simulating effective QED on quantum computers, Quantum **6**, 622 (2022).

[39] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, Rev. Mod. Phys. **86**, 153 (2014).

[40] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, K.-M. C. Fu, M. Greiner, K. R. Hazzard, R. G. Hulet, A. J. Kollár, B. L. Lev, M. D. Lukin, R. Ma, X. Mi, S. Misra, C. Monroe, K. Murch, Z. Nazario, K.-K. Ni, A. C. Potter, P. Roushan, M. Saffman, M. Schleier-Smith, I. Siddiqi, R. Simmonds, M. Singh, I. Spielman, K. Temme, D. S. Weiss, J. Vučković, V. Vuletić, J. Ye, and M. Zwierlein, Quantum simulators: Architectures and opportunities, PRX Quantum **2**, 017003 (2021).

[41] J. R. Stryker, Oracles for gauss's law on digital quantum computers, Phys. Rev. A **99**, 042301 (2019).

[42] Z. Davoudi, I. Raychowdhury, and A. Shaw, Search for efficient formulations for Hamiltonian simulation of non-abelian lattice gauge theories, Phys. Rev. D **104**, 074505 (2021).

[43] A. Ciavarella, N. Klco, and M. J. Savage, Trailhead for quantum simulation of SU(3) Yang-

Mills lattice gauge theory in the local multiplet basis, Physical Review D **103**, 094501 (2021), arXiv:2101.10227.

[44] A. Roggero and J. Carlson, Dynamic linear response quantum algorithm, Phys. Rev. C **100**, 034610 (2019).

[45] A. Roggero, Spectral-density estimation with the gaussian integral transform, Phys. Rev. A **102**, 022409 (2020).

[46] V. Galitski, G. Juzeliunas, and I. B. Spielman, Artificial gauge fields with ultracold atoms, Physics Today **72**, 39 (2019).

[47] R. Dasgupta and I. Raychowdhury, Cold-atom quantum simulator for string and hadron dynamics in non-Abelian lattice gauge theory, Physical Review A **105**, 023322 (2022), arXiv:2009.13969.

[48] P. Kairys and T. S. Humble, Parametrized hamiltonian simulation using quantum optimal control, Phys. Rev. A **104**, 042602 (2021).

[49] Y. Tong, V. V. Albert, J. R. McClean, J. Preskill, and Y. Su, Provably accurate simulation of gauge theories and bosonic systems, arXiv preprint arXiv:2110.06942 (2021).

[50] N. H. Nguyen, M. C. Tran, Y. Zhu, A. M. Green, C. H. Alderete, Z. Davoudi, and N. M. Linke, Digital quantum simulation of the schwinger model and symmetry protection with trapped ions, arXiv preprint arXiv:2112.14262 (2021).

[51] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, Nature communications **8**, 1 (2017).

[52] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nature Physics **14**, 447 (2018).

[53] G. Torlai and R. G. Melko, Machine-learning quantum states in the nisq era, Annual Review of Condensed Matter Physics **11**, 325 (2020).

[54] D. Salamani, S. Gadatsch, T. Golling, G. A. Stewart, A. Ghosh, D. Rousseau, A. Hasib, and J. Schaarschmidt, Deep generative models for fast shower simulation in atlas, in *2018 IEEE 14th International Conference on e-Science (e-Science)* (IEEE, 2018) pp. 348–348.

[55] R. Di Sipio, M. F. Giannelli, S. K. Haghighat, and S. Palazzo, DijetGAN: a generative-adversarial network approach for the simulation of QCD dijet events at the LHC, Journal of high energy physics **2019**, 1 (2019).

[56] Y. Alanazi, N. Sato, T. Liu, W. Melnitchouk, P. Ambrozewicz, F. Hauenstein, M. P. Kuchera, E. Pritchard, M. Robertson, R. Strauss, *et al.*, Simulation of electron-proton scattering events by a feature-augmented and transformed generative adversarial network (fat-gan), arXiv preprint arXiv:2001.11103 (2020).

[57] A. Hariri, D. Dyachkova, and S. Gleyzer, Graph generative models for fast detector simulations in high energy physics, arXiv preprint arXiv:2104.01725 (2021).

