

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

Lawrence Berkeley National Laboratory

LBL Publications

Title

Quantum Computing Technology Roadmaps and Capability Assessment for Scientific Computing - An analysis of use cases from the NERSC workload

Permalink

<https://escholarship.org/uc/item/0xb5w33t>

Authors

Camps, Daan
Rrapaj, Ermal
Klymko, Katherine
[et al.](#)

Publication Date

2025-09-11

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at <https://creativecommons.org/licenses/by-nc-sa/4.0/>



Report No. LBNL-2001699

Quantum Computing Technology Roadmaps and Capability Assessment for Scientific Computing

An analysis of use cases from the NERSC workload

Daan Camps Ermal Rrapaj Katherine Klymko Hyeongjin Kim Kevin Gott
Siva Darbha Jan Balewski Brian Austin Nicholas J. Wright

National Energy Research Scientific Computing Center (NERSC)
Lawrence Berkeley National Laboratory (LBNL)
Berkeley, CA, USA

September 11, 2025

Abstract

The National Energy Research Scientific Computing Center (NERSC), as the high-performance computing (HPC) facility for the Department of Energy’s Office of Science, recognizes the essential role of quantum computing in its future mission. In this report, we analyze the NERSC workload and identify materials science, quantum chemistry, and high-energy physics as the science domains and application areas that stand to benefit most from quantum computers. These domains jointly make up over 50% of the current NERSC production workload, which is illustrative of the impact quantum computing could have on NERSC’s mission going forward. We perform an extensive literature review and determine the quantum resources required to solve classically intractable problems within these science domains. This review also shows that the quantum resources required have consistently decreased over time due to algorithmic improvements and a deeper understanding of the problems. At the same time, public technology roadmaps from a collection of ten quantum computing companies predict a dramatic increase in capabilities over the next five to ten years. Our analysis reveals a significant overlap emerging in this time frame between the technological capabilities and the algorithmic requirements in these three scientific domains. We anticipate that the execution time of large-scale quantum workflows will become a major performance parameter and propose a simple metric, the Sustained Quantum System Performance (SQSP), to compare system-level performance and throughput for a heterogeneous workload.

Contents

Executive Summary	2
1 Introduction	4
2 NERSC Workload	7
3 A Primer on Quantum Computing	9
3.1 On qubits, gates, and circuits	9
3.2 Handling errors: quantum error mitigation and correction	10
3.3 Overview of quantum algorithms	12
4 Resource Estimates for the Quantum Relevant Workload	14
4.1 Condensed matter physics and materials science	15
4.2 Quantum chemistry	16
4.3 Nuclear and high energy physics	19
4.4 Other application areas	20
4.5 Summary	23
5 Vendor technology roadmaps	25
5.1 Overview of individual vendor roadmaps	26
5.1.1 Qubits based on superconducting circuits	27
5.1.2 Qubits based on trapped ions	30
5.1.3 Qubits based on neutral atoms	32
5.1.4 Vendors without public technology roadmaps	32
5.2 Overview of all vendor roadmaps	34
5.3 Vendor roadmaps versus application requirements	35
6 Execution times and system throughput	36
6.1 Execution times	36
6.2 Workload and system throughput	38
6.2.1 Prototypical workload	39
6.2.2 Prototypical quantum computers	39
6.2.3 Execution times of benchmark workload	40
6.2.4 Sustained quantum system performance	41
7 Conclusion	43
Acknowledgments	45
List of Acronyms	46
Glossary of Terms	48
Bibliography	52

Executive Summary

The National Energy Research Scientific Computing Center (NERSC) is charged with accelerating scientific discovery at the U.S. Department of Energy (DOE) Office of Science (SC) through High Performance Computing (HPC) and data analysis. As the DOE SC mission HPC user facility, NERSC pushes the envelope of HPC by strategically deploying new supercomputers on approximately a five year life cycle.

Quantum computers may fundamentally alter the HPC landscape and profoundly advance the DOE SC mission by the end of the decade. Quantum Computing (QC) is expected to *(i)* lead to exponential speedups for important scientific problems that classical HPC cannot solve, *(ii)* do so in an energy-efficient manner that is decoupled from Moore’s law, and *(iii)* offer opportunities and challenges to develop new paradigms for HPC-QC integration. The analysis in this report shows that **at least 50% of the DOE SC production workload, which aggregates the computational needs of more than 12,000 NERSC users across the DOE landscape, is allocated to solving problems for which exponential improvements in terms of problem complexity, time-to-solution, and/or solution accuracy are anticipated with the advent of QC.**

We focus on **three major scientific domains** of strategic relevance to NERSC’s mission — *materials science, quantum chemistry, and high energy physics*— and collect over 140 end-to-end resource estimates for benchmark problems from the scientific literature. We draw **three main conclusions** from this analysis. First, **Hamiltonian simulation** of quantum mechanical systems, also known as real-time evolution, is *the* key quantum algorithmic primitive upon which many known applications and speedups rely. Second, the **end-to-end resource estimates differ significantly by scientific domain**. Condensed matter physics offers a prime candidate for early quantum advantage as model problems relevant to materials science are more easily mapped to quantum computers and thus the least quantum resources are required at a given problem scale. Problems in quantum chemistry, including electronic structure, require an intermediate number of resources as the encoding overhead is larger. High energy physics problems, including lattice gauge theories, show the most dramatic scaling due to the encoding overhead encountered in including both fermionic and gauge degrees of freedom. Third, we show that, while asymptotically optimal algorithms exist, constant factor **algorithmic improvements over the past five years have reduced the quantum resources** needed to compute the ground state energy of a strongly-correlated molecule **by orders of magnitude**. We expect this trend to continue going forward and advance applications in all aforementioned domains. We show that current estimates of quantum resources required for scientific quantum advantage start at about 50 to a 100 qubits and a million quantum gate operations and go up from there.

In addition, we analyzed and combined **public technology roadmaps** from ten different quantum vendors and observe that all vendors **predict an exponential increase in quantum computer performance over the next decade by up to nine orders of magnitude**. When realized, this will unlock unprecedented capabilities that will accelerate scientific discovery at DOE SC and enable breakthroughs that were previously out of reach.

Figure 1 summarizes the vendor roadmap milestones by showing the highest performance QC systems expected at the end of 2025, in five years and ten years from today. The anticipated technological progress shows excellent overlap with the application requirements for known scientific quantum advantage of interest to DOE SC mission (area shown in blue), with potential for early scientific quantum utility within the next two years (area shown in white).

Finally, we discuss execution time estimates for quantum applications at scientific advantage scale and describe eight timescales —from less than a second up to more than a year— and their relevance




Figure 1: Schematic overview of quantum resources in terms of quantum gate count (y-axis) and qubit count (x-axis). The figure shows 4 different regions of quantum circuit volume. (1) quantum circuits in the light yellow region are considered classically simulable as they are either too narrow (i.e. at low qubit count) and can be simulated exactly, or too shallow (i.e. low quantum gate depth) and can be simulated with approximate methods such as tensor networks. (2) The region in blue starting at 50 qubits and 1 million quantum gates is the area where we expect scientific quantum advantage of relevance to NERSC based on the resource estimate data we collected. (3) The intermediate region shown in white might offer early quantum advantage but we have not identified resource estimates in this area. (4) We label quantum gate depths larger than 10^{14} in red to indicate that we expect this region to be unfeasible in the next ten years based on impractical quantum execution time estimates. The lines demarcate the highest performance quantum computers expected by the end of 2025, in five years from today, and in ten years from today.

to end users and system administrators. Quantum computer clock speed estimates presented in the literature range from sub-kHz up to GHz speeds and will determine the execution time. Based on the available data, **we construct a simple model, called the SQSP, to estimate the throughput** for nine possible quantum computer specifications on a prototypical workload of eight applications. Our model shows that the number of applications that can be run in a one year window can be as low as one and go up to tens of millions, with **most scenarios predicting a throughput of hundreds to (tens of) thousands of applications per year**. Our methodology can be helpful to **inform decision-making processes for QC system procurements**.

Chapter 1

Introduction

NERSC [177], located at Lawrence Berkeley National Laboratory, is the U.S. Department of Energy (DOE) Office of Science (SC) mission High Performance Computing (HPC) facility. An ever-growing community, currently made up of over 12,000 scientists across the DOE research complex, makes use of leading-edge compute resources deployed at NERSC. NERSC has an outstanding track record of successfully bringing advanced compute technologies to a large user community by supporting pathfinding projects such as the NERSC Science Acceleration Program (NESAP) [178], and by providing comprehensive training and support.

To remain at the leading-edge of computing, NERSC has strategically deployed new supercomputers on approximately a five year life cycle. The current *Perlmutter* system [179], deployed in 2021, is the ninth generation NERSC supercomputer and the first to combine a heterogeneous architecture offering both CPU and GPU compute nodes. Perlmutter has advanced NERSC’s ability to support numerical modeling, simulation, and data analysis. Furthermore, Perlmutter has unlocked new capabilities for Artificial Intelligence (AI) for scientific applications, a trend that NERSC expects to expand further with its tenth generation system. The NERSC-10 system, named *Doudna* [180] in honor of Dr. Jennifer Doudna, was announced in May 2025 and is expected to become available to users in the 2026 time frame. Doudna will provide at least 10× the performance of Perlmutter, and fundamentally expand the AI and workflow-based compute capabilities at NERSC. Doudna will offer real-time data processing to experimental facilities that rely on NERSC for their compute needs, seamlessly linking to DOE’s experimental and observational scientific user facilities as part of the Integrated Research Infrastructure (IRI) [113].

To ensure that NERSC continues to fulfill its central role in the DOE SC mission as we move toward the 2030s, it is essential that NERSC understands the applicability of quantum computing to its mission. Concurrently, a 2023 DOE workshop report [153] identifies critical research priorities for advancing quantum technology toward practical utility in scientific applications. Drawing from a decade of DOE investments in Quantum Computing (QC) software and hardware, the report outlines a grand challenge to “*demonstrate a rigorously quantifiable, end-to-end quantum advantage relative to state-of-the-art classical counterparts, particularly for problems with practical or scientific significance for which asymptotic exponential quantum advantages have been established,*” along with five priority research directions (respectively in software, algorithms, benchmarking, error handling, and networks) that require collaborative efforts among computer scientists, mathematicians, and physicists to achieve advances across the full technology stack. The workshop emphasized “*the need for end-to-end demonstrations of disruptive quantum advantages for scientific applications, positioning QC as a transformative technology that will complement and potentially revolutionize HPC capabilities for DOE’s scientific mission areas.*”

Building on this research-focused report, DOE proposed a Quantum Information Science (QIS) Applications Roadmap [20] in 2024 that outlines a strategic vision for QC development that positions DOE at the forefront of this transformative technology. The roadmap charts a clear progression through four distinct eras over an estimated 20 year window: from current Noisy Intermediate Scale Quantum (NISQ) devices and small demonstrations of Quantum Error Correction (QEC) (0-5 years), to small quantum computers with QEC (5-10 years), large quantum computers with QEC (10-20

years), and finally very large Fault-Tolerant Quantum Computing (FTQC) (20+ years). Key technology milestones include demonstrating 1,000 physical qubits with error rates ten times below code threshold, scaling to 10,000 physical qubits while maintaining low error rates, and developing quantum interconnects for modular architectures. The roadmap highlights promising application areas including chemistry and materials science, where quantum simulations could eliminate uncontrolled approximations that afflict classical methods and enable revolutionary scientific discoveries. The report says that within the next five years, calculations on specific molecules or materials that are on the boundary of classical tractability might cross that threshold, potentially augmented by HPC resources. DOE’s national laboratories and user facilities are uniquely positioned to lead this effort through interdisciplinary collaborative research, software stack development, and the establishment of widely accessible quantum user facilities that will accelerate the development and deployment of quantum computers for these breakthrough scientific applications.

Consequently, NERSC sees QC as a crucial and rapidly growing pillar of its ten year strategic plan for 2024-2034 [176]. These reports, roadmaps and strategies are motivated by the observation that quantum computers offer a fundamentally different paradigm of computing. By leveraging the principles of quantum mechanics, they have the potential to solve certain problems asymptotically faster than any computer based on the principles of classical physics is expected to. An important caveat is that proving that a quantum algorithm is faster than any classical algorithm is an extremely challenging task. Instead, researchers typically compare quantum algorithms against the best classical methods *currently known* – arguably, the strongest evidence of a exponential speedup using this approach exists for Shor’s prime factoring algorithm [223, 225]. One significant drawback of this approach is that if a superior classical algorithm is discovered later, claimed quantum advantages may disappear. Key areas (besides Shor’s algorithm) for which exponential asymptotic quantum speedups are expected to be exponential include simulation of quantum system dynamics [6, 62, 250, 257], computing static properties of quantum systems [124], and solving linear systems of equations [104]. Other areas, such as search and optimization could see low-degree polynomial speedups on quantum computers in certain regimes [100]. We refer to [68, 118] for an overview of known quantum algorithms and applications.

In the context of these trends, this report examines how QC is expected to impact NERSC’s mission and workload over the next ten years. We synthesize findings from academic literature with the forward-looking projections of QC companies. We observe that the projections and roadmaps that have been proposed by industry follow an accelerated trajectory compared to DOE’s roadmap [20] and predict very large-scale systems to become available within ten years. We remark that our endeavor necessitates various assumptions and approximations, which, compounded by the inherent uncertainty of technological forecasting, leads to a significant degree of uncertainty in our estimates. As this field is characterized by rapid evolution, we will continue to monitor its development, refine our analysis and adjust our strategy accordingly. Therefore, we intend this to be a living document, and plan to release updated versions when appropriate.

The remainder of this report is structured as follows. We begin by reviewing the workload that NERSC users are running on the current Perlmutter system in Chapter 2 and identify the three main scientific domains of interest for which quantum computers will deliver speedups. Chapter 3 provides a primer on key QC concepts that are mentioned throughout the report. Chapter 4 collects and reviews the resources required to solve more than 100 target problems across the scientific domains of interest to NERSC that were identified in Chapter 2. We base our analysis on estimates available in the scientific literature that go beyond the asymptotic speedups and account for the end-to-end resource requirements. We distill these results into a single figure of merit which captures the space-time cost of a quantum application. Next, in Chapter 5, we analyze various quantum technology roadmaps that different vendors have put forward for the next decade. We uniformize the different vendor milestones with respect to the previously introduced figure of merit. This allows us to bound and compare the capabilities of forthcoming quantum computers. We also compare these bounds to the resources required to solve applications in the various scientific domains and observe a significant overlap emerging in the next five to ten years. We comment on the execution time as a key constraint for large-scale quantum applications in Chapter 6 and present a simple model, called the Sustained

Quantum System Performance (SQSP) metric, that allows us to compare system-level performance and throughput for a heterogeneous workload. Chapter 7 concludes the report by summarizing the main findings. A List of Acronyms and Glossary of Terms is included after the conclusion.

Chapter 2

NERSC Workload

The workload at NERSC is inherently varied and complex with a large community of users coming from six different program offices in the U.S. Department of Energy (DOE) Office of Science (SC), including Advanced Scientific Computing Research (ASCR), Basic Energy Sciences (BES), Biological and Environmental Research (BER), Fusion Energy Science (FES), High Energy Physics (HEP), and Nuclear Physics (NP).

NERSC users develop and run a breadth of scientific computing codes at a large range of scales and GPU-readiness levels. This means that NERSC compute resources must be versatile as well as highly performant. The Perlmutter system provided a CPU as well as a GPU partition in an era when multiple relevant scientific codes were still in the process of being ported to GPUs. At the same time, NERSC provided extensive support to enable many more codes to benefit from GPU acceleration.

Figure 2.1 summarizes the resource allocation for various workloads on NERSC systems. Thanks to NERSC’s ongoing efforts, for example through the aforementioned NERSC Science Acceleration Program (NESAP) program, most of Perlmutter compute performance is stemming from the GPU partition.




Figure 2.1: Distribution of the compute cycles used at NERSC across different categories of algorithms and application areas.

Density Functional Theory (DFT) stands as a cornerstone of computational **materials science** and **quantum chemistry** and takes more than 30% of the yearly cycles at NERSC. It is a computational method to approximate the electronic structure of complex atomic and molecular systems, and achieves this at reasonable computational cost [128, 191]. Its strengths lie in its efficiency, allowing for the study of large systems, and its ability to provide valuable insights into ground-state properties. However, the accuracy of DFT is often limited by the choice of exchange-correlation functional, which can fail to capture essential electron correlation effects required for describing van der Waals interactions [127], electronic band gaps [169], and charge-transfer excitations [78]. These limitations

can be partially addressed by using more sophisticated exchange-correlation functionals [193] which incorporate exact exchange or perturbative correlation to improve accuracy, though at a significantly higher computational cost. However, even with these enhancements, DFT remains insufficient for solving the most challenging problems in quantum chemistry, such as strongly correlated systems or multi-reference electronic states, where more advanced wavefunction-based methods are typically required [64]. Quantum chemistry codes that go beyond DFT make up another 2.6% of the cycles. This is where quantum computers hold significant promise. By leveraging quantum mechanics, these machines could potentially tackle the exponentially complex many-body problem inherent in electronic structure calculations, offering the prospect of more accurate exchange-correlation functionals and enabling the simulation of strongly correlated materials at the precision of full configuration interaction (CI), thus overcoming key weaknesses of traditional DFT while scaling exponentially better than classical approaches.

Nuclear and high energy physics make up over 18% of the workload. Most of their computational time is dedicated to Lattice Quantum Chromodynamics (LQCD) runs. This is a computational method for studying the strong force and its associated particles, hadrons, from first principles. Its strength lies in its ability to non-perturbatively calculate properties of hadrons, such as their masses and decay constants, by discretizing spacetime onto a lattice. However, LQCD is both computationally demanding and constrained to imaginary time, requiring vast resources to simulate realistic quark masses and large volumes, thereby limiting the precision and scope of calculations. The demands are particularly severe for computing complex observables like scattering amplitudes and real-time dynamics, since the computation is performed on Euclidean space, and thus cannot access these quantities directly. Additionally, the sign problem poses a significant barrier to the applicability of LQCD to dense systems at finite chemical potential [98, 174]. Quantum computers promise to revolutionize this field by efficiently tackling these computational challenges. Their inherent ability to simulate quantum systems could overcome the sign problem and enable simulations at physical quark masses and large volumes, ultimately leading to more accurate and comprehensive understanding of strongly-interacting matter.

In addition, there are smaller pieces of the workload, like **N-body molecular dynamics and fusion using Particle-In-Cell (PIC) methods**, that also stand to benefit from quantum computers. As such, more than half of the computational resources on Perlmutter are dedicated to solving quantum many-body problems, highlighting the enormous potential impact and benefit quantum computers can have on NERSC’s mission. Collectively, we refer to all DOE SC mission applications that stand to benefit from large-scale Quantum Computing (QC) as the *Quantum Relevant Workload (QRW)*.

We summarize this section with the following finding:

Finding 1. At least 50% of NERSC compute resources are spent on solving quantum mechanical problems relevant to *materials science*, *quantum chemistry*, and *nuclear and high energy physics*, primarily using DFT and LQCD codes. Quantum computers naturally have the potential to accelerate scientific discovery in these areas as they do not have to rely on the same approximations as classical algorithms to solve computational problems. We refer to these applications as the QRW.

Chapter 3

A Primer on Quantum Computing

Quantum Computing (QC), like many scientific fields, uses its own jargon. While this enables efficient exchange of ideas among experts, it raises the bar for non-experts to understand those same ideas. In an effort to make this report broadly accessible, we provide a primer on the key QC concepts we refer to in the remainder of this text. Our goal is to offer a broad, high-level perspective, not a precise or comprehensive account. For detailed overviews, we refer the interested reader to [182, 248]. Readers already familiar with QC may consider skipping this chapter and move directly on to Chapter 4.

3.1 On qubits, gates, and circuits

QC represents a profound shift in how information can be processed, harnessing the principles of quantum mechanics to tackle problems beyond the reach of classical machines. At its foundation is the qubit, the quantum analogue of the classical bit. While a classical bit exists strictly as 0 or 1, a qubit can occupy a superposition of both states, $\alpha|0\rangle + \beta|1\rangle$. When qubits become entangled, the state of one is intrinsically correlated with the state of another, no matter the distance between them—a property that underpins much of the power of QC.

Computation in a quantum system is achieved through quantum gates, which are reversible operations that transform qubit states. These gates are assembled into quantum circuits, sequences of gate operations that together implement a computation. The performance and complexity of such circuits are often described in terms of quantum circuit volume, a metric that accounts for both the number of qubits n_Q and the number of gates n_G of the computation. A common characteristic of all quantum circuits is that at the end of the quantum computation (i.e. after running a quantum circuit), some or all of the qubits are measured and classical information is extracted in the form of a bitstring that corresponds to the states in which the qubits were measured. This procedure is often referred to as taking a shot of a quantum circuit. Figure 3.1 summarizes the notions of (qu)bits, quantum gates and quantum circuits (incl. quantum measurements) using a common circuit diagram notation [182].

In practical hardware, qubits are physical devices prone to errors. Current quantum computers have up to a few hundred physical qubits and offer imperfect control over those qubits. Consequently, various sources of noise, including due to State Preparation and Measurement (SPAM), imperfect gate operations and decoherence, limit the performance of quantum computers. This class of quantum systems has previously been dubbed Noisy Intermediate Scale Quantum (NISQ) devices [199]. When using a NISQ device, one typically maps the qubits and quantum gates directly onto error-prone physical qubits. In the NISQ paradigm, gates that entangle two qubits, such as the CNOT gate (see Fig. 3.1), typically incur the largest error and consequently are a dominant factor determining quantum circuit performance. The capabilities of NISQ hardware and applications can be extended by leveraging Quantum Error Mitigation (QEM) techniques [47] (more details in Section 3.2) which typically requires taking additional samples on the quantum computer (which is also known as the sampling overhead for QEM) to improve the quality of the result. However, there are fundamental limits on the level of improvement that is achievable and the sampling overhead typically scales exponentially [234].

For large-scale, reliable computation, sufficiently good physical qubits can be grouped into logical




Figure 3.1: (a) A bit can be either 0 or 1; while (b) a qubit can be in a linear combination (or superposition) of $|0\rangle$ and $|1\rangle$ and its state space can be visualized by a Bloch sphere; (c) quantum gates modify the state of one or more qubits, the top panel shows a single-qubit Hadamard gate applied to a qubit in the $|0\rangle$ state and transforms it to $|+\rangle$, the bottom panel shows a two-qubit Controlled-NOT (CNOT) gate that leaves the control qubit unchanged and applies a NOT to the target qubit conditioned on the state of the control; and finally (d) quantum circuits are combinations of quantum gates and measurement operations, the figure shows a two-qubit circuit that consists of a Hadamard gate and CNOT gate and prepares a state known as a Bell state.

qubits – error-protected units formed by encoding information across many physical qubits. This process is central to Fault-Tolerant Quantum Computing (FTQC) in which computations can proceed reliably despite the presence of noise due to an approach known as Quantum Error Correction (QEC) [143, 240]. Fault tolerance can be implemented more easily for some gates than for others. Clifford gates –including the Hadamard, Phase, and aforementioned CNOT gates– form a set that is relatively straightforward to realize. Non-Clifford gates, such as the T gate or the three-qubit Toffoli gate (a quantum analog of the classical AND gate), extend the Clifford set to achieve universal quantum computation but are significantly more resource-intensive in an FTQC setting. Implementing non-Clifford gates typically relies on the preparation of high-fidelity *resource states*¹ also known as magic states. Examples include $|T\rangle$ - and $|CCZ\rangle$ -states which respectively can be used to implement T and Toffoli gates in a fault-tolerant manner using a Clifford circuit [263]. High-fidelity magic states are costly to prepare, requiring a procedure known as magic state distillation and generated by magic state factories. Recent progress has significantly reduced the cost of producing magic states over a substantial range of errors (see: magic state cultivation [89]), nonetheless it remains a major cost for fault-tolerant quantum computation.

Analyzing (and optimizing) the end-to-end cost of running a computation on a FTQC system forms the topic of the field known as quantum resource estimation. In light of our previous discussion, the number of non-Clifford gates (e.g. T , Toffoli or continuous-angle rotation gates) is typically the most important metric considered in quantum resource estimation. In Chapter 4, we collect and present numerous resource estimates for quantum computations within the physical sciences.

3.2 Handling errors: quantum error mitigation and correction

QEM and QEC share the goal of increasing the robustness of a quantum computer to errors but go about reaching this goal in different ways. QEM accepts the inevitability of noise and errors, and compensates for these errors by applying a probabilistic correction to the *distribution* of measurements. This comes at a cost of increasing the sampling overhead (i.e. multiplying the number of shots needed to reach a prescribed error target). QEC, on the other hand, protects the quantum information against the introduction of noise by detecting and correcting for errors within *individual* shots. This comes at a cost of increasing the encoding overhead (i.e. multiplying the number of physical qubits used to redundantly store the state of a single logical qubit). The distinction between QEM and QEC is depicted graphically in Fig. 3.2

For the remainder of this section, we will introduce the principles guiding QEM and QEC protocols by focusing respectively on Probabilistic Error Cancellation (PEC) and the surface code as prototypical

¹These are called resource states as they are consumed during the computation.