[58] P.-L. Dallaire-Demers and N. Killoran, Quantum generative adversarial networks, Physical Review A **98**, 012324 (2018).

[59] C. Zoufal, A. Lucchi, and S. Woerner, Quantum generative adversarial networks for learning and loading random distributions, npj Quantum Information **5**, 1 (2019).

[60] S. Chakrabarti, H. Yiming, T. Li, S. Feizi, and X. Wu, Quantum wasserstein generative adversarial networks, Advances in Neural Information Processing Systems **32** (2019).

[61] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, Quantum boltzmann machine, Physical Review X **8**, 021050 (2018).

[62] C. Zoufal, A. Lucchi, and S. Woerner, Variational quantum boltzmann machines, Quantum Machine Intelligence **3**, 1 (2021).

[63] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, npj Quantum Information **5**, 1 (2019).

[64] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The Born supremacy: quantum advantage and training of an Ising Born machine, npj Quantum Information **6**, 1 (2020).

[65] K. Nakaji and N. Yamamoto, Expressibility of the alternating layered ansatz for quantum computation, Quantum **5**, 434 (2021).

[66] M. L. Dahlhauser and T. S. Humble, Modeling noisy quantum circuits using experimental characterization, Phys. Rev. A **103**, 042603 (2021).

[67] M. L. Dahlhauser and T. S. Humble, Benchmarking characterization methods for noisy quantum circuits, arXiv preprint arXiv:2201.02243 (2022).

[68] A. Lowe, M. H. Gordon, P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, Unified approach to data-driven quantum error mitigation, Physical Review Research **3**, 033098 (2021).

[69] K. E. Hamilton and R. C. Pooser, Error-mitigated data-driven circuit learning on noisy quantum hardware, Quantum Machine Intelligence **2**, 1 (2020).

[70] K. E. Hamilton, T. Kharazi, T. Morris, A. J. McCaskey, R. S. Bennink, and R. C. Pooser, Scalable quantum processor noise characterization, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2020) pp. 430–440.

[71] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, Scalable mitigation of measurement errors on quantum computers, PRX Quantum **2**, 040326 (2021).

[72] S. Dasgupta and T. S. Humble, Characterizing the reproducibility of noisy quantum circuits, Entropy **24**, 244 (2022).

[73] A. Stakia, T. Dorigo, G. Banelli, D. Bortoletto, A. Casa, P. de Castro, C. Delaere, J. Donini, L. Finos, M. Gallinaro, A. Giammanco, A. Held, F. J. Morales, G. Kotkowski, S. P. Liew, F. Maltoni, G. Menardi, I. Papavergou, A. Saggio, B. Scarpa, G. C. Strong, C. Tosciri, J. Varela, P. Vischia, and A. Weiler, Advances in multi-variate analysis methods for new physics searches at the large hadron collider, Reviews in Physics **7**, 100063 (2021).

[74] V. Mikuni and F. Canelli, Point cloud transformers applied to collider physics, Machine Learning: Science and Technology **2**, 035027 (2021).

[75] B. Nachman and J. Thaler, Learning from many collider events at once, Physical Review D **103**, 10.1103/physrevd.103.116013 (2021).

[76] B. Ostdiek, A. D. Rivero, and C. Dvorkin, Image segmentation for analyzing galaxy-galaxy strong lensing systems, Astronomy & Astrophysics **657**, L14 (2022).

[77] F. List, N. L. Rodd, and G. F. Lewis, Extracting the galactic center excess' source-count distribution with neural nets, Physical Review D **104**, 10.1103/physrevd.104.123022 (2021).

[78] J. Brehmer and K. Cranmer, Simulation-based inference methods for particle physics (2020), arXiv:2010.06439 [hep-ph].

[79] S. Carrazza, J. Cruz-Martinez, and T. R. Rabemananjara, Compressing PDF sets using generative adversarial networks, The European Physical Journal C **81**, 10.1140/epjc/s10052-021-09338-8 (2021).

[80] Y. Iiyama, G. Cerminara, A. Gupta, J. Kieseler, V. Loncar, M. Pierini, S. R. Qasim, M. Rieger, S. Summers, G. V. Onsem, K. A. Wozniak, J. Ngadiuba, G. D. Guglielmo, J. Duarte, P. Harris, D. Rankin, S. Jindariani, M. Liu, K. Pedro, N. Tran, E. Kreinar, and Z. Wu, Distance-weighted graph neural networks on FPGAs for real-time particle reconstruction in high energy physics, Frontiers in Big Data **3**, 10.3389/fdata.2020.598927 (2021).

[81] D. Magano, A. Kumar, M. Kālis, A. Locāns, A. Glos, S. Pratapsi, G. Quinta, M. Dimitrijevs, A. Rivošs, P. Bargassa, J. Seixas, A. Ambainis, and Y. Omar, Quantum speedup for track reconstruction in particle accelerators, arXiv:2104.11583 (2021).

[82] A. Zlokapa, A. Anand, J.-R. Vlimant, J. M. Duarte, J. Job, D. Lidar, and M. Spiropulu, Charged particle tracking with quantum annealing optimization, Quantum Machine Intelligence **3**, 10.1007/s42484-021-00054-w (2021).

[83] F. Bapst, W. Bhimji, P. Calafiura, H. Gray, W. Lavrijsen, L. Linder, and A. Smith, A pattern recognition algorithm for quantum annealers, Computing and Software for Big Science **4**, 10.1007/s41781-019-0032-5 (2019).

[84] I. Shapoval and P. Calafiura, Quantum associative memory in hep track pattern recognition, EPJ Web of Conferences **214**, 01012 (2019).

[85] S. Das, A. J. Wildridge, S. B. Vaidya, and A. Jung, Track clustering with a quantum annealer for primary vertex reconstruction at hadron colliders, arXiv:1903.08879 (2020).

[86] C. Tüysüz, C. Rieger, K. Novotny, B. Demirköz, D. Dobos, K. Potamianos, S. Vallecorsa, J.-R. Vlimant, and R. Forster, Hybrid quantum classical graph neural networks for particle track reconstruction, Quantum Machine Intelligence **3**, 10.1007/s42484-021-00055-9 (2021).

[87] C. Tüysüz, F. Carminati, B. Demirköz, D. Dobos, F. Fracas, K. Novotny, K. Potamianos, S. Vallecorsa, and J.-R. Vlimant, Particle track reconstruction with quantum algorithms, EPJ Web of Conferences **245**, 09013 (2020).

[88] G. Quiroz, L. Ice, A. Delgado, and T. S. Humble, Particle track classification using quantum associative memory, NIM-A **1010**, 165557 (2021).

[89] A. Y. Wei, P. Naik, A. W. Harrow, and J. Thaler, Quantum algorithms for jet clustering, Physical Review D **101**, 10.1103/physrevd.101.094015 (2020).

[90] D. Pires, Y. Omar, and J. Seixas, Adiabatic quantum algorithm for multijet clustering in high energy physics, arXiv:2012.14514 (2020).

[91] D. Pires, P. Bargassa, J. Seixas, and Y. Omar, A digital quantum algorithm for jet clustering in high-energy physics (2021).

[92] A. Gianelle, P. Koppenburg, D. Lucchesi, D. Nicotra, E. Rodrigues, L. Sestini, J. de Vries, and D. Zuliani, Quantum machine learning for *b*-jet identification (2022), arXiv:2202.13943 [hep-ex].

[93] C. Bravo-Prieto, J. Baglio, M. Cè, A. Francis, D. M. Grabowska, and S. Carrazza, Style-based quantum generative adversarial networks for Monte Carlo events, arXiv:2110.06933 (2021).

[94] S. Y. Chang, S. Herbert, S. Vallecorsa, E. F. Combarro, and R. Duncan, Dual-parameterized quantum circuit GAN model in high energy physics, EPJ Web of Conferences **251**, 03050 (2021).

[95] A. Pérez-Salinas, J. Cruz-Martinez, A. A. Alhajri, and S. Carrazza, Determining the proton content with a quantum computer, Physical Review D **103**, 034027 (2021).