Figure 3.2: (a) Quantum Error Mitigation aims to improve the result by repeating and combining multiple noisy measurement results into an improved measurement result with lower error. (b) Quantum Error Correction encodes quantum information over multiple physical qubits. When an error event occurs, as indicated by the data qubit in red, it is detected through measurement using (blue) ancilla qubits (red plaquettes) and corrective action can be taken. The entire diagram represents a single logical qubit encoded in a $d = 5$ surface code.

examples of each. For our discussion, we assume that (1) quantum gate errors are the dominant sources of errors in the system, thus ignoring SPAM errors, and, (2) errors occur independently and as discrete events. If we consider ϵ as the probability an error occurs each time a quantum gate is applied, then the overall fault rate λ of a quantum circuit becomes,

$$\lambda = n_G \epsilon, \quad (3.1)$$

and therefore the quantum circuit fault rate becomes of order unity when $n_G \approx \epsilon^{-1}$ [47].

QEM is most often considered in the context of producing accurate expectation values $\langle O \rangle$ of an observable O on a NISQ device. PEC [81, 236] is a widely adopted QEM method and here we consider it as the prototypical QEM procedure. PEC can provide bias-free estimates of expectation values by (1) characterizing a noise channel, Λ_i , associated with a quantum gate operation, (2) expanding it in an over-complete basis of (noisy) gates that can be executed on the quantum system, and (3) quasi-probabilistically implementing the inverse channel Λ_i^{-1} before each quantum gate operation to cancel out the noise. The trade off in reducing or eliminating the bias in the estimator for the observable, $\hat{O}_{\text{PEC}}$, is that the variance of the estimator is increased compared to the noisy estimator $\hat{O}_{\text{noise}}$ obtained from computing the sample mean of the noisy quantum circuit runs. The ratio is called the sampling overhead and it has been shown that for PEC it scales approximately as [31, 47]

$$C_{\text{PEC}} \approx e^{4\lambda} = e^{4n_{G,\text{PEC}}\epsilon}, \quad (3.2)$$

from which it follows that the error of the PEC estimator $\hat{O}_{\text{PEC}}$ compared to the exact observable $\langle O \rangle$ can be made arbitrarily small at a sampling overhead that scales exponentially in both the number of quantum gate operations in the quantum circuit and the error rate per quantum gate. Rewriting Equation (3.2) as

$$n_{G,\text{PEC}} \approx \frac{\ln C_{\text{PEC}}}{4\epsilon} = \frac{\ln C_{\text{PEC}}}{4} n_{G,\text{noise}}, \quad (3.3)$$

we see that the number of gates that can be run using PEC remains inversely proportional to the error rate but PEC can improve the pre-factor, $\ln C_{\text{PEC}}/4$, at an exponential cost in the sampling overhead.

QEC is the hallmark of FTQC. The prevailing QEC scheme combines multiple physical qubits to redundantly encode information in a single logical qubit. Within a logical qubit, some of the physical

qubits are designated as ‘data’ qubits to be used to reliably store quantum information. Others are designated as ‘ancilla’ qubits to be used for detecting error events that would distort the quantum information. This is loosely analogous to the use of parity bits for error correcting codes (ECC) in classical memory systems. Measurements of the ancillary qubits can be used to detect errors without collapsing the logical wave function.

Robust and efficient encoding of information in logical qubits is essential for QEC. A $[[n, k, d]]$ -code encodes k logical qubits in the state of n physical qubits in such a manner that any error event that maps a valid encoded state to another valid encoded state must act on at least d physical qubits. The value d is known as the code distance and quantifies the robustness of the code against errors; it is analogous to the number of bit errors that can be detected by a Hamming codes in classical ECC. Typically, reaching larger code distances, and thus an increased resilience against errors, requires using more physical qubits per logical qubit.

QEC cannot be effective unless the physical error rate is below the code threshold. For physical error rates below the code threshold, logical errors can be (exponentially) suppressed by increasing the code distance. For physical error rates above code threshold, increasing the code distance will instead worsen the logical errors due to the increased number of error-prone operations involved.

FTQC using a QEC code typically progresses in *code cycles* which involve measuring a subset of the physical qubits that make up the logical qubit(s) in order to extract the error syndrome which can then be used to correct any logical errors and implement quantum gate operations on the logical qubits.

The surface code [39, 72, 125] is a prototypical topological error correcting code that encodes a logical qubit in a planar, square lattice of physical qubits that only require nearest-neighbor connectivity. For a single logical qubit with code distance d , the encoding requires $2d^2 - 1$ physical qubits per logical qubit, consisting of d^2 data qubits and $d^2 - 1$ ancilla qubits for error syndrome extraction via parity check measurements. In summary, the surface code is a $[[n = 2d^2 - 1, k = 1, d]]$ -code. A $d = 5$ surface code logical qubit is sketched in Fig. 3.2(b). A commonly used rule-of-thumb relating the logical error rate ϵ_L to the physical device error rate ϵ_P for a code distance d surface code is [85]

$$\epsilon_L = 0.1(100\epsilon_P)^{(d+1)/2}, \quad (3.4)$$

assuming the physical error rate is below the code threshold. Small examples of the surface code have already been demonstrated experimentally by several groups [3, 132, 161, 235, 261]. The surface code offers many advantages: it is well understood, allows for a planar topology, has a high code threshold, and has been used to propose universal quantum computer architectures [145]. However, it also is notoriously resource-intensive since the number of qubits scales quadratically as a function of code distance. Naturally, more efficient encodings have been proposed in the literature. Prominent examples include Quantum Low Density Parity Check (qLDPC) codes [42, 97, 108, 109, 138, 139, 189, 190, 241, 242]; for example, qLDPC codes based on expander graphs can in principle achieve scaling $[[n, k = \Theta(n), d = \Theta(n)]]$, which is quadratically more efficient than the surface code.


QEC can in principle be combined with QEM approaches. Indeed, it has recently been shown that PEC can be beneficial in combination with QEC [233], and has the potential to allow for a reduction of the code distance by five while maintaining similar error rates at the cost of a sample overhead factor $C \leq 100$ and more intensive post-processing.

3.3 Overview of quantum algorithms

On top of these hardware and architectural foundations sit the algorithms that define QC’s potential [118]. One of the most important classes of quantum algorithmic kernels for the physical sciences is Hamiltonian simulation [83, 151], in which a quantum computer models the evolution of a quantum system for which the total energy is described by a Hamiltonian operator. This capability is invaluable for chemistry, materials science, and fundamental physics, where classical simulation quickly becomes intractable. Early approaches use product formulas also known as Trotter-Suzuki decomposition [232], which break the time evolution into a sequence of simpler steps to approximate the real-time evolution as illustrated in Fig. 3.3(a). More recent methods, such as qubitization [155], encode the Hamiltonian

operator into a structured form, and achieve optimal computational complexity: (1) the number of queries to the Hamiltonian data (oracle) scales linearly in the simulation time and as $\log(1/\epsilon)$ in the spectral error; (2) the number of qubits scales as $\log(N) + m$ with N the dimension of the Hamiltonian and m some comparatively small number of extra ancilla qubits. This is an exponential reduction in resources compared to a brute-force classical implementation. Quantum Signal Processing (QSP) [154] builds on this by enabling polynomial transformations of quantum operators through carefully orchestrated rotations, and the Quantum Singular Value Transformation (QSVT) [91] generalizes the technique to manipulate arbitrary operators directly—unlocking efficient algorithms for solving linear systems, performing principal component analysis, and other tasks relevant to scientific computing. Collectively, these constitute a unified framework for quantum algorithm development and discovery that is applicable to many known quantum algorithms exhibiting computational speedup [162]. Additionally, Quantum Phase Estimation (QPE) [124] and its modern variants [74, 76, 77, 107, 144, 181, 246, 260], shown in Fig. 3.3(b), use Hamiltonian simulation to compute ground state energies of quantum systems, a key primitive underpinning much of the expected QC speedups in quantum chemistry and condensed matter physics (see Chapter 4).

(a) **Hamiltonian Simulation via Trotterization**



(b) **QPE via Qubitization**




Figure 3.3: (a) When implementing Hamiltonian simulation via Trotterization, the time evolution operator e^{-iHt} (orange) is approximated as a product of simpler time evolutions resulting in structured circuits with single-qubit gates (yellow) and two-qubit gates (blue) (b) Implementing Quantum Phase Estimation to compute the ground state of a Hamiltonian H using qubitization begins with a state preparation phase χ_m , followed by a sequence of (controlled) *walk operators* $\mathcal{W}$, which encode the spectrum of H via qubitization, and reflection oracles $\mathcal{R}$, and ends with an inverse Quantum Fourier Transform. Figure adapted from [18].

While these advanced algorithms promise significant long-term benefits, today’s quantum hardware is still in the noisy intermediate-scale quantum (NISQ) era, where qubit counts are limited and error rates are high. To make progress in the near term, researchers have developed variational algorithms, which blend quantum and classical computation. In these hybrid approaches, a quantum processor prepares a parameterized quantum state, and a classical optimizer iteratively updates the parameters to minimize a cost function. Examples include the Variational Quantum Eigensolver (VQE) [195], used to estimate molecular ground-state energies, and the Quantum Approximate Optimization Algorithm (QAOA), designed for combinatorial optimization. These methods are more tolerant of noise and shorter circuit depths, making them practical stepping stones toward the large-scale, fault-tolerant algorithms that will eventually unlock QC’s full potential.

Chapter 4

Resource Estimates for the Quantum Relevant Workload

In this section, we collect and compare resource estimates from the literature for more than 140 target applications across the three scientific domains identified in the previous section that make up the Quantum Relevant Workload (QRW). We discuss condensed matter physics and materials science in Section 4.1, quantum chemistry in Section 4.2, and nuclear and high energy physics in Section 4.3. Section 4.5 provides an overview of the three domains and Section 4.4 discusses the state of affairs for certain other quantum algorithms and application areas of interest to the U.S. Department of Energy (DOE) Office of Science (SC) mission, including linear algebra, optimization, simulating differential equations, and quantum machine learning. For these other areas, we do not include resource estimates.

For each resource estimate, we report a commonly used space-time figure of merit [200] for the quantum circuit consisting of the number of qubits (n_Q), i.e., the space complexity, and the number of quantum gates (n_G), i.e., the time complexity. We name this quantity the $\mathcal{P}$ -vector (for performance) and define it as follows,

$$\mathcal{P} = (n_Q, n_G). \quad (4.1)$$

We will often visualize the $\mathcal{P}$ -vector as an xy-coordinate on a 2D plot. The $\mathcal{P}$ -vector corresponds to an estimate of the volume of the quantum circuit required to solve the application. Multiple factors can influence the $\mathcal{P}$ -vector required to solve an application. These factors include, but are not limited to:

1. *Space versus time trade-offs.* Most quantum algorithms and applications allow a trade-off between a reduction in the number of qubits for an increase in the number of quantum gates (or vice versa). A well-known example is Hamiltonian simulation, a key quantum-algorithmic kernel used in the majority of applications we discuss. Implementing a Hamiltonian simulation using a Trotter decomposition typically requires significantly fewer qubits and significantly more quantum gates compared to a method that combines qubitization with Quantum Signal Processing (QSP) at a similar error level.
2. *Dominant quantum gate complexity.* Most resource estimates only report the leading-order of gates that require the most resources to execute, such as T gates, Toffoli gates, or continuous-angle rotation gates, and thus are the dominant factor determining the execution time. This assumes a cost model that is typical for the (early) Fault-Tolerant Quantum Computing (FTQC) rather than Noisy Intermediate Scale Quantum (NISQ) era. However, the type of quantum gate that is most expensive might change in the future and depend on the system architecture. For example, recent work on magic state cultivation [89] has significantly reduced the complexity to implement a T gate [146]; over a certain parameter range, the cost of T gates is now much closer to that of a Clifford gate like a CNOT. More breakthroughs in this area can further alter this picture and in practice the dominant quantum gate cost might differ based on the quantum computer hardware architecture. For our analysis, we express all resource estimates in terms of T gates, which are the dominant resource constraint for many algorithms. In some cases,

previously published resource estimates have been reported in terms of different gate sets. In those instances, we convert the estimates into an equivalent T gate cost using the following conversion rates:

- 1 Toffoli gate = 4 T gates [117],
 - 1 rotation gate = 100 T gates [220].
3. *Topology or connectivity.* The two- and multi-qubit gates present in a quantum circuit determine the connectivity and topology requirements on the quantum computer architecture to run the quantum circuit without needing additional transpilation and the introduction of SWAP operations. The worst-case transpilation overhead on the number of gates to map a quantum circuit requiring all-to-all connectivity to 2D nearest-neighbor connectivity scales as $\mathcal{O}(\sqrt{n_Q})^1$. For the remainder of this report, we do not consider the effect of topology on the $\mathcal{P}$ -vectors reported in the literature.
 4. *Circuit versus active volume.* The quantum circuit volume (and hence the $\mathcal{P}$ -vector) might differ based on the quantum computer architecture. For example, *active volume* architectures have been proposed for quantum computers based on photonics [147]. In this setting, qubits that are idle at certain points in the quantum circuit do not count towards the overall quantum circuit volume. It has been estimated that an active volume architecture can reduce the number of gates to run Shor’s prime factoring algorithm by an order of magnitude [147] and it can also reduce the resources required for quantum chemistry applications [46]. Other such architectural designs that improve the $\mathcal{P}$ -vector might be developed in the future.
 5. *Critical depth and gate parallelization.* It is (often implicitly) assumed throughout this report that the number of gates reported represents the critical depth of the quantum circuit. Due to the anticipated serial or small-batch generation of the resource states necessary for implementing non-Clifford gates (i.e. T gates), early scalable FTQC are not expected to support significant parallelization for non-Clifford gates. Specifically, in our discussion on execution time estimates, we do not consider the possible impact of gate parallelization.

We conclude, based on the considerations discussed above, that the $\mathcal{P}$ -vector estimates should be considered as an upper bound. Furthermore, $\mathcal{P}$ -vectors are also expected to reduce further by future algorithmic advances and improved implementations. We will motivate this expectation by presenting a case study from quantum chemistry in Section 4.2.

4.1 Condensed matter physics and materials science

Condensed matter physics and materials science pose some of the most promising problems for quantum simulation. Strongly correlated electrons, emergent phases of matter, and exotic excitations underlie phenomena such as high-temperature superconductivity, magnetism, and topological order, areas where classical computational methods often reach their limits due to exponential scaling [10, 88, 129]. Accurate quantum simulations of these systems would not only provide insight into fundamental physics [15, 249], but also accelerate the discovery of practical materials for energy storage [13, 208], catalysis [52, 94], and quantum technologies [27, 164]. As a result, condensed matter and materials science are widely viewed as domains where quantum computers could deliver some of the earliest impacts [12, 35].

We collected resource estimates for a total of 80 applications in condensed matter physics and materials science that span a spectrum of model complexity and are considered out of reach for classical solvers. These include simpler spin models, such as the Kitaev and Heisenberg models that can help scientists study new and interesting physical phenomena and prepare exotic phases of matter (e.g. spin glasses and quantum chaos) [17, 24, 61, 84, 255]; more complicated lattice models such as the Fermi-Hubbard (FH) model [18, 29, 48, 84, 255] that can serve as a proxy to study superconducting materials;

¹Mapping to a linear topology would scale as $\mathcal{O}(n_Q)$ but almost all fault-tolerant architectures require at least a 2D square grid connectivity.

and realistic models of materials that take into account the relevant electronic structure [126]. As an example of the latter category, Refs. [71, 222] propose Quantum Phase Estimation (QPE)+qubitization using first-quantized plane-wave basis to simulate materials for lithium-ion batteries such as $\text{Li}_2\text{FeSiO}_4$. Subsequent work [33] shows that first quantization outperforms second quantization for materials simulations of a Lithium-Nickel-Oxide (LNO) material. The $\mathcal{P}$ -vectors of the resource estimates we collected are shown in Figure 4.1. The estimates at lower quantum gate count correspond to the simulations of spin and FH models, while the estimates for materials simulations in first quantization show comparatively larger $\mathcal{P}$ -vectors.




Figure 4.1: Overview of various resource estimates for computational problems related to materials science in terms of number of logical qubits (n_Q) and number of gates (n_G). Different marker types indicate the different quantum algorithms being used.

We observe that resource estimates start at about 10^6 gates and 10^2 qubits and go up from there in both quantum gate depth and qubit count. We do expect condensed matter physics and materials science to be among the first scientific fields to benefit from quantum computers since many of these problems are formulated in terms of spins, which map naturally to qubits, or can be encoded in qubits with minimal resource overhead. This is the primary reason we consider them to be the earliest candidates for scientific quantum advantage among the DOE QRW.

4.2 Quantum chemistry

The field of quantum chemistry addresses a number of important scientific and technological challenges, from designing efficient catalysts for sustainable fuels and industrial processes to developing advanced materials for batteries, photovoltaics, and quantum information technologies [27, 52, 208]. Quantum chemistry stands out as a leading target for Quantum Computing (QC), as QC has the potential to revolutionize our ability to simulate molecular behavior. Within quantum chemistry, a key distinction often lies between static and dynamic problems. Static problems focus on characterizing a molecule in equilibrium, providing a snapshot of its properties, for example determining the ground-state energy, electronic structure, and other time-independent properties. In contrast, dynamic problems explore how molecules evolve over time, such as chemical reactions, photoexcitation, charge and energy transfer, and vibrational dynamics. QC holds the potential to provide accurate solutions to both static and dynamic quantum chemistry problems, however our survey of the literature revealed that static chemistry problems, and in particular Ground State Energy Estimation (GSEE), have been studied in much greater detail than dynamic ones.

The field of quantum resource estimation for GSEE problems in quantum chemistry has matured over the past five years compared to other application areas. This is reflected by the observation that a few standardized benchmark problems have been established in the literature. This includes

the FeMoco complex [209] and the Cytochrome P450 system [93]. These target problems have been re-analyzed as new algorithms for GSEE have been proposed and as such form an excellent yardstick to track progress in quantum algorithms for quantum chemistry. The reduction in resources required for GSEE of the FeMoco complex is visualized in Figure 4.2 for both the original, smaller (54e, 54o) active space system, using the def2-TZVP basis [209], and the larger (113e, 76o) active space system, using the TZP-DKH basis for Fe, S, and Mo and the def2-SVP basis for the other atoms [142]. The dimension of the FCI space for the larger system is on the order of 10^{35} for the spin $S = 3/2$ ground state. The quantum resources required for running QPE on the smaller system, shown in Figure 4.2a, were originally estimated using a Trotterization method for the time evolution [209]. These estimates correspond to the data points labeled ‘2017’ and require fewer qubits than more recent estimates at the cost of an order of magnitude more gates. Recent approaches all use a qubitization of the electronic structure Hamiltonian, and distinguish themselves primarily by how they factorize and compress the two-electron integral tensor. The main approaches studied include the single factorization method [34], the double factorization method [43], and Tensor Hypercontraction (THC) [46, 137, 212]. In addition to these factorization strategies, spectrum amplification techniques have recently been developed to further reduce resource costs, most notably the sum-of-squares spectral amplification (SOSSA) framework [123, 156], which builds on spectrum amplification ideas to accelerate simulations for Hamiltonians with efficiently computable sum-of-squares representations. We observe from Figure 4.2 that these algorithmic improvements have significantly reduced the computational resources required to solve the GSEE problem leading to roughly a $1000\times$ reduction in the number of gates and a $5\times$ reduction in qubit count. Similar observations hold for other standardized benchmark problems such as the Cytochrome P450 system.




Figure 4.2: Overview of resource estimates for ground state energy estimation of two versions of the FeMoco system: (*left*) a (54e, 54o) active space [209] model, and (*right*) a (113e, 76o) active space [142] model.

While GSEE has received the bulk of attention, there is growing interest in developing quantum algorithms for dynamics problems, which aim to capture the time-dependent behavior of molecular systems. Classical approaches to these problems, such as Born-Oppenheimer molecular dynamics, Ehrenfest dynamics, and ab initio molecular dynamics based on time-dependent Density Functional Theory (DFT), are computationally intensive and often struggle with nonadiabatic effects or strongly correlated electrons [163, 244]. Quantum computers offer the potential to simulate such processes natively by evolving quantum states under a time-dependent or time-independent molecular Hamiltonian and by directly computing forces, energy gradients, and other observables needed for nuclear propagation in molecular dynamics simulations. For example, algorithms have been developed to compute molecular energy derivatives, including forces, more efficiently on quantum devices, achieving lower circuit repetition costs and even Heisenberg-limited scaling in the fault-tolerant limit [184]. Variational quantum algorithms have also been suggested for calculating forces within ab initio molecular dynamics frameworks [228]. Algorithmic strategies for real-time evolution include Trotterized and qubitized time evolution [155], variational quantum dynamics using a time-dependent Variational Quantum Eigensolver (VQE) [256], and approaches based on linear combination of unitaries [60] and QSP [154].

Ref. [186] reviews quantum algorithms for simulating molecular quantum dynamics, highlighting both representations using first quantization and second quantization, and comparing variational and Trotter-based real-time evolution methods as promising routes for capturing electron-nuclear dynamics. A number of papers [56, 213] have estimated the resource requirements for quantum chemistry dynamics problems and found a need for significantly higher gate counts than GSEE. While these studies represent important first steps in assessing the feasibility of quantum chemistry dynamics on quantum hardware, the level of effort devoted to reducing their resource requirements remains far behind that for ground-state estimation, where years of algorithmic refinements have dramatically lowered cost estimates. Consequently, we do not plot resource estimates for dynamics problems. Bridging this gap remains both a challenge and an opportunity for advancing quantum simulations of chemically relevant processes.

In addition to time evolution, significant progress has been made in developing quantum algorithms for excited state calculations, which, while not dynamical themselves, provide necessary inputs for simulating chemical dynamics. Access to accurate excitation energies and electronic couplings enables the construction of potential energy surfaces, which are needed for modeling processes like photoisomerization and energy transfer. Techniques such as state-averaged VQE [253], equation-of-motion approaches [187], and subspace-search methods [175, 221] have been proposed to efficiently target excited states on quantum hardware. These developments will be useful for quantum simulations of nonadiabatic molecular dynamics, particularly in regions near conical intersections where potential energy surfaces of different electronic states intersect and traditional approximations fail [75].

Bosonic QC platforms, particularly those using continuous-variable (CV) encodings, have also emerged as promising tools for simulating chemical dynamics. These platforms naturally represent vibrational degrees of freedom using bosonic modes, enabling compact and efficient simulations of molecular vibrations and vibronic transitions [79]. Recent theoretical work has proposed using CV-based quantum processors for real-time wavepacket evolution and vibronic dynamics in anharmonic molecular potentials [160], and hybrid bosonic-qubit [136, 149] schemes are being explored to model coupled electron-vibration dynamics more efficiently. Despite these recent theoretical developments, detailed quantitative resource estimates, such as required number of bosonic modes, effective squeezing levels, gate counts, or coherence time budgets, are not yet available in the literature. Similarly, hardware roadmaps from CV-focused vendors, such as Xanadu (photonic architectures) and Alice&Bob (cat qubits), describe platforms capable of simulating bosonic systems but as of yet provide no application specific resource benchmarks.

We expect resource reductions and algorithmic speedups to become pervasive across all domains as quantum algorithms continue to develop and mature with future scientific advances. Indeed, this is a larger trend within algorithmic development which can be seen in the field of classical algorithms over the past decades [69, 87, 90, 159, 243, 245], and has historically been a major factor in guiding investments at DOE SC [122].

Figure 4.3 displays all 33 target problems we identified in the literature. These include ground state preparation and estimation problems, such as FeMoco and Cytochrome P450, as well as computational catalysis for CO₂ fixation [43] and other strongly-correlated problems [36, 80, 141]. The figure shows that solving scientifically relevant, hard quantum chemistry problems will require the capability to run at least 10^8 gates on about 10^3 qubits, with estimates going up to 10^{12} gates and beyond, depending on the specific application and algorithm under consideration.

We summarize our findings for this section as follows:

Finding 2. Resource estimation for computing ground state energies for quantum chemistry systems is among the most mature disciplines in quantum algorithm development. The workhorse algorithm has become QPE+qubitization. Significant progress over the last five years in terms of improved factorizations and compression of the two-electron integrals has reduced the resources required to solve key benchmark problems by orders of magnitude. We expect that future improvements to quantum algorithms will continue to reduce the resources required for applications from all domains in the DOE QRW.




Figure 4.3: Overview of various resource estimates for computational problems related to quantum chemistry in terms of number of logical qubits (n_Q) and number of gates (n_G). QPE+Trotter based applications are marked by circles, and QPE+qubitization by squares.

4.3 Nuclear and high energy physics

QC holds significant potential for impact in nuclear and high energy physics. Despite the great success of classical lattice field theory, many phenomena remain out of reach for euclidean path-integral Monte Carlo simulations. For instance, real time dynamics and out-of-equilibrium properties can not be currently studied with classical methods. One of the most important field theories that could take advantage of QC is the strong force. Interactions between quarks and gluons that make up hadrons are described by the theory of Quantum Chromodynamics (QCD). The resulting strong force is one of the four primary computing focus areas delineated by the DOE in the quantum information science applications roadmap [20]. Hadrons are responsible for almost all the mass of matter, but the details of how hadrons form and more generally how nuclei form remain a mystery. This lack of knowledge has profound implications on our understanding of the evolution of the universe, and many experimental facilities like the Large Hadron Collider aim to help increase our knowledge for fundamental physics. These experimental searches require a strong theoretical and computational footing for progress to be made. As already mentioned, classical computing is limited to time-independent studies for only a subset of the strong force phenomenology, as first principle calculations are hindered by the exponential resource requirements [28, 119]. QC will provide scientists with access to real-time dynamics and finite density, thus increasing the reach to all the phenomenology of the strong force. In field theories there is an additional complication arising from the encoding of both fermionic and bosonic degrees of freedom, which can increase the overhead and lead to non-local interactions. Other areas of interest in High Energy Physics (HEP) that could benefit from QC are particle scattering [55] and particle dynamic properties such as collective neutrino flavor oscillations. For a detailed summary we refer the reader to recent review articles [73].

Here we provide a collection of applications aimed at summarizing the current state of research in adopting QC. The Schwinger model is one of the simplest gauge theories. A topological term in the model leads to the infamous sign problem in the classical Monte Carlo method, which means classical computational resources needed scale exponentially. As there is no sign problem for QC, this model is a good use case for early investigations. In [216], the authors simulate time evolution by combining a block encoding of the Hamiltonian, after gauge fixing and mapping with the Jordan-Wigner transform, with the Quantum Singular Value Transformation (QSVT). In [211], the authors develop protocols for time evolution for $U(1)$, $SU(2)$, and $SU(3)$ lattice gauge theories using a variety of methods including Linear Combination of Unitaries (LCU) + qubitization / Trotterization. For geometrically local Hamiltonians, one can take advantage of the HHKL algorithm [101] to further reduce the cost of simulating time evolution [211]. Related results on QC for lattice gauge theories

can also be found in [70, 120].

Other HEP applications included in this work are collective neutrino flavor oscillations, which influence the dynamics of core-collapse supernovae and neutron star mergers, and thus terrestrial detection of these events. Neutrino processes are among the most challenging aspects of current High Performance Computing (HPC) numerical simulations of these environments, a situation complicated by the inherent quantum many-body dynamics in the problem due to neutrino-neutrino interactions, which demand a quantum treatment. There have been recent studies on both qubit and qutrit processors [229] where the time evolution has been performed with first order Trotter steps. For recent resource estimations from Los Alamos National Laboratory see [24, 102]. Here, we combine the latest estimates for the Trotter error [229] to provide conservative bounds on the quantum resources required. The number of Trotter steps required for a given simulation time t and target overall error ϵ is $\mathcal{O}\left(\frac{t^2 \mu N}{\epsilon}(\Delta\omega_{\max} + \mu)\right)$, where μ is the strength of the two-body interaction and $\Delta\omega_{\max}$ is the maximal difference between the Fourier mode frequencies. For our analysis, we estimate $\Delta\omega_{\max} \approx 100\mu$, and $t = \mathcal{O}(\mu^{-1})$. For each Trotter step, the number of T gates is $n_G = 50(N + N(N - 1)/2)$. Using this number of gates per Trotter step, we provide the total number of gates required for different system sizes of interest under a total error budget of 0.01 for the full time evolution.

We conclude by providing an overview of the resource estimates for 32 data points from both applications discussed here in Figure 4.4. Compared to Figures 4.1 and 4.3, we see observe that the number of gates grows significantly faster as a function of qubit count (for an easier visual comparison see Figure 4.5), which we attribute to the larger encoding overhead for applications in HEP, and to a lesser degree to the extent to which the topic has been studied so far. Indeed, in comparison to materials science and chemistry, QC for HEP is less developed and recent works, as outlined here, have been focused primarily on resource estimation for simplified field theory models. We expect more progress to come in the near future as algorithmic approaches are currently being developed to reduce the quantum resources needed for both encoding the problems and for computing the dynamics [63, 207].




Figure 4.4: Overview of various resource estimates for computational problems related to nuclear and high energy physics in terms of number of logical qubits (n_Q) and number of gates (n_G). The inset shows all 32 $\mathcal{P}$ -vectors including some at very large quantum gate and/or qubit count.

4.4 Other application areas

Beyond materials, quantum chemistry, and high energy physics, which form the core of the DOE QRW, there are multiple other areas of computational science where quantum computers can impact the DOE SC mission. Here we summarize the state of affairs for solving problems in linear algebra, differential equations, optimization, and machine learning with quantum computers. We do not include resource

estimates for these categories as they often lack standardized benchmark problems, rely on heuristics, do not describe an application with an end-to-end speedup (including input/readout), or have only examined speedups using oracle queries.

Linear algebra. The field of quantum linear algebra originates from the work by Harrow, Hassidim, and Lloyd (HHL) [104] on solving linear systems of equations on quantum computers with an exponential scaling improvement on the problem size at the cost of a polynomial increase in the dependence on problem conditioning. Subsequent work further improved the dependence on the precision [57] and problem conditioning [9, 188]. QSVT [91] has become the de facto standard for solving various kinds of linear algebra problems, including linear systems and matrix functions. A key bottleneck is the block encoding input model, which can be highly resource efficient for structured data [51, 231] but scales exponentially for unstructured data [49]. Other approaches reduce the complexity of the input model by trading off for an increased sample complexity [247]. We conclude that quantum linear algebra kernels hold great potential, mainly due to the exponential improvement in the dependence on the problem size, and can have considerable implications for a vast number of scientific codes. However, end-to-end complexities and speedups have yet to be developed and will ultimately hinge on problem characteristics (structure) and output requirements.

Simulation of differential equations. The simulation of differential equations is ubiquitous in the physical sciences and engineering, ranging from topics like fluid dynamics and plasma physics, to structural mechanics and computational materials science. Consequently, many applications that are run at NERSC solve a system of differential equations in some form or another. Classical algorithms for solving differential equations form a mature subfield of computational sciences and many codes and solvers have been developed. At the same time, it remains challenging to accurately simulate certain physical phenomena like turbulence, where due to the immense range of interacting lengths and time scales present, extraordinarily fine grids and small time steps are necessary and lead to computationally prohibitive costs.

Research into quantum algorithms for simulating differential equations has seen steady progress since it was first proposed [140] shortly after the introduction of the celebrated HHL algorithm for linear systems [104]. Theoretical results can be categorized in various ways, for example algorithms for Ordinary Differential Equation (ODE) systems [133] and for Partial Differential Equation (PDE) systems [59], or alternatively algorithms for linear differential equations [116] and for nonlinear systems of differential equations [148].

Many approaches convert the system of differential equations to an equivalent problem for which a quantum algorithm exists. This includes mapping the system of differential equations to a system of linear equations [32, 58, 59, 152], to a Hamiltonian simulation problem by introducing auxiliary variables [116, 252], or to Fourier space using the quantum Fourier transform (QFT) that leads to simple quantum circuit designs [157]. Nonlinear problems require careful consideration as solving them on a quantum computer is intrinsically difficult due to the linear nature of quantum mechanics (see Ref. [238] for a review). One approach is to use linearization [148, 237] to transform a nonlinear system to a linear one. Another method is to use a hybrid classical-quantum workflow [134, 158, 197] that translates the differential equation to a minimization problem, which can be evaluated on a quantum computer and variationally minimized using a classical computer. Similar to the case of quantum linear algebra, simulating differential equations on quantum computers holds great potential, but showing evidence of end-to-end speedups remains an open problem.

Optimization and search. Quantum algorithms for combinatorial optimization have been widely studied after the initial introduction of Grover’s search algorithm [100]. Grover’s algorithm provides an asymptotic quadratic speedup for unstructured search problems. Recent work studying Max- k -SAT [45] suggests that practical end-to-end quantum speedups may be challenging to achieve once constant factors are taken into consideration. Furthermore, the overhead required for achieving fault-tolerance using a surface code encoding may outweigh the advantages a quadratic speedup offers [19]. This does not rule out that heuristic and approximate quantum approaches [82, 173] to optimization can

offer substantial benefits. Moreover, recent work suggests that certain optimization problems may be amenable to beyond-quadratic speedups, for example [218] suggests a quartic speedup for the planted noisy k XOR problem.

Scheduling algorithms, specifically load-balancing algorithms, are a category of optimization problems of great relevance to the HPC community. Modern HPC distributed memory problems have increased substantially in complexity, both technically and scientifically. Multi-threading, GPU performance, code coupling, multi-physics, multi-scale, I/O methodologies, digital twin techniques, and workflow applications have all increased the complexity of optimizing the use of HPC compute resources for large-scale, dynamically variable simulations. As Moore’s Law reaches its end and “re-compile and go faster” performance gains disappear, it’s critical for distributed memory HPC applications to use available resources as efficiently as possible to maintain scientific scalability with next-generation architectures.

Classical load-balancers use approximation algorithms due to time constraints — dynamic in-situ load-balancing cannot substantially increase the time-to-solution. So, simple load balancers typically use a broad overview of application performance and aim to do “good enough, fast enough”. High-dimensional optimization solutions have the potential to provide more accurate descriptions of applications, and therefore create more accurate solutions that greatly improve resource utilization. Quantum optimization algorithms that can quickly solve a complex scheduling problem would be of board interest to a variety of distributed memory HPC simulations.

This category of quantum algorithm research has started to be explored using quantum annealing [259] as has the similar field of task scheduling [44].

Machine learning and artificial intelligence. The integration of QC with Machine Learning (ML) and Artificial Intelligence (AI) is poised to overcome several fundamental computational barriers faced by classical methods. Quantum Processing Units (QPUs) offer the potential to address key deficiencies in classical ML, particularly in the efficient handling of high-dimensional feature spaces, combinatorial optimization, and probabilistic sampling. For example, quantum kernel methods can map data into exponentially large Hilbert spaces, enabling the creation of classifiers and regressors that are classically infeasible [219]. In the context of optimization – a central task in AI and ML – Variational Quantum Algorithms (VQAs) leverage quantum parallelism to explore complex solution landscapes, providing a pathway to escape local minima that often trap classical heuristics. This has implications for the training of generative models, solving large-scale combinatorial problems, and optimizing policies in reinforcement learning [54]. Additionally, quantum algorithms can accelerate sampling from complex probability distributions, directly impacting the scalability of Bayesian inference and other probabilistic modeling tasks [150]. As quantum hardware matures, these algorithmic advances could enable ML and AI models to tackle problems previously out of reach for classical computation.

A key challenge, however, is the quantum data loading problem: encoding large classical datasets into quantum states can require resources that scale polynomially with dataset size, potentially eliminating any quantum speedup. As a result, most practical approaches currently focus on hybrid quantum-classical models, where classical computers handle data-intensive tasks and QPUs are used for select, computationally hard subroutines [54, 219]. Another challenge lies in the large regions in the parameter space of quantum circuits where the cost function gradient becomes exponentially small, making optimization extremely difficult. This phenomenon is commonly referred to barren plateaus [86, 165].

The relationship between ML and quantum computing is also reciprocal, as ML techniques are crucial for advancing quantum hardware. Machine learning methods, especially neural networks, are increasingly used for Quantum Error Correction (QEC), learning to decode and correct errors in quantum systems more efficiently than traditional methods [103]. Furthermore, reinforcement learning agents have demonstrated success in automating the calibration and control of qubits, leading to high-fidelity quantum gate operations and improved system stability [183, 227]. This interplay is accelerating progress toward scalable FTQC.

4.5 Summary

We summarize the qualitative differences in $\mathcal{P}$ -vectors for condensed matter/materials science, quantum chemistry, and HEP in Figure 4.5. We observe that in general the condensed matter problems require comparatively the fewest number of gates for a given qubit count (which we use as a proxy for problem size) and that they exhibit the most advantageous scaling behavior. We attribute this to the fact that (1) these problems either map 1-to-1 to qubits or with minimal encoding overhead, and (2) the fields of condensed matter physics and materials science have spent significant efforts to study, develop, and optimize quantum algorithms and applications. The $\mathcal{P}$ -vectors of quantum chemistry applications fall in the intermediate region in between materials and HEP. This is due to the greater qubit encoding overhead compared to condensed matter physics and the greater maturity of the field compared to HEP. Finally, resource estimates for HEP require the largest $\mathcal{P}$ -vectors and the poorest scaling with problem size as they have both a high encoding overhead and the applications studied are not as mature. Based on these observations, we expect that condensed matter physics and materials science problems will be among the first domains of interest to DOE SC mission to benefit from quantum computers with chemistry and HEP likely following later. We summarize this with the following finding.

Finding 3. Condensed matter and materials science applications are expected to be impacted first as they map naturally from spin to qubit systems and require the fewest resources. Quantum chemistry applications require an intermediate amount of resources and have seen steady algorithmic improvements. HEP applications require the most resources among the problems that were considered and likely offer opportunities for significant algorithmic improvements.




Figure 4.5: Qualitative overview of the different regions and scaling of $\mathcal{P}$ -vectors for applications in condensed matter and materials science (yellow), quantum chemistry (blue), and nuclear and high energy physics (red).

With regard to algorithms, we expect Hamiltonian simulation, either through Trotterization or QSVT, and determination of ground state properties to dominate the span across the QRW.

Finding 4. All known quantum applications within the physical sciences for which evidence of exponential speedups exist fall in two categories: dynamical properties (correlation functions, response functions, adiabatic state preparation, ...) and static properties (GSEE, excited states, ...). Consequently, the main algorithmic primitives are Hamiltonian simulation, either using Trotterization (minimal qubit requirements) or qubitization/QSVT (optimal quantum gate requirements), and QPE which relies on Hamiltonian simulation as a subroutine. For quantum chemistry in particular, the majority of resource estimates focus on GSEE with QPE; dynamical problems are comparatively understudied.

Our analysis has centered around the $\mathcal{P}$ -vector which represents the volume of a quantum circuit. As we will discuss in Chapter 5, the $\mathcal{P}$ -vector determines which quantum computer can run which application. However, depending on the nature of the end-to-end application, a different number of circuit repetitions or shots may be required. We write the number of shots as N_s and also refer to this quantity as the sampling complexity. The combination of the $\mathcal{P}$ -vector and the number of shots N_s will determine our execution time estimate presented in Chapter 6.

Table 4.1: Three prototypical sampling complexities for different use cases considered in this report.

Application	N_s
Ground state energy via QPE	10
Measuring an observable $\langle\psi O \psi\rangle$ (VQE, Trotter, ...)	10^3
Dynamical simulations requiring multiple configurations	10^5

Table 4.1 lists three different prototypical sampling complexities that we use for the remainder of this report, and in particular in Chapter 6. For applications that use QPE to determine a ground state energy, multiple trials may be required when the reference state is not the exact ground state. However, we expect QPE to be used only when a good reference state is available and use $N_s = 10$. For applications that require measuring an observable (e.g. magnetization) $\langle\psi|O|\psi\rangle$, with $|\psi\rangle$ prepared using VQE, Trotter, or some other state preparation procedure, we assume a larger $N_s = 10^3$ as many samples are required to reduce the variance of the estimator according to the standard quantum limit. Finally, we contemplate a workflow where the quantum computer is used to support dynamical simulations of some quantum many-body system which requires simulating many different configurations of the many-body system. An example includes calculating the binding affinity of a drug to a target protein which can demand millions of energy calculations using different molecular configurations [185]. Optimistically, we assume $N_s = 10^5$ for a similar but smaller workflow.

Chapter 5

Vendor technology roadmaps

In this section, we collect and unify the public technology roadmaps of ten vendors active in the quantum computing space. We focus on the three technological platforms that are currently leading development efforts: superconducting qubits, trapped ions, and neutral atoms. The different vendors report their roadmap milestones according to different metrics, including:

- The number of qubits in the system, n_Q (physical qubits and/or logical qubits);
- The error rate of the system, ϵ ;
- The number of quantum gates the system can reliably execute, n_G ;
- The quantum circuit depth the system can reliably execute, d_C .

Based on the milestone data provided, we uniformize the roadmaps against the volumetric $\mathcal{P}$ -vector introduced in Equation (4.1). We use the following rules-of-thumb to convert, respectively, error rates and quantum circuit depths to an equivalent number of gates:

$$n_G = \epsilon^{-1}, \quad \text{and} \quad n_G = n_Q \cdot d_C / 2. \quad (5.1)$$

For vendors that report a maximum quantum circuit depth d_C , we convert to quantum gate depth by assuming a maximally dense quantum circuit of $3n_Q/4$ gates per circuit layer. We stress that our rules-of-thumb are a first order approach to obtain estimates that can be reasonably compared to each other. A more precise analysis would require more detailed information that is often publicly unavailable or perhaps still unknown.

The different vendor roadmaps report the number of qubits, the first metric, at their technology milestones either by specifying the number of physical qubits, logical qubits, or both, or remaining ambiguous. The distinction between physical qubits and logical qubits becomes relevant as the field matures from Noisy Intermediate Scale Quantum (NISQ) systems [199] to (early) Fault-Tolerant Quantum Computing (FTQC) [224]. In a NISQ system the (noisy) physical qubits are used directly as a platform for logical (but noisy) quantum operations, while in a FTQC many imperfect physical qubits work together in a carefully orchestrated manner to represent fewer logical qubits with improved performance. We observe that different milestones either specify that the proposed system is a NISQ system or a (early) FTQC system, or do not specify which category the system falls into. In the latter case, we assume that systems that can run 10^4 gates or fewer are NISQ systems (labeled N), systems that can reliably run between 10^4 to 10^6 gates are early FTQC systems (labeled EF), and systems that can reliably run more than 10^6 gates are FTQC systems (labeled F). The number of qubits n_Q that we report in the $\mathcal{P}$ -vector is the number of physical qubits for NISQ systems and the number of logical qubits for (early) FTQC systems. The $\mathcal{P}$ -vector then estimates an upper bound on the performance of the technology that each vendor is developing. Any quantum application that requires at most n_Q qubits and at most n_G gates will be able to run on the forthcoming system. We denote this as $\mathcal{P}^{\text{app}} \leq \mathcal{P}^{\text{qc}}$.

It follows from our discussion on Quantum Error Mitigation (QEM) and Quantum Error Correction (QEC) in Section 3.2 that these protocols can improve the figures of merit and performance bounds

if we allow to take more samples and increase the classical post processing cost. For the purposes of our analysis, we assume that QEM can at most double the number of gates that can be executed on a NISQ device, i.e. $n_{G,QEM} = 2n_G$. In the case of Probabilistic Error Cancellation (PEC), we get from Equation (3.3) that this requires a sampling overhead of $C_{PEC} = e^8 \approx 3,000$, which is at the limit of what we deem practical. For roadmap milestones pertaining to NISQ devices, we present two values for the number of gates: the value based on the information provided by the vendor, and twice this value when assuming QEM is used. The second, higher number $\mathcal{P}_{QEM}^{qc} = (n_Q, 2n_G)$ will be considered the *region of extended capabilities* of the NISQ device and will be plotted as a light shaded area. We remark that for NISQ milestones there is no leeway in the n_Q parameter of the figure of merit as every individual physical qubit is directly used for computation. For FTQC milestones, we consider that QEC and QEM may be used in conjunction and allow for a reduction of the code distance d by five while maintaining the same error rate at the cost of a sample overhead factor $C \leq 100$ and more intensive post-processing [233]. This leads to a figure of merit $\mathcal{P}_{QEM+QEC,1}^{qc} = (\hat{n}_Q, n_G)$ with $\hat{n}_Q > n_Q$ and the exact increase depends on the error correction code and code distance. Similarly, at fixed code distance, one could reach error rates similar to code distance $d + 5$ leading to a figure of merit $\mathcal{P}_{QEM+QEC,2}^{qc} = (n_Q, \hat{n}_G)$ with $\hat{n}_G > n_G$. We use these assumptions to estimate a region of extended capabilities for FTQC milestones that will be plotted as a light shaded area. In this case, both n_Q (by reducing the code distance) and n_G (by keeping the code distance constant but decreasing the effective error rate) can be improved using QEC+QEM.

As vendors support different underlying physical qubits, their error correction strategies will likely differ as well. By default, we will assume surface codes for error correction of operations and PEC for error mitigation. Whenever more detailed information about the future architecture is available, we aim to incorporate it in our estimates. Depending on the hardware characteristics, and the sometimes partial information provided by vendors, we fill in the gaps based on our assumptions stated above.

5.1 Overview of individual vendor roadmaps

In this section we provide an overview of all public roadmap data we identified at the time of writing. We categorize the vendors according to the core qubit technology of their hardware platforms:

- Section 5.1.1: superconducting circuits, which include transmon, coaxmon, and cat qubit types;
- Section 5.1.2: trapped ions;
- Section 5.1.3: neutral atoms;
- Section 5.1.4: other technologies and vendors that have not released official roadmaps.

Within each section, vendors are listed in alphabetical order. We refer to [210] for a different overview of quantum vendor roadmaps.

As our goal is to glean the expected capabilities of near-term quantum systems that the industry is pursuing, we strive to cover as many public roadmaps as possible across different qubit modalities, but we do not claim that our overview is complete. Furthermore, as the field of quantum computing is rapidly progressing, existing roadmaps are updated regularly and new roadmaps are announced frequently. The data we collected reflects a moment in time and is subject to rapid change.

For each vendor, we provide a brief overview of the key concepts and ideas behind their technology and approach to building a scalable quantum computer. We summarize each roadmap in a two panel figure, focusing on recent and future milestones. On the left hand side, a table shows data provided by the vendor in black font and inferred from our assumptions in blue font; the table columns indicate the year for each milestone, the name of the milestone where available, the type of milestone (N for NISQ, EF for early FTQC, or F for FTQC), the number of qubits n_Q (both physical qubits and logical qubits), and a metric that can be related to the number of gates (i.e. the error, the number of gates, or the quantum circuit depth). On the right hand side, a scatter plot shows the $\mathcal{P}$ -vector for each milestone, a dark shaded bounding box highlights the milestone with the largest $\mathcal{P}$ -vector, and a light shaded area indices the region of extended performance, based on QEM for NISQ milestones and QEC+QEM for FTQC milestones.

5.1.1 Qubits based on superconducting circuits

Alice & Bob. Alice & Bob is a start-up based out of Paris, France that is developing quantum computers based on *cat qubits* [168]. Cat qubits represent an innovative approach to building more stable quantum computers by encoding information in superpositions of classical-like states of microwave light, analogous to Schrödinger’s cat being in two states at once. The physical realization of cat states is typically achieved within superconducting microwave resonators, which act as cavities to confine microwave photons that make up the coherent states. Josephson junctions introduce a non-linearity in the superconducting circuit that is required to control the state and perform quantum operations.

An advantage of cat qubits is an inherent, hardware-level protection against bit-flip errors, one of the two main types of quantum errors; this protection becomes exponentially stronger as the size (average photon number $\bar{n}$) of the cat state increases. While this comes at the cost of a smaller linear increase in phase-flip errors, the resulting biased noise profile – where one error type is far less likely than the other – significantly simplifies the complex task of QEC. The QEC reduces from a 2D problem, possibly requiring a surface code to correct both bit- and phase-flips, to a (nearly) 1D problem which can be stabilized by a $[[2d - 1, 1, d]]$ repetition code leading to a linear reduction in qubits compared to a surface code of similar code distance. In this case, we can estimate the logical phase-flip error rate as [215],

$$\epsilon_L = 0.056 \times \left(\frac{\bar{n}^{0.86} k_1/k_2}{0.013} \right)^{(d+1)/2} + (d-1)e^{-2\bar{n}}, \quad (5.2)$$

where k_1 is the single photon loss rate, and k_2 is the two-photon dissipation rate. Based on [215], we assume that $k_1/k_2 \approx 10^{-4}$ and $\bar{n} = 11$, that the encoding ratio of logical qubits to physical qubits is $2d - 1$, and that there is hardware-level protection against bit-flips leading to a logical bit-flip error rate smaller than 10^{-9} for average photon number $\bar{n} = 11$.




Figure 5.1: Roadmap data for: Alice & Bob [8] – *Cat qubits*. *Left*: Table with technology milestones. The Lithium and Beryllium systems have no target year on the roadmap and we interpolate the dates based on the target years for the Helium and Graphene systems. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

Figure 5.1 summarizes the roadmap Alice & Bob released in December 2024 [8]. It lays out 5 technology milestones starting with the Boson 4 system released in 2024, which consists of a single cat qubit and supports a limited (non-universal) set of single-qubit gates, and culminates in the Graphene system in 2030 which is projected to have 100 logical qubits at an error rate of 10^{-6} . Based on the available data, we assume Alice & Bob to pursue an architecture inspired by a repetition code to suppress phase flip errors as their main strategy until the logical phase-flip error rate becomes comparable to the bit-flip error. We thus use Equation (5.2) to estimate the region of extended performance. Around the 2030 milestone, Alice & Bob’s roadmap suggests they will start using sparse LDPC codes to further reduce the logical error rates.