[96] A. Mott, J. Job, J.-R. Vlimant, D. Lidar, and M. Spiropulu, Solving a Higgs optimization problem with quantum annealing for machine learning, Nature **550**, 375 (2017).

[97] A. Zlokapa, A. Mott, J. Job, J.-R. Vlimant, D. Lidar, and M. Spiropulu, Quantum adiabatic machine learning by zooming into a region of the energy surface, Physical Review A **102**, 062405 (2020), DOI:10.1103/PhysRevA.102.062405.

[98] M. Kim, P. Ko, J. hyeon Park, and M. Park, Leveraging quantum annealer to identify an event-topology at high energy colliders, arXiv:2111.07806 (2021).

[99] J. Caldeira, J. Job, S. H. Adachi, B. Nord, and G. N. Perdue, Restricted Boltzmann machines for galaxy morphology classification with a quantum annealer (2020).

[100] V. Belis, S. González-Castillo, C. Reissel, S. Vallecorsa, E. F. Combarro, G. Dissertori, and F. Reiter, Higgs analysis with quantum classifiers, EPJ Web of Conferences **251**, 03070 (2021).

[101] K. Terashi, M. Kaneda, T. Kishimoto, M. Saito, R. Sawada, and J. Tanaka, Event classification with quantum machine learning in high-energy physics, Computing and Software for Big Science **5**, 10.1007/s41781-020-00047-7 (2021).

[102] A. A. Armenakas and O. K. Baker, Implementation and analysis of quantum computing application to Higgs boson reconstruction at the large hadron collider, Scientific Reports **11**, 10.1038/s41598-021-01552-4 (2021).

[103] P. Bargassa, T. Cabos, A. C. O. Choi, T. Hessel, and S. Cavinato, Quantum algorithm for the classification of supersymmetric top quark events, Physical Review D **104**, 10.1103/physrevd.104.096004 (2021).

[104] K. T. Matchev, P. Shyamsundar, and J. Smolinsky, A quantum algorithm for model independent searches for new physics, arXiv preprint: 2003.02181 (2020).

[105] A. Blance and M. Spannowsky, Quantum machine learning for particle physics using a variational quantum classifier, Journal of High Energy Physics **2021**, 10.1007/jhep02(2021)212 (2021).

[106] A. Blance and M. Spannowsky, Unsupervised event classification with graphs on classical and photonic quantum computers, Journal of High Energy Physics **2021**, 10.1007/jhep08(2021)170 (2021).

[107] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, Quantum convolutional neural networks for high energy physics data analysis, arXiv preprint: 2012.12177 (2020).

[108] J. Heredge, C. Hill, L. Hollenberg, and M. Sevior, Quantum support vector machines for continuum suppression in b meson decays, Computing and Software for Big Science 5, 10.1007/s41781-021-00075-x (2021).

[109] P. Date, D. Arthur, and L. Pusey-Nazzaro, Qubo formulations for training machine learning models, Scientific reports 11, 10029 (2021).

[110] P. Date and T. Potok, Adiabatic quantum linear regression, Scientific reports 11, 21905 (2021).

[111] D. Arthur and P. Date, Balanced k-means clustering on an adiabatic quantum computer, Quantum Information Processing 20, 1 (2021).

[112] P. Date, C. Schuman, R. Patton, and T. Potok, A classical-quantum hybrid approach for unsupervised probabilistic machine learning, in *Future of Information and Communication Conference* (Springer, 2019) pp. 98–117.

[113] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, Annals of Physics 298, 210 (2002).

[114] C. Derby, J. Klassen, J. Bausch, and T. Cubitt, Compact fermion to qubit mappings, Phys. Rev. B 104, 035118 (2021).

[115] K. Setia, S. Bravyi, A. Mezzacapo, and J. D. Whitfield, Superfast encodings for fermionic quantum simulation, Phys. Rev. Research 1, 033033 (2019).

[116] A. J. Landahl and B. C. A. Morrison, Logical majorana fermions for fault-tolerant quantum simulation (2021), arXiv:2110.10280 [quant-ph].