Google Quantum AI. Google Quantum AI is building quantum processors based on superconducting transmon qubits arranged in a square grid topology where each qubit has four nearest-neighbors. This naturally matches the connectivity requirements for a surface code [3]. Google’s current 105 qubit system named Willow was announced in late 2024 [2], while the previous system named Sycamore was originally announced in 2019 with 53 qubits [11] and upgraded to 70 qubits in 2023 [170]. Google Quantum AI presents their technology roadmap [95] in terms of 6 milestones and we summarize them in Figure 5.2. The first milestone, labeled *beyond classical*, was achieved in 2019 on the Sycamore device [11]. Milestone 2, labeled *Quantum Error Correction*, was achieved in 2023 on the second generation Sycamore-2 device [3]. Future milestones 3 through 6 continuously improve both the scale and the quality of the system and Google reports the expected number of physical qubits and the logical error rates in their roadmap. Based on this information, we estimate the code distance using Equation (3.4), which in turns allows us to estimate the number of logical qubits. As their milestones have no estimates on when they will become available, we allow for a 3 year period between consecutive milestones, extrapolating from the announcement of Willow. We assume a physical error rate of 10^{-3} up to milestone 4, and 5×10^{-4} afterwards.




Figure 5.2: Roadmap data for: Google Quantum AI [95] – *Superconducting qubits*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

IBM Quantum. IBM Quantum is building superconducting transmon quantum processors with the qubits arranged in a heavy-hex lattice in recent systems. The company has most recently updated their roadmap [110, 111] in June 2025 and Figure 5.3 summarizes the data in our standardized format. Up until 2028, IBM Quantum reports the number of physical qubits and gates. Starting from the Starling system in 2029, IBM Quantum is expecting a FTQC device and the company reports the number of logical qubits and number of gates for the FTQC systems on their roadmap.

Recent research on FTQC architectures from IBM Quantum suggests that the company is pursuing a QEC strategy based on Quantum Low Density Parity Check (qLDPC) codes [40, 67]. More precisely, a modular architecture based on bivariate bicycle codes, dubbed the *gross code* [254] appears to be a leading candidate for future systems on the IBM Quantum roadmap. This type of architecture will require modifications to the current heavy-hex topology by the introduction of (long-range) couplers [41]. The company has presented first experimental results that demonstrate entanglement generation across such a coupler [106]. Based on [254], we use a $10\times$ reduction in the ratio of physical qubits to logical qubits over a surface code logical qubit with the same error rate as a rule of thumb. For Starling and Blue Jay time frames, we assume a physical error rate of approximately 10^{-3} .

IQM. IQM, founded in 2018, is a company headquartered in Finland that is developing quantum computers based on superconducting transmon qubits. They have released a public roadmap which includes milestones through 2033 [115]. At present, IQM develops systems in two different topologies: a *crystal* topology implementing a 2D square lattice compatible with a surface code encoding and a

Year	Name	Type	n_Q		n_G
			Phys.	Log.	
2020	Falcon	N	27	—	—
2022	Eagle	N	127	—	—
2024	Heron	N	133	—	5,000
2025	Nighthawk	N	120	—	5,000
2026	Nighthawk	N	120	—	7,500
2027	Nighthawk	N	120	—	10,000
2028	Nighthawk	N	120	—	15,000
2029	Starling	F	8,000?	200	10^8
2033+	Blue Jay	F	58,000?	2,000	10^9




Figure 5.3: Roadmap data for: IBM Quantum [111] – *Superconducting qubits*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

star topology with one central qubit connected to all other qubits. According to their roadmap, these two technologies will be combined into a single approach in 2027. We conservatively assume a surface code encoding for their later milestones given the high number of physical qubits to logical qubits. IQM’s roadmap is summarized in Figure 5.4.

Year	Type	n_Q		error
		Phys.	Log.	
2024	N	54	—	1×10^{-3}
2025	N	150	—	8×10^{-4}
2026	N	300	—	6×10^{-4}
2027	EF	10^3	36	10^{-5}
2028	EF	5×10^3	180	10^{-6}
2030	F	4×10^4	720	10^{-7}
2031	F	10^5	1,800	10^{-8}
2033	F	10^6	7,200	10^{-9}




Figure 5.4: Roadmap data for: IQM [115] – *Superconducting qubits*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

Oxford Quantum Circuits. Oxford Quantum Circuits is a quantum computing startup founded in 2017 and headquartered in the United Kingdom. They are building quantum computers based on superconducting qubits with a coaxmon chip design, a 3D architecture in which the qubit and resonator are fabricated on opposing sides of the chip, allowing the separation of the control and readout wiring to different sides of the chip [206]. The company released a technology roadmap in June 2025 [65], which we summarize in Figure 5.5. Oxford Quantum Computing is actively developing multi-mode qubits [251] and in particular dimon qubits [105], where a single dimon qubit is made up of two Josephson junctions and allows for a dual rail encoding of the quantum information [239] which is potentially more error resistant, in particular against erasure errors. We remark that Oxford Quantum Circuits is one of the few companies reporting clock speeds for their future Titan, Athena, and Atlas

systems.




Figure 5.5: Roadmap data for: Oxford Quantum Computing [65] – *Superconducting qubits*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

Rigetti Computing. Rigetti Computing is a full-stack quantum computing company founded in 2013, headquartered in Berkeley, California, and developing superconducting quantum processors based on transmons. While Rigetti Computing has not released a multi-year technology roadmap, their investor presentation from March 2025 [66] contained details about recent systems as well as a system expected later in 2025. The milestones are summarized in Figure 5.6; we remark that all listed devices are NISQ devices and we do not infer an error correction strategy.




Figure 5.6: Roadmap data for: Rigetti Computing [66] – *Superconducting qubits*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

5.1.2 Qubits based on trapped ions

IonQ. IonQ, founded in 2015 and headquartered in College Park, Maryland, is building quantum computers based on linear ion traps. An advantage of this approach is that two-qubit gates between two ions are mediated through the collective motion of the chain which allows for a native all-to-all connectivity within a trap.

In a recent press release from June 13, 2025 [114], IonQ stated updated milestones out to 2030 which include plans for early and large-scale FTQC systems. We summarize these most recent projections in Figure 5.7. IonQ is pursuing a modular architecture which links together multiple ion traps via photonic interconnects. For its 2027 milestone, IonQ wants to achieve 10,000 physical qubits on

a single chip before moving to a modular system in 2028 where two chips will be interconnected leading to a system with 20,000 physical qubits. IonQ envisions their architecture to be compatible with multiple error correction schemes. However, no further details are provided about their error correction strategy but given (1) the high degree of connectivity within a trap and (2) a modest ratio of 12.5 up to 25 physical qubits per logical qubits for their technology milestones, one could expect an approach based on sparse LDPC codes.

Year	Type	n_Q		error
		Phys.	Log.	
2025	N	64	—	10^{-4}
2026	EF	256	12	10^{-7}
2027	EF	10,000	800	10^{-7}
2028	EF	20,000	1,600	10^{-7}
2029	F	200,000	8,000	10^{-12}
2030	F	2,000,000	80,000	10^{-12}




Figure 5.7: Roadmap data for: IonQ [114] – *Trapped ions*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

Quantinuum. Quantinuum is a quantum computing company headquartered in Cambridge, United Kingdom and Broomfield, Colorado that was formed after the merger of Cambridge Quantum and Honeywell Quantum Solutions. Quantinuum is building trapped ion quantum computers based on Quantum Charge-Coupled Device (QCCD) architecture. Their H-series systems are currently industry-leading [92] in terms of the quantum volume benchmark for NISQ systems. The company has announced their roadmap [202] of ion trap devices up to the Apollo system in 2029 and specifies milestones in terms of qubits and error rates. Note that Quantinuum reports the Apollo system to have “hundreds” of logical qubits; we choose 500 as a proxy. The Quantinuum roadmap is summarized in Figure 5.8. One could expect qLDPC or many-hypercube codes [96] as their strategy for error correction due to high qubit connectivity.

Year	Name	Type	n_Q		error
			Phys.	Log.	
2025	Helios	EF	96	50	1×10^{-4}
2027	Sol	EF	192	100	1×10^{-5}
2029	Apollo	F	5,000	500 [?]	1×10^{-10}




Figure 5.8: Roadmap data for: Quantinuum [202] – *Trapped ions*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

5.1.3 Qubits based on neutral atoms

Infleqtion. Infleqtion[30, 112], based out of Boulder, Colorado, is building quantum computers out of dual-species arrays of neutral atoms and reports the number of logical qubits and a combination of error rates and quantum circuit depths on their roadmap that is summarized in Figure 5.9. Their leading QEC strategy revolves around qLDPC codes, and they have recently introduced an open source software library for constructing and analyzing qLDPC codes [194]. On the experimental and architecture side, Infleqtion is exploring techniques for non-destructive readout and tightly focused laser beams that can address single atoms to reduce quantum gate execution times by removing the need to shuttle atoms [205].




Figure 5.9: Roadmap data for: Infleqtion [112] – *Neutral atoms*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

Pasqal. Pasqal is a quantum computing startup based in France that is developing quantum computers based on arrays of Rubidium atoms. The company has recently demonstrated the capability to load large arrays of over a 1,000 atoms [196] and announced their technical roadmap in June of 2025 [192] which we summarize in Fig. 5.10. They aim to combine analog with FTQC approaches in future hybrid systems such as Centaurus and Lyra. Although we were not able to identify specific details on this hybrid approach, recent research [166] suggest that analog mode may be used for preparing an initial many-body state that is subsequently used by a quantum algorithm running in digital mode. We remark that Pasqal is one of the few companies reporting clock speeds for their future systems.

5.1.4 Vendors without public technology roadmaps

In this section we provide an overview of vendors active in the quantum computing space that have not released a technology roadmap that allows us to infer a $\mathcal{P}$ -vector based on our prior assumptions. We do not claim that our overview is exhaustive, since the quantum computing space is vast and evolves constantly. We do remark that while we arguably covered the most mature qubit modalities in the first part of this section (superconducting, ions, and neutral atoms), many different approaches are currently under investigation and development; some alternative technologies include photonic quantum computers, topological quantum computing, silicon spin qubits, and color centers in diamonds.

Alpine Quantum Technologies. Alpine Quantum Technologies is a company focused on building quantum computers based on trapped-ion technology. Alpine Quantum Technologies is based out of Innsbruck, Austria and currently offers a 20 qubit rack-mounted system. At the time of writing, Alpine Quantum Technologies has not released a technology roadmap that allows us to infer a $\mathcal{P}$ -vector.




Figure 5.10: Roadmap data for: Pasqal [192] – *Neutral atoms*. *Left*: Table with technology milestones. *Right*: Plot with equivalent technology milestones on number of qubits versus number of gates axes. The light shaded area sketches the performance boundary using a quantum error mitigation approach on the largest-scale system.

Atom Computing. Atom Computing is a startup based in Berkeley, California and was founded in 2018. The company is developing quantum computers based on atomic arrays of optically-trapped neutral atoms. They encode the qubit state in nuclear spin states of the atom, which leads to qubits that interact weakly with their environment and have a long coherence time [22]. Recent news stories by the company include the announcement of the second generation system with over a 1,000 physical qubits [14], a demonstration of universal gates on a neutral atom system [172], and an announcement of a collaboration with Microsoft [258]. Atom Computing has not released a technology roadmap that allows us to infer a $\mathcal{P}$ -vector.

Equal1. Equal1 is a quantum startup founded in 2017 in Dublin, Ireland. Their technology is rack-mounted silicon based spin qubits [25, 26, 230]. One of their main goals is to make it easy to integrate with High Performance Computing (HPC)-class environments. One of the key advantages of this technology lies in the compatibility with standard semiconductor manufacturing, enabling, among other capabilities, co-integration with control and readout electronics, thus improving scalability and reducing interconnect complexity [131]. The company has not released a public technology roadmap.

Microsoft. Microsoft is pursuing quantum computers based on hardware-protected topological qubits. The company recently announced the Majorana 1 chip [4, 167] and has proposed plans to scale up quantum computational systems based on topological qubits to large-scale fault-tolerant systems [1]. At the time of writing, Microsoft has not released technology milestones that allow us to infer a $\mathcal{P}$ -vector.

PsiQuantum. The mission statement of PsiQuantum is to build and deploy the first useful quantum computers. To this end, the company is pursuing an approach based on photonics and leveraging the semiconductor manufacturing industry [7]. PsiQuantum aims to implement their physical qubits via a dual-rail encoding based on single photons, relying on single photon sources, single photon detectors, and optical switches. Recent research suggests that they are pursuing an approach known as *fusion-based quantum computing* which relies on entangling measurements (or fusions) on the qubits of constant-sized resource states [23], and can be used in an *active volume* architecture [147] which holds the potential to significantly reduce the $\mathcal{P}$ -vector for many scientifically relevant applications. At the time of writing, PsiQuantum has not released a public technology roadmap.

Quantum Circuits Inc. Quantum Circuits Inc. was founded in 2015 out of research by Yale University. The company is building quantum computers based on superconducting circuitry. Their qubits use a dual-rail encoding which allows for native, high-fidelity detection of erasure errors due to

single photon loss, a known, dominant error channel for superconducting qubits [203]. At the time of writing, Quantum Circuits has not released a public technology roadmap.

QuEra Computing. QuEra Computing is a startup based in Boston, Massachusetts that is building neutral atom quantum computers. Their current roadmap [204] only specifies milestones in terms of number of qubits and does not provide any information from which we can infer the number of gates and we thus do not provide estimates for a $\mathcal{P}$ -vector. Their milestones include: (i) ten logical qubits on at least 256 physical qubits in 2024, (ii) 30 logical qubits on at least 3,000 physical qubits in 2025, and (iii) 100 logical qubits on at least 10,000 physical qubits in 2026. In light of recent publications [262], we can expect qLDPC codes to be their target quantum error correction strategy.

Xanadu. Xanadu is a quantum computing company founded in 2016 and headquartered in Toronto, Canada. Xanadu is building quantum computers based on photonics. Their qubits are encoded in optical states called Gottesman-Kitaev-Preskill states which offer a universal quantum gate set based on Gaussian operations [38]. Recent experimental results demonstrate the generation of such GKP states [135] and a prototype of a modular, networked photonic quantum computer [5]. The company has not released a public technology roadmap.

5.2 Overview of all vendor roadmaps

Next, we combine and consolidate the milestones from all ten vendor roadmaps in Figure 5.11. The markers show individual milestones sorted by qubit type, the grayed-out area marks the performance bounds of current quantum systems, and the solid lines show the *expected* performance bounds by the end of the current calendar year, and in five and ten years from today. On the y-axis, we list quantum gate depths up to 10^4 that are achievable by QEM methods and quantum gate depths larger than 10^6 that are expected to require QEC approaches. The intermediate region $[10^4, 10^6]$ will likely be unreachable by QEM only, but perhaps Quantum Error Detection (QED) with post-selection can bridge this region. Nonetheless, to achieve the most ambitious five year technology milestones, which predict an increase in quantum gate depth by *9 orders of magnitude* and in qubit count by *3 orders of magnitude*, it is indisputable that FTQC will be necessary.




Figure 5.11: Overview of all technology milestones presented in Section 5.1. Purple circles indicate milestones for superconducting quantum computers, green squares indicate trapped ion systems, and red triangles correspond to neutral atom systems. The current limits on quantum computer performance correspond to the area in gray. The dark blue line shows the expected performance by the end of 2025, the medium blue line shows the expected performance in five years from the time of writing, and the light blue line shows the expected performance ten years from today.

We summarize our conclusions from evaluating the various roadmaps in the following finding:

Finding 5. The different vendor roadmaps outlining their projected technological progress over the next five to ten years show reasonable alignment in terms of the (n_Q, n_G) figure-of-merit. All vendors project substantial improvements in the figure-of-merit, often by many orders of magnitude. We do foresee that the transition to FTQC will be necessary within the next five years to realize this upscaling.

5.3 Vendor roadmaps versus application requirements

Finally, in Figure 5.12 we combine the results displayed in Figures 4.5 and 5.11 and plot the observed scaling of scientific applications in materials science, chemistry, and High Energy Physics (HEP) together with the hardware performance bounds we derived from vendor projections. Figure 5.12 is a more detailed version of Figure 1.




Figure 5.12: Overview of vendor milestones expected by the end of 2025, in five years from today, and in ten years from today (without the individual milestones shown in Figure 5.11) together with the regions of application resource requirements for high energy physics in red, quantum chemistry in blue, and materials science in yellow shown previously in Figure 4.5. We observe a significant overlap between the capabilities of future quantum computers listed on the vendor roadmaps and requirements for scientific quantum advantage relevant to U.S. Department of Energy (DOE) Office of Science (SC) mission.

As the plots show, with time progression, we can expect the hardware technology, coupled with error mitigation, to provide support for many scientific applications in the Quantum Relevant Workload (QRW) within the next five years. This finding is summarized below.

Finding 6. Under our assumptions, vendor technology roadmaps and DOE SC application resource requirements to support the QRW are expected to converge, meet, and overlap within the next five to ten years.

Chapter 6

Execution times and system throughput

Our analysis thus far has centered around the volumetric figure of merit $\mathcal{P}$ introduced at the beginning of the report. The $\mathcal{P}$ -vector only indicates whether a given quantum computer can run a particular application. It does not provide an indication on the execution time of said application. Execution times are crucial both in terms of the end user experience and to understand the value of a quantum computer in terms of the rate at which it produces scientific results. An improved understanding of execution times is essential to make informed procurement, planning and operational decisions when hosting Quantum Computing (QC) systems.

In Section 6.1 we discuss the expected clock speeds for (early) Fault-Tolerant Quantum Computing (FTQC) systems based on results in the literature and describe key assumptions made. Based on this discussion we define eight regions of feasibility for quantum application execution time. Next, we proceed to introduce a simple model, called Sustained Quantum System Performance (SQSP), for estimating system throughput in Section 6.2. Our model uses a benchmark set of six applications which represent the Quantum Relevant Workload (QRW) and nine different quantum computer configurations, informed by Chapters 4 and 5 respectively. This conceptual model allows us to reason quantitatively about a range of different scenarios.

6.1 Execution times

Execution times on FTQC systems will be highly relevant, both for economical and practical considerations. For example, if an application has a runtime of 1 month and the quantum system has a production life cycle of 3 years, or 36 months, then the value of solving this application should be at least $1/36^{\text{th}}$, or roughly 3%, of the total cost of ownership of the quantum system. On the other hand, increasingly long execution times rapidly become impractical for end users, certainly if the quantum computation is part of a larger computational workflow which requires multiple iterations or shots.

Accurately estimating execution times requires deep insights into the quantum computer architecture and structure of the application. The result of quantum circuit compilation, which is an NP-complete problem [37], can strongly influence the quantum gate depth of a circuit when mapped to hardware. Other factors that play an important role include:

- The duration for implementing each quantum gate from a universal quantum gate set (including both the leading order type of quantum gates as well as the lower order ones). This step is typically rate-limited by the measurement time required to extract the error syndrome;
- The speed of the Quantum Error Correction (QEC) decoder that analyzes the error syndrome data and plans corrective actions can slow down the quantum computer [50];
- The overhead incurred by topology, routing, and qubit shuttling when mapping an application to a quantum computer.

Collectively, these determine the *logical clock speed* or *frequency* f of the quantum computer. As discussed in Section 5.1, only two vendors, Oxford Quantum Computing (1–10 MHz) and Pasqal (10–100+ Hz), provide data on the logical clock speeds as part of their technology roadmaps.

As detailed performance projections are scarce, we provide a high-level overview of execution times. Our discussion is founded on the two data points available from vendors and supplemented with results from the literature, a level of detail we consider adequate for the purpose of reasoning about our findings.

We use the following ranges of clock speed for each technology type:

- *Superconducting*: $f \in [100 \text{ kHz}, 10 \text{ MHz}]$. This range is based on clock speed estimates from Oxford Quantum Computing and [41, 53, 121, 171].
- *Trapped ions*: $f \in [1 \text{ kHz}, 1 \text{ MHz}]$. This range is based on clock speed estimates from [21, 217].
- *Neutral atoms*: $f \in [10 \text{ Hz}, 100 \text{ kHz}]$. This range is based on clock speed estimates from Pasqal and [99, 205, 226].
- *Photons*: $f \in [100 \text{ MHz}, 10 \text{ GHz}]$. This range is based on clock speed estimates from [214].

We observe that these clock speeds span nine orders of magnitude and it is natural to assume that execution times and throughput metrics will span a comparably large range.

We consider the logical clock speed as the rate at which quantum gates can be executed. So an application which requires n_G quantum gates of the dominant type per run and is repeated for N_s times incurs a total quantum gate cost $\mathcal{C}_{\text{tot}} = N_s n_g$ and has an estimated execution time of,

$$t_{\text{exec}} = \mathcal{C}_{\text{tot}} / f. \quad (6.1)$$

We remark that according to this model, the execution time estimate is independent of the number of logical qubits required by an application. This assumption may approximately be satisfied if resource states to implement a dominant gate can be generated and applied to any qubits in the circuit in a time that does not depend on n_Q . More detailed models must take into account more fine-grained aspects of the FTQC architecture, which may vary significantly between different vendors and is left for future analyses.

Under these assumptions, Figure 6.1 shows eight different regions of feasibility for the range of clock speeds f starting at 1 Hz and going up to 1 THz (informed by the previous discussion) and for a total cost $\mathcal{C}_{\text{tot}}$ up to 10^{20} (informed by the application requirements discussion in Chapter 4). These color-coded regions include execution times of:

- *Less than 1 second*: in this case the application runs in near real-time from an end user perspective, it is easy to run applications in interactive mode, and it is practical to integrate the quantum application into a larger High Performance Computing (HPC) workflow with limited scheduling concerns.
- *Less than 1 minute*: in this case the time to execute the application will become noticeable for the end user, running the application in an interactive environment starts to become cumbersome, and scheduling concerns are relevant for integrated workflows.
- *Less than 1 hour*: in this case it is impractical to run the application in interactive mode and an approach using batch jobs is desired, co-scheduling HPC resources in concert with the quantum system will be necessary for workflow-style jobs.
- *Less than 1 day*: this is approaching the limit of what we deem practical for a single application at NERSC; at this timescale scheduling might be less of a concern as multiple other jobs can be run while the quantum system is active.
- *Less than 1 week*: this is at the limit of what we deem practical, this timescale should be reserved for valuable production runs only.

- *Less than 1 month*: this timescale should be considered only in extraordinary circumstances for important applications, for example when it resolves a scientific question of significant value (3% of cost of ownership of the quantum computer); it is unlikely that this will become a relevant use case at NERSC.
- *Less than 1 year*: this timescale should be considered only in the most extraordinary circumstances for high-impact applications where the value of the solution is on par with the total cost of ownership of the quantum computer; it is unlikely that this will become a relevant use case at NERSC.
- *More than 1 year*: we do not consider this to be a useful timescale.

The current time limit on jobs submitted to the Perlmutter system at NERSC is 48h. Due to the different nature and speed of FTQC, it is plausible that longer running jobs will be supported on future quantum systems (i.e. up to 1 week) for valuable production runs. We estimate that within a 1 week period an FTQC system can run an application with a total cost of about 10^7 for the slowest clock speed estimates up to more than 10^{16} for the fastest clock speed estimates.




Figure 6.1: Execution time as a function of logical clock speed f and total cost $C_{\text{tot}} = N_s n_g$. The execution time is estimated as C_{tot}/f . Regions where the execution times are less than 1 second, 1 minute, 1 hour, 1 day, and 1 week are shown in gradually darker shades of green. Execution times less than 1 month are shown in yellow, less than 1 year in orange, and more than 1 year in red. The horizontal strips at the top of the figure show the range of clock frequencies for each technology type as described in the main text. The strips are stacked vertically in order to visually distinguish them.

6.2 Workload and system throughput

In this section, we use data collected thus far to construct a simple quantum workload model that we subsequently optimize for 9 different quantum system specifications. Our model is conceptual in nature and lacks sufficient detail both on the applications and quantum system specifications to draw precise conclusions. Nonetheless, we consider this an important stepping stone towards modeling the value and performance of FTQC systems. System level performance models have historically been used by NERSC as part of procurement processes for HPC systems. Examples include the Sustained System Performance (SSP) [130] and Sustained System Improvement (SSI) [16] metrics which are respectively absolute and relative system-level metrics for a benchmark workload relevant to NERSC.

6.2.1 Prototypical workload

To define a quantum version of a system level performance model, we start by proposing a collection of prototypical benchmark applications that represent the QRW. As there is no historical quantum workload data at NERSC, we use a set of six applications based on the discussion in Chapter 4. The six applications (A–F) are summarized in Table 6.1. The values for the number of shots N_s are based on Table 4.1.

Table 6.1: Six quantum computing applications (A–F), determined by number of qubits n_Q , number of gates n_G , number of samples/repetitions (N_s), total cost (C_{tot}), and prototypical use case.

App.	n_Q	n_G	N_s	C_{tot}	Prototypical use case
■ A	200	10^6	10^1	10^7	Condensed matter physics (static)
■ B	200	10^6	10^3	10^9	Condensed matter physics (dynamics)
■ C	1,000	10^9	10^1	10^{10}	Quantum chemistry (static)
■ D	1,000	10^9	10^5	10^{14}	Quantum chemistry (dynamics)
■ E	2,000	10^{11}	10^1	10^{12}	Quantum chemistry (static, large system)
■ F	250	10^{12}	10^1	10^{13}	High energy physics (static)

Applications A and B listed in Table 6.1 are proxies for condensed matter physics and materials science problems discussed in Section 4.1 and both have a gate depth of 10^6 which was the minimal gate depth identified for this category. Application A assumes ten repetitions, which we deem a reasonable number for a static problem such as Ground State Energy Estimation (GSEE) provided that the reference wavefunction has sufficient overlap with the ground state. Application B on the other hand assumes 10^3 repetitions which represents the sampling overhead required to measure the observable of interest after a quantum dynamics simulation.

Applications C, D and E are all inspired by the quantum chemistry discussion in Section 4.2; C and E both represent static problems like GSEE for respectively a smaller and larger system, they both assume $N_s = 10$ but differ in gate depth between 10^9 and 10^{11} . Application D on the other hand represents some dynamics problem. As was mentioned in Section 4.2, resource data for dynamics problems is scarce in the literature. Hence, we use the same n_G as in Application C, but with drastically increased sampling overhead of $N_s = 10^5$ to reflect the larger number of measurements typically required for a dynamics problem.

Finally, Application F represents the High Energy Physics (HEP) applications discussed in Section 4.3 as demonstrated by the highest gate depth of 10^{12} among all six problems.

Within each scientific domain, we erred on selecting applications on the cheaper side since (1) faster execution times do lead to higher throughput and thus we expect our results to reflect an optimistic scenario, and (2) we expect a general trend of resource reduction through algorithmic advances. Overall, the suite of benchmark applications in Table 6.1 spans seven orders of magnitude from 10^7 up to 10^{14} in total cost and six orders of magnitude in gate depth. We remark that, while we list data for n_Q for these benchmark applications, the execution time estimates are independent of n_Q based on our prior assumptions and modeling approach.

6.2.2 Prototypical quantum computers

Table 6.2 describes nine different possible quantum system specifications varying in maximal gate depth $n_g \in [10^6, 10^9, 10^{12}]$ (in line with the vendor roadmaps), also dubbed megaquop [198] (for million quantum operations), gigaquop and teraquop FTQC systems, and clock speed $f \in [1 \text{ kHz}, 1 \text{ MHz}, 1 \text{ GHz}]$ (in line with Fig. 6.1). We remark that we do not list n_Q for these reference systems as the execution time estimate is independent of n_Q and we make the blanket assumption that the number of qubits is never the constraining factor that determines whether or not an application can run.

Table 6.2: Nine different quantum computer configurations as determined by the maximum number of gates they can reliably run (n_G) and clock frequency f .

n_G	f		
	1 kHz	1 MHz	1 GHz
10^6 (Megaquop)	①	②	③
10^9 (Gigaquop)	④	⑤	⑥
10^{12} (Teraquop)	⑦	⑧	⑨

6.2.3 Execution times of benchmark workload

Figure 6.2 shows both the $\mathcal{P} = (n_G, n_Q)$ and $(n_G, \mathcal{C}_{\text{tot}})$ vectors for the six benchmark applications introduced in Table 6.1 together with the domain-specific regions identified in Fig. 4.5. This shows that within each domain, we selected $\mathcal{P}$ -vectors on the lower end of their respective range. When we consider $\mathcal{C}_{\text{tot}}$ instead of n_G , the cost increases multiplicatively by a factor of N_s but this does not affect whether a given system is able to run an application reliably, it merely increases the total execution time on the quantum system.




Figure 6.2: Domain specific resource requirements introduced in Figure 4.5 together with $\mathcal{P}$ -vectors for the six benchmark applications (A–F) introduced in Table 6.1 as markers with black edges. The corresponding $(n_Q, \mathcal{C}_{\text{tot}})$ -vector for each application is plotted as an edge-less marker of the same type connected by a vertical line to the $\mathcal{P}$ -vector. The solid, dashed, and dotted horizontal lines show the largest application that can be run within a day period on respectively 1 kHz, 1 MHz, and 1 GHz devices considered in Table 6.2

Table 6.3 summarizes the execution time estimates based on Eq. (6.1) for each combination of application (A–F) and quantum system (1–9) introduced in the previous section. Execution time estimates are listed in terms of the eight ranges introduced in Fig. 6.1. Applications C–F and E–F cannot run on respectively the megaquop and gigaquop systems as their number of gates n_G lies beyond the limits of the devices.

The main conclusion we draw from Table 6.3 is that there is a clear separation into three categories of applications that impose increasingly demanding requirements on the quantum systems:

1. Application A and B: These have the lowest total cost $\mathcal{C}_{\text{tot}}$ at respectively 10^7 and 10^9 and can run with acceptable performance on each of the nine quantum systems considered. Naturally, faster clock speeds lead to faster quantum calculations, but even at a frequency of 1 kHz applications of this size become feasible.

2. Application C and E: These have an intermediate total cost of respectively 10^{10} and 10^{12} and require a processing speed of 1 MHz to achieve feasible execution times. Application E has a gate depth beyond 10^9 and thus is only deemed feasible on systems 8 and 9.
3. Application D and F: These have the highest total cost at respectively 10^{14} and 10^{13} . At this level, only the GHz-rate systems result in a feasible execution time estimate. Application F has a gate depth beyond 10^9 and thus is only deemed feasible on system 9.

Table 6.3: Execution times for quantum applications A–F on nine quantum computing systems. Colors indicate time ranges: ■ < 1s, ■ < 1min, ■ < 1h, ■ < 1day, ■ < 1wk, ■ < 1mo, ■ < 1yr, ■ > 1yr. Invalid combinations (i.e. when the number of gates in the application is larger than the number of gates the system can reliably run) are marked with $\times$.

System Type	System	f	Quantum Applications					
			■ A	■ B	■ C	■ D	■ E	■ F
Megaquop $n_G = 10^6$	(1)	1 kHz	< 1day	< 1mo	$\times$	$\times$	$\times$	$\times$
	(2)	1 MHz	< 1min	< 1h	$\times$	$\times$	$\times$	$\times$
	(3)	1 GHz	< 1s	< 1min	$\times$	$\times$	$\times$	$\times$
Gigaquop $n_G = 10^9$	(4)	1 kHz	< 1day	< 1mo	< 1yr	> 1yr	$\times$	$\times$
	(5)	1 MHz	< 1min	< 1h	< 1day	> 1yr	$\times$	$\times$
	(6)	1 GHz	< 1s	< 1min	< 1min	< 1wk	$\times$	$\times$
Teraquop $n_G = 10^{12}$	(7)	1 kHz	< 1day	< 1mo	< 1yr	> 1yr	> 1yr	> 1yr
	(8)	1 MHz	< 1min	< 1h	< 1day	> 1yr	< 1mo	< 1yr
	(9)	1 GHz	< 1s	< 1min	< 1min	< 1wk	< 1h	< 1day

6.2.4 Sustained quantum system performance

In this section, we combine the data generated in Table 6.3 into a single system level metric that we call the SQSP metric and is an extension of the SSP metric previously introduced by NERSC [130] to quantum systems. The SSP metric for a supercomputer is defined as the geometric mean of the flop rate of each application in a properly defined benchmark suite on said supercomputer. In absence of the capability to run a representative benchmark suite on a large-scale quantum computer, we use the execution time estimates in Table 6.3 to compute the expected throughput T_i of application i over a one year time window, i.e.,

$$T_i = 1\text{yr} / t_{\text{exec},i}. \quad (6.2)$$

The one year time period is chosen as it corresponds to the typical cycle at which projects are allocated computer resources at NERSC.

We then define SQSP as the geometric mean of the throughput values of every application in a suite of n benchmarks:

$$\text{SQSP} = (\prod_{i=1}^n T_i)^{1/n}. \quad (6.3)$$

Applications that cannot run on a particular quantum system ($\times$ in Table 6.3) have infinite execution time and zero throughput, and the SQSP for such a system will be zero. The value of SQSP metric thus depends critically on the suite of benchmarks, and SQSP values for different suites cannot be compared to each other.

Table 6.4 shows the results of the SQSP metric based on the data collected in Table 6.3 for the nine quantum systems we consider. The SQSP for each system type is evaluated using a different benchmark

suite consisting of only the subset of applications from Table 6.3 that can run on that architecture. We do observe that the SQSP metric spans about eight orders of magnitude. We conclude from our analysis that (1) faster FTQC devices naturally lead to high SQSP results, (2) at megaquop-scale all systems lead to SQSP of at least $\mathcal{O}(10^3)$ megaquop-jobs/year, (3) at gigaquop-scale, the throughput of the kHz system drops below 10 gigaquop-jobs/year, indicating that higher frequencies will be a necessity, and (4) at teraquop-scale, clock speeds will need to reach at least MHz-range in order to achieve meaningful throughput.

Table 6.4: SQSP for a one year window for every system defined in Table 6.2 with respect to a benchmark suite which is a subset of the workload defined in Table 6.1.

System Type	Suite	System	f	SQSP
Megaquop $n_G = 10^6$	A-B	①	1 kHz	
		②	1 MHz	
		③	1 GHz	
Gigaquop $n_G = 10^9$	A-D	④	1 kHz	
		⑤	1 MHz	
		⑥	1 GHz	
Teraquop $n_G = 10^{12}$	A-F	⑦	1 kHz	
		⑧	1 MHz	
		⑨	1 GHz	
				

The rudimentary performance model expressed by Eq. (6.1) does not account for details of the FTQC architecture (such as qubit connectivity, parallel gate operations, or multi-qubit gates) that have the potential to dramatically alter system performance. Practical use of SQSP would avoid this limitation by relying on measured runtimes or more detailed device-specific models. Nevertheless, the current analysis is capable of summarizing anticipated application-level capabilities across the QC industry.

Chapter 7

Conclusion

NERSC is exploring and evaluating alternative compute technologies that can deliver increased performance in a more power-efficient manner in the 2030's time-frame for the next-next HPC generation. One key direction in this regard is represented by quantum computers that can lead to exponential speedups for important scientific problems that classical HPC cannot solve. Based on the Quantum Information Science (QIS) Applications Roadmap of the U.S. Department of Energy (DOE) [20], we have focused in this report on materials science, quantum chemistry, and high energy physics, and collected over 140 end-to-end resource estimates for benchmark problems from the scientific literature. These domains of science cover over 50% of the DOE Office of Science (SC) production workload, which aggregates the computational needs of more than 12,000 NERSC users across the DOE landscape.

The space of quantum hardware development has seen major improvement in recent years, especially through noise reduction and early tests of error correction [2, 3, 132, 161, 235, 261] across different technology propositions such as neutral atoms, ions, and superconducting qubits. Given these developments, we see very ambitious roadmaps across the whole sector, as this is a high-risk high-reward environment when considering the disruptive power of quantum computing. On the government side, we notice a similar fast approach as can be seen from the DARPA Quantum Benchmarking Initiative (QBI), which targets Fault-Tolerant Quantum Computing (FTQC) in under 10 years [201]. This bullish approach will require risk mitigation as the technology is not yet mature and recent successes, while very promising, have been in the realm of experimental laboratories and the production chain is currently under development. The market is exploring many hardware technologies with superconducting, neutral atoms, and ions currently taking the lead. Alternative technologies like photonics, spin qubits, topological qubits, etc are also currently under exploration by various vendors. At the present, it remains too early to assess which approach will stand the test of time, and the coming years will be crucial in technology development and the transition from laboratory to production scale.

We have superimposed the resources needed for the scientific applications collected from the literature with the vendors roadmaps and see convergence in the near future, particularly so considering algorithmic developments. Indeed, on the algorithmic front, we notice resource requirements being greatly reduced with every development (see FeMoco in quantum chemistry). The combination of these two trends promises early fault tolerant quantum relevant workloads in the next five to ten years. Additionally, there is a high likelihood that new and heuristic impactful applications will be developed as better quantum systems become available to researchers.

Whether this utility will cover a reasonable fraction of High Performance Computing (HPC) workloads (e.g. around 50%), or become a dedicated resource for certain workloads, remains to be seen, and further algorithmic development and hardware validation are needed in this direction. To get an initial estimate of timing and throughput, we selected a few applications from condensed matter, chemistry, and HEP and computed their runtime for quantum hardware clock speeds between 1 KHz and 1 GHz. We proposed the Sustained Quantum System Performance (SQSP) as a metric to compare system-level throughput for a quantum system. Perhaps not surprisingly, the number of tasks completed varies wildly between different hardware specifications. This is an indication of the need for careful selection of applications for early FTQC, where resources are expected to be limited.

In addition, the classical resources needed for error correction and hybrid algorithms with both

quantum and classical resources in the loop, might require tight coupling with low latencies between HPC and Quantum Computing (QC) [50]. This inevitably adds another layer of complication to the proposed vendor roadmaps. A tight collaboration between the public and private sector is key in overcoming these challenges and unlocking the quantum computing potential.

Acknowledgments

We thank Akel Hashim for feedback on the manuscript. This research was supported by the U.S. Department of Energy (DOE) under Contract No. DEAC02-05CH11231, through the National Energy Research Scientific Computing Center (NERSC), an Office of Science User Facility located at Lawrence Berkeley National Laboratory.

List of Acronyms

AI	Artificial Intelligence 4, 22
ASCR	Advanced Scientific Computing Research 7
BER	Biological and Environmental Research 7
BES	Basic Energy Sciences 7
CNOT	Controlled-NOT 9, 10
CPU	Central Processing Unit 4, 7
DFT	Density Functional Theory 7, 8, 17
DOE	U.S. Department of Energy 2, 4, 5, 7, 8, 14, 16, 18–20, 23, 35, 43
FES	Fusion Energy Science 7
FTQC	Fault-Tolerant Quantum Computing 5, 10–12, 14, 15, 22, 25, 26, 28, 30, 32, 34–39, 42, 43
GPU	Graphics Processing Unit 4, 7, 22
GSEE	Ground State Energy Estimation 16–18, 24, 39
HEP	High Energy Physics 7, 19, 20, 23, 35, 39
HPC	High Performance Computing 2, 4, 5, 20, 22, 33, 37, 38, 43, 44
IRI	Integrated Research Infrastructure 4
LCU	Linear Combination of Unitaries 19
LQCD	Lattice Quantum Chromodynamics 8
ML	Machine Learning 22
NERSC	National Energy Research Scientific Computing Center 2–5, 7, 8, 21, 37–39, 41, 43
NESAP	NERSC Science Acceleration Program 4, 7
NISQ	Noisy Intermediate Scale Quantum 4, 9, 11, 14, 25, 26, 30, 31
NP	Nuclear Physics 7
ODE	Ordinary Differential Equation 21
PDE	Partial Differential Equation 21
PEC	Probabilistic Error Cancellation 10, 11, 26
PIC	Particle-In-Cell 8

QC Quantum Computing 2–5, 8, 9, 12, 13, 16, 18–20, 22, 36, 42, 44

QCCD Quantum Charge-Coupled Device 31

QCD Quantum Chromodynamics 19

QEC Quantum Error Correction 4, 10–12, 22, 25–28, 32, 34, 36

QED Quantum Error Detection 34

QEM Quantum Error Mitigation 9–12, 25, 26, 34

QIS Quantum Information Science 4, 43

qLDPC Quantum Low Density Parity Check 12, 28, 31, 32, 34

QPE Quantum Phase Estimation 13, 16–19, 24

QPU Quantum Processing Unit 22

QRW Quantum Relevant Workload 8, 14, 16, 18, 20, 23, 35, 36, 39

QSP Quantum Signal Processing 13, 14, 17

QSVT Quantum Singular Value Transformation 13, 19, 21, 23, 24

SC Office of Science 2, 4, 7, 8, 14, 18, 20, 23, 35, 43

SPAM State Preparation and Measurement 9, 11

SQSP Sustained Quantum System Performance 1, 3, 5, 36, 41–43

SSI Sustained System Improvement 38

SSP Sustained System Performance 38, 41

THC Tensor Hypercontraction 17

VQA Variational Quantum Algorithm 22

VQE Variational Quantum Eigensolver 13, 17, 18, 24

Glossary of Terms

- Code Distance** A measure of a quantum error correction code’s robustness, defined as the minimum number of physical qubit errors required to cause a logical error. A code with distance d can detect up to $d - 1$ errors and correct up to $\lfloor (d - 1)/2 \rfloor$ errors. Higher code distances provide stronger error protection but require more physical qubits per logical qubit. 12, 26–28
- Code Threshold** The critical physical error rate below which quantum error correction becomes beneficial, allowing logical error rates to decrease exponentially as more resources are added. Above this threshold, adding more error correction actually increases the logical error rate. For surface codes, the threshold is approximately 1% error rate per operation. 5, 12
- Controlled-NOT** An controlled-NOT gate (also known as a controlled- X gate) is a two-qubit entangling gate. It applies a NOT (or Pauli- X) operation on the target qubit when the control qubit is in the $|1\rangle$ -state. Its action is summarized as $|a\rangle|b\rangle \mapsto |a\rangle|b \oplus a\rangle$. 9, 10, 46
- Error Syndrome** The pattern of measurement outcomes from ancilla qubits that indicates which errors have occurred in a quantum system without directly measuring the data qubits. Syndrome extraction allows errors to be identified and corrected without collapsing the quantum state, forming the basis of quantum error correction protocols. 12, 36
- Fault-Tolerant Quantum Computing** A quantum computing paradigm where logical operations can be performed reliably despite the presence of errors in physical components. This requires quantum error correction codes, fault-tolerant gate implementations, and error rates below the code threshold to ensure that errors do not propagate faster than they can be corrected. 5, 10–12, 14, 15, 22, 25, 26, 28, 30, 32, 34–39, 42, 43, 46
- First Quantization** A formulation of quantum mechanics where particles are treated as quantum mechanical objects with wavefunctions, while fields remain classical. In this approach, the Schrödinger equation describes the evolution of N -particle wavefunctions in configuration space, with the wavefunction amplitude giving the probability of finding particles at specific positions. While conceptually straightforward, first quantization becomes computationally challenging for many-body systems due to the exponential growth of the configuration space with particle number. 16, 18
- Ground State Energy Estimation** A fundamental quantum computing application that determines the lowest energy eigenvalue of a quantum system’s Hamiltonian. This problem is central to quantum chemistry and materials science, as ground state energies determine molecular properties and reaction rates. Quantum algorithms like variational quantum eigensolver (VQE) and quantum phase estimation can solve this problem with potential advantages over classical methods. 16–18, 24, 39, 46
- Linear Combination of Unitaries** An important algorithm in developing block encoding for arbitrary operators by first decomposing them in terms of unitary operators, such as the Pauli basis. In a second stage, the applications of prepare, select, and unprepare primitives allows for the construction of the block encoding of the target operator. This algorithm can serve as the preparatory step before QSVT algorithm is applied. 19, 46

Logical Qubit An encoded qubit protected by quantum error correction, formed from multiple physical qubits working together to store quantum information redundantly. Logical qubits can maintain coherence far longer than individual physical qubits by detecting and correcting errors, enabling reliable quantum computation at the cost of increased resource requirements. 9, 12, 16, 19, 20, 25–29, 31, 32, 34, 37

Noisy Intermediate Scale Quantum Noisy Intermediate-Scale Quantum devices, referring to current-generation quantum computers with 50-1000 qubits that lack full error correction. NISQ devices can potentially demonstrate quantum advantage for specific problems but are limited by noise, short coherence times, and restricted circuit depths, requiring specialized algorithms and error mitigation techniques. 4, 9, 11, 14, 25, 26, 30, 31, 46

Physical Qubit The fundamental hardware unit of quantum information in a quantum computer, such as a superconducting circuit, trapped ion, or quantum dot. Physical qubits are subject to decoherence and operational errors, typically maintaining quantum states for microseconds to seconds depending on the platform, necessitating error correction for practical quantum computing. 5, 9, 10, 12, 25–31, 33, 34

Probabilistic Error Cancellation A quantum error mitigation technique that reduces the effect of noise by running multiple circuit variants with inserted recovery operations. The results are combined with appropriate weights to cancel out error contributions statistically. While this method can significantly improve accuracy, it incurs an exponential sampling overhead that limits its applicability to moderate circuit depths. 10, 11, 26, 46

Quantum Circuit A sequence of quantum gates applied to qubits, representing a quantum algorithm or computation. Circuits are typically expressed as directed acyclic graphs where nodes represent gates and edges indicate qubit connections. Circuit depth, width, and gate count are key metrics that determine resource requirements and feasibility on a quantum computer. 3, 9, 11, 14, 15, 21, 24–26, 32, 36

Quantum Error Correction The process of protecting quantum information from decoherence and operational errors by encoding logical qubits into entangled states of multiple physical qubits. Error correction codes detect and correct errors without directly measuring the protected information, enabling fault-tolerant quantum computation when error rates fall below the code threshold. 4, 10–12, 22, 25–28, 32, 34, 36, 47

Quantum Error Detection The process of identifying when errors have occurred in a quantum system without necessarily correcting them. Error detection requires fewer resources than full error correction and can be used to post-select successful runs or trigger re-computation. This approach is particularly useful for quantum communication and certain NISQ applications where modest error rates are acceptable. 34, 47

Quantum Error Mitigation Techniques to reduce the impact of errors in NISQ devices without the full overhead of quantum error correction. Methods include zero-noise extrapolation, probabilistic error cancellation, and symmetry verification. While these approaches cannot eliminate errors entirely, they can significantly improve computation accuracy at the cost of increased sampling or classical post-processing. 9–12, 25, 26, 34, 47

Quantum Gate A reversible operation on one or more qubits that forms the basic building block of quantum circuits. Common gates include single-qubit rotations (X, Y, Z, Hadamard) and two-qubit entangling gates (CNOT, CZ). Gate fidelity, duration, and connectivity constraints determine the practical implementation of quantum algorithms on physical hardware. 2, 3, 9, 11, 12, 14, 16, 20, 22, 24, 25, 32, 34, 36, 37

Quantum Information Science The interdisciplinary field studying how quantum mechanical systems can be used to store, process, and transmit information. QIS encompasses quantum computing, quantum communication, quantum cryptography, and quantum sensing, exploring both

fundamental limits of information processing and practical applications. The field combines physics, computer science, mathematics, and engineering to develop quantum technologies that can surpass classical limitations in computation, communication security, and measurement precision. 4, 43, 47

Quantum Low Density Parity Check A class of quantum error correction codes characterized by sparse parity check matrices, where each qubit participates in only a few syndrome measurements. These codes can achieve better encoding rates than surface codes while maintaining good error correction properties, though they typically require non-local connectivity that poses implementation challenges on current hardware. 12, 28, 31, 32, 34, 47

Quantum Phase Estimation A fundamental quantum algorithm that extracts eigenvalue information from unitary operators by encoding phase differences into measurable probability amplitudes. This algorithm underlies many quantum applications including Shor’s factoring algorithm and quantum chemistry simulations, offering exponential speedup for certain problems but requiring deep circuits. 13, 16–19, 24, 47

Quantum Signal Processing A framework for designing quantum algorithms by constructing polynomial transformations through sequences of rotation gates and signal operators. QSP provides optimal quantum algorithms for a wide range of problems by encoding desired polynomial functions into alternating sequences of signal operators and signal processing operators. This technique forms the foundation for quantum singular value transformation and achieves optimal query complexity for tasks including Hamiltonian simulation, matrix inversion, and amplitude amplification. 13, 14, 17, 47

Quantum Singular Value Transformation A unified framework for quantum algorithms that transforms singular values of block-encoded matrices through polynomial approximations. This technique generalizes many quantum algorithms including phase estimation, amplitude amplification, and Hamiltonian simulation, providing optimal query complexity for a broad class of matrix transformation problems. 13, 19, 21, 23, 24, 47

Qubit The fundamental unit of quantum information, analogous to a classical bit but capable of existing in superposition states $|0\rangle$, $|1\rangle$, or any linear combination $\alpha|0\rangle + \beta|1\rangle$. Qubits enable quantum parallelism and entanglement, the key resources for quantum computational advantage, but require careful isolation from environmental noise to maintain coherence. 2, 3, 9, 14, 16, 17, 20, 23, 24, 34, 37

Qubitization A quantum algorithm technique that encodes the eigenvalues of a Hamiltonian into the phases of a unitary operator, enabling efficient quantum simulation and eigenvalue estimation. This method block-encodes the Hamiltonian into a larger unitary matrix, allowing quantum phase estimation to extract eigenvalues with optimal query complexity. Qubitization often provides better scaling than Trotterization for Hamiltonian simulation problems. 12, 14, 16–19, 24

Qutrit A three-level quantum system that extends the qubit concept to a superposition of three basis states $|0\rangle$, $|1\rangle$, $|2\rangle$. Qutrits can provide advantages in certain quantum algorithms and error correction schemes by offering higher information density and more efficient gate decompositions, though they are generally more challenging to control and maintain than qubits. 20