[117] Y. Li, Fault-tolerant fermionic quantum computation based on color code, Phys. Rev. A 98, 012336 (2018).

[118] Z. Jiang, J. McClean, R. Babbush, and H. Neven, Majorana loop stabilizer codes for error mitigation in fermionic quantum simulations, Phys. Rev. Applied 12, 064041 (2019).

[119] N. Klco and M. J. Savage, Hierarchical qubit maps and hierarchically implemented quantum error correction, Phys. Rev. A 104, 062425 (2021).

[120] A. Rajput, A. Roggero, and N. Wiebe, Quantum error correction with gauge symmetries (2021), arXiv:2112.05186 [quant-ph].

[121] M. Kreshchuk, W. M. Kirby, G. Goldstein, H. Beauchemin, and P. J. Love, Quantum Simulation of Quantum Field Theory in the Light-Front Formulation, arXiv:2002.04016 (2020), arXiv: 2002.04016.

[122] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, Light-Front Field Theory on Current Quantum Computers, Entropy 23, 597 (2021).

[123] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. J. Savage, Quantum-classical computation of Schwinger model dynamics using quantum computers, Physical Review A 98, 032331 (2018).

[124] L. Bittel and M. Kliesch, Training variational quantum algorithms is np-hard, Physical Review Letters 127, 120502 (2021).

[125] G. Verdon, J. Pye, and M. Broughton, A universal training algorithm for quantum deep learning (2018), arXiv:1806.09729 [quant-ph].

[126] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, Training deep quantum neural networks, Nature communications 11, 1 (2020).

[127] A. W. Harrow and J. C. Napp, Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms, Phys. Rev. Lett. 126, 140502 (2021).

[128] J.-G. Liu and L. Wang, Differentiable learning of quantum circuit Born machines, Physical Review A 98, 062324 (2018).

[129] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, Physical Review A 98, 032309 (2018).

[130] K. E. Hamilton, E. F. Dumitrescu, and R. C. Pooser, Generative model benchmarks for superconducting qubits, Physical Review A 99, 062323 (2019).

[131] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Brecque, *et al.*, Training of quantum circuits on a hybrid quantum computer, Science advances 5, eaaw9918 (2019).

[132] J. C. Spall *et al.*, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, IEEE transactions on automatic control 37, 332 (1992).

[133] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nature communications 9, 1 (2018).

[134] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, Nature communications 12, 1 (2021).

[135] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, Nature communications 12, 1 (2021).

[136] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, Absence of barren plateaus in quantum convolutional neural networks, Physical Review X 11, 041011 (2021).

[137] A. Uvarov and J. D. Biamonte, On barren plateaus and cost function locality in variational quantum algorithms, Journal of Physics A: Mathematical and Theoretical 54, 245301 (2021).

[138] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Physical review letters 119, 180509 (2017).

[139] J. Preskill, Quantum computing in the nisq era and beyond, Quantum 2, 79 (2018).

[140] A. Wack, H. Paik, A. Javadi-Abhari, P. Jurcevic, I. Faro, J. M. Gambetta, and B. R. Johnson, Quality, speed, and scale: three key attributes to measure the performance of near-term quantum computers, arXiv preprint arXiv:2110.14108 (2021).

[141] A. McCaskey, E. Dumitrescu, D. Liakh, and T. Humble, Hybrid programming for near-term quantum computing systems, in *2018 IEEE International Conference on Rebooting Computing (ICRC)* (IEEE, 2018) pp. 1–12.

[142] QED-C: Quantum economic development consortium, https://quantumconsortium.org/, accessed: 2022-02-21.

[143] QIR alliance, https://github.com/qir-alliance, accessed: 2022-02-21.

[144] C. Lattner, M. Amini, U. Bondhugula, A. Cohen, A. Davis, J. Pienaar, R. Riddle, T. Shpeisman, N. Vasilache, and O. Zinenko, MLIR: A compiler infrastructure for the end of Moore's law, arXiv preprint arXiv:2002.11054 (2020).

[145] T. Nguyen, A. Santana, T. Kharazi, D. Claudino, H. Finkel, and A. McCaskey, Extending C++ for Heterogeneous Quantum-Classical Computing, arXiv e-prints , arXiv:2010.03935 (2020), arXiv:2010.03935 [quant-ph].