Second Quantization A formulation of quantum mechanics using creation and annihilation operators to describe many-body quantum systems in terms of occupation numbers rather than individual particle coordinates. This approach naturally handles variable particle numbers and identical particle statistics (bosons/fermions), making it essential for quantum chemistry and condensed matter physics. Second quantization provides the foundation for efficient quantum simulation algorithms, as fermionic or bosonic operators can be mapped to qubit operations through transformations like Jordan-Wigner or Bravyi-Kitaev. 16, 18

Shot A single execution of a quantum circuit on a quantum computer, producing one measurement outcome from the probabilistic quantum state. Since quantum measurements are inherently probabilistic, multiple shots are typically required to estimate expectation values or probability distributions. The number of shots needed scales with desired statistical accuracy, with standard error decreasing as $\mathcal{O}(N_s^{-1/2})$ for N_s shots. 9, 10, 24, 36

Surface Code The leading quantum error correction code for near-term implementation, arranging physical qubits on a 2D lattice where data and measurement qubits alternate in a checkerboard pattern. Surface codes require only nearest-neighbor interactions, tolerate error rates up to $\sim 1\%$, and scale favorably, making them the primary choice for many quantum computing architectures despite their high qubit overhead. 12, 21, 26–29

Trotter A method for simulating quantum time evolution by decomposing the exponential of a sum of non-commuting operators into a product of simpler exponentials. While conceptually simple and widely used, Trotterization typically requires more gates than advanced methods like qubitization for achieving the same accuracy, though it has the advantage of requiring fewer ancilla qubits, making it more suitable for near-term quantum devices with limited qubit counts. 14, 17–20, 23, 24

Bibliography

- [1] D. Aasen et al. *Roadmap to fault tolerant quantum computation using topological qubit arrays*. 2025. DOI: [10.48550/ARXIV.2502.12252](https://arxiv.org/abs/2502.12252). URL: <https://arxiv.org/abs/2502.12252>.
- [2] R. Acharya et al. “Quantum error correction below the surface code threshold”. In: *Nature* 638.8052 (Dec. 2024), pp. 920–926. DOI: [10.1038/s41586-024-08449-y](https://doi.org/10.1038/s41586-024-08449-y). URL: <http://dx.doi.org/10.1038/s41586-024-08449-y>.
- [3] R. Acharya et al. “Suppressing quantum errors by scaling a surface code logical qubit”. In: *Nature* 614.7949 (Feb. 2023), pp. 676–681. DOI: [10.1038/s41586-022-05434-1](https://doi.org/10.1038/s41586-022-05434-1). URL: <https://doi.org/10.1038/s41586-022-05434-1>.
- [4] M. Aghaee et al. “Interferometric single-shot parity measurement in InAs–Al hybrid devices”. In: *Nature* 638.8051 (Feb. 2025), pp. 651–655. DOI: [10.1038/s41586-024-08445-2](https://doi.org/10.1038/s41586-024-08445-2). URL: <http://dx.doi.org/10.1038/s41586-024-08445-2>.
- [5] H. Aghaee Rad et al. “Scaling and networking a modular photonic quantum computer”. In: *Nature* 638.8052 (Jan. 2025), pp. 912–919. DOI: [10.1038/s41586-024-08406-9](https://doi.org/10.1038/s41586-024-08406-9). URL: <http://dx.doi.org/10.1038/s41586-024-08406-9>.
- [6] D. Aharonov and A. Ta-Shma. *Adiabatic Quantum State Generation and Statistical Zero Knowledge*. 2003. arXiv: [quant-ph/0301023](https://arxiv.org/abs/quant-ph/0301023) [quant-ph]. URL: <https://arxiv.org/abs/quant-ph/0301023>.
- [7] K. Alexander et al. “A manufacturable platform for photonic quantum computing”. In: *Nature* 641.8064 (Feb. 2025), pp. 876–883. DOI: [10.1038/s41586-025-08820-7](https://doi.org/10.1038/s41586-025-08820-7). URL: <http://dx.doi.org/10.1038/s41586-025-08820-7>.
- [8] Alice&Bob. *Alice&Bob’s Quantum Computing Roadmap*. <https://alice-bob.com/wp-content/uploads/2024/12/Alice-Bob-Roadmap-Long-Version.pdf>. [Accessed 03-03-2025]. Dec. 2024.
- [9] A. Ambainis. “Variable time amplitude amplification and quantum algorithms for linear algebra problems”. In: *29th International Symposium on Theoretical Aspects of Computer Science (STACS 2012)*. Ed. by C. Dürr and T. Wilke. Vol. 14. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2012, pp. 636–647. DOI: [10.4230/LIPIcs.STACS.2012.636](https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.STACS.2012.636). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.STACS.2012.636>.
- [10] P. W. Anderson. “The resonating valence bond state in La_2CuO_4 and superconductivity”. In: *Science* 235.4793 (1987), pp. 1196–1198.
- [11] F. Arute et al. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (Oct. 2019), pp. 505–510. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5). URL: <http://dx.doi.org/10.1038/s41586-019-1666-5>.
- [12] A. Aspuru-Guzik and P. Walther. “Photonic quantum simulators”. In: *Nat. Phys.* 4 (2008), pp. 437–439.
- [13] A. Aspuru-Guzik et al. “Simulated quantum computation of molecular energies”. In: *Science* 309.5741 (2005), pp. 1704–1707.

- [14] atomcomputing. *Quantum startup Atom Computing first to exceed 1,000 qubits - Atom Computing* — *atom-computing.com*. <https://atom-computing.com/quantum-startup-atom-computing-first-to-exceed-1000-qubits/>. [Accessed 17-06-2025]. 2023.
- [15] A. Auerbach. *Interacting Electrons and Quantum Magnetism*. Springer, 1994.
- [16] B. Austin et al. “A Metric for Evaluating Supercomputer Performance in the Era of Extreme Heterogeneity”. In: *2018 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. 2018, pp. 63–71. DOI: [10.1109/PMBS.2018.8641549](https://doi.org/10.1109/PMBS.2018.8641549).
- [17] R. Babbush, D. W. Berry, and H. Neven. “Quantum simulation of the Sachdev-Ye-Kitaev model by asymmetric qubitization”. In: *Phys. Rev. A* 99 (4 Apr. 2019), p. 040301. DOI: [10.1103/PhysRevA.99.040301](https://doi.org/10.1103/PhysRevA.99.040301). URL: <https://link.aps.org/doi/10.1103/PhysRevA.99.040301>.
- [18] R. Babbush et al. “Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity”. In: *Physical Review X* 8.4 (Oct. 2018). DOI: [10.1103/physrevx.8.041015](https://doi.org/10.1103/physrevx.8.041015). URL: <http://dx.doi.org/10.1103/PhysRevX.8.041015>.
- [19] R. Babbush et al. “Focus beyond Quadratic Speedups for Error-Corrected Quantum Advantage”. In: *PRX Quantum* 2 (1 Mar. 2021), p. 010103. DOI: [10.1103/PRXQuantum.2.010103](https://doi.org/10.1103/PRXQuantum.2.010103). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.010103>.
- [20] A. Baczewski et al. *Quantum Information Science Applications Roadmap*. https://www.quantum.gov/wp-content/uploads/2024/12/DOE_QIS_Roadmap_Final.pdf. 2024.
- [21] C. J. Ballance et al. “High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits”. In: *Phys. Rev. Lett.* 117 (6 Aug. 2016), p. 060504. DOI: [10.1103/PhysRevLett.117.060504](https://doi.org/10.1103/PhysRevLett.117.060504). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.060504>.
- [22] K. Barnes et al. “Assembly and coherent control of a register of nuclear spin qubits”. In: *Nature Communications* 13.1 (May 2022). DOI: [10.1038/s41467-022-29977-z](https://doi.org/10.1038/s41467-022-29977-z). URL: <http://dx.doi.org/10.1038/s41467-022-29977-z>.
- [23] S. Bartolucci et al. “Fusion-based quantum computation”. In: *Nature Communications* 14.1 (Feb. 2023). DOI: [10.1038/s41467-023-36493-1](https://doi.org/10.1038/s41467-023-36493-1). URL: <http://dx.doi.org/10.1038/s41467-023-36493-1>.
- [24] A. Bäertschi et al. *Potential Applications of Quantum Computing at Los Alamos National Laboratory*. 2024. DOI: [10.48550/ARXIV.2406.06625](https://doi.org/10.48550/ARXIV.2406.06625). URL: <https://arxiv.org/abs/2406.06625>.
- [25] I. Bashir et al. “A Single-Electron Injection Device for CMOS Charge Qubits Implemented in 22-nm FD-SOI”. In: *IEEE Solid-State Circuits Letters* 3 (2020), pp. 206–209. DOI: [10.1109/LSSC.2020.3010822](https://doi.org/10.1109/LSSC.2020.3010822).
- [26] I. Bashir et al. “Monolithically Integrated Quantum Dots in a 22-nm Fully Depleted Silicon-on-Insulator Process Operating at 3 K”. In: *International Journal of Circuit Theory and Applications* 53.7 (2025), pp. 3902–3911. DOI: <https://doi.org/10.1002/cta.4350>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/cta.4350>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cta.4350>.
- [27] B. Bauer et al. “Quantum algorithms for quantum chemistry and quantum materials science”. In: *Chem. Rev.* 120.22 (2020), pp. 12685–12717.
- [28] C. W. Bauer et al. “Quantum Simulation for High-Energy Physics”. In: *PRX Quantum* 4 (2 May 2023), p. 027001. DOI: [10.1103/PRXQuantum.4.027001](https://doi.org/10.1103/PRXQuantum.4.027001). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.4.027001>.
- [29] A. J. Bay-Smidt et al. *Fault-tolerant quantum simulation of generalized Hubbard models*. 2025. arXiv: [2501.10314](https://arxiv.org/abs/2501.10314) [quant-ph]. URL: <https://arxiv.org/abs/2501.10314>.
- [30] M. J. Bedalov et al. *Fault-Tolerant Operation and Materials Science with Neutral Atom Logical Qubits*. Dec. 2024. arXiv: [2412.07670](https://arxiv.org/abs/2412.07670) [quant-ph].

- [31] E. van den Berg et al. “Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors”. In: *Nature Physics* 19.8 (May 2023), pp. 1116–1121. DOI: [10.1038/s41567-023-02042-2](https://doi.org/10.1038/s41567-023-02042-2). URL: <http://dx.doi.org/10.1038/s41567-023-02042-2>.
- [32] D. W. Berry. “High-order quantum algorithm for solving linear differential equations”. In: *Journal of Physics A: Mathematical and Theoretical* 47.10 (Feb. 2014), p. 105301. DOI: [10.1088/1751-8113/47/10/105301](https://doi.org/10.1088/1751-8113/47/10/105301). URL: <http://dx.doi.org/10.1088/1751-8113/47/10/105301>.
- [33] D. W. Berry et al. “Quantum simulation of realistic materials in first quantization using non-local pseudopotentials”. In: *npj Quantum Information* 10.1 (Dec. 2024). DOI: [10.1038/s41534-024-00896-9](https://doi.org/10.1038/s41534-024-00896-9). URL: <http://dx.doi.org/10.1038/s41534-024-00896-9>.
- [34] D. W. Berry et al. “Qubitization of Arbitrary Basis Quantum Chemistry Leveraging Sparsity and Low Rank Factorization”. In: *Quantum* 3 (Dec. 2019), p. 208. DOI: [10.22331/q-2019-12-02-208](https://doi.org/10.22331/q-2019-12-02-208). URL: <https://doi.org/10.22331/q-2019-12-02-208>.
- [35] K. Bharti et al. “Noisy intermediate-scale quantum (NISQ) algorithms”. In: *Rev. Mod. Phys.* 94.1 (2022), p. 015004.
- [36] N. S. Blunt et al. “Perspective on the Current State-of-the-Art of Quantum Computing for Drug Discovery Applications”. In: *Journal of Chemical Theory and Computation* 18.12 (Nov. 2022), pp. 7001–7023. DOI: [10.1021/acs.jctc.2c00574](https://doi.org/10.1021/acs.jctc.2c00574). URL: <http://dx.doi.org/10.1021/acs.jctc.2c00574>.
- [37] A. Botea, A. Kishimoto, and R. Marinescu. “On the Complexity of Quantum Circuit Compilation”. In: *Proceedings of the International Symposium on Combinatorial Search* 9.1 (Sept. 2021), pp. 138–142. DOI: [10.1609/socs.v9i1.18463](https://doi.org/10.1609/socs.v9i1.18463). URL: <https://ojs.aaai.org/index.php/SOCS/article/view/18463>.
- [38] J. E. Bourassa et al. “Blueprint for a Scalable Photonic Fault-Tolerant Quantum Computer”. In: *Quantum* 5 (Feb. 2021), p. 392. DOI: [10.22331/q-2021-02-04-392](https://doi.org/10.22331/q-2021-02-04-392). URL: <https://doi.org/10.22331/q-2021-02-04-392>.
- [39] S. Bravyi and A. Kitaev. *Quantum codes on a lattice with boundary*. 1998. eprint: [quant-ph/9811052](https://arxiv.org/abs/quant-ph/9811052). URL: <https://arxiv.org/abs/quant-ph/9811052>.
- [40] S. Bravyi et al. “High-threshold and low-overhead fault-tolerant quantum memory”. In: *Nature* 627.8005 (Mar. 2024), pp. 778–782. DOI: [10.1038/s41586-024-07107-7](https://doi.org/10.1038/s41586-024-07107-7). URL: <https://doi.org/10.1038/s41586-024-07107-7>.
- [41] S. Bravyi et al. “The future of quantum computing with superconducting qubits”. In: *Journal of Applied Physics* 132.16 (Oct. 2022). DOI: [10.1063/5.0082975](https://doi.org/10.1063/5.0082975). URL: <http://dx.doi.org/10.1063/5.0082975>.
- [42] N. P. Breuckmann and J. Eberhardt. “Balanced Product Quantum Codes”. In: *IEEE Transactions on Information Theory* 67.10 (Oct. 2021), pp. 6653–6674.
- [43] V. von Burg et al. “Quantum computing enhanced computational catalysis”. In: *Physical Review Research* 3.3 (July 2021). DOI: [10.1103/physrevresearch.3.033055](https://doi.org/10.1103/physrevresearch.3.033055). URL: <http://dx.doi.org/10.1103/PhysRevResearch.3.033055>.
- [44] A. Butko et al. *TIGER: Topology-aware Assignment using Ising machines Application to Classical Algorithm Tasks and Quantum Circuit Gates*. 2020. arXiv: [2009.10151 \[cs.ET\]](https://arxiv.org/abs/2009.10151). URL: <https://arxiv.org/abs/2009.10151>.
- [45] C. Cade et al. “Quantifying Grover speed-ups beyond asymptotic analysis”. In: *Quantum* 7 (Oct. 2023), p. 1133. DOI: [10.22331/q-2023-10-10-1133](https://doi.org/10.22331/q-2023-10-10-1133). URL: <https://doi.org/10.22331/q-2023-10-10-1133>.
- [46] A. Caesura et al. *Faster quantum chemistry simulations on a quantum computer with improved tensor factorization and active volume compilation*. 2025. DOI: [10.48550/ARXIV.2501.06165](https://arxiv.org/abs/2501.06165). URL: <https://arxiv.org/abs/2501.06165>.

- [47] Z. Cai et al. “Quantum error mitigation”. In: *Rev. Mod. Phys.* 95 (4 Dec. 2023), p. 045005. DOI: [10.1103/RevModPhys.95.045005](https://doi.org/10.1103/RevModPhys.95.045005). URL: <https://link.aps.org/doi/10.1103/RevModPhys.95.045005>.
- [48] E. T. Campbell. “Early fault-tolerant simulations of the Hubbard model”. In: *Quantum Science and Technology* 7.1 (Nov. 2021), p. 015007. DOI: [10.1088/2058-9565/ac3110](https://doi.org/10.1088/2058-9565/ac3110). URL: <http://dx.doi.org/10.1088/2058-9565/ac3110>.
- [49] D. Camps and R. Van Beeumen. “FABLE: Fast Approximate Quantum Circuits for Block-Encodings”. In: *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2022, pp. 104–113. DOI: [10.1109/QCE53715.2022.00029](https://doi.org/10.1109/QCE53715.2022.00029).
- [50] D. Camps et al. “Evaluation of the Classical Hardware Requirements for Large-Scale Quantum Computations”. In: *ISC High Performance 2024 Research Paper Proceedings (39th International Conference)*. 2024, pp. 1–12. DOI: [10.23919/ISC.2024.10528937](https://doi.org/10.23919/ISC.2024.10528937).
- [51] D. Camps et al. “Explicit Quantum Circuits for Block Encodings of Certain Sparse Matrices”. In: *SIAM Journal on Matrix Analysis and Applications* 45.1 (2024), pp. 801–827. DOI: [10.1137/22M1484298](https://doi.org/10.1137/22M1484298).
- [52] Y. Cao et al. “Quantum Chemistry in the Age of Quantum Computing”. In: *Chem. Rev.* 119.19 (2019), pp. 10856–10915.
- [53] L. Caune et al. *Demonstrating real-time and low-latency quantum error correction with superconducting qubits*. 2024. arXiv: [2410.05202 \[quant-ph\]](https://arxiv.org/abs/2410.05202). URL: <https://arxiv.org/abs/2410.05202>.
- [54] M. Cerezo et al. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3.9 (2021), pp. 625–644.
- [55] Y. Chai et al. “Fermionic wave packet scattering: a quantum computing approach”. In: *Quantum* 9 (Feb. 2025), p. 1638. DOI: [10.22331/q-2025-02-19-1638](https://doi.org/10.22331/q-2025-02-19-1638). URL: <https://doi.org/10.22331/q-2025-02-19-1638>.
- [56] H. H. S. Chan et al. “Grid-based methods for chemistry simulations on a quantum computer”. In: *Science Advances* 9.9 (2023), eabo7484.
- [57] A. M. Childs, R. Kothari, and R. D. Somma. “Quantum Algorithm for Systems of Linear Equations with Exponentially Improved Dependence on Precision”. In: *SIAM Journal on Computing* 46.6 (2017), pp. 1920–1950. DOI: [10.1137/16M1087072](https://doi.org/10.1137/16M1087072).
- [58] A. M. Childs and J.-P. Liu. “Quantum Spectral Methods for Differential Equations”. In: *Communications in Mathematical Physics* 375.2 (Feb. 2020), pp. 1427–1457. DOI: [10.1007/s00220-020-03699-z](https://doi.org/10.1007/s00220-020-03699-z). URL: <http://dx.doi.org/10.1007/s00220-020-03699-z>.
- [59] A. M. Childs, J.-P. Liu, and A. Ostrander. “High-precision quantum algorithms for partial differential equations”. In: *Quantum* 5 (Nov. 2021), p. 574. DOI: [10.22331/q-2021-11-10-574](https://doi.org/10.22331/q-2021-11-10-574). URL: <https://doi.org/10.22331/q-2021-11-10-574>.
- [60] A. M. Childs and N. Wiebe. “Hamiltonian simulation using linear combinations of unitary operations”. In: *Quantum Information & Computation* 12.11–12 (2012), pp. 901–924. arXiv: [1202.5822](https://arxiv.org/abs/1202.5822).
- [61] A. M. Childs et al. “Toward the first quantum simulation with quantum speedup”. In: *Proceedings of the National Academy of Sciences* 115.38 (2018), pp. 9456–9461. DOI: [10.1073/pnas.1801723115](https://doi.org/10.1073/pnas.1801723115). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1801723115>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1801723115>.
- [62] A. M. Childs and E. H. Farhi. “Quantum information processing in continuous time”. AAI0806375. PhD thesis. USA, 2004.
- [63] A. N. Ciavarella, I. M. Burbano, and C. W. Bauer. *Efficient Truncations of $SU(N_c)$ Lattice Gauge Theory for Quantum Simulation*. 2025. arXiv: [2503.11888 \[hep-lat\]](https://arxiv.org/abs/2503.11888). URL: <https://arxiv.org/abs/2503.11888>.

- [64] A. J. Cohen, P. Mori-Sánchez, and W. Yang. “Insights into current limitations of density functional theory”. In: *Science* 321.5890 (2008), pp. 792–794.
- [65] O. Q. Computing. *The OQC Roadmap*. <https://oqc.tech/technical-roadmap>. [Accessed 12-06-2025]. 2025.
- [66] R. Computing. *Investor Presentation*. <https://investors.rigetti.com/static-files/fbac3801-223f-4f0f-a207-47d25084a1d7>. [Accessed 04-06-2025]. 2025.
- [67] A. Cross et al. *Improved QLDPC Surgery: Logical Measurements and Bridging Codes*. 2024. arXiv: [2407.18393](https://arxiv.org/abs/2407.18393) [quant-ph]. URL: <https://arxiv.org/abs/2407.18393>.
- [68] A. M. Dalzell et al. *Quantum algorithms: A survey of applications and end-to-end complexities*. 2023. arXiv: [2310.03011](https://arxiv.org/abs/2310.03011) [quant-ph]. URL: <https://arxiv.org/abs/2310.03011>.
- [69] E. Darve. “The Fast Multipole Method I: Error Analysis and Asymptotic Complexity”. In: *SIAM Journal on Numerical Analysis* 38.1 (2000), pp. 98–128. URL: <https://doi.org/10.1137/S0036142999330379>.
- [70] Z. Davoudi, A. F. Shaw, and J. R. Stryker. “General quantum algorithms for Hamiltonian simulation with applications to a non-Abelian lattice gauge theory”. In: *Quantum* 7 (Dec. 2023), p. 1213. DOI: [10.22331/q-2023-12-20-1213](https://doi.org/10.22331/q-2023-12-20-1213). URL: <https://doi.org/10.22331/q-2023-12-20-1213>.
- [71] A. Delgado et al. “Simulating key properties of lithium-ion batteries with a fault-tolerant quantum computer”. In: *Physical Review A* 106.3 (Sept. 2022). DOI: [10.1103/PhysRevA.106.032428](https://doi.org/10.1103/PhysRevA.106.032428). URL: <http://dx.doi.org/10.1103/PhysRevA.106.032428>.
- [72] E. Dennis et al. “Topological quantum memory”. In: *Journal of Mathematical Physics* 43.9 (Sept. 2002), pp. 4452–4505. DOI: [10.1063/1.1499754](https://doi.org/10.1063/1.1499754). URL: <http://dx.doi.org/10.1063/1.1499754>.
- [73] A. Di Meglio et al. “Quantum Computing for High-Energy Physics: State of the Art and Challenges”. In: *PRX Quantum* 5 (3 Aug. 2024), p. 037001. DOI: [10.1103/PRXQuantum.5.037001](https://doi.org/10.1103/PRXQuantum.5.037001). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.5.037001>.
- [74] Z. Ding and L. Lin. “Even Shorter Quantum Circuit for Phase Estimation on Early Fault-Tolerant Quantum Computers with Applications to Ground-State Energy Estimation”. In: *PRX Quantum* 4 (2 May 2023), p. 020331. DOI: [10.1103/PRXQuantum.4.020331](https://doi.org/10.1103/PRXQuantum.4.020331). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.4.020331>.
- [75] W. Domcke, D. R. Yarkony, and H. Köppel. *Conical Intersections: Electronic Structure, Dynamics & Spectroscopy*. World Scientific, 2004. DOI: [10.1142/5553](https://doi.org/10.1142/5553).
- [76] Y. Dong, L. Lin, and Y. Tong. “Ground-State Preparation and Energy Estimation on Early Fault-Tolerant Quantum Computers via Quantum Eigenvalue Transformation of Unitary Matrices”. In: *PRX Quantum* 3 (4 Oct. 2022), p. 040305. DOI: [10.1103/PRXQuantum.3.040305](https://doi.org/10.1103/PRXQuantum.3.040305). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.3.040305>.
- [77] U. Dorner et al. “Optimal Quantum Phase Estimation”. In: *Phys. Rev. Lett.* 102 (4 Jan. 2009), p. 040403. DOI: [10.1103/PhysRevLett.102.040403](https://doi.org/10.1103/PhysRevLett.102.040403). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.040403>.
- [78] A. Dreuw and M. Head-Gordon. “Failure of Time-Dependent Density Functional Theory for Long-Range Charge-Transfer Excited States: The Zincbacteriochlorin-Bacteriochlorin and Bacteriochlorophyll-Spheroidene Complexes”. In: *Journal of the American Chemical Society* 126.12 (2004), pp. 4007–4016. DOI: [10.1021/ja039556n](https://doi.org/10.1021/ja039556n).
- [79] R. Dutta et al. “Simulating chemistry on bosonic quantum devices”. In: *Journal of Chemical Theory and Computation* 20.15 (2024), pp. 6426–6441.
- [80] V. E. Elfving et al. *How will quantum computers provide an industrially relevant computational advantage in quantum chemistry?* 2020. arXiv: [2009.12472](https://arxiv.org/abs/2009.12472) [quant-ph]. URL: <https://arxiv.org/abs/2009.12472>.

- [81] S. Endo, S. C. Benjamin, and Y. Li. “Practical Quantum Error Mitigation for Near-Future Applications”. In: *Phys. Rev. X* 8 (3 July 2018), p. 031027. DOI: [10.1103/PhysRevX.8.031027](https://doi.org/10.1103/PhysRevX.8.031027). URL: <https://link.aps.org/doi/10.1103/PhysRevX.8.031027>.
- [82] E. Farhi, J. Goldstone, and S. Gutmann. *A Quantum Approximate Optimization Algorithm*. 2014. arXiv: [1411.4028](https://arxiv.org/abs/1411.4028) [quant-ph]. URL: <https://arxiv.org/abs/1411.4028>.
- [83] R. P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6–7 (June 1982), pp. 467–488. DOI: [10.1007/bf02650179](https://doi.org/10.1007/bf02650179). URL: <http://dx.doi.org/10.1007/BF02650179>.
- [84] S. Flannigan et al. “Propagation of errors and quantitative quantum simulation with quantum advantage”. In: *Quantum Science and Technology* 7.4 (Aug. 2022), p. 045025. DOI: [10.1088/2058-9565/ac88f5](https://doi.org/10.1088/2058-9565/ac88f5). URL: <https://dx.doi.org/10.1088/2058-9565/ac88f5>.
- [85] A. G. Fowler and C. Gidney. *Low overhead quantum computation using lattice surgery*. 2018. DOI: [10.48550/ARXIV.1808.06709](https://arxiv.org/abs/1808.06709). URL: <https://arxiv.org/abs/1808.06709>.
- [86] R. J. Garcia et al. “Barren plateaus from learning scramblers with local cost functions”. In: *Journal of High Energy Physics* 2023.1 (Jan. 2023), p. 90. DOI: [10.1007/JHEP01\(2023\)090](https://doi.org/10.1007/JHEP01(2023)090). URL: [https://doi.org/10.1007/JHEP01\(2023\)090](https://doi.org/10.1007/JHEP01(2023)090).
- [87] J. Garcke and M. Griebel. *Sparse Grids and Applications*. Springer Publishing Company, Incorporated, 2012.
- [88] A. Georges et al. “Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions”. In: *Rev. Mod. Phys.* 68 (1996), pp. 13–125.
- [89] C. Gidney, N. Shutty, and C. Jones. *Magic state cultivation: growing T states as cheap as CNOT gates*. Sept. 2024. arXiv: [2409.17595](https://arxiv.org/abs/2409.17595) [quant-ph].
- [90] M. B. Giles. *Multilevel Monte Carlo methods*. 2013. arXiv: [1304.5472](https://arxiv.org/abs/1304.5472) [math.NA]. URL: <https://arxiv.org/abs/1304.5472>.
- [91] A. Gilyén et al. “Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 193–204. DOI: [10.1145/3313276.3316366](https://doi.org/10.1145/3313276.3316366). URL: <https://doi.org/10.1145/3313276.3316366>.
- [92] *GitHub - CQCL/quantinuum-hardware-quantum-volume: Repository for sharing Quantinuum’s Quantum Volume data — github.com*. <https://github.com/CQCL/quantinuum-hardware-quantum-volume>. [Accessed 08-07-2025].
- [93] J. J. Goings et al. “Reliably assessing the electronic structure of cytochrome P450 on today’s classical computers and tomorrow’s quantum computers”. In: *Proceedings of the National Academy of Sciences* 119.38 (2022), e2203533119. DOI: [10.1073/pnas.2203533119](https://doi.org/10.1073/pnas.2203533119). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2203533119>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2203533119>.
- [94] J. J. Goings et al. “Reliably determining reaction mechanisms on quantum computers”. In: *J. Phys. Chem. Lett.* 13.21 (2022), pp. 4798–4809.
- [95] Google Quantum AI. *Roadmap — Google Quantum AI — quantumai.google*. <https://quantumai.google/roadmap>. [Accessed 03-03-2025].
- [96] H. Goto. “High-performance fault-tolerant quantum computing with many-hypercube codes”. In: *Sci. Adv.* 10.36 (2024), adp6388. DOI: [10.1126/sciadv.adp6388](https://doi.org/10.1126/sciadv.adp6388). arXiv: [2403.16054](https://arxiv.org/abs/2403.16054) [quant-ph].
- [97] D. Gottesman. “Fault-tolerant quantum computation with constant overhead”. In: *Quantum Info. Comput.* 14.15–16 (Nov. 2014), pp. 1338–1372.

- [98] V. A. Goy et al. “Sign problem in finite density lattice QCD”. In: *Progress of Theoretical and Experimental Physics* 2017.3 (Mar. 2017), p. 031D01. DOI: [10.1093/ptep/ptx018](https://academic.oup.com/ptep/article-pdf/2017/3/031D01/11150731/ptx018.pdf). eprint: <https://academic.oup.com/ptep/article-pdf/2017/3/031D01/11150731/ptx018.pdf>. URL: <https://doi.org/10.1093/ptep/ptx018>.
- [99] T. M. Graham, E. Oh, and M. Saffman. “Multiscale architecture for fast optical addressing and control of large-scale qubit arrays”. In: *Appl. Opt.* 62.12 (Apr. 2023), pp. 3242–3251. DOI: [10.1364/AO.484367](https://opg.optica.org/ao/abstract.cfm?URI=ao-62-12-3242). URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-62-12-3242>.
- [100] L. K. Grover. “A fast quantum mechanical algorithm for database search”. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC ’96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 212–219. DOI: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866). URL: <https://doi.org/10.1145/237814.237866>.
- [101] J. Haah et al. “Quantum Algorithm for Simulating Real Time Evolution of Lattice Hamiltonians”. In: *SIAM Journal on Computing* 52.6 (2023), FOCS18-250-FOCS18–284. DOI: [10.1137/18M1231511](https://doi.org/10.1137/18M1231511). eprint: <https://doi.org/10.1137/18M1231511>. URL: <https://doi.org/10.1137/18M1231511>.
- [102] B. Hall et al. “Simulation of collective neutrino oscillations on a quantum computer”. In: *Phys. Rev. D* 104 (6 Sept. 2021), p. 063009. DOI: [10.1103/PhysRevD.104.063009](https://link.aps.org/doi/10.1103/PhysRevD.104.063009). URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.063009>.
- [103] R. Harper and S. T. Flammia. “Fast decoding of the surface code with machine learning”. In: *Physical Review Letters* 127.13 (2021), p. 130504.
- [104] A. W. Harrow, A. Hassidim, and S. Lloyd. “Quantum Algorithm for Linear Systems of Equations”. In: *Phys. Rev. Lett.* 103 (15 Oct. 2009), p. 150502. DOI: [10.1103/PhysRevLett.103.150502](https://link.aps.org/doi/10.1103/PhysRevLett.103.150502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.103.150502>.
- [105] S. Hazra et al. “Engineering cross resonance interaction in multi-modal quantum circuits”. In: *Applied Physics Letters* 116.15 (Apr. 2020), p. 152601. DOI: [10.1063/1.5143440](https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.5143440/13021480/152601_1_online.pdf). eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.5143440/13021480/152601_1_online.pdf. URL: <https://doi.org/10.1063/1.5143440>.
- [106] K. Heya et al. *Randomized benchmarking of a high-fidelity remote CNOT gate over a meter-scale microwave interconnect*. 2025. arXiv: [2502.15034](https://arxiv.org/abs/2502.15034) [quant-ph]. URL: <https://arxiv.org/abs/2502.15034>.
- [107] B. L. Higgins et al. “Entanglement-free Heisenberg-limited phase estimation”. In: *Nature* 450.7168 (Nov. 2007), pp. 393–396. DOI: [10.1038/nature06257](http://dx.doi.org/10.1038/nature06257). URL: <http://dx.doi.org/10.1038/nature06257>.
- [108] O. Higgott and N. P. Breuckmann. “Constructions and Performance of Hyperbolic and Semi-Hyperbolic Floquet Codes”. In: *PRX Quantum* 5 (4 Nov. 2024), p. 040327. DOI: [10.1103/PRXQuantum.5.040327](https://link.aps.org/doi/10.1103/PRXQuantum.5.040327). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.5.040327>.
- [109] O. Higgott and N. P. Breuckmann. “Subsystem Codes with High Thresholds by Gauge Fixing and Reduced Qubit Overhead”. In: *Phys. Rev. X* 11 (3 Aug. 2021), p. 031039. DOI: [10.1103/PhysRevX.11.031039](https://link.aps.org/doi/10.1103/PhysRevX.11.031039). URL: <https://link.aps.org/doi/10.1103/PhysRevX.11.031039>.
- [110] IBM Quantum. *IBM Quantum Development & Innovation Roadmap 2025*. <https://www.ibm.com/downloads/documents/us-en/131cf87ab63319bf>. [Accessed 06-12-2025].
- [111] IBM Quantum. *IBM Quantum Roadmap*. <https://www.ibm.com/roadmaps/quantum/>. [Accessed 06-12-2025].
- [112] Inflection. *Inflection Unveils 5-year Quantum Computing Roadmap, Advancing Plans to Commercialize Quantum*. <https://www.youtube.com/watch?v=23X6-AQ9Cnc>. [Accessed 03-03-2025].
- [113] *Integrated Research Infrastructure Architecture Blueprint Activity*. 2023. URL: <https://www.osti.gov/servlets/purl/1984466/>.

- [114] IonQ. *IonQ's Accelerated Roadmap: Turning Quantum Ambition into Reality*. <https://ionq.com/blog/ionqs-accelerated-roadmap-turning-quantum-ambition-into-reality>. [Accessed 21-07-2025]. June 2025.
- [115] IQM. *IQM Roadmap*. <https://www.meetiqm.com/technology/roadmap>. [Accessed 03-03-2025].
- [116] S. Jin, N. Liu, and Y. Yu. “Quantum Simulation of Partial Differential Equations via Schrödingerization”. In: *Phys. Rev. Lett.* 133 (23 Dec. 2024), p. 230602. DOI: [10.1103/PhysRevLett.133.230602](https://doi.org/10.1103/PhysRevLett.133.230602). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.133.230602>.
- [117] C. Jones. “Low-overhead constructions for the fault-tolerant Toffoli gate”. In: *Phys. Rev. A* 87 (2 Feb. 2013), p. 022328. DOI: [10.1103/PhysRevA.87.022328](https://doi.org/10.1103/PhysRevA.87.022328). URL: <https://link.aps.org/doi/10.1103/PhysRevA.87.022328>.
- [118] S. P. Jordan. *Quantum Algorithm Zoo*. <https://quantumalgorithmzoo.org>.
- [119] S. P. Jordan, K. S. M. Lee, and J. Preskill. “Quantum Algorithms for Quantum Field Theories”. In: *Science* 336.6085 (2012), pp. 1130–1133. DOI: [10.1126/science.1217069](https://doi.org/10.1126/science.1217069). eprint: <https://www.science.org/doi/pdf/10.1126/science.1217069>. URL: <https://www.science.org/doi/abs/10.1126/science.1217069>.
- [120] A. Kan and Y. Nam. *Lattice Quantum Chromodynamics and Electrodynamics on a Universal Quantum Computer*. 2022. arXiv: [2107.12769](https://arxiv.org/abs/2107.12769) [quant-ph]. URL: <https://arxiv.org/abs/2107.12769>.
- [121] A. Kandala et al. “Demonstration of a High-Fidelity cnot Gate for Fixed-Frequency Transmons with Engineered ZZ Suppression”. In: *Phys. Rev. Lett.* 127 (13 Sept. 2021), p. 130501. DOI: [10.1103/PhysRevLett.127.130501](https://doi.org/10.1103/PhysRevLett.127.130501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.130501>.
- [122] D. Keyes. *Building a Science-Based Case for Large-Scale Simulation*. 2003. URL: <https://wgropp.cs.illinois.edu/bib/reports/scales2-2005.pdf>.
- [123] R. King et al. “Quantum simulation with sum-of-squares spectral amplification”. In: *arXiv preprint arXiv:2505.01528* (2025).
- [124] A. Y. Kitaev. *Quantum measurements and the Abelian Stabilizer Problem*. 1995. arXiv: [quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026) [quant-ph]. URL: <https://arxiv.org/abs/quant-ph/9511026>.
- [125] A. Kitaev. “Fault-tolerant quantum computation by anyons”. In: *Annals of Physics* 303.1 (Jan. 2003), pp. 2–30. DOI: [10.1016/S0003-4916\(02\)00018-0](https://doi.org/10.1016/S0003-4916(02)00018-0). URL: [http://dx.doi.org/10.1016/S0003-4916\(02\)00018-0](http://dx.doi.org/10.1016/S0003-4916(02)00018-0).
- [126] I. D. Kivlichan et al. “Improved Fault-Tolerant Quantum Simulation of Condensed-Phase Correlated Electrons via Trotterization”. In: *Quantum* 4 (July 2020), p. 296. DOI: [10.22331/q-2020-07-16-296](https://doi.org/10.22331/q-2020-07-16-296). URL: <https://doi.org/10.22331/q-2020-07-16-296>.
- [127] J. Klimeš, D. R. Bowler, and A. Michaelides. “Chemical accuracy for the van der Waals density functional”. In: *Journal of Physics: Condensed Matter* 22.2 (2009), p. 022201.
- [128] W. Kohn and L. J. Sham. “Self-consistent equations including exchange and correlation effects”. In: *Physical review* 140.4A (1965), A1133.
- [129] G. Kotliar et al. “Electronic structure calculations with dynamical mean-field theory”. In: *Rev. Mod. Phys.* 78 (2006), pp. 865–951.
- [130] W. Kramer, J. Shalf, and E. Strohmaier. *The NERSC Sustained System Performance (SSP) Metric*. Tech. rep. LBNL-57155. Lawrence Berkeley National Laboratory, 2005. URL: <https://escholarship.org/uc/item/4f5621q9>.
- [131] I. Kriekouki et al. “Silicon spin qubits: A scalable solution for quantum computing”. In: *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*. 2025, pp. 1–5. DOI: [10.1109/ISCAS56072.2025.11043641](https://doi.org/10.1109/ISCAS56072.2025.11043641).

- [132] S. Krinner et al. “Realizing repeated quantum error correction in a distance-three surface code”. In: *Nature* 605.7911 (May 2022), pp. 669–674. DOI: [10.1038/s41586-022-04566-8](https://doi.org/10.1038/s41586-022-04566-8). URL: <https://doi.org/10.1038/s41586-022-04566-8>.
- [133] H. Krovi. “Improved quantum algorithms for linear and nonlinear differential equations”. In: *Quantum* 7 (Feb. 2023), p. 913. DOI: [10.22331/q-2023-02-02-913](https://doi.org/10.22331/q-2023-02-02-913). URL: <https://doi.org/10.22331/q-2023-02-02-913>.
- [134] O. Kyriienko, A. E. Paine, and V. E. Elfving. “Solving nonlinear differential equations with differentiable quantum circuits”. In: *Phys. Rev. A* 103 (5 May 2021), p. 052416. DOI: [10.1103/PhysRevA.103.052416](https://link.aps.org/doi/10.1103/PhysRevA.103.052416). URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.052416>.
- [135] M. V. Larsen et al. “Integrated photonic source of Gottesman–Kitaev–Preskill qubits”. In: *Nature* 642.8068 (June 2025), pp. 587–591. DOI: [10.1038/s41586-025-09044-5](https://doi.org/10.1038/s41586-025-09044-5). URL: <http://dx.doi.org/10.1038/s41586-025-09044-5>.
- [136] J. Lee et al. “Fault-tolerant quantum computation by hybrid qubits with bosonic cat code and single photons”. In: *PRX Quantum* 5.3 (2024), p. 030322.
- [137] J. Lee et al. “Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction”. In: *PRX Quantum* 2 (3 July 2021), p. 030305. DOI: [10.1103/PRXQuantum.2.030305](https://link.aps.org/doi/10.1103/PRXQuantum.2.030305). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.030305>.
- [138] A. Leverrier, J.-P. Tillich, and G. Zémor. “Quantum Expander Codes”. In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE, 2015, pp. 810–824.
- [139] A. Leverrier and G. Zémor. “Quantum Tanner codes”. In: Feb. 2022. DOI: [10.1109/FOCS54457.2022.00117](https://arxiv.org/abs/2202.13641). arXiv: [2202.13641 \[quant-ph\]](https://arxiv.org/abs/2202.13641).
- [140] S. K. Leyton and T. J. Osborne. *A quantum algorithm to solve nonlinear differential equations*. 2008. arXiv: [0812.4423 \[quant-ph\]](https://arxiv.org/abs/0812.4423). URL: <https://arxiv.org/abs/0812.4423>.
- [141] Z. Li and G. K.-L. Chan. “Spin-Projected Matrix Product States: Versatile Tool for Strongly Correlated Systems”. In: *Journal of Chemical Theory and Computation* 13.6 (May 2017), pp. 2681–2695. DOI: [10.1021/acs.jctc.7b00270](https://doi.org/10.1021/acs.jctc.7b00270). URL: <http://dx.doi.org/10.1021/acs.jctc.7b00270>.
- [142] Z. Li et al. “The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations”. In: *The Journal of Chemical Physics* 150.2 (Jan. 2019), p. 024302. DOI: [10.1063/1.5063376](https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.5063376/14715992/024302_1_online.pdf). eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.5063376/14715992/024302_1_online.pdf. URL: <https://doi.org/10.1063/1.5063376>.
- [143] D. A. Lidar and T. A. Brun, eds. *Quantum Error Correction*. Cambridge University Press, 2013. DOI: [10.1017/CBO9781139034807](https://doi.org/10.1017/CBO9781139034807).
- [144] L. Lin and Y. Tong. “Heisenberg-Limited Ground-State Energy Estimation for Early Fault-Tolerant Quantum Computers”. In: *PRX Quantum* 3 (1 Feb. 2022), p. 010318. DOI: [10.1103/PRXQuantum.3.010318](https://link.aps.org/doi/10.1103/PRXQuantum.3.010318). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.3.010318>.
- [145] D. Litinski. “A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery”. In: *Quantum* 3 (Mar. 2019), p. 128. DOI: [10.22331/q-2019-03-05-128](https://doi.org/10.22331/q-2019-03-05-128). URL: <https://doi.org/10.22331/q-2019-03-05-128>.
- [146] D. Litinski. “Magic State Distillation: Not as Costly as You Think”. In: *Quantum* 3 (Dec. 2019), p. 205. DOI: [10.22331/q-2019-12-02-205](https://doi.org/10.22331/q-2019-12-02-205). URL: <https://doi.org/10.22331/q-2019-12-02-205>.
- [147] D. Litinski and N. Nickerson. *Active volume: An architecture for efficient fault-tolerant quantum computers with limited non-local connections*. 2022. DOI: [10.48550/ARXIV.2211.15465](https://arxiv.org/abs/2211.15465). URL: <https://arxiv.org/abs/2211.15465>.

- [148] J.-P. Liu et al. “Efficient quantum algorithm for dissipative nonlinear differential equations”. In: *Proceedings of the National Academy of Sciences* 118.35 (2021), e2026805118. DOI: [10.1073/pnas.2026805118](https://doi.org/10.1073/pnas.2026805118). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2026805118>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2026805118>.
- [149] Y. Liu et al. “Hybrid oscillator-qubit quantum processors: Instruction set architectures, abstract machine models, and applications”. In: *arXiv preprint arXiv:2407.10381* (2024).
- [150] Y. Liu, S. Arunachalam, and K. Temme. “A rigorous and robust quantum speed-up in supervised machine learning”. In: *Nature Physics* 17.9 (2021), pp. 1013–1017.
- [151] S. Lloyd. “Universal Quantum Simulators”. In: *Science* 273.5278 (Aug. 1996), pp. 1073–1078. DOI: [10.1126/science.273.5278.1073](https://doi.org/10.1126/science.273.5278.1073). URL: <http://dx.doi.org/10.1126/science.273.5278.1073>.
- [152] S. Lloyd et al. *Quantum algorithm for nonlinear differential equations*. 2020. arXiv: [2011.06571](https://arxiv.org/abs/2011.06571) [quant-ph]. URL: <https://arxiv.org/abs/2011.06571>.
- [153] P. Lougovski et al. *Report for the ASCR Workshop on Basic Research Needs in Quantum Computing and Networking - 2023*. Tech. rep. Sandia National Lab. (SNL-NM), Albuquerque, NM (United States), Aug. 2024. DOI: [10.2172/2430035](https://doi.org/10.2172/2430035). URL: <https://www.osti.gov/biblio/2430035>.
- [154] G. H. Low and I. L. Chuang. “Optimal Hamiltonian simulation by quantum signal processing”. In: *Physical review letters* 118.1 (2017), p. 010501.
- [155] G. H. Low and I. L. Chuang. “Hamiltonian simulation by qubitization”. In: *Quantum* 3 (2019), p. 163. DOI: [10.22331/q-2019-07-12-163](https://doi.org/10.22331/q-2019-07-12-163).
- [156] G. H. Low et al. *Fast quantum simulation of electronic structure by spectrum amplification*. 2025. DOI: [10.48550/ARXIV.2502.15882](https://doi.org/10.48550/ARXIV.2502.15882). URL: <https://arxiv.org/abs/2502.15882>.
- [157] M. Lubasch et al. *Quantum circuits for partial differential equations in Fourier space*. 2025. arXiv: [2505.16895](https://arxiv.org/abs/2505.16895) [quant-ph]. URL: <https://arxiv.org/abs/2505.16895>.
- [158] M. Lubasch et al. “Variational quantum algorithms for nonlinear problems”. In: *Phys. Rev. A* 101 (1 Jan. 2020), p. 010301. DOI: [10.1103/PhysRevA.101.010301](https://doi.org/10.1103/PhysRevA.101.010301). URL: <https://link.aps.org/doi/10.1103/PhysRevA.101.010301>.
- [159] M. W. Mahoney. “Randomized Algorithms for Matrices and Data”. In: *Found. Trends Mach. Learn.* 3.2 (Feb. 2011), pp. 123–224. DOI: [10.1561/22000000035](https://doi.org/10.1561/22000000035). URL: <https://doi.org/10.1561/22000000035>.
- [160] S. Malpathak, S. D. Kallullathil, and A. F. Izmaylov. “Simulating Vibrational Dynamics on Bosonic Quantum Devices”. In: *The Journal of Physical Chemistry Letters* 16.8 (2025), pp. 1855–1864.
- [161] J. F. Marques et al. “Logical-qubit operations in an error-detecting surface code”. In: *Nature Physics* 18.1 (Jan. 2022), pp. 80–86. DOI: [10.1038/s41567-021-01423-9](https://doi.org/10.1038/s41567-021-01423-9). URL: <https://doi.org/10.1038/s41567-021-01423-9>.
- [162] J. M. Martyn et al. “Grand Unification of Quantum Algorithms”. In: *PRX Quantum* 2 (4 Dec. 2021), p. 040203. DOI: [10.1103/PRXQuantum.2.040203](https://doi.org/10.1103/PRXQuantum.2.040203). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.040203>.
- [163] D. Marx and J. Hutter. *Ab initio molecular dynamics: basic theory and advanced methods*. Cambridge University Press, 2009.
- [164] S. McArdle et al. “Quantum computational chemistry”. In: *Rev. Mod. Phys.* 92.1 (2020), p. 015003.
- [165] J. R. McClean et al. “Barren plateaus in quantum neural network training landscapes”. In: *Nature Communications* 9.1 (Nov. 2018), p. 4812. DOI: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4). URL: <https://doi.org/10.1038/s41467-018-07090-4>.