[146] A. W. Cross, A. Javadi-Abhari, T. Alexander, N. de Beaudrap, L. S. Bishop, S. Heidel, C. A. Ryan, J. Smolin, J. M. Gambetta, and B. R. Johnson, OpenQASM 3: A broader and deeper quantum assembly language, arXiv preprint arXiv:2104.14722 (2021).

[147] J. van de Wetering, Zx-calculus for the working quantum computer scientist, arXiv preprint arXiv:2012.13966 (2020).

[148] A. Kissinger and J. van de Wetering, Reducing the number of non-clifford gates in quantum circuits, Phys. Rev. A 102, 022406 (2020).

[149] R. Iten, R. Moyard, T. Metger, D. Sutter, and S. Woerner, Exact and practical pattern matching for quantum circuit optimization, ACM Transactions on Quantum Computing 3, 1 (2022).

[150] 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, 1 (2018).

[151] M. Amy and V. Gheorghiu, staq—a full-stack quantum processing toolkit, Quantum Science and Technology 5, 034016 (2020).

[152] G. Li, Y. Ding, and Y. Xie, Tackling the qubit mapping problem for nisq-era quantum devices, in Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (2019) pp. 1001–1014.

[153] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers, in Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (2019) pp. 1015–1029.

[154] N. Hatano and M. Suzuki, Finding exponential product formulas of higher orders, in Quantum Annealing and Other Optimization Methods, edited by A. Das and B. K. Chakrabarti (Springer Berlin Heidelberg, Berlin, Heidelberg, 2005) pp. 37–68.

[155] A. M. Childs, A. Ostrander, and Y. Su, Faster quantum simulation by randomization, Quantum 3, 182 (2019).

[156] 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, 1218 (2019).

[157] D. Camps, E. Kökcü, L. Bassman, W. A. de Jong, A. F. Kemper, and R. Van Beeumen, An algebraic quantum circuit compression algorithm for Hamiltonian simulation, arXiv preprint arXiv:2108.03283 (2021).

[158] T. Nguyen, D. Lyakh, R. C. Pooser, T. S. Humble, T. Proctor, and M. Sarovar, Quantum circuit transformations with a multi-level intermediate representation compiler, arXiv preprint arXiv:2112.10677 (2021).

[159] T. S. Humble, A. McCaskey, D. I. Lyakh, M. Gowrishankar, A. Frisch, and T. Monz, Quantum computers for high-performance computing, IEEE Micro 41, 15 (2021).

[160] K. A. Britt and T. S. Humble, High-performance computing with quantum processing units, ACM Journal on Emerging Technologies in Computing Systems (JETC) 13, 1 (2017).

[161] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, Journal of Mathematical Physics 43, 4452 (2002).

[162] S. Bravyi, M. Suchara, and A. Vargo, Efficient algorithms for maximum likelihood decoding in the surface code, Phys. Rev. A 90, 032326 (2014).

[163] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A 86, 032324 (2012).

[164] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, et al., Quantum supremacy using a programmable superconducting processor, Nature 574, 505 (2019).

[165] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, and S. Mandrà, Establishing the quantum supremacy frontier with a 281 pflop/s simulation, Quantum Science and Technology 5, 034003 (2020).

[166] A. McCaskey, E. Dumitrescu, M. Chen, D. Lyakh, and T. Humble, Validating quantum-classical programming models with tensor network simulations, PloS one 13, e0206704 (2018).

[167] K. A. Britt, F. A. Mohiyaddin, and T. S. Humble, Quantum accelerators for high-performance computing systems, in 2017 IEEE International Conference on Rebooting Computing (ICRC) (IEEE, 2017) pp. 1–7.

[168] T. Nguyen, L. Bassman, P. C. Lotshaw, D. Lyakh, A. McCaskey, V. Leyton-Ortega, R. Pooser, W. Elwasif, T. S. Humble, and W. A. de Jong, Quasimo: A composable library to program hybrid workflows for quantum simulation, IET Quantum Communication 2, 160 (2021).