- [166] A. Michel et al. “Blueprint for a digital-analog variational quantum eigensolver using Rydberg atom arrays”. In: *Phys. Rev. A* 107 (4 Apr. 2023), p. 042602. DOI: [10.1103/PhysRevA.107.042602](https://doi.org/10.1103/PhysRevA.107.042602). URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.042602>.
- [167] *Microsoft Quantum — Microsoft Quantum Hardware — quantum.microsoft.com*. <https://quantum.microsoft.com/en-us/solutions/microsoft-quantum-hardware>. [Accessed 17-06-2025].
- [168] M. Mirrahimi. “Cat-qubits for quantum computation”. In: *Comptes Rendus Physique* 17.7 (2016). Quantum microwaves / Micro-ondes quantiques, pp. 778–787. DOI: <https://doi.org/10.1016/j.crhy.2016.07.011>. URL: <https://www.sciencedirect.com/science/article/pii/S1631070516300627>.
- [169] P. Mori-Sánchez, A. J. Cohen, and W. Yang. “Localization and Delocalization Errors in Density Functional Theory and Implications for Band-Gap Prediction”. In: *Physical review letters* 100.14 (2008), p. 146401.
- [170] A. Morvan et al. “Phase transitions in random circuit sampling”. In: *Nature* 634.8033 (Oct. 2024), pp. 328–333. DOI: [10.1038/s41586-024-07998-6](https://doi.org/10.1038/s41586-024-07998-6). URL: <http://dx.doi.org/10.1038/s41586-024-07998-6>.
- [171] I. N. Moskalenko et al. “High fidelity two-qubit gates on fluxoniums using a tunable coupler”. In: *npj Quantum Information* 8.1 (Nov. 2022). DOI: [10.1038/s41534-022-00644-x](https://doi.org/10.1038/s41534-022-00644-x). URL: <http://dx.doi.org/10.1038/s41534-022-00644-x>.
- [172] J. A. Muniz et al. “High-Fidelity Universal Gates in the ^{171}Yb Ground-State Nuclear-Spin Qubit”. In: *PRX Quantum* 6.2 (May 2025). DOI: [10.1103/prxquantum.6.020334](https://doi.org/10.1103/prxquantum.6.020334). URL: <http://dx.doi.org/10.1103/PRXQuantum.6.020334>.
- [173] H. Munoz-Bauza and D. Lidar. “Scaling Advantage in Approximate Optimization with Quantum Annealing”. In: *Phys. Rev. Lett.* 134 (16 Apr. 2025), p. 160601. DOI: [10.1103/PhysRevLett.134.160601](https://doi.org/10.1103/PhysRevLett.134.160601). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.134.160601>.
- [174] K. Nagata. “Finite-density lattice QCD and sign problem: Current status and open problems”. In: *Progress in Particle and Nuclear Physics* 127 (2022), p. 103991. DOI: <https://doi.org/10.1016/j.pnpnp.2022.103991>. URL: <https://www.sciencedirect.com/science/article/pii/S0146641022000497>.
- [175] K. M. Nakanishi, K. Mitarai, and K. Fujii. “Subspace-search variational quantum eigensolver for excited states”. In: *Physical Review Research* 1.3 (2019), p. 033062.
- [176] National Energy Research Scientific Computing Center. *Strategic Plan FY2024 – 2034*. <https://www.nersc.gov/assets/Annual-Reports/2024-2034-NERSC-10-yr-Strategic-Plan-v2.pdf>. [Accessed 03-03-2025]. 2024.
- [177] NERSC. *National Energy Research Scientific Computing Center*. <https://www.nersc.gov>. [Accessed 03-03-2025].
- [178] NERSC. *NERSC Science Acceleration Program*. <https://www.nersc.gov/what-we-do/support-for-scientists/nersc-science-acceleration-program>. [Accessed 07-01-2025].
- [179] NERSC. *Perlmutter Architecture*. <https://docs.nersc.gov/systems/perlmutter/architecture/>. [Accessed 03-03-2025].
- [180] NERSC. *Perlmutter Architecture*. <https://www.nersc.gov/what-we-do/computing-for-science/doudna-system>. [Accessed 07-01-2025].
- [181] H. Ni, H. Li, and L. Ying. “On low-depth algorithms for quantum phase estimation”. In: *Quantum* 7 (Nov. 2023), p. 1165. DOI: [10.22331/q-2023-11-06-1165](https://doi.org/10.22331/q-2023-11-06-1165). URL: <https://doi.org/10.22331/q-2023-11-06-1165>.
- [182] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. 10th Anniversary. Cambridge University Press, 2010.

- [183] M. Y. Niu et al. “Universal quantum control through deep reinforcement learning”. In: *npj Quantum Information* 5.1 (2019), p. 33.
- [184] T. E. O’Brien et al. “Efficient quantum computation of molecular forces and other energy gradients”. In: *Physical Review Research* 4.4 (2022), p. 043210.
- [185] T. E. O’Brien et al. “Efficient quantum computation of molecular forces and other energy gradients”. In: *Phys. Rev. Res.* 4 (4 Dec. 2022), p. 043210. DOI: [10.1103/PhysRevResearch.4.043210](https://doi.org/10.1103/PhysRevResearch.4.043210). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.4.043210>.
- [186] P. J. Ollitrault, A. Miessen, and I. Tavernelli. “Molecular quantum dynamics: A quantum computing perspective”. In: *Accounts of Chemical Research* 54.23 (2021), pp. 4229–4238.
- [187] P. J. Ollitrault et al. “Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor”. In: *Physical Review Research* 2.4 (2020), p. 043140.
- [188] D. Orsucci and V. Dunjko. “On solving classes of positive-definite quantum linear systems with quadratically improved runtime in the condition number”. In: *Quantum* 5 (Nov. 2021), p. 573. DOI: [10.22331/q-2021-11-08-573](https://doi.org/10.22331/q-2021-11-08-573). URL: <https://doi.org/10.22331/q-2021-11-08-573>.
- [189] P. Pantelev and G. Kalachev. “Asymptotically good Quantum and locally testable classical LDPC codes”. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2022. Rome, Italy: Association for Computing Machinery, 2022, pp. 375–388. DOI: [10.1145/3519935.3520017](https://doi.org/10.1145/3519935.3520017). URL: <https://doi.org/10.1145/3519935.3520017>.
- [190] P. Pantelev and G. Kalachev. “Quantum LDPC Codes With Almost Linear Minimum Distance”. In: *IEEE Transactions on Information Theory* 68.1 (2022), pp. 213–229. DOI: [10.1109/TIT.2021.3119384](https://doi.org/10.1109/TIT.2021.3119384).
- [191] R. G. Parr. “Density functional theory of atoms and molecules”. In: *Horizons of Quantum Chemistry: Proceedings of the Third International Congress of Quantum Chemistry Held at Kyoto, Japan, October 29-November 3, 1979*. Springer. 1989, pp. 5–15.
- [192] Pasqal. *Inflection Unveils 5-year Quantum Computing Roadmap, Advancing Plans to Commercialize Quantum*. <https://www.pasqal.com/technology/roadmap/>. [Accessed 03-03-2025].
- [193] J. P. Perdew and K. Schmidt. “Jacob’s ladder of density functional approximations for the exchange-correlation energy”. In: *AIP conference proceedings*. Vol. 577. 1. American Institute of Physics. 2001, pp. 1–20.
- [194] M. A. Perlin. *qLDPC*. <https://github.com/Inflection/qLDPC>. 2023.
- [195] A. Peruzzo et al. “A variational eigenvalue solver on a photonic quantum processor”. In: *Nature Communications* 5.1 (July 2014). DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213). URL: <http://dx.doi.org/10.1038/ncomms5213>.
- [196] G. Pichard et al. “Rearrangement of individual atoms in a 2000-site optical-tweezer array at cryogenic temperatures”. In: *Phys. Rev. Appl.* 22 (2 Aug. 2024), p. 024073. DOI: [10.1103/PhysRevApplied.22.024073](https://doi.org/10.1103/PhysRevApplied.22.024073). URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.22.024073>.
- [197] A. J. Pool et al. “Nonlinear dynamics as a ground-state solution on quantum computers”. In: *Phys. Rev. Res.* 6 (3 Sept. 2024), p. 033257. DOI: [10.1103/PhysRevResearch.6.033257](https://doi.org/10.1103/PhysRevResearch.6.033257). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.6.033257>.
- [198] J. Preskill. “Beyond NISQ: The Megaquop Machine”. In: *ACM Transactions on Quantum Computing* 6.3 (Apr. 2025), pp. 1–7. DOI: [10.1145/3723153](https://doi.org/10.1145/3723153). URL: <http://dx.doi.org/10.1145/3723153>.
- [199] J. Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (Aug. 2018), p. 79. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79). URL: <https://doi.org/10.22331/q-2018-08-06-79>.
- [200] T. Proctor et al. “Benchmarking quantum computers”. In: *Nature Reviews Physics* 7.2 (2025), pp. 105–118.

- [201] *QBI: Quantum Benchmarking Initiative*. 2025. URL: <https://www.darpa.mil/research/programs/quantum-benchmarking-initiative>.
- [202] Quantinuum. *Quantinuum Unveils Accelerated Roadmap to Achieve Universal, Fully Fault-Tolerant Quantum Computing by 2030*. <https://www.quantinuum.com/press-releases/quantinuum-unveils-accelerated-roadmap-to-achieve-universal-fault-tolerant-quantum-computing-by-2030>. [Accessed 03-03-2025].
- [203] Quantum Circuits Inc. *Technology* — *quantumcircuits.com*. <https://quantumcircuits.com/technology/>. [Accessed 17-07-2025].
- [204] QuEra Computing. *Error-Corrected Computers*. <https://www.quera.com/qec>. [Accessed 03-03-2025].
- [205] A. G. Radnaev et al. *A universal neutral-atom quantum computer with individual optical addressing and non-destructive readout*. 2025. arXiv: [2408.08288](https://arxiv.org/abs/2408.08288) [quant-ph]. URL: <https://arxiv.org/abs/2408.08288>.
- [206] J. Rahamim et al. “Double-sided coaxial circuit QED with out-of-plane wiring”. In: *Applied Physics Letters* 110.22 (May 2017), p. 222602. DOI: [10.1063/1.4984299](https://doi.org/10.1063/1.4984299). eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.4984299/14500580/222602_1_online.pdf. URL: <https://doi.org/10.1063/1.4984299>.
- [207] I. Raychowdhury and J. R. Stryker. “Loop, string, and hadron dynamics in SU(2) Hamiltonian lattice gauge theories”. In: *Phys. Rev. D* 101 (11 June 2020), p. 114502. DOI: [10.1103/PhysRevD.101.114502](https://doi.org/10.1103/PhysRevD.101.114502). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.114502>.
- [208] M. Reiher et al. “Elucidating reaction mechanisms on quantum computers”. In: *Proc. Natl. Acad. Sci. USA* 114.29 (2017), pp. 7555–7560.
- [209] M. Reiher et al. “Elucidating reaction mechanisms on quantum computers”. In: *Proceedings of the National Academy of Sciences* 114.29 (July 2017), pp. 7555–7560. DOI: [10.1073/pnas.1619152114](https://doi.org/10.1073/pnas.1619152114). URL: <http://dx.doi.org/10.1073/pnas.1619152114>.
- [210] I. Reppel. *Quantum computer roadmaps for various qubit modalities and vendors, including references to official sources* — *ianreppel.org*. <https://ianreppel.org/quantum.html>. [Accessed 05-06-2025].
- [211] M. L. Rhodes, M. Kreshchuk, and S. Pathak. “Exponential Improvements in the Simulation of Lattice Gauge Theories Using Near-Optimal Techniques”. In: *PRX Quantum* 5.4 (Dec. 2024). DOI: [10.1103/prxquantum.5.040347](https://doi.org/10.1103/prxquantum.5.040347). URL: <http://dx.doi.org/10.1103/PRXQuantum.5.040347>.
- [212] D. Rocca et al. “Reducing the Runtime of Fault-Tolerant Quantum Simulations in Chemistry through Symmetry-Compressed Double Factorization”. In: *Journal of Chemical Theory and Computation* 20.11 (May 2024), pp. 4639–4653. DOI: [10.1021/acs.jctc.4c00352](https://doi.org/10.1021/acs.jctc.4c00352). URL: <http://dx.doi.org/10.1021/acs.jctc.4c00352>.
- [213] N. C. Rubin et al. “Quantum computation of stopping power for inertial fusion target design”. In: *Proceedings of the National Academy of Sciences* 121.23 (2024), e2317772121.
- [214] T. Rudolph. *What is the logical gate speed of a photonic quantum computer?* June 2023. URL: <https://quantumfrontiers.com/2023/06/21/what-is-the-logical-gate-speed-of-a-photonic-quantum-computer/>.
- [215] D. Ruiz et al. “LDPC-cat codes for low-overhead quantum computing in 2D”. In: *Nature Communications* 16.1 (Jan. 2025), p. 1040. DOI: [10.1038/s41467-025-56298-8](https://doi.org/10.1038/s41467-025-56298-8). URL: <https://doi.org/10.1038/s41467-025-56298-8>.
- [216] K. Sakamoto et al. “End-to-end complexity for simulating the Schwinger model on quantum computers”. In: *Quantum* 8 (Sept. 2024), p. 1474. DOI: [10.22331/q-2024-09-17-1474](https://doi.org/10.22331/q-2024-09-17-1474). URL: <http://dx.doi.org/10.22331/q-2024-09-17-1474>.

- [217] V. M. Schäfer et al. “Fast quantum logic gates with trapped-ion qubits”. In: *Nature* 555.7694 (Mar. 2018), pp. 75–78. DOI: [10.1038/nature25737](https://doi.org/10.1038/nature25737). URL: <http://dx.doi.org/10.1038/nature25737>.
- [218] A. Schmidhuber et al. “Quartic Quantum Speedups for Planted Inference”. In: *Phys. Rev. X* 15 (2 June 2025), p. 021077. DOI: [10.1103/PhysRevX.15.021077](https://doi.org/10.1103/PhysRevX.15.021077). URL: <https://link.aps.org/doi/10.1103/PhysRevX.15.021077>.
- [219] M. Schuld and N. Killoran. “Is quantum advantage the right question?” In: *PRX Quantum* 3.3 (2022), p. 030101.
- [220] P. Selinger. “Efficient Clifford+T approximation of single-qubit operators”. In: *Quantum Info. Comput.* 15.1–2 (Jan. 2015), pp. 159–180.
- [221] Y. Shen et al. “Efficient Measurement-Driven Eigenenergy Estimation with Classical Shadows”. In: *arXiv preprint arXiv:2409.13691* (2024).
- [222] M. Shokrian Zini et al. “Quantum simulation of battery materials using ionic pseudopotentials”. In: *Quantum* 7 (July 2023), p. 1049. DOI: [10.22331/q-2023-07-10-1049](https://doi.org/10.22331/q-2023-07-10-1049). URL: <https://doi.org/10.22331/q-2023-07-10-1049>.
- [223] P. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 124–134. DOI: [10.1109/SFCS.1994.365700](https://doi.org/10.1109/SFCS.1994.365700).
- [224] P. Shor. “Fault-tolerant quantum computation”. In: *Proceedings of 37th Conference on Foundations of Computer Science*. 1996, pp. 56–65. DOI: [10.1109/SFCS.1996.548464](https://doi.org/10.1109/SFCS.1996.548464).
- [225] P. W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In: *SIAM Journal on Computing* 26.5 (1997), pp. 1484–1509. DOI: [10.1137/S0097539795293172](https://doi.org/10.1137/S0097539795293172). eprint: <https://doi.org/10.1137/S0097539795293172>. URL: <https://doi.org/10.1137/S0097539795293172>.
- [226] J. Sinclair et al. “Fault-tolerant optical interconnects for neutral-atom arrays”. In: *Phys. Rev. Res.* 7 (1 Mar. 2025), p. 013313. DOI: [10.1103/PhysRevResearch.7.013313](https://doi.org/10.1103/PhysRevResearch.7.013313). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.7.013313>.
- [227] V. V. Sivak et al. “Model-Free Quantum Control with Reinforcement Learning”. In: *Physical Review X* 12.1 (2022), p. 011059. DOI: [10.1103/PhysRevX.12.011059](https://doi.org/10.1103/PhysRevX.12.011059).
- [228] I. O. Sokolov et al. “Microcanonical and finite-temperature ab initio molecular dynamics simulations on quantum computers”. In: *Physical Review Research* 3.1 (2021), p. 013125.
- [229] L. Spagnoli et al. “Collective neutrino oscillations in three flavors on qubit and qutrit processors”. In: *Phys. Rev. D* 111 (10 May 2025), p. 103054. DOI: [10.1103/gjr1-lf8s](https://doi.org/10.1103/gjr1-lf8s). URL: <https://link.aps.org/doi/10.1103/gjr1-lf8s>.
- [230] R. B. Staszewski et al. “Cryogenic Controller for Electrostatically Controlled Quantum Dots in 22-nm Quantum SoC”. In: *IEEE Open Journal of the Solid-State Circuits Society* 2 (2022), pp. 103–121. DOI: [10.1109/OJSSCS.2022.3213528](https://doi.org/10.1109/OJSSCS.2022.3213528).
- [231] C. Sünderhauf, E. Campbell, and J. Camps. “Block-encoding structured matrices for data input in quantum computing”. In: *Quantum* 8 (Jan. 2024), p. 1226. DOI: [10.22331/q-2024-01-11-1226](https://doi.org/10.22331/q-2024-01-11-1226). URL: <https://doi.org/10.22331/q-2024-01-11-1226>.
- [232] M. Suzuki. “General theory of fractal path integrals with applications to many-body theories and statistical physics”. In: *Journal of Mathematical Physics* 32.2 (Feb. 1991), pp. 400–407. DOI: [10.1063/1.529425](https://doi.org/10.1063/1.529425). URL: <http://dx.doi.org/10.1063/1.529425>.
- [233] Y. Suzuki et al. “Quantum Error Mitigation as a Universal Error Reduction Technique: Applications from the NISQ to the Fault-Tolerant Quantum Computing Eras”. In: *PRX Quantum* 3.1 (Mar. 2022). DOI: [10.1103/prxquantum.3.010345](https://doi.org/10.1103/prxquantum.3.010345). URL: <http://dx.doi.org/10.1103/PRXQuantum.3.010345>.

- [234] R. Takagi et al. “Fundamental limits of quantum error mitigation”. In: *npj Quantum Information* 8.1 (Sept. 2022). DOI: [10.1038/s41534-022-00618-z](https://doi.org/10.1038/s41534-022-00618-z). URL: <http://dx.doi.org/10.1038/s41534-022-00618-z>.
- [235] M. Takita et al. “Demonstration of Weight-Four Parity Measurements in the Surface Code Architecture”. In: *Phys. Rev. Lett.* 117 (21 Nov. 2016), p. 210505. DOI: [10.1103/PhysRevLett.117.210505](https://doi.org/10.1103/PhysRevLett.117.210505). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.210505>.
- [236] K. Temme, S. Bravyi, and J. M. Gambetta. “Error mitigation for short-depth quantum circuits”. In: *Physical review letters* 119.18 (2017), p. 180509.
- [237] F. Tennie and L. Magri. *Solving nonlinear differential equations on Quantum Computers: A Fokker-Planck approach*. 2024. arXiv: [2401.13500](https://arxiv.org/abs/2401.13500) [quant-ph]. URL: <https://arxiv.org/abs/2401.13500>.
- [238] F. Tennie et al. “Quantum computing for nonlinear differential equations and turbulence”. In: *Nature Reviews Physics* 7.4 (Jan. 2025), pp. 220–230. DOI: [10.1038/s42254-024-00799-w](https://doi.org/10.1038/s42254-024-00799-w). URL: <http://dx.doi.org/10.1038/s42254-024-00799-w>.
- [239] J. D. Teoh et al. “Dual-rail encoding with superconducting cavities”. In: *Proceedings of the National Academy of Sciences* 120.41 (2023), e2221736120. DOI: [10.1073/pnas.2221736120](https://doi.org/10.1073/pnas.2221736120). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2221736120>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2221736120>.
- [240] B. M. Terhal. “Quantum error correction for quantum memories”. In: *Reviews of Modern Physics* 87.2 (Apr. 2015), pp. 307–346. DOI: [10.1103/revmodphys.87.307](https://doi.org/10.1103/revmodphys.87.307). URL: <http://dx.doi.org/10.1103/RevModPhys.87.307>.
- [241] J.-P. Tillich and G. Zémor. “Quantum LDPC Codes with Positive Rate and Minimum Distance Proportional to the Square Root of the Blocklength”. In: *IEEE Transactions on Information Theory* 60.2 (Feb. 2014), pp. 1193–1202.
- [242] M. A. Tremblay, N. Delfosse, and M. E. Beverland. “Constant-Overhead Quantum Error Correction with Thin Planar Connectivity”. In: *Phys. Rev. Lett.* 129 (5 July 2022), p. 050504. DOI: [10.1103/PhysRevLett.129.050504](https://doi.org/10.1103/PhysRevLett.129.050504). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.050504>.
- [243] U. Trottenberg and A. Schuller. *Multigrid*. USA: Academic Press, Inc., 2000.
- [244] J. C. Tully. “Molecular dynamics with electronic transitions”. In: *The Journal of Chemical Physics* 93.2 (1990), pp. 1061–1071.
- [245] J. S. Walker. *Fast Fourier transforms*. Boca Raton, FL: CRC Press, 1996.
- [246] K. Wan, M. Berta, and E. T. Campbell. “Randomized Quantum Algorithm for Statistical Phase Estimation”. In: *Phys. Rev. Lett.* 129 (3 July 2022), p. 030503. DOI: [10.1103/PhysRevLett.129.030503](https://doi.org/10.1103/PhysRevLett.129.030503). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.030503>.
- [247] S. Wang, S. McArdle, and M. Berta. “Qubit-Efficient Randomized Quantum Algorithms for Linear Algebra”. In: *PRX Quantum* 5 (2 Apr. 2024), p. 020324. DOI: [10.1103/PRXQuantum.5.020324](https://doi.org/10.1103/PRXQuantum.5.020324). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.5.020324>.
- [248] J. Watrous. *Understanding Quantum Information and Computation*. 2025. arXiv: [2507.11536](https://arxiv.org/abs/2507.11536) [quant-ph]. URL: <https://arxiv.org/abs/2507.11536>.
- [249] X.-G. Wen. *Quantum Field Theory of Many-body Systems*. Oxford University Press, 2004.
- [250] S. Wiesner. *Simulations of Many-Body Quantum Systems by a Quantum Computer*. 1996. arXiv: [quant-ph/9603028](https://arxiv.org/abs/quant-ph/9603028) [quant-ph]. URL: <https://arxiv.org/abs/quant-ph/9603028>.
- [251] J. Wills. “Multi-mode coaxial transmon qubits for quantum computing and sensing”. PhD thesis. University of Oxford, 2022.
- [252] L. Wright et al. “Noisy intermediate-scale quantum simulation of the one-dimensional wave equation”. In: *Phys. Rev. Res.* 6 (4 Nov. 2024), p. 043169. DOI: [10.1103/PhysRevResearch.6.043169](https://doi.org/10.1103/PhysRevResearch.6.043169). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.6.043169>.

- [253] S. Yalouz et al. “A state-averaged orbital-optimized hybrid quantum–classical algorithm for a democratic description of ground and excited states”. In: *Quantum Science and Technology* 6.2 (2021), p. 024004.
- [254] T. J. Yoder et al. *Tour de gross: A modular quantum computer based on bivariate bicycle codes*. 2025. arXiv: [2506.03094](https://arxiv.org/abs/2506.03094) [quant-ph]. URL: <https://arxiv.org/abs/2506.03094>.
- [255] N. Yoshioka et al. “Hunting for quantum-classical crossover in condensed matter problems”. In: *npj Quantum Information* 10.1 (Apr. 2024). DOI: [10.1038/s41534-024-00839-4](https://doi.org/10.1038/s41534-024-00839-4). URL: <http://dx.doi.org/10.1038/s41534-024-00839-4>.
- [256] X. Yuan et al. “Theory of variational quantum simulation”. In: *Quantum* 3 (2019), p. 191.
- [257] C. Zalka. “Simulating quantum systems on a quantum computer”. In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 454.1969 (Jan. 1998), pp. 313–322. DOI: [10.1098/rspa.1998.0162](https://doi.org/10.1098/rspa.1998.0162). URL: <http://dx.doi.org/10.1098/rspa.1998.0162>.
- [258] J. Zander. *Microsoft announces the best performing logical qubits on record and will provide priority access to reliable quantum hardware in Azure Quantum - The Official Microsoft Blog* — [blogs.microsoft.com](https://blogs.microsoft.com/blog/2024/09/10/microsoft-announces-the-best-performing-logical-qubits-on-record-and-will-provide-priority-access-to-reliable-quantum-hardware-in-azure-quantum/). <https://blogs.microsoft.com/blog/2024/09/10/microsoft-announces-the-best-performing-logical-qubits-on-record-and-will-provide-priority-access-to-reliable-quantum-hardware-in-azure-quantum/>. [Accessed 17-06-2025]. 2024.
- [259] J. Zawalska et al. “Leveraging Hybrid Classical-Quantum Methods for Efficient Load Rebalancing in HPC”. In: *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 2024, pp. 1713–1722. DOI: [10.1109/SCW63240.2024.00214](https://doi.org/10.1109/SCW63240.2024.00214).
- [260] R. Zhang, G. Wang, and P. Johnson. “Computing Ground State Properties with Early Fault-Tolerant Quantum Computers”. In: *Quantum* 6 (July 2022), p. 761. DOI: [10.22331/q-2022-07-11-761](https://doi.org/10.22331/q-2022-07-11-761). URL: <https://doi.org/10.22331/q-2022-07-11-761>.
- [261] Y. Zhao et al. “Realization of an Error-Correcting Surface Code with Superconducting Qubits”. In: *Phys. Rev. Lett.* 129 (3 July 2022), p. 030501. DOI: [10.1103/PhysRevLett.129.030501](https://doi.org/10.1103/PhysRevLett.129.030501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.030501>.
- [262] H. Zhou et al. *Algorithmic Fault Tolerance for Fast Quantum Computing*. June 2024. arXiv: [2406.17653](https://arxiv.org/abs/2406.17653) [quant-ph].
- [263] X. Zhou, D. W. Leung, and I. L. Chuang. “Methodology for quantum logic gate construction”. In: *Phys. Rev. A* 62 (5 Oct. 2000), p. 052316. DOI: [10.1103/PhysRevA.62.052316](https://doi.org/10.1103/PhysRevA.62.052316). URL: <https://link.aps.org/doi/10.1103/PhysRevA.62.052316>.