[169] T. Nguyen, L. Bassman, D. Lyakh, P. C. Lotshaw, A. McCaskey, R. S. Bennink, V. Leyton-Ortega, R. C. Pooser, T. S. Humble, and W. A. de Jong, Scalable programming workflows for validation of quantum computers, in 2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS) (IEEE, 2021) pp. 80–87.

[170] I. Arrazola, J. S. Pedernales, L. Lamata, and E. Solano, Digital-Analog Quantum Simulation of Spin Models in Trapped Ions, Scientific Reports 6, 30534 (2016).

[171] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, Light-front field theory on current quantum computers, Entropy 23, 597 (2021).

[172] M. Echevarria, I. Egusquiza, E. Rico, and G. Schnell, Quantum simulation of light-front parton correlators, Physical Review D 104, 10.1103/physrevd.104.014512 (2021).

[173] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, Science 374, 1242 (2021).

[174] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, Nature 595, 233 (2021).

[175] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, S. Choi, M. Serbyn, M. Greiner, V. Vuletic, and M. D. Lukin, Controlling quantum many-body dynamics in driven Rydberg atom arrays, Science **371**, 1355 (2021).

[176] P. de Forcrand, Simulating QCD at finite density (2010), arXiv:1005.0539 [hep-lat].

[177] O. Philipsen, Lattice calculations at non zero chemical potential, in *Proc. VIIIth Conf. Quark Confinement and the Hadron Spectrum* (Sissa Medialab, 2012).

[178] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution (2000), arXiv:quant-ph/0001106 [quant-ph].

[179] T. Albash and D. A. Lidar, Adiabatic quantum computation, Reviews of Modern Physics **90**, 10.1103/revmodphys.90.015002 (2018).

[180] M. Marvian, D. A. Lidar, and I. Hen, On the computational complexity of curing non-stoquastic Hamiltonians, Nature Communications **10** (2019).

[181] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, Physical Review E **58**, 5355 (1998).

[182] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, Quantum annealing with manufactured spins, Nature **473**, 194 (2011).

[183] K. Ikeda, Universal computation with quantum fields, Quantum Information Processing **19** (2020).

[184] V. Kuvshinov, V. Shaparau, and E. Bagashov, Quantum Optics Properties of QCD Vacuum, EPJ Web Conf. **164**, 07030 (2017).

[185] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, and J.-R. Vlimant, Quantum machine learning in high energy physics, Machine Learning: Science and Technology **2**, 011003 (2021).

[186] M. Grossi, L. Crippa, A. Aita, G. Bartoli, V. Sammarco, E. Picca, N. Said, F. Tramonto, and F. Mattei, A serverless cloud integration for quantum computing, arXiv preprint arXiv:2107.02007 (2021).

[187] Quantum Computing User Program, https://www.olcf.ornl.gov/olcf-resources/compute-systems/quantum-computing-user-program.

[188] C. Amsler, D. Barna, H. Breuker, S. Chesnevskaya, G. Costantini, R. Ferragut, M. Giammarchi, A. Gligorova, H. Higaki, M. Hori, E. D. Hunter, Y. Kanai, V. Kletzl, V. Kraxberger, N. Kuroda, A. Lanz, M. Leali, V. Mäckel, G. Maero, C. Malbrunot, V. Mascagna, Y. Matsuda, S. Migliorati, D. J. Murtagh, Y. Nagata, A. Nanda, L. Nowak, E. Pasino, W. Pirkl, M. Romé, M. C. Simon, M. Tajima, V. Toso, S. Ulmer, U. Uggerhøj, L. Venturelli, A. Weiser, E. Widmann, T. Wolz, Y. Yamazaki, and J. Zmeskal (ASACUSA Collaboration), *Status report of the ASACUSA experiment - progress in 2020 and plans for 2021*, Tech. Rep. (CERN, Geneva, 2021).

[189] A. Di Meglio, M. Doser, B. Frisch, D. M. Grabowska, M. Pierini, and S. Vallecorsa (CERN QTI Collaboration), *CERN Quantum Technology Initiative Strategy and Roadmap*, Tech. Rep. (2021).